

Release Notes

Published
2024-08-08

Junos OS Release 22.4R1®

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX. These release notes accompany Junos OS Release 22.4R1. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can find release notes for all Junos OS releases at https://www.juniper.net/documentation/product/us/en/junos-os#cat=release_notes.

Table of Contents

Key Features in Junos OS Release 22.4 | 1

Junos OS Release Notes for ACX Series

What's New | 3

- Class of Service | 3
- EVPN | 4
- Junos Telemetry Interface | 4
- MPLS | 5
- OpenConfig | 5
- Precision Time Protocol (PTP) | 6
- Routing Protocols | 6
- Additional Features | 6

What's Changed | 7

Known Limitations | 8

Open Issues | 8

Resolved Issues | 10

Migration, Upgrade, and Downgrade Instructions | 13

- Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 14

Junos OS Release Notes for cRPD

What's New | 15

- Routing Protocols | 15
- Additional Features | 16

What's Changed | 16

Known Limitations | 16

Open Issues | 17

Resolved Issues | 17

Junos OS Release Notes for cSRX

What's New | 19

Network Address Translation (NAT) | 19

VPNs | 19

What's Changed | 19

Known Limitations | 19

Open Issues | 20

Resolved Issues | 20

Junos OS Release Notes for EX Series

What's New | 20

Authentication and Access Control | 21

EVPN | 22

J-Web | 24

Juniper Extension Toolkit (JET) | 25

Junos OS API and Scripting | 25

Junos Telemetry Interface | 25

Layer 2 Features | 26

Network Management and Monitoring | 27

Security | 27

Additional Features | 27

What's Changed | 28

Known Limitations | 29

Open Issues | 29

Resolved Issues | 33

Migration, Upgrade, and Downgrade Instructions | 37

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 37

Junos OS Release Notes for JRR Series

What's New | 39

Routing Protocols | 39

What's Changed | 40

Known Limitations | 40

Open Issues | 40

Resolved Issues | 40

Migration, Upgrade, and Downgrade Instructions | 41

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 41

Junos OS Release Notes for Juniper Secure Connect

What's New | 43

What's Changed | 43

Known Limitations | 43

Open Issues | 43

Resolved Issues | 43

Junos OS Release Notes for Junos Fusion for Enterprise

What's New | 44

What's Changed | 44

Known Limitations | 44

Open Issues | 45

Resolved Issues | 45

Migration, Upgrade, and Downgrade Instructions | 45

Junos OS Release Notes for Junos Fusion for Provider Edge

[What's New | 51](#)

[What's Changed | 52](#)

[Known Limitations | 52](#)

[Open Issues | 52](#)

[Resolved Issues | 52](#)

[Migration, Upgrade, and Downgrade Instructions | 52](#)

Junos OS Release Notes for MX Series

[What's New | 62](#)

[Class of Service | 63](#)

[EVPN | 64](#)

[Flow-based and Packet-based Processing | 65](#)

[High Availability | 66](#)

[Interfaces | 66](#)

[Junos Telemetry Interface | 66](#)

[Layer 2 VPN | 68](#)

[Licensing | 68](#)

[MPLS | 71](#)

[Network Address Translation \(NAT\) | 73](#)

[Network Management and Monitoring | 73](#)

[OpenConfig | 73](#)

[Precision Time Protocol \(PTP\) | 74](#)

[Routing Protocols | 74](#)

[Services Applications | 77](#)

[Software Installation and Upgrade | 77](#)

[Source Packet Routing in Networking \(SPRING\) or Segment Routing | 78](#)

Subscriber Management and Services | 78

VPNs | 80

Additional Features | 80

What's Changed | 83

Known Limitations | 85

Open Issues | 87

Resolved Issues | 96

Migration, Upgrade, and Downgrade Instructions | 115

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 116

Junos OS Release Notes for NFX Series

What's New | 117

What's Changed | 117

Known Limitations | 117

Open Issues | 118

Resolved Issues | 118

Resolved Issues | 119

Migration, Upgrade, and Downgrade Instructions | 119

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 120

Junos OS Release Notes for PTX Series

What's New | 121

Authentication and Access Control | 122

EVPN | 122

High Availability | 123

Junos Telemetry Interface | 123

MPLS | 124

OpenConfig | 125

Routing Policy and Firewall Filters | 125

Routing Protocols | 125

Source Packet Routing in Networking (SPRING) or Segment Routing | 128

Additional Features | 128

What's Changed | 128

Known Limitations | 130

Open Issues | 131

Resolved Issues | 132

Migration, Upgrade, and Downgrade Instructions | 134

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 138

Junos OS Release Notes for QFX Series

What's New | 139

Authentication and Access Control | 140

Ethernet Switching and Bridging | 140

EVPN | 141

High Availability | 143

Interfaces | 143

Juniper Extension Toolkit (JET) | 143

Junos OS API and Scripting | 143

Junos Telemetry Interface | 144

MPLS | 144

Routing Protocols | 145

Additional Features | 146

What's Changed | 147

Known Limitations | 149

Open Issues | 149

Resolved Issues | 152

Migration, Upgrade, and Downgrade Instructions | 158

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 171

Junos OS Release Notes for SRX Series

What's New | 172

EVPN | 172

High Availability | 173

Interfaces | 174

Juniper Extension Toolkit (JET) | 174

Junos OS API and Scripting | 174

Network Address Translation (NAT) | 175

Content Security (UTM) | 175

VPNs | 175

What's Changed | 176

Known Limitations | 179

Open Issues | 180

Resolved Issues | 181

Migration, Upgrade, and Downgrade Instructions | 186

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 186

Junos OS Release Notes for vMX

What's New | 188

EVPN | 188

Junos Telemetry Interface | 189

What's Changed | 190

Known Limitations | 191

Open Issues | 191

Resolved Issues | 192

Upgrade Instructions | 192

Junos OS Release Notes for vRR

What's New | 193

| **Routing Protocols | 193**

What's Changed | 194

Known Limitations | 194

Open Issues | 194

Resolved Issues | 194

Junos OS Release Notes for vSRX

What's New | 195

| **EVPN | 195**

| **Interfaces | 196**

| **Network Address Translation (NAT) | 196**

| **VPNs | 196**

What's Changed | 197

Known Limitations | 198

Open Issues | 199

Resolved Issues | 199

Migration, Upgrade, and Downgrade Instructions | 201

| **Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 207**

Licensing | 208

Finding More Information | 209

Requesting Technical Support | 210

Revision History | 211

Key Features in Junos OS Release 22.4

Start here to learn about the key features in Junos OS Release 22.4. For more information about a feature, click the link in the feature description.

- **Pure EVPN Type 5 routes with EVPN-VXLAN (SRX Series and vSRX)**—Starting in Junos OS Release 22.4R1, you can configure pure Type 5 routes in an Ethernet VPN–Virtual Extensible LAN (EVPN–VXLAN) environment. These devices use EVPN Type 5 routes to advertise IP prefixes for intersubnet connectivity within and across data centers.

[See [Understanding EVPN Pure Type 5 Routes](#) and [ip-prefix-routes](#).]

- **Source NAT port overload (cSRX, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 22.4R1, We've updated the hash algorithm to allow for improved distribution of network traffic, when using the port overloading capability. Enabling better utilization per IP, as appropriate to the type of network traffic.

The hash algorithm uses the reverse traffic from the server, matches the existing sessions, and reuses the same Network Address Translation (NAT) resources.

You can configure the updated hash algorithm using the `enhanced-port-overloading-algorithm` statement at the `[security nat source pool pool-name port]` and `[security nat source interface]` hierarchy levels.

[See [pool \(Security Source NAT\)](#) and [source \(Security Source NAT\)](#).]

- **Source NAT preserve range support (SRX Series)**—Starting in Junos OS Release 22.4R1, we support a preserve range for the source NAT. You can assign a port within the same range as the incoming port, either 0 through 1023 or 1024 through 65,535.

To enable the preserve range, configure the `preserve-range` statement at the `[security nat source pool pool-name port]` hierarchy level.

[See [pool \(Security Source NAT\)](#) and [preserve-range](#).]

- **Support for 1-Gbps speed on SRX5K-IOC4-10G card (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 22.4R1, you can change the speed configuration of a 10-Gbps port to operate at 1-Gbps. You can make this change by configuring the speed value as 1-Gbps in the `set interfaces <intf-name> gigether-options speed 1g` command. After you commit the configuration, the operating speed of the 10-Gbps port changes to 1 Gbps.

[See [Port Speed on SRX5K-IOC4-MRATE](#).]

- **Support for PTP over Ethernet (hybrid mode) (ACX7509)**—Starting in Junos OS Release 22.4R1, ACX7509 routers support Precision Time Protocol (PTP over Ethernet (hybrid mode) with the G.8275.1 telecom profile.

PTP over Ethernet effectively implements the packet-based technology. This feature helps operators deliver synchronization services on packet-based mobile backhaul (MBH) networks.

[See [PTP Profiles](#) and [Precision Time Protocol Overview](#).]

- **Telemetry support for interfaces and chassis (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-MP, EX4100-F, EX4300-MP, EX4400-MP, EX4400, EX4650, QFX5110, QFX5120, QFX5200, and QFX5210)**—Junos OS Release 22.4R1 introduces support for streaming operational state statistics and counters for chassis and interfaces using OpenConfig sensor paths. We also support the following new Junos-specific sensor paths for statistics that are unsupported in OpenConfig:

- `/state/chassis/`
- `/state/interfaces/`

[See [Telemetry Sensor Explorer](#).]

- **Telemetry support for PoE (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-MP, EX4100-F, EX4300-MP, EX4400-MP, EX4400, and EX4650)**—Junos OS Release 22.4R1 introduces support for streaming operational state statistics and counters for PoE using the new Junos-specific sensor path `/state/poe/`.

[See [Telemetry Sensor Explorer](#).]

- **View or log out the LNS L2TP subscriber sessions associated with a routing instance (MX Series)**—Starting in Junos OS Release 22.4R1, we've introduced the following two L2TP operational commands for MX Series devices that support BNG L2TP functionality. Use the following new operational commands to view or log out all the L2TP subscriber sessions simultaneously.

- `show service l2tp session routing-instance name`
- `clear service l2tp session routing-instance name`

[See [show service l2tp session routing-instance](#) and [clear service l2tp session routing-instance](#).]

Junos OS Release Notes for ACX Series

IN THIS SECTION

 [What's New | 3](#)

- What's Changed | 7
- Known Limitations | 8
- Open Issues | 8
- Resolved Issues | 10
- Migration, Upgrade, and Downgrade Instructions | 13

What's New

IN THIS SECTION

- Class of Service | 3
- EVPN | 4
- Junos Telemetry Interface | 4
- MPLS | 5
- OpenConfig | 5
- Precision Time Protocol (PTP) | 6
- Routing Protocols | 6
- Additional Features | 6

Learn about new features introduced in this release for ACX Series routers.

Class of Service

- **Support for hierarchical class of service (HCoS) on EVPN Interfaces (ACX5448, ACX5448-D, and ACX5448-M)**—Starting in Junos OS Release 22.4R1, you can apply up to four levels of hierarchical traffic scheduling and shaping features to EVPN interfaces on ACX 5448, ACX 5448-D, and ACX 5448-M devices. The four levels of hierarchical scheduling are:
 - Physical interfaces
 - Logical interface sets
 - Logical interfaces

- Queues

[See [Hierarchical Class of Service in ACX Series Routers.](#)]

- **Support for hierarchical class of service (HCoS) (ACX710)**—Starting in Junos OS Release 22.4R1, you can apply up to four levels of hierarchical traffic scheduling and shaping features to the following interfaces on ACX710 devices:
 - EVPN
 - Layer 3 VPN (L3VPN)
 - Multichassis link aggregation group (MC-LAG)
 - Multicast

The four levels of hierarchical scheduling are:

- Physical interfaces
- Logical interface sets
- Logical interfaces
- Queues

[See [Hierarchical Class of Service in ACX Series Routers.](#)]

EVPN

- **EBGP and IBGP support over IRB interfaces across PE-CE links for EVPN-MPLS (ACX5448)**—Starting in Junos OS Release 22.4R1, you can configure external BGP (EBGP) and internal BGP (IBGP) over an integrated routing and bridging (IRB) interface which spans a provider edge (PE) device to customer edge (CE) device (PE-CE) link.

[See [EVPN with IRB Solution Overview.](#)]

Junos Telemetry Interface

- **Event-driven streaming of sensor data for MPLS LSP record route objects (ACX5448, ACX7100, MX204, MX240, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX1000, and vMX)**—Junos OS Release 22.4R1 introduces ON_CHANGE notification for streaming MPLS label-switched path (LSP) record route object statistics. Using ON_CHANGE mode, data values are not streamed but sent only when data values change. Support includes leaf nodes under the resource path `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/sessions/session/record-route-objects/record-route-object/state/`.

[See [Telemetry Sensor Explorer.](#)]

- **OpenConfig OSPF configuration and operational state sensors (ACX5448, ACX7100, MX150, MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, and PTX1000)**—Junos OS Release 22.4R1 introduces support for the OpenConfig OSPF data model `openconfig-ospfv2.yang` (v.0.3.1). We now support configuration and streaming of operational state data under the resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.

To learn about OpenConfig configuration mappings, see [Mapping OpenConfig OSPF Commands to Junos Configuration](#). For state sensors, see [Telemetry Sensor Explorer](#).

- **System health reporting sensors support on gRPC (ACX5448 and ACX710 routers; MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10008, and MX10016 routers; and PTX10002 routers)**—Starting in Junos OS Release 22.4R1, Junos telemetry interface (JTI) Junos telemetry interface (JTI) supports data model `openconfig-system.yang` using gRPC remote procedure calls (gRPC) and provides new health-monitoring sensors.

[See [Telemetry Sensor Explorer](#).]

MPLS

- **Support for LDP tunneling over SR-TE (ACX5448, MX480, MX960, and MX2010)**—Starting in Junos OS Release 22.4R1, you can tunnel LDP label-switched paths (LSPs) over segment routing-traffic engineering (SR-TE) in OSPF networks. Tunneling LDP over SR-TE provides consistency and coexistence of both LDP LSPs and SR-TE LSPs.

To configure LDP tunneling over SR-TE, include the `tunnel-source-protocol` configuration statement at the `[edit protocols ospf traffic-engineering]` and `ldp-tunneling` configuration statement at the `[edit protocols ospf source-packet-routing source-routing-path]` hierarchy levels.

[See [Tunneling LDP over SR-TE](#).]

OpenConfig

- **OpenConfig OSPF configuration and operational state sensors (ACX5448, ACX7100, MX150, MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, and PTX1000)**—Starting in Junos OS Release 22.4R1, we support the OpenConfig OSPF data model `openconfig-ospfv2.yang` (version 0.3.1). We also support configuration and streaming of operational state data under the resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.

See [Mapping OpenConfig OSPF Commands to Junos Configuration](#) for OpenConfig configuration mappings. See [Telemetry Sensor Explorer](#) for state sensors.

Precision Time Protocol (PTP)

- **Support for PTP over Ethernet (hybrid mode) (ACX7509)**—Starting in Junos OS Release 22.4R1, ACX7509 routers support Precision Time Protocol (PTP over Ethernet (hybrid mode) with the G.8275.1 telecom profile.

PTP over Ethernet effectively implements the packet-based technology. This feature helps operators deliver synchronization services on packet-based mobile backhaul (MBH) networks.

[See [PTP Profiles](#) and [Precision Time Protocol Overview](#).]

Routing Protocols

- **BMP local RIB monitoring support for all RIBs with sharding (ACX Series, cRPD, PTX Series, QFX Series, and vRR)**—Starting in Junos OS Release 22.4R1, you can configure a policy to monitor routing information bases also known as routing table (RIBs) of virtual routers and virtual routing and forwarding instances (VRF). You can specify two separate sets of RIBs in the BGP Monitoring Protocol (BMP), one for monitoring and the other for reporting. With this feature, BMP can filter traffic based on the routes and routing instances.

[See [BGP Monitoring Protocol](#), [loc-rib](#), and [rib-list](#).]

- **OSPF FAPM and interarea support (ACX5448, MX204, MX240, MX480, MX960, MX10003, MX10008, MX2008, MX2010, MX2020, PTX1000, and QFX10002)**—Starting with Junos OS Release 22.4R1, the Flexible Algorithm Prefix Metric (FAPM) is defined to allow an optimal end-to-end path for an inter-area prefix. The Area Border Router (ABR) *must* include the FAPM when advertising the prefix between areas that areas reachable in that given Flex-Algorithm. When a prefix is unreachable, the ABR *must not* include that prefix in the Flex-Algorithm when advertising between areas. The defined FAPM provides inter-area support.

[See [Understanding OSPF Flexible Algorithm for Segment Routing](#).]

[See [show ospf database](#), [show route table](#), [show ted database](#)

- **Flex algo and FAPM leaking across IS-IS multi-instance (ACX5448, MX480, MX960, MX2010)**—Starting in Junos OS Release 22.4R1, we've added support to readvertise flexible algorithm (flex algo) prefix-segment identifiers (SIDs) and Flexible Algorithm Prefix Metrics (FAPMs) across interior gateway protocol (IGP) instances. We have also added support to readvertise other protocol prefixes and assign flex algo prefix-SIDs via policy to those prefixes.

Additional Features

Support for the following features has been extended to these platforms.

- **SNMP MIB support for the timing Synchronous Ethernet feature (ACX710)**. You can get the Synchronous Ethernet , Precision Time Protocol (PTP) defect and event notifications. We've introduced timing MIB to support this feature. The trap and alarm notifications are disabled by

default. To enable trap and alarm notifications for the timing feature, include the `timing-events` statement at the `[edit snmp trap-group group-name categories]` hierarchy level.

[See [Configuring SNMP Trap Groups](#) and [Enterprise-Specific SNMP MIBs Supported by Junos OS.](#)]

What's Changed

IN THIS SECTION

- [EVPN | 7](#)
- [Network Management and Monitoring | 7](#)
- [User Interface and Configuration | 8](#)

Learn about what changed in this release for ACX Series routers.

EVPN

- Flow-label configuration status for EVPN ELAN services. The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

Network Management and Monitoring

- **Junos YANG modules for RPCs include the `junos:command extension` statement (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command extension` statement in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the

YANG schemas on the local device, the YANG modules for RPCs include the `junos:command` extension statement.

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml` filter or the `| display json` filter to display command output in XML or JSON format in any release. We've removed these options from the CLI.

Known Limitations

There are no known limitations in hardware or software in this release for ACX Series routers.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

IN THIS SECTION

- [General Routing | 9](#)
- [Routing Protocols | 10](#)

Learn about open issues in this release for ACX Series routers.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- On WRL8 based VMHost platforms (ACX6360), there is no log rotation for resild log and temperature sensor info is incorrectly written into resild log which could result in continuous logs in resild log file. The disk usage might keep increasing due to this issue. The disk usage could be eventually full which could cause system to hang and reboot. [PR1480217](#)
- Due to the BRCM KBP issue, route lookup might fail. [PR1533513](#)
- On the ACX500 devices, service MIC does not work. [PR1569103](#)
- On ACX5448 and ACX710 platforms, all types of delegated BFD sessions configured on routing-instance other than the default routing-instance might not come up. [PR1633395](#)
- Interop for 1G interfaces between EX4100 SKUs and ACX5448/ACX5448-M/D or MX480 does not work. [PR1657766](#)
- For ACX5448 devices, if a non-default ssh port is configured for system login, after upgrade to 21.4 release, the FPC is stuck in offline. To avoid such issue please use default SSH port and use protect Routing Engine filter to only allow the access from the trusted source. [PR1660446](#)
- When TCP Main and TCP remaining attached together on IFD its observed that Improper Scheduler MAP is getting configured on HQoS IFD for some corner scenarios. This is a sequence issue from CoSD(RE) which not guaranteed at Packet Forwarding Engine side. And this applicable for all platforms. [PR1664785](#)
- On ACX5000 devices, in VPLS MH cases, the standby UNI ifl in the backup router gets programmed in the Disable state by adding the UNI interface to invalid the VPN ID in hardware. During switch over, the UNI ifl gets deleted and added under the VPLS instance VPN ID. In issue case, UNI interface added under the invalid VPN ID in the backup router tries to get deleted by passing the VPLS instance VPN ID, causing the issue. [PR1665178](#)
- The Queue statistics might show constant PPS / bps after interface is disabled. The statistics does not increment and remain same when the interface goes down. [PR1685344](#)
- Reserved buffers might be shown as 0. But internally reserved buffers do get used to queue and transmit traffic on the queue. [PR1689183](#)

- The aggregate Ethernet statistics might show 0 bps for output traffic. It is a CLI output display issue. It does not impact the traffic output. [PR1689185](#)
- dc-pfe: HEAP malloc(0) detected! when a VPLS instance is deactivated in ACX5048. [PR1692400](#)
- Convergence time can be more than 60ms for OSPF TILFA Node protection testing. [PR1695292](#)

Routing Protocols

- SRTE secondary LSP should be only standby in forwarding table, however, it is also active and forwarding traffic due to the wrong metric calculation. [PR1696598](#)

Resolved Issues

IN THIS SECTION

- [Class of Service | 10](#)
- [EVPN | 11](#)
- [General Routing | 11](#)
- [Network Management and Monitoring | 13](#)
- [Routing Protocols | 13](#)
- [User Interface and Configuration | 13](#)

Learn about the issues fixed in this release for ACX Series routers.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service

- The default code-point aliases and respective CoS value Bit patterns are inconsistent with Junos OS. [PR1667404](#)

EVPN

- Junos OS and Junos OS Evolved: RPD core upon receipt of a specific EVPN route by a BGP route reflector in an EVPN environment (CVE-2022-22219). [PR1675054](#)

General Routing

- The Layer 3 packets with the destination as IPv6 Link Local address will not work. [PR1638642](#)
- If a firewall has a log action and it is applied on physical interface or lo0, the LDP cannot go up. [PR1648968](#)
- BGP Sensor `/bgp-rib/afi-safis/afi-safi/ipv4-unicast/loc-rib/` not available as a periodic sensor. [PR1649529](#)
- HTTP(S) file download hangs over EVPN-ETREE. [PR1653531](#)
- On ACX5448 platform, physical interfaces of FPC remain up even though it lost communication with Routing Engine. [PR1659949](#)
- Packet count might occasionally be 0 for some interfaces in **monitor interface traffic**. [PR1661617](#)
- The Layer 2 circuit backup might not get reverted to primary in rare condition. [PR1661802](#)
- After activating EVPN-ETREE service on ACX5448 Packet Forwarding Engine might crash. [PR1662686](#)
- The rpd core might be seen when there is a synchronization issue. [PR1663050](#)
- Multicast upstream interface does not change to back up link when PIM neighbor is removed or flapped and causes a traffic impact. [PR1663271](#)
- Transit traffic drop was seen for BGP-LU(Border Gateway Protocol-Labeled Unicast) prefix on ACX5448 or ACX710 when BGP-LU label routes has ECMP(Equal-Cost Multipath) forwarding path. [PR1663563](#)
- Adding an empty interface to an aggregate Ethernet bundle causes traffic drop. [PR1663651](#)
- FXPC core might be observed when deactivating a child member link from aggregate Ethernet bundle. [PR1665511](#)
- In the SRTE scenario, sensors are wrongly populated for colored tunnel BSID routes when uncolored tunnels are enabled. [PR1665943](#)

- Inline BFDv6 Sessions might go DOWN and stay in that state on ACX5448 and ACX710 platforms. [PR1666746](#)
- Traffic loss is observed when the VRRP is configured over the aggregate Engine interface. [PR1666853](#)
- On ACX710 and ACX5448 its variants Packet Forwarding Engine might crash due to configuration of BFD. [PR1667129](#)
- Shutting the CE interface and bringing back up causes traffic (going towards the core) drop. [PR1667724](#)
- LLDP neighborship might fail if the chassis-id format of the LLDP packet is xx:xx:xx:XX:XX:xx. [PR1669677](#)
- Chassis alarms for smart errors not set or cleared. [PR1669968](#)
- ACX710 : Log related to resources reported after EVPN RI are deactivated / activated multiple times : ACX_BD_ERR: dnx_bd_alloc_l2_svlan: System reached L3 IFL and BD limit(12286). [PR1670683](#)
- MX-SPC3 PIC core dump is seen when a CPCD service is modified. [PR1675990](#)
- The LLDP packets will not be transmitted over Layer 2 circuit on the ACX platform. [PR1678752](#)
- Memory leak is seen on ACX710 and ACX5448 when the core link flaps. [PR1681980](#)
- ACX5448 : RIO DNX PFE wrongly spelled as QUMARN instead of Qumran. [PR1682819](#)
- The traffic drop would be observed with inter-vlan configuration when deactivating and activating the EVPN routing instance. [PR1683321](#)
- On the ACX710 device, the IEEE 802.1p priority and DEI values in the locally generated VLAN-based IP packets might be changed when sourced from the IRB interface. [PR1683770](#)
- ACX5448:ACX710 Layer 2 circuit traffic drop with control-word enabled or control-word configuration change. [PR1683900](#)
- Auto-mdix is not working in ACX710 devices. [PR1685431](#)
- [acx710-22.4]: Mc-lag down in odin post bringup. [PR1688958](#)
- EVPN traffic is classified in the wrong queue. [PR1689604](#)
- Packet forwarding fails on specific ACX Junos OS platforms due to flapping of core interface member link in the MPLS-EVPN environment. [PR1690590](#)

Network Management and Monitoring

- The snmpd core might be observed with filter-duplicates configuration. [PR1669510](#)

Routing Protocols

- Ipv6 inline BFD sessions are down when neighbor is not resolved. [PR1650677](#)
- Routing Process Daemon (rpd) crashes and restarts when a specific timing condition is hit with BGP configuration. [PR1659441](#)
- MCSNOOPD will be restarted and will again learn the states after core. [PR1672488](#)
- The process rpd (route process daemon) crashes with BGP VPN (Border gateway protocol - Virtual Private Network) configuration, while ebgp (external bgp) routes exported into ibgp (internal bgp) core with vrf (virtual route forward) configured. [PR1675893](#)

User Interface and Configuration

- Commit failure when changing BGP well-known community attributes. [PR1669375](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 14](#)

This section contains the upgrade and downgrade support policy for Junos OS for ACX Series routers. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

For information about software installation and upgrade, see the https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/software-installation-and-upgrade/software-installation-and-upgrade.html Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 1: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for cRPD

IN THIS SECTION

- [What's New | 15](#)
- [What's Changed | 16](#)
- [Known Limitations | 16](#)
- [Open Issues | 17](#)
- [Resolved Issues | 17](#)

What's New

IN THIS SECTION

- [Routing Protocols | 15](#)
- [Additional Features | 16](#)

Learn about new features introduced in this release for cRPD.

Routing Protocols

- **BMP local RIB monitoring support for all RIBs with sharding (ACX Series, cRPD, PTX Series, QFX Series, and vRR)**—Starting in Junos OS Release 22.4R1, you can configure a policy to monitor routing information bases also known as routing table (RIBs) of virtual routers and virtual routing and forwarding instances (VRF). You can specify two separate sets of RIBs in the BGP Monitoring Protocol (BMP), one for monitoring and the other for reporting. With this feature, BMP can filter traffic based on the routes and routing instances.

[See [BGP Monitoring Protocol](#), [loc-rib](#), and [rib-list](#).]

- **Support for bootstrapping route-validation database from a local file (cRPD, JRR200, MX204, PTX10008, and QFX10008)**—Starting in Junos OS Release 22.4R1, we support the ability to read validation records from a local binary file and install into the specified named route-validation

databases within RPD. This feature implements syntactic and semantic checks on the content of the file to ensure that it is a well-specified set of validation records. If the syntactic and semantic checks fail, the entire file is rejected as a source of validation records. Use the `source-file` statement at the `[edit routing-options validation]` hierarchy level to source route-validation records from a local file source. You can use the `show validation source-file` command to display the properties of a local validation source file.

[See [validation](#).]

- **MVPN feature support with sharding (cRPD, JRR200, MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support the following features:
 - Multicast virtual private network (MVPN) inactive route query from the main thread to shards
 - Extranet and auto-export support with sharding
 - Interact functions with RT-proxy client and server
 - New data structure to store the inactive route data on the main thread
 - Asynchronous route processing on the main thread

You can use `show mvpn c-multicast` to display the inactive route data stored on the main thread.

[See [rib-sharding](#) and [show mvpn c-multicast](#) .]

Additional Features

Support for the following features has been extended to these platforms.

- **OS upgrade(cRPD)**. We've upgraded the operating system (OS) for the containerized routing protocol process (daemon) (cRPD) docker image from Ubuntu version 18 to 22.

What's Changed

There are no changes in behavior and syntax in this release for cRPD.

Known Limitations

There are no known limitations in hardware or software in this release for cRPD.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

There are no known issues in hardware or software in this release for cRPD.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

IN THIS SECTION

- [Platform and Infrastructure](#) | 17

Learn about the issues fixed in this release for cRPD.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Platform and Infrastructure

- In cRPD, the following command issues are fixed:
 - clear arp all
 - monitor interface traffic
 - monitor interface *interface-name*
 - show log user
 - request system process terminate
 - show system processes extensive
 - show system processes detail
 - show system processes summary

- show system processes brief
- show system virtual-memory
- show system users
- show system statistics
- show system memory
- request system storage cleanup
- show system software
- In cRPD, the following commands are deprecated:
 - request system reboot
 - request system halt
 - show route localization
 - show system snapshot
 - traceroute routing-instance
 - clear interfaces statistics
 - show system khms-stats

[PR1672670](#)

Junos OS Release Notes for cSRX

IN THIS SECTION

- [What's New | 19](#)
- [What's Changed | 19](#)
- [Known Limitations | 19](#)
- [Open Issues | 20](#)
- [Resolved Issues | 20](#)

What's New

IN THIS SECTION

- [Network Address Translation \(NAT\) | 19](#)
- [VPNs | 19](#)

Learn about new features introduced in this release for cSRX.

Network Address Translation (NAT)

VPNs

- **PSK-based IPsec VPN support on cSRX**—Starting in Junos OS Release 22.4R1, we support Internet Key Exchange (IKE) with Pre-shared key authentication for IPsec VPN on cSRX instances using traffic selector-based tunnels. IPsec can be used to encrypt traffic between devices protected by cSRX and a remote IPsec gateway.

[See [Internet Key Exchange](#) and [Internet Key Exchange \(IKE\) for IPsec VPN.](#)]

What's Changed

There are no changes in behavior and syntax in this release for cSRX.

Known Limitations

There are no known limitations in hardware or software in this release for cSRX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

There are no known issues in hardware or software in this release for cSRX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

There are no resolved issues in this release for cSRX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Junos OS Release Notes for EX Series

IN THIS SECTION

- [What's New | 20](#)
- [What's Changed | 28](#)
- [Known Limitations | 29](#)
- [Open Issues | 29](#)
- [Resolved Issues | 33](#)
- [Migration, Upgrade, and Downgrade Instructions | 37](#)

What's New

IN THIS SECTION

- [Authentication and Access Control | 21](#)

- [EVPN | 22](#)
- [J-Web | 24](#)
- [Juniper Extension Toolkit \(JET\) | 25](#)
- [Junos OS API and Scripting | 25](#)
- [Junos Telemetry Interface | 25](#)
- [Layer 2 Features | 26](#)
- [Network Management and Monitoring | 27](#)
- [Security | 27](#)
- [Additional Features | 27](#)

Learn about new features introduced in this release for cSRX.

Authentication and Access Control

- **802.1X authentication with EVPN-VXLAN (EX4650, QFX5120)**—Starting in Junos OS Release 22.4R1, the EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM switches that act as access switches can use 802.1X authentication to protect an EVPN-VXLAN network from unauthorized end devices.

These switches support the following 802.1X authentication features on access and trunk ports:

- Access ports: single, single-secure, and multiple supplicant modes
- Trunk ports: single and single-secure supplicant modes
- Guest VLAN
- Server fail
- Server reject
- Dynamic VLAN
- Dynamic firewall filters
- RADIUS accounting
- Port bounce with Change of Authorization (CoA) requests
- MAC RADIUS client authentication
- Central Web Authentication (CWA) with redirect URL

- Captive portal client authentication
- Flexible authentication with fallback scenarios

[See [802.1X Authentication](#)

EVPN

- **EVPN-MPLS E-LAN flow-aware transport (FAT) label load balancing (MX Series, EX9200, vMX)** – Starting in Junos OS Release 22.4R1, you can configure provider edge (PE) devices to use FAT labels in an Ethernet VPN-MPLS (EVPN-MPLS) routing instance, according to Request for Comments (RFC) 6391. PE devices use these labels to load-balance EVPN-MPLS unicast packets across equal-cost multipaths (ECMPs) without performing deep packet inspection of the MPLS payload. This feature supports emulated LAN (ELAN) with single-homing and multi-homing active/standby and active/active topologies and supports the VLAN-based, VLAN-bundle, and VLAN-aware bundle EVPN-MPLS variants.

NOTE: This feature does not support MX Series devices with Advanced Forwarding Toolkit (AFT) cards.

NOTE: On MX Series devices, a configuration where the local PE has a static-flow-label and the remote PE does not have a static-flow-label, the remote PE can process packets without dropping any traffic.

Enabling Load Balancing Using Fat Labels for EVPN Routing Instances:



WARNING: Configuring a flow label or deleting a flow label with the following CLI commands causes a catastrophic event for the routing instance. As a best practice, perform these CLI commands during a maintenance period to avoid network disruptions.

- Configure the flow-label-static statement at the [edit routing-instances routing-instance-name protocols evpn] hierarchy level on PE devices to insert FAT flow labels into pseudowire packets sent to remote PE devices.
- Configure the flow-label statement at the [edit routing-instances routing-instance-name protocols evpn] hierarchy level on PE devices to signal flow-label capability in the EVPN Layer 2 Attributes Extended Community by setting the flow-label (F) bit in the EVPN Type 3 route.

[See [flow-label](#) and [flow-label-static](#).]

- **EVPN-VXLAN to EVPN-VXLAN seamless stitching for EVPN Type 5 routes (EX4100-24T, EX4400-24MP, EX4400-24P, EX4400-48F, EX4650, MX204, MX240, QFX10002-60C, QFX10002, QFX10008, QFX10016, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 22.4R1, you can configure Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) to EVPN-VXLAN seamless stitching with EVPN Type 5 (IP prefix) routes between two interconnected data centers or between two points of delivery (pods) in a data center.

In the EVPN-VXLAN fabric, border leaf or border spine devices act as interconnection gateways. You enable EVPN Type 5 routes in virtual routing and forwarding (VRF) instances on both sides of the interconnection. For each VRF instance, the server leaf devices in the first data center create VXLAN tunnels for Type 5 routes (with corresponding virtual network identifiers [VNIs]) toward their local gateway devices. The gateway devices map those VXLAN tunnels to an interconnection tunnel (with a new route distinguisher [RD], route target, and VNI) toward the second data center. The gateway devices in the second data center re-create the Type 5 VXLAN tunnels using their local RD.

We support one-to-one mapping of Type 5 VRF instances across the interconnection.

- **Support for VXLAN group-based policy with ingress and egress configuration (EX4100, EX4400, EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting in Junos OS Release 22.4R1, we've added enhancements to the group-based policy (GBP) feature and made some changes to the CLIs.

The enhancements are:

- You can enforce the policy on the ingress endpoint or the egress tunnel endpoint. Ingress enforcement optimizes the network bandwidth. To configure policy enforcement at the ingress endpoint, use the `set forwarding-options evpn-vxlan gbp ingress-enforcement` command.
- We support these match conditions for GBP tagging:
 - `interface <interface_name>`
 - `mac-address <mac address>`
 - `vlan-id <vlan id>`
 - `ip-version ipv4 <ip address> or <prefix-list>`
 - `ip-version ipv6 <ip address> or <prefix-list>`

[See [Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN.](#)]

- **Routing protocols on EVPN-VXLAN overlay IRB interfaces in the default routing instance (EX4400, EX4650, EX9200, EX9253, QFX5110, QFX5120, QFX10002, QFX1008, and QFX10016)**—Starting in Junos OS Release 22.4R1, you can run routing protocols on Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) overlay integrated routing and bridging (IRB) interfaces in the IPv4 or IPv6 default routing instance associated with the underlay (default.inet.0 or default.inet6.0). To perform this task,

you can set and export a policy with the `install-next-hop except overlay-vxlan-interfaces` policy qualifier option. The policy configuration avoids routing loops that can happen if the device uses overlay IRB routes for underlay VTEP reachability. To support this use case in releases prior to 22.4R1, you can configure the IRB interface in a routing instance of type `vrf` instead of in the default routing instance.

[See [install-next-hop](#).]

- **Protect core support for EVPN-VXLAN (EX4300-MP, EX4400-48MP, EX4650, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 22.4R1, you can configure the protect core feature in an EVPN-VXLAN environment. You can use the protect core feature to install a route in the forwarding table for use as an alternative path when an existing route fails or if connectivity is lost.

[See [protect core](#).]

- **Overlay and CE-IP ping and traceroute support for EVPN-VXLAN (EX4300-MP, EX4400, EX4650, QFX5110, QFX5120, QFX5200, QFX10002, QFX10002-60C, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.4R1, you can perform ping and traceroute operations within an EVPN-VXLAN overlay or to a specific customer edge (CE) device IP address (CE-IP) across an EVPN-VXLAN overlay. You can use ping, traceroute, CE-IP ping, and CE-IP traceroute utilities to detect and isolate faults in overlay networks.

[See [Understanding Overlay Ping and Traceroute Packet Support](#).]

- **Persistent MAC learning (sticky MAC) with EVPN-VXLAN (EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T)**— Starting in Junos OS Release 22.4R1, you can enable network interfaces to retain dynamically learned MAC addresses when the switch is restarted or when an interface goes down and comes back up again.

NOTE: We don't support persistent MAC learning on virtual tunnel endpoint (VTEP) interfaces.

[See [Understanding and Using Persistent MAC Learning](#).]

J-Web

- **Support for EX4100 and EX4100-F switches (EX Series)**—Starting in Junos OS Release 22.4R1, you can configure, monitor, and manage EX4100 and EX4100-F switches by using J-Web. To configure the EX4100 and EX4100-F switches, you must connect the Ethernet cable from the PC's Ethernet port to the port labeled **MGMT** on the switch's rear panel. The chassis viewer on the Dashboard page supports both the standalone device view and the Virtual Chassis configuration view (graphical view of each member switch).

[See [Dashboard for EX Series Switches](#) and [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#).]

Juniper Extension Toolkit (JET)

- Prevent script execution based on current system memory usage (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4300, EX4300-MP, EX4300-VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9208, EX9251, EX9253, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, SRX4100, SRX4200, and SRX4600)—Starting in Junos OS Release 22.4R1, you can configure the system memory usage threshold above which the device prevents the execution of Juniper Extension Toolkit (JET) scripts. You can configure the start start-options mem-factor statement for individual JET scripts or all JET scripts. The device doesn't execute the script if the system's memory usage exceeds the configured value at the time the script is invoked. This configuration ensures that a device executes only essential scripts when system resources are limited, thereby enabling the device to continue performing all critical network functions.

[See [Configure Script Start Options.](#)]

Junos OS API and Scripting

- Prevent script execution based on current system memory usage (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4300, EX4300-MP, EX4300-VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9208, EX9251, EX9253, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, SRX4100, SRX4200, and SRX4600)—Starting in Junos OS Release 22.4R1, you can configure the system memory usage threshold above which the device prevents the execution of certain op, event, or SNMP scripts. You can configure the start start-options mem-factor statement for individual scripts or all scripts of a given type. The device doesn't execute the script if the system's memory usage exceeds the configured value at the time the script is invoked. This configuration ensures that a device executes only essential scripts when system resources are limited, thereby enabling the device to continue performing all critical network functions.

[See [Configure Script Start Options.](#)]

Junos Telemetry Interface

- Telemetry support for interfaces and chassis (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-MP, EX4100-F, EX4300-MP, EX4400-MP, EX4400, EX4650, QFX5110, QFX5120, QFX5200, and QFX5210)—Junos OS Release 22.4R1 introduces support for streaming

operational state statistics and counters for chassis and interfaces using OpenConfig sensor paths. We also support the following new Junos-specific sensor paths for statistics that are unsupported in OpenConfig:

- `/state/chassis/`
- `/state/interfaces/`

[See [Telemetry Sensor Explorer](#).]

- **Telemetry support for PoE (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-MP, EX4100-F, EX4300-MP, EX4400-MP, EX4400, and EX4650)**—Junos OS Release 22.4R1 introduces support for streaming operational state statistics and counters for PoE using the new Junos-specific sensor path `/state/poe/`.

[See [Telemetry Sensor Explorer](#).]

Layer 2 Features

- **Configure MAC learning priority and enable persistent MAC learning on trunk interfaces (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T)**—Starting in Junos OS Release 22.4R1, you can configure MAC learning priority on the interfaces so that the high-priority interface always learns the MAC addresses. Configuring MAC learning priority ensures:
 - Traffic sent to the high-priority interface that learns the MAC address is accepted.
 - Traffic with the same MAC address sent to the low-priority interface is dropped.

By default, the switch discards the traffic if you do not configure an explicit action.

MAC address move is allowed when you configure the interfaces with the same MAC learning priority. When interfaces are not configured with MAC learning priority, then the default priority for each interface is 4. By default, discard action is taken if an explicit action is not configured. To configure MAC learning priority, use the `mac-learning-priority` configuration statement at the `[edit switch-options interface interface-name]` hierarchy level.

To accept traffic on the low-priority interface, you need to configure persistent MAC learning on the high-priority interface.

NOTE: Do not configure both MAC learning priority and persistent MAC learning on the same interface.

From Junos OS Release 22.4R1 onward, you can enable persistent MAC learning on the trunk or vlan-tagged interfaces so that traffic is accepted on the low-priority interface.

[See [Configuring MAC Learning Priority](#).]

Network Management and Monitoring

- **1:N port mirroring for sending a source packet to multiple Layer 2 destinations (EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 switches)**—Starting in Junos OS Release 22.4R1, you can use the 1:N port mirroring feature to mirror traffic to multiple Layer 2 destinations, by configuring one or both of the following configurations:
 - A port-mirroring instance that is based on a firewall filter. Use the configuration statements in the [edit forwarding-options port-mirroring instance] hierarchy.
 - A native analyzer. Use the configuration statements in the [edit forwarding-options analyzer] hierarchy.

For both configuration methods, you must also configure next-hop groups with a group type of layer-2 to direct the mirrored packets to their destinations.

Security

- **Source MAC filtering on aggregated Ethernet interfaces (EX4100 and EX4400 switches)**—Starting in Junos OS Release 22.4R1, you can configure source media access control (MAC) filtering on an aggregated Ethernet interface on EX4100 and EX4400 switches. Ingress packets are matched on the source MAC address list that you configure at the accept-source-mac mac-address hierarchy level on the logical interface of the aggregated Ethernet interface.

[See [Understanding MAC Limiting on Layer 3 Routing Interfaces](#) and [accept-source-mac](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **MACsec-bounded delay protection** (EX4100, EX4100-F, and EX4100-Multigigabit)

[See [Configuring Bounded Delay Protection](#).]

- **MAC-VRF instances with EVPN-VXLAN** (MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX2010, and MX2020 with MPC10E line cards)

We support vlan-based, vlan-aware, and vlan-bundle service types for Ethernet VPN (EVPN) unicast and multicast traffic.

[See [MAC-VRF Routing Instance Type Overview](#), [mac-vrf](#), and [service-type](#).]

- **Precision Time Protocol (PTP) transparent clock** (EX4400-48P and EX4400-48T)

[See [PTP Transparent Clocks](#) , [e2e-transparent](#), and [show ptp global-information](#).]

What's Changed

IN THIS SECTION

- [EVPN | 28](#)
- [Network Management and Monitoring | 28](#)
- [User Interface and Configuration | 29](#)

Learn about what changed in this release for EX Series switches.

EVPN

- **Flow-label configuration status for EVPN ELAN services**—The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

Network Management and Monitoring

- **Junos YANG modules for RPCs include the `junos:command` extension statement (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command` extension statement in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the YANG schemas on the local device, the YANG modules for RPCs include the `junos:command` extension statement.

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type identityref (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type identityref, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the identityref value in different modules.
- **The file copy command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The file copy command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The file copy command does not support using the | display xml filter or the | display json filter to display command output in XML or JSON format in any release. We've removed these options from the CLI.

Known Limitations

There are no known limitations in hardware or software in this release for EX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 30](#)
- [EVPN | 30](#)
- [General Routing | 30](#)
- [Interfaces and Chassis | 31](#)
- [Layer 2 Ethernet Services | 32](#)
- [Layer 2 Features | 32](#)

- Platform and Infrastructure | 32
- Virtual Chassis | 32

Learn about open issues in this release for EX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- On all Junos OS platforms, in a scaled scenario when some of the ge/xe/et interfaces are members of aggregated Ethernet (AE) and the Class of Service (CoS) forwarding-class-set configuration is applied with a wildcard for all the physical interfaces and aggregated Ethernet interfaces, it might trigger a Flexible PIC Concentrators (FPC) crash which leads to traffic loss. [PR1688455](#)

EVPN

- After Routing Engine switchover, a momentary traffic loss may be observed with EVPN VxLAN on EX4400 switches. [PR1659315](#)

General Routing

- runt, fragment and jabber counters are not incrementing on EX4300-MPs. [PR1492605](#)
- On EX2300, EX3400, EX4300-48MP and EX4300, pause frames counters does not get incremented when pause frames are sent. [PR1580560](#)
- In rare circumstances when doing Routing Engine switchover, the routing protocol daemon in former active Routing Engine (new backup Routing Engine) might restart generating a core file while in process of being terminated. [PR1589432](#)
- EX4100-24mp, 48mp, 24p/t, 48p/t, F-24p/t, F-48-p/t: In an interop scenario, when using 1G SFP Optic on PIC-2, auto-negotiation should be disabled on the peer. [PR1657766](#)

- EX4100 MACsec interface statistics of encrypted/decrypted bytes do not increment further after reaching a 40-bit limit (1099511627775). [PR1658584](#)
- If MVRP is enabled on an MSTP enabled interface, the interface will be made part of all the existing instances on the switch, Therefore, if there are two interfaces between R1 and R2 as below:
R1(et-0/0/1 and et-0/0/2)=====(et-0/0/1 and et-0/0/2)R2. And one interface is MVRP enabled (say et-0/0/1), and et-0/0/2 is not MVRP enabled. By configuration et-0/0/1 is part of MSTI-1 and et-0/0/2 is part of MSTI-2. MSTI-1 is running on vlan-100 and MSTI-2 is running on Vlan-200. R2 in this case, is advertising only vlan-100. The MVRP enabled interface will become part of all the MSTIs(MSTI-1 and MSTI-2 both) configured on the device and it will take part in the FSM of all the MSTIs. Although et-0/0/1 is not member interface of vlan-200 (corresponding to MSTI-2). This potentially can cause a problem where et-0/0/1 although not a vlan-200 member, will go into FWD state and et-0/0/2, genuine member of vlan-200 goes into BLK state for MSTI-2. Therefore, when traffic is received in vlan-200 it will be sent out of et-0/0/1, an it will be dropped.[PR1686596](#)
- Traffic loss could be seen in case configuration changes lead to switching from fallback to primary or vice-versa are committed, while SAK rollover from current live session is in progress. The issue is dependent on sequence of event at specific time. Example - MACsec session is live with primary key and at non-keyserver CAK for primary is changed this will lead to switching to fallback session, in case at same instance SAK rollover was triggered by Key-server then traffic loss will be observed.[PR1698687](#)
- On EX Series, QFX5000, and MX Series platforms having persistent binding for Dynamic Host Configuration Protocol (DHCP) snooping configured might cause Ternary Content Addressable Memory (TCAM) space exhaustion for DHCP snooping learning on the trusted port after the device reboot.[PR1699777](#)
- On EX4400 platforms in the EVPN-VXLAN environment, the overlay equal-cost multipath (ECMP) route does not get programmed in hardware resulting in traffic drops when hierarchal overlay ECMP is configured.[PR1704470](#)

Interfaces and Chassis

- On Junos OS platforms such as EX4600 and QFX5100 line of switches configured with Virtual Chassis (VC), if a master member is unplugged or forced to power off, the unicast traffic is dropped due to mac-persistence-timer expiry there is a difference in mac addresses between logical aggregated parent interface and member aggregated ethernet(ae) interface.[PR1695663](#)

Layer 2 Ethernet Services

- When an EX3400 Virtual Chassis (VC) member is zeroized or if it powered on for the first time after halt, set `chassis auto-image-upgrade` configuration is removed during the process of VC formation. Absence of this configuration will not allow the user to download configuration and images via ZTP. [PR1694952](#)

Layer 2 Features

- The memory leak might occur because of the `eswd` daemon on EX Series platforms. A message like the following is displayed in the system log: `eswd[1330]: JTASK_OS_MEMHIGH: Using 212353 KB of memory, 158 percent of available /kernel: KERNEL_MEMORY_CRITICAL: System low on free memory, notifying init (#2). /kernel: Process (1254,eswd) has exceeded 85% of RLIMIT_DATA: used 114700 KB Max 131072 KB.` [PR1262563](#)

Platform and Infrastructure

- On EX4300 platform, if encapsulation `ethernet-bridge` is configured, the interface is getting programmed as trunk instead of access in VLAN membership. This leads to untagged traffic drop. [PR1665785](#)
- On EX4300-24T, EX4300-48P, EX4300-VC, EX430024P, EX430032F and EX430048T platforms, when a VLAN Spanning Tree Protocol (VSTP) Bridge Protocol Data Unit (BPDU) arrives with a VLAN ID that is not configured in the switch, but that matches with an hardware token of any other configured VLAN, the VLAN ID of the BPDU will be changed to the VLAN ID corresponding to the matched hardware token and flooded. This disrupts STP convergence on the configured VLAN because some ports can incorrectly go into blocking state. [PR1673000](#)

Virtual Chassis

- On Junos OS EX4600 Virtual Chassis (VC), the master Routing Engine reboot and all-members reboot lead to the Packet Forwarding Engine manager hogging logs when SFP-T pluggable is installed in. The Packet Forwarding Engine manager hogging logs has no functionality impact (PR 1641556).[PR1685067](#)
- When you execute `request system reboot all members` command on EX4600 Virtual Chassis, one of the FPCs may disconnect and join the Virtual Chassis back post reboot. The FPC reboots with reason

"FXPC_RENESAUS_RETRY_FAILURE: fxc_fpga_fpc_ideeprom_read: FAILED retry to Renesaus chip". This issue may be hit approximately once in 10-12 attempts of executing the CLI command request system reboot all members. [PR1700133](#)

Resolved Issues

IN THIS SECTION

- [Forwarding and Sampling | 33](#)
- [General Routing | 33](#)
- [Infrastructure | 35](#)
- [Layer 2 Ethernet Services | 36](#)
- [Network Management and Monitoring | 36](#)
- [Platform and Infrastructure | 36](#)
- [Routing Protocols | 36](#)
- [User Interface and Configuration | 36](#)
- [Virtual Chassis | 37](#)

Learn about the issues fixed in this release for EX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Forwarding and Sampling

- Traffic drop seen and filter do not hit as expected for match condition traffic class with FLT option configured. [PR1573350](#)

General Routing

- On the EX4650 device, the filter action to change VLAN does not work. [PR1362609](#)

- DHCP traffic might be dropped when DHCP-security and RTG are enabled. [PR1647209](#)
- EX4100: Junos Telemetry Interface FAN and power supply names do not match with CLI. [PR1648739](#)
- EX4100: Wrong trap generated after removal of fan0 in FPC4. [PR1652388](#)
- EVPN-VXLAN: With EVPN type 5 and NSR configuration, a Routing Engine switchover might result in momentary traffic drop for about 2-3 seconds. [PR1655052](#)
- The egress traffic is not tagged properly in a L2PT scenario. [PR1655511](#)
- Filter-based forwarding filter might not work as expected. [PR1656117](#)
- The interface might not come up on EX Series platforms. [PR1656540](#)
- DHCP packets might get looped in a VXLAN setup. [PR1657597](#)
- port/mac gbp tags may not be carried forward to the spine. [PR1659384](#)
- Packet count may occasionally be 0 for some interfaces in "monitor interface traffic". [PR1661617](#)
- LEDs on ports 0-35 are always lit on EX4400-48MP platforms. [PR1662288](#)
- The fxpc crash might be observed with the RPF check enabled. [PR1662508](#)
- In the EVPN-VXLAN scenario, the DHCP packets will get dropped when the DHCP relay agent is configured. [PR1662524](#)
- SSH traffic might be affected when filter log action is used. [PR1663126](#)
- MAC address learning failure and traffic loss might be observed on VXLAN enabled ports having native-VLAN configured. [PR1663172](#)
- MAC addresses learned on the RTG interface are not aging out. [PR1664955](#)
- MAC-IP bindings for IPv4 (ARP) and IPv6 (ND) may not be processed for IRB interfaces in an EVPN scenario. [PR1665828](#)
- High numbers of PDs connected may result in high CPU utilization. [PR1667564](#)
- Shaping-rate is not taking 20 bytes of overhead into account. [PR1667879](#)
- MAC RADIUS authentication without restrict option updates authenticated VLAN information before client authentication. [PR1668144](#)
- Traffic flow will be affected as interfaces will be removed from VLAN. [PR1675861](#)
- EX4100 and EX4100-F series: Traffic would not go through on management port at link speed 10 and 100M. [PR1676433](#)

- VLAN translation mapping gets deleted when one of the member interface removed from LAG. [PR1676772](#)
- Aggregated Ethernet interface will receive unknown unicast traffic on FPC3 reboot of a Virtual Chassis. [PR1678430](#)
- Firewall functions will not work as expected when egress firewall filter is configured. [PR1679574](#)
- DHCP binding will fail for the clients (Clients connected on an aggregated Ethernet interface with 2 or more VLANs) on a VLAN where DHCP security is not configured. [PR1679094](#)
- On EX2300 and EX3400, set system ports console log-out-on-disconnect does not allow user to log in via console. [PR1680408](#)
- Multicast traffic loss is seen with igmp-snooping running on EX4100. [PR1681478](#)
- EX4100-24mp/48mp/48p/48t/24p/24t: Activity LED is lit on some ports if 1G optic is inserted without link being present or up. [PR1682633](#)
- The I2cpd process crash might be observed when disabling RSTP on an interface. [PR1684072](#)
- Licenses on the device might become invalid when the device is upgraded from a legacy licensing-based release to an agile licensing-based release. [PR1684842](#)
- MAC address learning might not happen on specific EX Series and QFX Series platforms. [PR1685938](#)
- The I2ald core file seen after zeroize. [PR1686097](#)
- On EX4100 and EX4400 platform, alarm 'PEM is not supported' might be seen. [PR1690674](#)
- The factory default configuration does not have xe-0/2/0. [PR1691174](#)
- Few uplink ports of EX2300-48MP are not coming up. [PR1692579](#)
- The dot1x reauthentication will not work for a port with VoIP VLAN. [PR1693640](#)
- On a PVLAN with DAI ARP packets will be forwarded between isolated ports. [PR1694800](#)
- The dot1x authentication will not be enabled on interfaces with specific configuration combination. [PR1696906](#)
- Dot1x authentication do not occur with EVPN VXLAN end-to-end configuration. [PR1697995](#)

Infrastructure

- The auxiliary serial port (of type USB-C on the front panel) does not show any output. [PR1616315](#)

- On EX4100, if a live vmcore is attempted to be created, the DUT might get stuck and reboot. [PR1656625](#)
- On EX4400 upgrade fails when upgrading through a USB drive. [PR1681783](#)

Layer 2 Ethernet Services

- phone-home and SZTP may fail if phone-home daemon restarts. [PR1693124](#)

Network Management and Monitoring

- Observed memory leak in eventd leak during GRES. [PR1602536](#)
- The "snmpd" process might crash if SNMP timeout occur. [PR1666548](#)

Platform and Infrastructure

- On EX4300 platform, high CPU is seen with generation of log message `"/kernel: %KERN-3: i802_3_slow_rcv_input:oam/esmc PDU dropped"`. [PR1661332](#)
- EX9000 and MX Series platforms do not relay a DHCP offer with a broadcast flag under EVPN-VXLAN scenario. [PR1670923](#)
- The fxpc process might crash on EX4300 and EX4300-VC platforms. [PR1675977](#)

Routing Protocols

- High CPU will be seen due to frequent triggering of SPF for IS-IS. [PR1667575](#)

User Interface and Configuration

- Commit and commit check fails when the `interface-range` statement is configured. [PR1656565](#)

Virtual Chassis

- On Junos OS EX4600, EX4650, and QFX5000 Virtual Chassis platforms, line card might be disconnected from Virtual Chassis post master Routing Engine reboot. [PR1669241](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 37](#)

This section contains the upgrade and downgrade support policy for Junos OS for EX Series switches. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases - 21.1, 21.2 and 21.3 or downgrade to the previous three releases - 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases - 21.2 and 21.4 or downgrade to the previous two EEOL releases - 20.2 and 19.4.

Table 2: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for JRR Series

IN THIS SECTION

- [What's New | 39](#)
- [What's Changed | 40](#)
- [Known Limitations | 40](#)
- [Open Issues | 40](#)
- [Resolved Issues | 40](#)
- [Migration, Upgrade, and Downgrade Instructions | 41](#)

What's New

IN THIS SECTION

- [Routing Protocols](#) | 39

Learn about new features introduced in this release for JRR Series Route Reflectors.

Routing Protocols

- **Support for bootstrapping route-validation database from a local file (cRPD, JRR200, MX204, PTX10008, and QFX10008)**—Starting in Junos OS Release 22.4R1, we support the ability to read validation records from a local binary file and install into the specified named route-validation databases within RPD. This feature implements syntactic and semantic checks on the content of the file to ensure that it is a well-specified set of validation records. If the syntactic and semantic checks fail, the entire file is rejected as a source of validation records. Use the `source-file` statement at the `[edit routing-options validation]` hierarchy level to source route-validation records from a local file source. You can use the `show validation source-file` command to display the properties of a local validation source file.

[See [validation](#).]

- **MVPN feature support with sharding (cRPD, JRR200, MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support the following features:
 - Multicast virtual private network (MVPN) inactive route query from the main thread to shards
 - Extranet and auto-export support with sharding
 - Interact functions with RT-proxy client and server
 - New data structure to store the inactive route data on the main thread
 - Asynchronous route processing on the main thread

You can use `show mvpn c-multicast` to display the inactive route data stored on the main thread.

[See [rib-sharding](#) and [show mvpn c-multicast](#) .]

What's Changed

There are no changes in behavior and syntax in this release for JRR Series Route Reflectors.

Known Limitations

There are no known limitations in hardware or software in this release for JRR Series Route Reflectors.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

There are no known issues in hardware or software in this release for JRR Series Route Reflectors.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

IN THIS SECTION

- [General Routing | 41](#)

Learn about the issues fixed in this release for JRR Series Route Reflectors.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- IS-IS packet drops for packets with GRE over FTI-VXLAN header. [PR1676912](#)
- A 802.1Q tagged Ethernet traffic with an expected VLAN ID and with a non-zero 802.1P value ingressing a JRR200 VLAN enabled interface drops. [PR1691694](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 41](#)

This section contains the upgrade and downgrade support policy for Junos OS for the JRR Series Route Reflector. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

For information about software installation and upgrade, see the [JRR200 Route Reflector Quick Start](#) and [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if

the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 3: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for Juniper Secure Connect

IN THIS SECTION

- [What's New | 43](#)
- [What's Changed | 43](#)
- [Known Limitations | 43](#)
- [Open Issues | 43](#)
- [Resolved Issues | 43](#)

What's New

There are no new features or enhancements to existing features in this release for Juniper Secure Connect.

What's Changed

There are no changes in behavior and syntax in this release for Juniper Secure Connect.

Known Limitations

There are no known limitations in hardware or software in this release for Juniper Secure Connect.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

There are no known issues in hardware or software in this release for Juniper Secure Connect.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

There are no resolved issues in this release for Juniper Secure Connect.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Junos OS Release Notes for Junos Fusion for Enterprise

IN THIS SECTION

- [What's New | 44](#)
- [What's Changed | 44](#)
- [Known Limitations | 44](#)
- [Open Issues | 45](#)
- [Resolved Issues | 45](#)
- [Migration, Upgrade, and Downgrade Instructions | 45](#)

What's New

There are no new features or enhancements to existing features in this release for Junos fusion for enterprise.

What's Changed

There are no changes in behavior and syntax in this release for Junos Fusion for enterprise.

Known Limitations

There are no known limitations in hardware or software in this release for Junos fusion for enterprise.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

There are no known issues in hardware or software in this release for Junos Fusion for enterprise.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

There are no resolved issues in this release for Junos Fusion for enterprise.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | 46
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | 47
- [Preparing the Switch for Satellite Device Conversion](#) | 48
- [Converting a Satellite Device to a Standalone Switch](#) | 50
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 50
- [Downgrading Junos OS](#) | 51

This section contains the procedure to upgrade or downgrade Junos OS and satellite software for a Junos fusion for enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos fusion for enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the `junos-install` package. Use other packages (such as the `jbundle` package) only when so instructed by a Juniper Networks support representative. For information about the contents of the `junos-install` package and details of the installation process, see the [Installation and Upgrade Guide](#).

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the [Junos OS Administration Library](#).

To download and install Junos OS:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos fusion** to find the software that you want to download.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list on the right of the page.
5. Select the **Software** tab.
6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new `junos-install` package on the aggregation device.

NOTE: We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

Customers in the United States and Canada, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot source/package-name.n.tgz
```

All other customers, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot source/package-name.n-limited.tgz
```

Replace *source* with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - *ftp://hostname/pathname*
 - *http://hostname/pathname*
 - *scp://hostname/pathname* (available only for Canada and U.S. version)

The `validate` option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Adding the `reboot` command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to minimize disrupting network operations as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos fusion for enterprise. See [Configuring or Expanding a Junos fusion for enterprise](#).

For satellite device hardware and software requirements, see [Understanding Junos fusion for enterprise Software and Hardware Requirements](#).

Use the following command to install Junos OS on a switch before converting it into a satellite device:

```
user@host> request system software add validate reboot source/package-name
```

NOTE: The following conditions must be met before a Junos switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The switch running Junos OS can be converted only to SNOS 3.1 and later.
- Either the switch must be set to factory-default configuration by using the `request system zeroize` command, or the following command must be included in the configuration: `set chassis auto-satellite-conversion`.

When the interim installation has completed and the switch is running a version of Junos OS that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device using the console port.

2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

NOTE: The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port port-number
```

For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 0
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 1
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 2
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 3
```

This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, or preconfiguration. See [Configuring or Expanding a Junos fusion for enterprise](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Switch

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove it from the Junos fusion topology. For more information, see [Converting a Satellite Device to a Standalone Device](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 4: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Downgrading Junos OS

Junos fusion for enterprise is first supported in Junos OS Release 16.1, although you can downgrade a standalone EX9200 switch to earlier Junos OS releases.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

To downgrade a Junos fusion for enterprise, follow the procedure for upgrading, but replace the junos-install package with one that corresponds to the appropriate release.

Junos OS Release Notes for Junos Fusion for Provider Edge

IN THIS SECTION

- [What's New | 51](#)
- [What's Changed | 52](#)
- [Known Limitations | 52](#)
- [Open Issues | 52](#)
- [Resolved Issues | 52](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

What's New

There are no new features or enhancements to existing features in this release for Junos Fusion for Provider Edge.

What's Changed

There are no changes in behavior and syntax in this release for Junos Fusion for provider edge.

Known Limitations

There are no known limitations in hardware or software in this release for Junos fusion for provider edge.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

There are no known issues in hardware or software in this release for Junos Fusion for provider edge.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

Learn about the issues fixed in this release for Junos Fusion for provider edge.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device](#) | 53
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | 56

- [Preparing the Switch for Satellite Device Conversion | 56](#)
- [Converting a Satellite Device to a Standalone Device | 58](#)
- [Upgrading an Aggregation Device | 60](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 60](#)
- [Downgrading from Junos OS Release 22.4 | 61](#)

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for Junos fusion for provider edge. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

Basic Procedure for Upgrading an Aggregation Device

When upgrading or downgrading Junos OS, always use the `jinstall` package. Use other packages (such as the `jbundle` package) only when so instructed by a Juniper Networks support representative. For information about the contents of the `jinstall` package and details of the installation process, see the [Installation and Upgrade Guide](#).

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful.

Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the [Installation and Upgrade Guide](#).

The download and installation process for Junos OS Release 22.4R1 is different from that for earlier Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

<https://www.juniper.net/support/downloads/>

2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos fusion** to find the software that you want to download.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the page.
5. Select the **Software** tab.
6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new `jinstall` package on the aggregation device.

NOTE: We recommend that you upgrade all software packages out-of-band using the console, because in-band connections are lost during the upgrade process.

Customers in the United States and Canada, use the following commands.

- For 64-bit software:

NOTE: We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot source/jinstall64-22.4R1.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot source/jinstall-22.4R1.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

NOTE: We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot source/jinstall64-22.4R1.SPIN-
export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot source/jinstall-22.4R1.SPIN-
export-signed.tgz
```

Replace *source* with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - *ftp://hostname/pathname*
 - *http://hostname/pathname*
 - *scp://hostname/pathname* (available only for the Canada and U.S. version)

The *validate* option validates the software package against the current configuration as a prerequisite for adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is for a different release.

Adding the *reboot* command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 22.4R1 *jinstall* package, you cannot return to the previously installed software by issuing the `request system software rollback` command. Instead, you must issue the `request system software add validate` command and specify the *jinstall* package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately as follows to minimize disrupting network operations:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see [Understanding Junos fusion Software and Hardware Requirements](#)

NOTE: The following conditions must be met before a standalone switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The switch can be converted to only SNOS 3.1 and later.
- Either the switch must be set to factory-default configuration by using the `request system zeroize` command, or the following command must be included in the configuration: `set chassis auto-satellite-conversion`.

Customers with EX4300 switches, use the following command:

```
user@host> request system software add validate reboot source/jinstall-ex-4300-14.1X53-D43.3-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add reboot source/jinstall-qfx-5-14.1X53-D43.3-domestic-signed.tgz
```

When the interim installation has completed and the switch is running a version of Junos and OS on one line that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device by using the console port.
2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

NOTE: The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device by using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose your connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port port-number
```

For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 0
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 1
```

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 2
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 3
```

This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See [Configuring Junos fusion for provider edge](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Device

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove the satellite device from the Junos fusion topology.

NOTE: If the satellite device is a QFX5100 switch, you need to install a PXE version of Junos OS. The PXE version of Junos OS is software that includes *pxe* in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D43 is named `install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz`. If the satellite device is an EX4300 switch, you install a standard `jinstall-ex-4300` version of Junos OS.

The following steps explain how to download software, remove the satellite device from Junos fusion, and install the Junos OS software image on the satellite device so that the device can operate as a standalone device.

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:

<https://www.juniper.net/support/downloads>
2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos fusion** from the drop-down list and select the switch platform series and model for your satellite device.
4. Select the Junos OS Release 14.1X53-D30 software image for your platform.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.

7. Copy the software to the routing platform or to your internal software distribution site.
8. Remove the satellite device from the automatic satellite conversion configuration.

If automatic satellite conversion is enabled for the satellite device's member number, remove the member number from the automatic satellite conversion configuration. The satellite device's member number is the same as the FPC slot ID.

```
[edit]
user@aggregation-device# delete chassis satellite-management auto-satellite-conversion
satellite member-number
```

For example, to remove member number 101 from Junos fusion:

```
[edit]
user@aggregation-device# delete chassis satellite-management auto-satellite-conversion
satellite 101
```

You can check the automatic satellite conversion configuration by entering the show command at the [edit chassis satellite-management auto-satellite-conversion] hierarchy level.

9. Commit the configuration.

To commit the configuration to both Routing Engines:

```
[edit]
user@aggregation-device# commit synchronize
```

Otherwise, commit the configuration to a single Routing Engine:

```
[edit]
user@aggregation-device# commit
```

10. Install the Junos OS software on the satellite device to convert the device to a standalone device.

```
[edit]
user@aggregation-device> request chassis satellite install URL-to-software-package fpc-slot
member-number
```

For example, to install a PXE software package stored in the `/var/tmp` directory on the aggregation device onto a QFX5100 switch acting as the satellite device using FPC slot 101:

```
[edit]
user@aggregation-device> request chassis satellite install /var/tmp/install-media-pxe-
qfx-5-14.1X53-D43.3-signed.tgz fpc-slot 101
```

For example, to install a software package stored in the `var/tmp` directory on the aggregation device onto an EX4300 switch acting as the satellite device using FPC slot 101:

```
[edit]
user@aggregation-device> request chassis satellite install /var/tmp/jinstall-
ex-4300-14.1X53-D30.3-domestic-signed.tgz fpc-slot 101
```

The satellite device stops participating in the Junos fusion topology after the software installation starts. The software upgrade starts after this command is entered.

11. Wait for the reboot that accompanies the software installation to complete.
12. When you are prompted to log back into your device, unconnect the device from the Junos fusion topology. See [Removing a Transceiver from a QFX Series Device](#) or [Remove a Transceiver](#), as needed. Your device has been removed from Junos fusion.

NOTE: The device uses a factory-default configuration after the Junos OS installation is complete.

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 22.4R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 5: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Downgrading from Junos OS Release 22.4

To downgrade from Release 22.4 to another supported release, follow the procedure for upgrading, but replace the 22.4 `jinstall` package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for MX Series

IN THIS SECTION

- [What's New | 62](#)
- [What's Changed | 83](#)
- [Known Limitations | 85](#)
- [Open Issues | 87](#)
- [Resolved Issues | 96](#)
- [Migration, Upgrade, and Downgrade Instructions | 115](#)

What's New

IN THIS SECTION

- [Class of Service | 63](#)
- [EVPN | 64](#)
- [Flow-based and Packet-based Processing | 65](#)
- [High Availability | 66](#)
- [Interfaces | 66](#)
- [Junos Telemetry Interface | 66](#)
- [Layer 2 VPN | 68](#)
- [Licensing | 68](#)

- MPLS | 71
- Network Address Translation (NAT) | 73
- Network Management and Monitoring | 73
- OpenConfig | 73
- Precision Time Protocol (PTP) | 74
- Routing Protocols | 74
- Services Applications | 77
- Software Installation and Upgrade | 77
- Source Packet Routing in Networking (SPRING) or Segment Routing | 78
- Subscriber Management and Services | 78
- VPNs | 80
- Additional Features | 80

Learn about new features introduced in this release for the MX Series routers.

Class of Service

- **Support for hierarchical class of service (HCoS) on network slices on aggregated Ethernet interfaces (MX480, MX960, MX10003, and MX2020)**—Starting in Junos OS Release 22.4R1, you can provision HCoS for network slices on aggregated Ethernet (ae-) interfaces. You can use either the default scale mode or the replicate mode for the ae- interfaces.

To enable HCoS for network slices on ae- interfaces, set `hierarchical-scheduler` at the `[interface ae-interface-name]` hierarchy level, then attach `output-traffic-control-profile tcp-name` at the `[class-of-service interface ae-interface-name slice slice-name]` hierarchy level.

[See [Hierarchical Class of Service for Network Slicing](#).]

- **Support for IPv6 class of service in SRv6 network programming (MX Series)**—Starting with Junos OS Release 22.4R1, MX Series devices that deploy Segment Routing for IPv6 (SRv6) network programming can support:
 - IPv6 classifiers on both transit and end-point nodes.
 - Rewrite rules
 - Policy maps

The support applies to both transit and end-point nodes.

[See [How to Enable SRv6 Network Programming in IS-IS Networks.](#)]

EVPN

- **EVPN-MPLS E-LAN flow-aware transport (FAT) label load balancing (MX Series, EX9200, vMX)** — Starting in Junos OS Release 22.4R1, you can configure provider edge (PE) devices to use FAT labels in an Ethernet VPN-MPLS (EVPN-MPLS) routing instance, according to Request for Comments (RFC) 6391. PE devices use these labels to load-balance EVPN-MPLS unicast packets across equal-cost multipaths (ECMPs) without performing deep packet inspection of the MPLS payload. This feature supports emulated LAN (ELAN) with single-homing and multi-homing active/standby and active/active topologies and supports the VLAN-based, VLAN-bundle, and VLAN-aware bundle EVPN-MPLS variants.

NOTE: This feature does not support MX Series devices with Advanced Forwarding Toolkit (AFT) cards.

NOTE: On MX Series devices, a configuration where the local PE has a static-flow-label and the remote PE does not have a static-flow-label, the remote PE can process packets without dropping any traffic.

Enabling Load Balancing Using Fat Labels for EVPN Routing Instances:



WARNING: Configuring a flow label or deleting a flow label with the following CLI commands causes a catastrophic event for the routing instance. As a best practice, perform these CLI commands during a maintenance period to avoid network disruptions.

- Configure the flow-label-static statement at the [edit routing-instances routing-instance-name protocols evpn] hierarchy level on PE devices to insert FAT flow labels into pseudowire packets sent to remote PE devices.
- Configure the flow-label statement at the [edit routing-instances routing-instance-name protocols evpn] hierarchy level on PE devices to signal flow-label capability in the EVPN Layer 2 Attributes Extended Community by setting the flow-label (F) bit in the EVPN Type 3 route.

[See [flow-label](#) and [flow-label-static](#).]

- **EVPN-VPWS over SRv6 underlay (MX240, MX304, MX480, MX960, MX10003, MX10008, MX2010, and MX2020)**—Starting in Junos OS Release 22.4R1, you can configure a single-active or

an all-active multihomed Ethernet VPN–virtual private wireless service (EVPN-VPWS) network using segment routing over a Segment Routing for IPv6 (SRv6) underlay.

[See [Overview of VPWS with EVPN Signaling Mechanisms.](#)]

- **EVPN-VXLAN to EVPN-VXLAN seamless stitching for EVPN Type 5 routes (EX4100-24T, EX4400-24MP, EX4400-24P, EX4400-48F, EX4650, MX204, MX240, QFX10002-60C, QFX10002, QFX10008, QFX10016, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 22.4R1, you can configure Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) to EVPN-VXLAN seamless stitching with EVPN Type 5 (IP prefix) routes between two interconnected data centers or between two points of delivery (pods) in a data center.

In the EVPN-VXLAN fabric, border leaf or border spine devices act as interconnection gateways. You enable EVPN Type 5 routes in virtual routing and forwarding (VRF) instances on both sides of the interconnection. For each VRF instance, the server leaf devices in the first data center create VXLAN tunnels for Type 5 routes (with corresponding virtual network identifiers [VNIs]) toward their local gateway devices. The gateway devices map those VXLAN tunnels to an interconnection tunnel (with a new route distinguisher [RD], route target, and VNI) toward the second data center. The gateway devices in the second data center re-create the Type 5 VXLAN tunnels using their local RD.

We support one-to-one mapping of Type 5 VRF instances across the interconnection.

- **Support for Microsoft load-balancing node's static ARP entries with unicast MAC addresses (EX9208, MX-Series, and VMX)**—Starting in Junos OS Release 22.4R1, you can configure a Microsoft load-balancing node's static Address Resolution Protocol (ARP) entries for unicast MAC addresses on integrated routing and bridging (IRB) interfaces. On your provider edge (PE) device, you can create a static ARP entry for the Microsoft load-balancing node's virtual IP address and its unicast virtual MAC address. This static ARP configuration enables your PE devices to flood traffic for the Microsoft load-balancing node's virtual IP address to the virtual MAC address in an EVPN Layer 2 domain or any other Layer 2 domain.

To enable unicast MAC addresses on IRB interfaces, enable the `flood-as-unknown-unicast` option in the `[edit interfaces irb unit <logical-interface-number> family inet address <local-ip-address>/<prefix-length> arp <MSLB-virtual IP address> mac <MSLB-unicast-VMAC>]` hierarchy. The `flood-as-unknown-unicast` option enables flooding of virtual IP addresses and virtual MAC traffic flows from a Microsoft load-balancing cluster.

[See [EVPN User Guide.](#)]

Flow-based and Packet-based Processing

- **Addition of new server on ECMP**—Starting in Junos OS Release 22.4R1, when a new equal-cost multipath (ECMP) is added dynamically, we avoid complete redistribution of flows. The flows mapped to the existing ECMP remain consistent where only a proportion of the flows are moved to the newly

added ECMP. This form of redistribution ensures minimal disruption to the existing ECMP flows and also maintains uniformity in the flow from each active ECMP to the newly added ECMP.

The MX Series routers with MPC10, MPC11, LC9600, and MX304 cards support multiple path failures in an ECMP group. This capability prevents reordering of flows to active paths in an ECMP group when one or more paths fail.

High Availability

- **MVPN NSR with BGP Sharding enabled (MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we've enabled multicast virtual private network (MVPN) nonstop active routing (NSR) for border gateway function (BGP) sharding.

[See [Understanding BGP RIB sharding and BGP Update IO thread.](#)]

Interfaces

- **Support for TLS-Hello in Traffic Load Balancer (MX240, MX480, and MX960)**—Starting in Junos OS Release 22.4R1, we've enhanced Traffic Load Balancer (TLB) to support TLS-Hello, a new health-check type. We've added a CLI knob in the network monitoring profile and you can use his knob to configure TLS-Hello probing type.

TLS v1.2 and v1.3 health checks are supported for TLS-Hello over TCP.

[See [Traffic Load Balancing.](#)]

- **Support for symmetric load balancing on MPC10 and MPC11 (MX2010 and MX2020)**— Starting in Junos OS Release 22.4R1, we support symmetric load balancing on MPC10 and MPC11.

[See [Load Balancing on Aggregated Ethernet Interfaces.](#)]

Junos Telemetry Interface

- **Enhanced support for FIB telemetry streaming (MX240, MX960, MX2020, PTX5000, and PTX1000)**—Junos OS Release 22.4R1 introduces enhanced support for forwarding information base (FIB) telemetry streaming based on the OpenConfig Abstract Forwarding Table (AFT) model. We now support the following sensor paths:

- `/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/origin-protocol`
- `/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/origin-protocol`
- `/network-instances/network-instance/afts/next-hops/next-hop/state/encapsulate-header`

[See [Telemetry Sensor Explorer.](#)]

- **Event-driven streaming of sensor data for MPLS LSP record route objects (ACX5448, ACX7100, MX204, MX240, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008,**

MX2010, MX2020, PTX1000, and vMX)—Junos OS Release 22.4R1 introduces ON_CHANGE notification for streaming MPLS label-switched path (LSP) record route object statistics. Using ON_CHANGE mode, data values are not streamed but sent only when data values change. Support includes leaf nodes under the resource path `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/sessions/session/record-route-objects/record-route-object/state/`.

[See [Telemetry Sensor Explorer](#).]

- **OpenConfig OSPF configuration and operational state sensors (ACX5448, ACX7100, MX150, MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, and PTX1000)**—Junos OS Release 22.4R1 introduces support for the OpenConfig OSPF data model `openconfig-ospfv2.yang (v.0.3.1)`. We now support configuration and streaming of operational state data under the resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.

To learn about OpenConfig configuration mappings, see [Mapping OpenConfig OSPF Commands to Junos Configuration](#). For state sensors, see [Telemetry Sensor Explorer](#).

- **System health reporting sensors support on gRPC (ACX5448 and ACX710 routers; MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10008, and MX10016 routers; and PTX10002 routers)**—Starting in Junos OS Release 22.4R1, Junos telemetry interface (JTI) Junos telemetry interface (JTI) supports data model `openconfig-system.yang` using gRPC remote procedure calls (gRPC) and provides new health-monitoring sensors.

[See [Telemetry Sensor Explorer](#).]

- **Telemetry sensor support (MX10004)**—Starting in Junos OS Release 22.4R1, Junos telemetry interface (JTI) supports sensors in the following telemetry statistics areas:
 - Transceiver diagnostics
 - Native optics
 - Native fabric statistics
 - Physical Ethernet interface
 - Chassis management error (cmerror) configuration and counters
 - Platform, interface, and alarms
 - Flexible PIC Concentrator (FPC) environment

[See [Telemetry Sensor Explorer](#) and [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

- **Support for gRPC tunnel sessions (MX204, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX-5000, PTX1000, PTX10002, VMX, and QFX5110)**—Starting with Junos OS Release 22.4R1, you can configure a gRPC tunnel session to

establish a connection between an external TCP client and a TCP server. The gRPC tunnel session establishes a reverse connection when a TCP client can't reach the TCP server.

To establish a gRPC tunnel session, include the `grpc-tunnel` configuration statement in the `[edit system services]` hierarchy.

[See [gRPC Tunnels Overview](#).]

Layer 2 VPN

- **Improved pseudowire failover convergence in an H-VPLS network (MX Series)**—Starting in Junos OS Release 22.4R1, we've improved the performance of hot-standby redundant pseudowires in an hierarchical VPLS (H-VPLS) hub-and-spoke network. The hot-standby option enables the network to quickly recover from a failure in the active pseudowire and reduces the amount of traffic that gets discarded during the transition.

[See [Configure Hot-Standby Pseudowire Redundancy in H-VPLS](#) .]

- **Preserving next hop hierarchy for Layer 2 services (MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX2010, and MX2020)**—Starting in Junos OS Release 22.4R1, Junos OS supports expanded next hop hierarchy when forwarding BUM traffic for Layer 2 services. To support expanded next hop hierarchy, set the `preserve-nexthop-hierarchy` option at the `[edit routing-options resolution]` hierarchy level.

[See [resolution](#).]

Licensing

- **Juniper Agile Licensing (MX10004)**—Starting in Junos OS Release 22.4R1, the MX10004 device supports Juniper Agile Licensing.

Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can use Juniper Agile Licensing to install and manage licenses for hardware and software features.

Juniper Agile Licensing supports soft enforcement and hard enforcement of hardware and software feature licenses.

- With soft enforcement, if you configure a feature without a license, Junos OS displays a warning when you commit the configuration. However, the feature remains operational. In addition, Junos OS generates periodic alarms indicating that you need the license to use the feature. You can see the list of alarms at [System Log Explorer](#).
- With hard enforcement, if you configure a feature without a license, Junos OS displays a warning when you commit the configuration. The feature is not operational until you install the license. In addition, Junos OS generates periodic syslog messages indicating that you need the license to use the feature. You can see the list of syslog messages at [System Log Explorer](#).

Table 6: Licensed Features on MX10004 Device

License Model	Use Case Examples or Solutions	Features	Scale
Standard	Basic Layer 2 features	Bridging with port and single level VLAN (Dot1Q), LAG, and xSTP	Not Applicable

Table 6: Licensed Features on MX10004 Device *(Continued)*

License Model	Use Case Examples or Solutions	Features	Scale
Advanced	Transport	<ul style="list-style-type: none"> • Includes standard features • IP routing, IGP (OSFP and IS-IS), IP-FRR, PIM variants, and IGMP • Internet eBGP peering, BGP multihoming (add path and multipath), EPE, and BGP PIC • BGP Flow specification • All Layer 2 services— E-LINE (Layer 2 VPNs, Layer 2 ckt, EVPN-VPWS, EVPN FXC), E-LAN (bridging, H-VPLS, EVPN, and IRB), E-Tree (H-VPLS, EVPN, and IRB), Layer 2 multicast (snooping included) • All MPLS transport— LDP, RSVP-TE, SR, SR-TE, and MPLS-FRR (including TI-LFA) • IP fabrics (MPLS-over-UDP, VXLAN, and IP-in-IP) • GRE 	<p>32 IP VPNs</p> <p>8 multicast VPNs</p>

Table 6: Licensed Features on MX10004 Device *(Continued)*

License Model	Use Case Examples or Solutions	Features	Scale
		<ul style="list-style-type: none"> • Streaming telemetry and SNMP • Policers, ACLs, J-Flow, port mirroring, and per VLAN queuing • PWHT for Layer 2 • Timing (all variants) • OAM—BFD, Ethernet CFM or LFM, MPLS or SR (ping and traceroute), services OAM, RPM, and TWAMP 	
Premium	Services	<ul style="list-style-type: none"> • Includes advanced features • High scale IP-VPNs • IP fabrics (SRv6 and SRm6) • PWHT for Layer 3 VPNs • Inline NAT and inline MDI • 1:1 inline J-Flow 	32+ IP VPNs 8+ multicast VPNs

[See [Flex Software License for MX](#) and [Managing Licenses](#).]

MPLS

- **Automatic Derivation of remote discriminator for S-BFD session in SR-TE (MX204, MX480, MX960, MX10008, and MX2008)**—Starting in Junos OS Release 22.4R1, you can use the automatically

derived remote discriminator for Seamless BFD (S-BFD) sessions on segment routing-traffic engineering (SR-TE) paths. With this feature, you don't need to configure a remote discriminator in the S-FBD configuration on the ingress or transit device and a matching local-discriminator on its respective endpoint. Instead, the egress provider edge device will now accept an IP address as a local discriminator.

To configure this feature, use the `set protocols bfd sbfd local-discriminator-ip` command.

Additionally, you can now use a common sBFD template with the S-FBD configurations on multiple controller-provisioned SR-TE policies. In these sBFD sessions, Junos OS automatically derives the remote discriminator from the tunnel endpoint for matching SR-TE policies.

[See [sbfd](#) and [Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution.](#)]

- **Support for LDP tunneling over SR-TE (ACX5448, MX480, MX960, and MX2010)**—Starting in Junos OS Release 22.4R1, you can tunnel LDP label-switched paths (LSPs) over segment routing-traffic engineering (SR-TE) in OSPF networks. Tunneling LDP over SR-TE provides consistency and coexistence of both LDP LSPs and SR-TE LSPs.

To configure LDP tunneling over SR-TE, include the `tunnel-source-protocol` configuration statement at the `[edit protocols ospf traffic-engineering]` and `ldp-tunneling` configuration statement at the `[edit protocols ospf source-packet-routing source-routing-path]` hierarchy levels.

[See [Tunneling LDP over SR-TE.](#)]

- **PCEP multipath support for SR-TE (MX480, PTX10008, and QFX5200)**—Starting in Junos OS Release 22.4R1, you can configure the multipath feature (primary or secondary paths) for Path Computation Element Protocol (PCEP) segment routing-traffic engineering (SR-TE) as defined in *draft-ietf-pce-multipath-06*. We support the following multipath capabilities:
 - When the PCEP multipath feature is enabled, you can configure multiple primary or secondary paths in a candidate path that you configure and control using Path Computation Client (PCC). Note that the PCEP multipath feature is enabled by default.
 - When the PCEP multipath feature is disabled, you can configure only one primary path in a candidate path. Note that a secondary path configuration is not allowed.

The PCEP multipath feature removes the `compute-profile` restriction of 1 on the maximum number of segment lists (`maximum-computed-segment-lists`).

NOTE: When PCEP multipath is enabled, PCCD will not send constraints for PCC-controlled candidate paths.

[See [Configuring Multiple Paths for Path Computation Element Protocol SR-TE Overview.](#)]

- **Support for ingress and transit-chained CNHs for LDP, RSVP, L2VPN, L3VPN, and static protocols (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2010, and MX2020)**—Starting in Junos OS Release 22.4R1, you can configure chained composite next hops (CNHs) for devices carrying ingress or transit traffic in the network. We now support the following options:
 - On the LDP ingress router, set `routing-options forwarding-table chained-composite-next-hop ingress ldp`
 - For the LDP, RSVP, L2VPN, L3VPN, and static protocols on the transit router:
 - `set routing-options forwarding-table chained-composite-next-hop transit ldp`
 - `set routing-options forwarding-table chained-composite-next-hop transit rsvp`
 - `set routing-options forwarding-table chained-composite-next-hop transit static`
 - `set routing-options forwarding-table chained-composite-next-hop transit l2vpn`
 - `set routing-options forwarding-table chained-composite-next-hop transit l3vpn`

We support class of service (CoS) and rewrite rules for ingress and transit-chained CNHs for these configuration commands.

[See [labeled-bgp](#), [chained-composite-next-hop](#), and [ingress \(Chained Composite Next Hop\)](#).]

Network Address Translation (NAT)

Network Management and Monitoring

- **Enhanced management traffic (MX Series)**—Starting in Junos OS Release 22.4R1, you can prioritize the management traffic per packet classification. This enhancement improves the performance of the internal traffic and management traffic. The enhancement also improves packet processing.

We've also introduced three operational commands that provide easily interpretable results for debugging packet path issues in the Routing Engine.

[See [priority](#), [show system statistics netisr per-priority-stats](#), [show system statistics netisr qdrops](#), and [show system statistics netisr socket-drops](#).]

OpenConfig

- **OpenConfig OSPF configuration and operational state sensors (ACX5448, ACX7100, MX150, MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, and PTX1000)**—Starting in Junos OS Release 22.4R1, we support the OpenConfig OSPF data model `openconfig-ospfv2.yang` (version 0.3.1). We also support configuration and streaming of operational state data under the resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.

See [Mapping OpenConfig OSPF Commands to Junos Configuration](#) for OpenConfig configuration mappings. See [Telemetry Sensor Explorer](#) for state sensors.

Precision Time Protocol (PTP)

- **G.8275.1 telecom profile and PTP over Ethernet encapsulation support (MX10008 with MX10K-LC9600)**—Starting in Junos OS Release 22.4R1, MX10008 routers with MX10K-LC9600 line cards support Precision Time Protocol (PTP) over Ethernet encapsulation and the G.8275.1 telecom profile.

The G.8275.1 profile supports the architecture that is defined in ITU-T G.8275. The profile enables the distribution of phase and time with full timing support. You must ensure that all the devices in the network operate in combined or hybrid mode, which means that PTP and Synchronous Ethernet are enabled on all devices.

PTP over Ethernet effectively implements the packet-based technology. This feature helps operators deliver synchronization services on packet-based mobile backhaul (MBH) networks.

[See [PTP Profiles](#) and [Precision Time Protocol Overview](#).]

- **Support for Synchronous Ethernet (MX10004)**—Starting in Junos OS Release 22.4R1, we support Synchronous Ethernet over LAG) with Ethernet Synchronization Message Channel (ESMC).

[See [Synchronous Ethernet](#) and [Ethernet synchronization Message Channel \(ESMC\)](#).]

- **Support for PTP G.8275.1 over LAG (MX10008)**—Starting in Junos OS Release 22.4R1, we support the Precision Time Protocol (PTP) G.8275.1 profile over LAG interfaces.

[See [PTP Overview](#).]

- **Support for PTP G.8275.1 (MX10004)**—Starting in Junos OS Release 22.4R1, we support PTP G.8275.1 with or without LAG interfaces.

[See [PTP Overview](#).]

Routing Protocols

- **Distribution of segment routing–traffic engineering (SR-TE) policies and state in the TED by using BGP link-state (MX10004)**—Starting in Junos OS 22.4R1 Release, you can export the traffic engineering policies that originate from the segment routing protocol to the traffic engineering database (TED)and in the BGP link-state. After the export, you can use the BGP link-state to collect the information about traffic engineering policies. External controllers can then perform actions such as path computation, reoptimization, and network visualization within and across domains.

To import the Source Packet Routing in Networking (SPRING) policy information to the TED, we now have the `spring-te-policy` statement at the `[edit protocols source-packet-routing]` hierarchy level.

[See [Link-State Distribution Using BGP, source-routing-path](#), and [show ted spring-te-policy](#).]

- **MD5 authentication key rotation with overlap for key transition for OSPF (MX204, MX480, MX10003, PTX1000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support advertising OSPF MD5 authentication with multiple active keys to send packets with a maximum limit of two keys per interface. Having multiple keys active at any one time at the interface enables the smooth transition from one key to another for OSPF. You can delete old keys without any impact on the OSPF session.

[See [Understanding OSPFv2 Authentication authentication](#).]

[See [show \(ospf | ospf3\) interface](#).]

- **ICMP TTL expiry with source address (MX240, MX480, MX960, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 22.4R1, you can configure your device to use an IPv4 address as the source address for ICMP time-to-live (TTL) expiry error messages. This means you can configure the loopback address as the source address in response to ICMP error packets. Doing this is useful when you cannot use the device address for traceroute purposes because you have duplicate IPv4 addresses in your network.

Specify the address using the `ttl-expired-source-address source-address` option at the `[edit system icmp]` hierarchy level. The source address must be an IPv4 address. This configuration applies only to ICMP TTL expiry messages. Other ICMP error reply messages continue to use the address of the ingress interface as the source address.

[See [icmp \(System\)](#) and [ICMP Features](#).]

- **OSPF FAPM and interarea support (ACX5448, MX204, MX240, MX480, MX960, MX10003, MX10008, MX2008, MX2010, MX2020, PTX1000, and QFX10002)**—Starting with Junos OS Release 22.4R1, the Flexible Algorithm Prefix Metric (FAPM) is defined to allow an optimal end-to-end path for an inter-area prefix. The Area Border Router (ABR) *must* include the FAPM when advertising the prefix between areas that areas reachable in that given Flex-Algorithm. When a prefix is unreachable, the ABR *must not* include that prefix in the Flex-Algorithm when advertising between areas. The defined FAPM provides inter-area support.

[See [Understanding OSPF Flexible Algorithm for Segment Routing](#).]

[See [show ospf database](#), [show route table](#), [show ted database](#)

- **Support for S-BFD over EPE SIDs (MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2010, MX2020, PTX5000, PTX1000, PTX10002, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.4R1, seamless BFD (S-BFD) running between ingress devices and autonomous system boundary routers (ASBRs) can track BGP egress peer engineering (EPE) segment identifiers (SIDs). With this feature, you can prevent null-route filtering if a BGP EPE SID goes down.

[See [sbfd](#).]

- **CLI support for BFD echo and echo-lite modes (MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2010, MX2020, PTX5000, PTX1000, PTX10002, PTX10008, and PTX10016)**—

Starting in Junos OS Release 22.4R1, you can configure BFD echo mode and echo-lite mode through the Junos OS CLI. When BFD echo mode is active, a neighbor device transmits and loops back BFD echo packets to ensure that a forwarding path is available. BFD echo mode requires both the local device and neighbor device to support the full BFD protocol. However, BFD echo-lite mode can function even if the neighbor device doesn't support BFD.

You can use the following new CLI configuration commands to configure BFD echo mode and echo-lite mode:

- **echo mode:** `set routing-options static route address bfd-liveness-detection echo minimum-interval interval`
- **echo-lite mode:** `set routing-options static route address bfd-liveness-detection echo-lite minimum-interval interval`

[See [bfd-liveness-detection](#).]

- **Flex algo and FAPM leaking across IS-IS multi-instance (ACX5448, MX480, MX960, MX2010)**—Starting in Junos OS Release 22.4R1, we've added support to readvertise flexible algorithm (flex algo) prefix-segment identifiers (SIDs) and Flexible Algorithm Prefix Metrics (FAPMs) across interior gateway protocol (IGP) instances. We have also added support to readvertise other protocol prefixes and assign flex algo prefix-SIDs via policy to those prefixes.
- **Support for bootstrapping route-validation database from a local file (cRPD, JRR200, MX204, PTX10008, and QFX10008)**—Starting in Junos OS Release 22.4R1, we support the ability to read validation records from a local binary file and install into the specified named route-validation databases within RPD. This feature implements syntactic and semantic checks on the content of the file to ensure that it is a well-specified set of validation records. If the syntactic and semantic checks fail, the entire file is rejected as a source of validation records. Use the `source-file` statement at the `[edit routing-options validation]` hierarchy level to source route-validation records from a local file source. You can use the `show validation source-file` command to display the properties of a local validation source file.

[See [validation](#).]

- **Support for BGP RIB sharding and update threading features (MX304, MX10003, MX10004, MX10008, MX10016, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.4R1, you can use a new CLI option to override the sharding and update-threading configuration that might be present either through platform defaults or through explicit configuration. To override the configuration, use the `no-rib-sharding` and `no-update-threading` options at the `[edit system processes routing bgp]` hierarchy level.

[See [bgp](#).]

- **MVPN feature support with sharding (cRPD, JRR200, MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support the following features:

- Multicast virtual private network (MVPN) inactive route query from the main thread to shards
- Extranet and auto-export support with sharding
- Interact functions with RT-proxy client and server
- New data structure to store the inactive route data on the main thread
- Asynchronous route processing on the main thread

You can use `show mvpn c-multicast` to display the inactive route data stored on the main thread.

[See [rib-sharding](#) and [show mvpn c-multicast](#) .]

Services Applications

- **Inline active flow monitoring support for abstracted fabric (af) interfaces between guest network functions (MX240, MX480, MX960, MX10004, MX10008, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 22.4R1, we now support ingress and egress sampling of IPv4, IPv6, and MPLS traffic on abstracted fabric (af) interfaces between guest network functions (GNFs) in a node slicing scenario, for both the IPFIX and version 9 export formats.

[See [Understand Inline Active Flow Monitoring](#) and [Abstracted Fabric Interface](#).]

- **Full reassembly of IPv4 and IPv6 packets for MAP-E (MX Series Routers)**—Starting in Junos OS Release 22.4R1, the line cards on MX304, MX960, and MX10008 routers support full reassembly of IPv4 and IPv6 packets for Mapping of Address and Port with Encapsulation (MAP-E). We are introducing the following enhancements:
 - Maximum supported IP fragment size is increased to 15900 bytes.
 - Maximum IP packet size that can be fully reassembled is increased to 15900 bytes.

[See [Understanding Mapping of Address and Port with Encapsulation \(MAP-E\)](#).]

- **UDP source port configuration for FTI (MX Series Routers)**—Starting in Junos OS Release 22.4R1, you can define the UDP source port range of the packets for VXLAN encapsulation at the flexible tunnel interface (FTI) logical interface level using the command `source-port-range (min min-port-number | max max-port-number)` at the `[edit interfaces name unit name tunnel encapsulation]` hierarchy level. If you do not specify the UDP source port range, it is randomly set based on the hash value calculated using various packet headers field.

[See [vxlan-gpe \(FTI\)](#).]

Software Installation and Upgrade

- **Operation, Administration, and Maintenance (OAM) partition removal on the SSD (MX240, MX480, and MX960)**—Starting in Junos OS Release 22.4R1, you cannot use the Operation, Administration,

and Maintenance (OAM) partition on the solid-state drive (SSD) of the RE-S-X6-64G Routing Engine as a snapshot partition. The virtual machine (VM) host snapshot is a superset of the OAM snapshot as they backup or restore both host and guest components. You must use the VM host snapshot commands instead of OAM recovery commands on the VM host platforms.

[See [request vmhost snapshot](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **BGP Classful Transport (CT) support for IPv6 and Segment Routing–Traffic Engineering (SR-TE) color-only support (MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2010, MX2020, PTX10008, PTX10016, VMX)**—Starting in Junos OS Release 22.4R1, we support BGP-CT with IPv6 and BGP service-routes with a color-only mapping community. We have also enhanced the transport-class configuration statement by enabling shards to provide strict resolution without falling back on best-effort tunnels.

[See [use-transport-class](#).]

Subscriber Management and Services

- **Subscriber management functionality using MPC10E-10C-MRATE or MPC10E-15C-MRATE line cards (MX960, MX480, and MX240)**—Starting in Junos OS Release 22.4R1, we provide the following support:
 - Basic and advanced class of service (CoS) and filters (IPv4 or dual stack) support for:
 - DVLAN with DHCP subscribers
 - DVLAN with Point-to-Point Protocol (PPP) subscribers
 - DVLAN and agent circuit identifier (ACI) with DHCP subscribers
 - DVLAN and ACI with PPP subscribers
 - Stacked DVLAN with DHCP subscribers
 - Stacked DVLAN with PPP subscribers
 - Pseudowire DVLAN with DHCP subscribers
 - Pseudowire DVLAN with PPP subscribers
 - DVLAN with LAC (IPv4) basic and advanced CoS and filters
 - DVLAN with LNS (IPv4 and dual stack) basic CoS and filters
 - Advanced CoS and filters (IPv4 or dual stack) support for:
 - IFLSET with DHCP subscribers

- IFLSET with PPP subscribers
- L2TP tunnels (configured for each line card and each chassis)
- Subscriber services (customer solutions test scripts) processing
- Scaling and performance for the following features:
 - DHCP subscribers with authenticated dynamic VLAN
 - DHCP subscribers with authenticated dynamic S-VLAN
 - LNS subscribers
 - LAC subscribers
 - CoS service
 - Firewall service

[See [Subscriber Management VLAN Architecture Overview](#).]

- **View or log out the LNS L2TP subscriber sessions associated with a routing instance (MX Series)**—Starting in Junos OS Release 22.4R1, we've introduced the following two L2TP operational commands for MX Series devices that support BNG L2TP functionality. Use the following new operational commands to view or log out all the L2TP subscriber sessions simultaneously.
 - `show service l2tp session routing-instance name`
 - `clear service l2tp session routing-instance name`

[See [show service l2tp session routing-instance](#) and [clear service l2tp session routing-instance](#).]

- **Wireless CUPS: Support for IPv6 Frame Routing, IPv6 PD, and AGF use cases on User Plane Function (MX240, MX480, MX960, and MX10003)**—Starting in Junos OS Release 22.4R1, we support the following User Plane Function (UPF) features:
 - IPv4 address allocation via DHCPv4
 - IPv6 address allocation via DHCPv6
 - IPv6 prefix delegation via DHCPv6
 - IP6PL User Plane Functional (UPF) features
 - Framed routing for IPv6

[See [Multi-Access User Plane/Wireless CUPS Overview](#).]

VPNs

- **Multiple certificate types support on IKEv2 (MX240, MX480, and MX960 in USF mode, SRX1500, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0 running IKED process)**—Starting in Junos OS Release 22.4R1, you can establish the IKEv2 and IPsec SA tunnels irrespective of the type of certificate used on an initiator and a responder.

To support the multiple certificate types, configure the authentication method as certificates using the `certificates` option at the `[security ike proposal proposal-name authentication-method]` hierarchy.

[See [proposal \(Security IKE\)](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **MAC-VRF instances with EVPN-VXLAN** (MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX2010, and MX2020 with MPC10E line cards)

We support `vlan-based`, `vlan-aware`, and `vlan-bundle` service types for Ethernet VPN (EVPN) unicast and multicast traffic.

[See [MAC-VRF Routing Instance Type Overview](#), [mac-vrf](#), and [service-type](#).]

- **Passive flow monitoring** (MX304; MX2010 and MX2020 with the MPC10 and MPC11 line cards; and MX10004, MX10008, and MX10016 with LC9600 line cards)

[See [Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers](#) and [Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers](#) .]

- **Support for Access Gateway Function** (MX10004). Starting with Junos OS Release 22.4R1, we support the Access Gateway Function on the MX10004 router.

[See [Access Gateway Function User Guide](#) .]

- **Encryption with TPM 1.2** (MX240, MX480, MX960, MX2010, MX2020, and MX10003). We use the Trusted Platform Module (TPM) device to secure sensitive data including the master password on MX Series devices.

[See [Master Password Encryption](#).]

- **Remote Integrity Verification (RIV)** (MX240, MX480, MX960, MX2008, MX2010, and MX2020). We use RIV to determine the authenticity of software on a particular device.

[See [Remote Integrity Verification](#) .]

- **Support for SRv6** (MX240, MX480, and MX960). We support Segment Routing for IPv6 (SRv6) with the following features:

- Network programming and Layer 3 services over Segment Routing for IPv6 (SRv6) in BGP
- Operation, Administration, and Maintenance (OAM)
- Topology-independent loop-free alternate (TI-LFA link) and Node protection
- Static segment routing-traffic engineering (SR-TE) policy

[See [How to Enable SRv6 Network Programming in IS-IS Networks.](#)]

- **Support for rewrite rules** (MX304 and MX240, MX480, MX960 with MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards). You can configure IEEE 802.1 class-of-service (CoS) rewrite rules for host outbound traffic on subscriber interface for line cards based on advanced forwarding toolkit (AFT).

[See [rewrite-rules \(CoS Host Outbound Traffic\).](#)]

- **Support for unicast IP-IP tunneling for IPv4 and IPv6 traffic signaled by BGP** (MX304 and MX10004; and MX10008 with MX10K-LC9600 line cards). We support an IP over IP (IP-IP) encapsulation to facilitate IP overlay construction over an IP transport network.

[See [Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation.](#)]

- **Support for Layer 2 features** (MX304). Starting in Junos OS Release 22.4R1, we support the following Layer 2 features:

- Layer 2 Protocol Tunneling (L2PT)

[See [Layer 2 Protocol Tunneling.](#)]

- Ethernet VPN (EVPN) and virtual private wire service (VPWS)

[See [Overview of VPWS with EVPN Signaling Mechanisms](#) and [Overview of Headend Termination for EVPN VPWS for Business Services.](#)]

- EVPN-MPLS and EVPN-VXLAN overlay networks

[See [Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment](#) and [Multicast Support in EVPN-VXLAN Overlay Networks.](#)]

- EVPN routing policies

[See [Routing policies for EVPN.](#)]

- Multiple VLAN registration protocol (MVRP) and Ethernet ring protection switching (ERPS)

[See [Understanding Multiple VLAN Registration Protocol \(MVRP\) for Dynamic VLAN Registration](#) and [Ethernet Ring Protection Switching Overview.](#)]

- VPLS over transport class tunnels

[See [Introduction to VPLS](#) and [BGP Classful Transport Planes Overview](#).]

- Headend termination of pseudowire services in a VPLS-enabled virtual switch

[See [Layer 2 Services on Pseudowire Service Interface Overview](#).]

- **Support for inline (ITU-T Y.1731) performance monitoring** (MX Series routers with MPC10E and MPC11E line cards)

[See [ITU-T Y.1731 Ethernet Service OAM Overview](#)]

- **Feature support** (MX304, MX10004, and MX10008)—We support the following features:
 - Link delay measurement and advertising in OSPF.
 - SRv6 locator summarisation, locator anycast, and service mapping.
 - Layer 3 VPN service interworking between SRv6 and MPLS.
 - Remote LFA support for LDP in IS-IS.
 - SRv6 support for static SR-TE policy.
 - Compute traffic-engineered path delays using delay metrics.
 - Support for CBF fallback inside the service group.
 - AFT-based trio line cards on your network devices support OSPFv3 and IS-IS BFD sessions that use IPv6 link-local addresses.
 - Enhancements to Bidirectional Forwarding Detection (BFD)-triggered fast reroute (FRR) for unicast next hops and session-id-change-limiter-indirect.

[See [Understanding Link Delay Measurement and Advertising in OSPF](#) , [How to Enable SRv6 Network Programming in IS-IS Networks](#), [Understanding Remote LFA over LDP Tunnels in IS-IS Networks](#), and [Bidirectional Forwarding Detection \(BFD\) for MPLS](#).]

- **Support for platform resiliency** (MX304)

[See [show system errors active](#).]

- **Support for asynchronous notification** (MPC10, MPC11, MX304, MX10004, MX10008, MX2008, MX2010, MX2020, and MX10008 with MX10K-LC9600)

[See [Configuring Gigabit Ethernet Notification of Link Down Alarm](#).]

- **Supported transceivers, optical interfaces, and DAC cables**Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update this tool and provide the first supported release information when the optic becomes available.

What's Changed

IN THIS SECTION

- [EVPN | 83](#)
- [MPLS | 83](#)
- [Network Management and Monitoring | 84](#)
- [Platform and Infrastructure | 84](#)
- [User Interface and Configuration | 84](#)

Learn about what changed in this release for MX Series routers.

EVPN

- **Flow-label configuration status for EVPN ELAN services**--The output for the `show evpn instance` extensive command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

MPLS

- **CSPF LSP resignaling uses new instance ID (MX480)**--A Constrained Shortest Path First (CSPF) LSP uses a new instance ID when attempting to resignal an LSP that is down. In earlier releases, the CSPF LSPs that went down were stuck in CSPF path computation stage. You had to manually clear the affected LSPs and recompute the paths for the LSPs to be up again.

[See [LSP Computation.](#)]

Network Management and Monitoring

- **Junos YANG modules for RPCs include the `junos:command extension statement` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command extension statement` in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper <https://github.com/Juniper/yang> yang GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the YANG schemas on the local device, the YANG modules for RPCs include the `junos:command extension statement`.

Platform and Infrastructure

- **Enhanced bandwidth and burst policer value (MX Series and EX9200 Series)**—We've updated the default bandwidth value from 20000 to 100 pps and burst policer value from 20000 to 100 packets. This enhancement avoids the CPU usage of `eventd` and `snmpd` reaching more than 100 percent. Earlier to this release, when the system receives a violated traffic for SNMP along with other protocols traffic, the CPU usage of `eventd` and `snmpd` was reaching more than 100 percent with an error.

[See [show ddos-protection protocols parameters](#).]

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml` filter or the `| display json` filter

to display command output in XML or JSON format in any release. We've removed these options from the CLI.

Known Limitations

IN THIS SECTION

- [General Routing | 85](#)
- [MPLS | 86](#)
- [Network Management and Monitoring | 86](#)
- [Platform and Infrastructure | 86](#)
- [Routing Protocols | 87](#)
- [Services Applications | 87](#)

Learn about known limitations in this release for MX Series routers.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- In a scaled setup with LDP over RSVP configuration and maximum-ecmp as 32 or 64, line card CPU usage can remain high for extended duration on link flap operation. In this duration, LACP might take 5 minutes or more to converge and the aggregated Ethernet interface bundle to be active. [PR1624219](#)
- The SDN-Telemetry process might crash during long running streaming telemetry collectors. Telemetry data loss occurs streamed from the line cards till the process comes up. [PR1647568](#)
- If proper gap is given between channelisation and dechannelisation the issue is not seen. Proper gap means allowing the system to complete the previous config before we load the new configuration. Recommendation is to if we give channelisation config commit wait for the links to come up or atleast the ifd's get created on both evo and RE side and then only revert the configuration to dechannlisation and vice versa. [PR1665625](#)

- VM host snapshot recovery is not enabled for RE-S-X6-128G-K. [PR1674091](#)
- Even though GRES is enabled, the `show system filesystem encryption status` command display information about the specific Routing Engine. [PR1674373](#)
- For IPv6 traffic that is ingressing into an Abstract Fabric (AF) interface via MPC11e card, and also sampled, the `OutputIntf` in the flow records might not be captured if you do not enable the `nexthop-learning` command. [PR1680873](#)
- MVRP on PVLAN promiscuous port is not supported. If you configure MVRP on promiscuous port, then hosts connected to secondary VLAN ports will not be able to reach external world through promiscuous port carrying primary VLAN tags. [PR1693345](#)

MPLS

- With local reversion on, there is a possibility of transit router not informing headend of RSVP disabled link when the link flaps more than once. As a workaround, remove the local-reversion configuration. [PR1576979](#)

Network Management and Monitoring

- Junos might translate the custom yang configuration even after disabling the custom Yang package. [PR1599107](#)

Platform and Infrastructure

- On MX devices, under Ethernet VPN (EVPN) environment, packets routed using IRB interface could not be fragmented due to media maximum transmission unit (MTU) problem. [PR1522896](#)
- When the `deactivate services rpm` and `deactivate routing-options rpm-tracking` commands are applied together and then committed, some of the rpm tracked added routes are not deleted from the routing table. Issue cannot be seen using the following steps:
 1. Deactivate routing-options rpm-tracking .
 2. Commit the configuration then all the rpm tracked routes will be deleted. If the RPM service needs to be deactivated.
 3. Deactivate services rpm.

4. Commit.

[PR1597190](#)

- On MX platforms, VPLS flood traffic loss is observed if flood composite next-hops are out-of-sync on ingress and egress FPCs during transport path reversion. [PR1656216](#)

Routing Protocols

- When you do not configure `routing-options transport-class fallback none`, you cannot configure more than 10 transport-classes or advertise more than 10 distinct colors in SRTE or FlexAlgo. [PR1648490](#)

Services Applications

- In Junos OS release 17.4 and later, subscriber sessions on the LNS that send an ICRQ that includes RFC5515 AVPs might fail to establish a session. The client might receive a `receive-icrq-avp-missing-random-vector` CDN error in response. [PR1493289](#)

Open Issues

IN THIS SECTION

- [EVPN | 88](#)
- [Forwarding and Sampling | 88](#)
- [General Routing | 88](#)
- [High Availability \(HA\) and Resiliency | 93](#)
- [Interfaces and Chassis | 93](#)
- [Layer 2 Ethernet Services | 94](#)
- [Layer 2 Features | 94](#)
- [MPLS | 94](#)
- [Network Management and Monitoring | 94](#)
- [Platform and Infrastructure | 95](#)

- Routing Protocols | 95
- Services Applications | 96

Learn about open issues in this release for MX Series routers.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

EVPN

- In PBB-EVPN (Provider Backbone Bridging - Ethernet VPN) environment, ARP suppression feature which is not supported by PBB might be enabled unexpectedly. This could cause MAC addresses of remote CEs not to be learned and hence traffic loss. [PR1529940](#)
- On all platforms, MAC-IP route deletion and addition are triggered when re-ARP (Address Resolution Protocol) on MH (Multihoming) device fails in the EVPN-MPLS multihoming scenario resulting in traffic drop. [PR1691132](#)

Forwarding and Sampling

- On all Junos dual-RE platforms, when performing activate or deactivate Graceful Routing Engine Switchover (GRES) multiple times synchronization issues are observed between the primary and backup dfwd process. [PR1697959](#)

General Routing

- When there is an input failure on one of the AC PEMs (low or high) its wrongly categorized as Mix of AC PEMs so instead of PEM input failure you will see Mix of AC PEMs alarm raised. [PR1315577](#)
- On WRL8 based VMHost platforms there is no log rotation for resild log and temperature sensor info is incorrectly written into resild log which could result in continuous logs in resild log file. The disk usage might keep increasing due to this issue. The disk usage could be eventually full which could cause system to hang and reboot. [PR1480217](#)

- When there are HW link errors occurred on all 32 links on an FPC 11. Because of these link errors, all FPCs reported destination errors towards FPC 11 and FPC 11 was taken offline with reason offlined due to unreachable destinations. [PR1483529](#)
- After backup Routing Engine halt, CB1 goes offline and comes back online; this leads to the backup Routing Engine booting up, and it shows the reboot reason as "0x1:power cycle/failure." This issue is only for the RE reboot reason, and there is no other functional impact of this. [PR1497592](#)
- In the platform using INH (indirect next hop, such as Unilist) as route next hop type for multiple paths scenario (such as BGP PIC or ECMP), the session fast-reroute might be enabled in Packet Forwarding Engines (PFEs). When the version-id of session-id of INH is above 256, the PFE might not respond to session update, which might cause the session-id permanently to be stuck with the weight of 65535 in PFE. It might lead PFE to have a different view of Unilist against load-balance selectors. Then either the BGP PIC or the ECMP-FRR might not work properly and traffic might be dropped or silently discarded. [PR1501817](#)
- Software defect that causes a 10GE interface to flap continuously when configuring with the WAN-PHY framing with the default "hold-down" timer (0). Once upgrading a router to an affected software release, the interface may flap continuously. This is not applicable to an interface with the default framing - LAN-PHY. [PR1508794](#)
- Due to BRCM KBP issue route lookup might fail. Need to upgrade KBP to address this issue, Due to high risk KBP SDK upgrade planned for Junos OS release21.1. [PR1533513](#)
- With IPsec PMI/fat-core enabled, the `show services sessions utilization` command does not display the right CPU utilization. [PR1557751](#)
- The Sync-E to PTP transient simulated by Calnex Paragon Test equipment is not real network scenario. In real network deployment model typically there will be two Sync-E sources (Primary and Secondary) and switchover happens from one source to another source. MPCE7 would pass real network SyncE switchover and associated transient mask. [PR1557999](#)
- VE and CE mesh groups are default mesh groups created for a given Routing instance. On vlan/bridge-domain add, flood tokens and routes are created for both VE and CE mesh-group/flood-group. Ideally, VE mesh-group doesn't require on a CE router where IGMP is enabled on CE interfaces. Trinity based CE boxes have unlimited capacity of tokens, so this would not be a major issue. [PR1560588](#)
- When deactivate or activate of security configuration is executed continuously, there are instances in which when gkmd process can core while the process exits. [PR1566044](#)
- When the active backup interface is deactivated, the PTP lock status is set to 'INITIALIZING' state in `show ptp lock-status` output for few seconds before BMCA chooses the next best backup interface. This is the day-1 behavior and there is no functional impact. [PR1585529](#)

- Pim VXlan not working on TD3 chipsets enabling VXLAN flexflow after Junos OS release 21.3R1. Customers Pim VXlan or data plane VXLAN can use the Junos OS release 21.3R1. [PR1597276](#)
- During Routing Engine switchover, if there is a burst of ICMP/BFD/SSH/FTP/TELNET/RSVP packets (around 18,000 pps) you might see new backup Routing Engine restarting. [PR1604299](#)
- On MX-VC (Virtual Chassis) platforms with MS-MPC or SPC3 service cards and AMS(Aggregated Multi-Service), traffic on the line card in the backup chassis may not be load-balanced properly due to timing conditions. This works well on the line card in the master chassis. There might be traffic loss when interfaces are not properly balanced. [PR1605284](#)
- When user tries to disable AMS ifd using config knob, the ipsec tunnels are not deleted. Deactivating the services will provide the desired result. [PR1613432](#)
- In some NAPT44 and NAT64 scenarios, duplicate SESSION_CLOSE syslog gets generated. [PR1614358](#)
- On all Junos OS platforms, the MAC address of the 17th aggregate Ethernet interface might be changed after the upgrade from 18.4+ to 20.4+ releases. It will lead to MAC based service interruption. [PR1629050](#)
- For a topology with VSTP and VRRP configured and IPV6 traffic, if VSTP bridge priority is changed a couple of times (to trigger toggling of root bridge), it is possible that IPv6 traffic drops on some of the streams. [PR1629345](#)
- The fabric statistics counters are not displayed in the output of the show snmp mib walk ascii jnxFabricMib command. [PR1634372](#)
- On all devices running Junos OS or Junos OS Evolved, where this is a high BGP scale with flapping route and the BGP Monitoring Protocol (BMP) collector/station is very slow, the rpd process might crash due to memory pressure. [PR1635143](#)
- Source MAC should not be configured on the underlying static interface on the UP for PPPoE login to work correctly. [PR1641495](#)
- On Junos platform, PTP does not lock when port speed is not configured under PIC hierarchy or port speed for some additional random ports are configured under the PIC hierarchy or perform PIC deactivate or activate. [PR1645562](#)
- When per-interface egress and per-sid egress SR sensor stats are configured using the CLI commands below, the (pushed) MPLS label length does not get included in the output/Tx octets field that gets exported from the sensor. set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress set protocols isis source-packet-routing sensor-based-stats per-sid egress This is a day-1 behavior on all Trio ASIC based FPCs on the MX platform. [PR1646799](#)
- Pressing N during the PXE install and the reboot prompt is supposed to abort the installation, not reboot, and provide a debug shell. On MX304, the shell is not spawned and the system hangs for a while and then reboots. [PR1647534](#)

- Core dump reported intermittently where random grpc stack crash is observed. The license service will auto restart and recover. [PR1656975](#)
- On Junos platforms, in the VPLS environment when having "routing-options resolution preserve-nexthop-hierarchy" configured results in the packet dropped at egress PE devices for multiple MPLS stack labels. [PR1658406](#)
- For MX204 and MX10003 devices, if a non-default SSH port is configured for system login, after upgrade to 21.4 release, the FPC is stuck in offline. To avoid such issue please use default SSH port and use protect Routing Engine filter to only allow the access from the trusted source. [PR1660446](#)
- The /telemetry-system/subscriptions/dynamic-subscriptions/ support on GNF will be from 22.4. [PR1661106](#)
- Not all MAC addresses are learnt for some VPLS instances after "clear vpls mac-table" command is executed. [PR1664694](#)
- With following configuration changes subscribers are coming up. Config changes: =====
set forwarding-options dhcp-relay overrides allow-snooped-clients set forwarding-options dhcp-relay group DHCP-FO overrides allow-snooped-clients set forwarding-options dhcp-relay group DHCP-FO overrides user-defined-option-82 100.112.77.66 deactivate forwarding-options dhcp-relay group DHCP-FO interface ae31.0 overrides. [PR1665499](#)
- You must not modify the locator attributes, instead locator, SIDs should be deleted and configured back. Otherwise it will lead to coredump. [PR1667320](#)
- On MX devices with MIC-MACSEC-20GE, FEB (Forwarding Engine Board) might go down while activating or deactivating GRES configuration. [PR1668983](#)
- Sometimes cores are reported on backup Routing Engine during init after a reboot etc. When the backup Routing Engine initialization is being done and system is busy, some commands executed in context of spmbpfe are taking more time to complete due to the initial heavy lifting by the kernel, In this stage, if in case the commands from spmbpfe process do not complete for less than 2.5 seconds, then there are chances of spmbpfe cores. This is a temporary issue seen on backup Routing Engine during init time only. This may not be impacting because if in case spmbpfe process crashes due to this, it would restart by itself and continue to init and run once the initial high CPU condition has passed. It should not cause any functionality or performance impact especially since it is reported only on backup Routing Engine. [PR1675268](#)
- There will be drop of syslog packets seen for RT_FLOW: RT_FLOW_SESSION_CREATE_USF logs until this is fixed. [PR1678453](#)
- The Queue stats may show constant PPS / bps after interface is disabled. The stats don't increment and remain same when the interface went down. [PR1685344](#)

- When you change the hostname configuration, the change is not reflected in the RIFT output. Also when changes are made to the REDIS configuration, they are not applied until rift is restarted through "restart rift-proxyd". [PR1686233](#)
- If MVRP is enabled on an MSTP enabled interface, the interface will be made part of all the existing instances on the switch, So, if there are two interfaces between R1 and R2 as below: R1(et-0/0/1 & et-0/0/2)=====(et-0/0/1 & et-0/0/2)R2 And one interface is MVRP enabled (say et-0/0/1), and et-0/0/2 is not MVRP enabled. By configuration et-0/0/1 is part of MSTI-1 and et-0/0/2 is part of MSTI-2. MSTI-1 is running on vlan-100 and MSTI-2 is running on Vlan-200. R2 in this case, is advertising only vlan-100. The MVRP enabled interface will become part of all the MSTIs(MSTI-1 and MSTI-2 both) configured on the device and it will take part in the FSM of all the MSTIs. Although et-0/0/1 is not member interface of vlan-200(corresponding to MSTI-2). This potentially can cause a problem where et-0/0/1 although not a vlan-200 member, will go into FWD state and et-0/0/2, genuine member of vlan-200 goes into BLK state for MSTI-2. So, when traffic is received in vlan-200 it will be sent out of et-0/0/1, and it will be dropped.[PR1686596](#)
- With sharding enabled, when BGP route is resolved over RSVP LSP, LSP name is not displayed in the output of the show route extensive command for inactive route. [PR1687890](#)
- When an interface's configuration is changed between channelized and non-channelized modes and the PIC requires a bounce, rapid toggles between the two may result in the interface(s) not coming up. It is recommended to give the PIC time to complete its bounce and initialize the interfaces - at least 15 seconds - before issuing another configuration change for that interface. [PR1688767](#)
- JNP10K-LC9600: G.8275.1: SyncE to PTP and SyncE to 1PPS Transient Response not meeting G.8273.2 mask. [PR1692202](#)
- JNP10K-LC9600: G.8275.1: PTP and PTP to PTP Noise transfer performance not meeting G.8273.2 mask. [PR1692272](#)
- With Sharding enabled, BGP advertised metric are not displayed for active route in ?show route prefix extensive output. This information can be seen using the show route prefix extensive rib-sharding shard-name command. [PR1692755](#)
- With Sharding enabled, BGP flags like the following are not displayed on Active route in the output of the show route extensive command: Accepted Multipath MultipathContrib MultiNextHop" Per shard view, using "show route extensive prefix rib-sharding shard-name" . [PR1693207](#)
- On all Junos platforms supporting MPC10/11, due to a mismatch of API signature, the IPsec packets are not handed over to the Services Processing Card (eg. SPC3) for decryption and get dropped.[PR1694942](#)
- When MX10004 and MX10008 chassis has a malfunctioning Line card plugged in, Upstream SyncE source interface will get stuck in abort state. [PR1695156](#)

- We have traffic drop seen for some streams in intra-as srte color only ipv6 tunneling shard test on VMX10008/VMX304 with IPv6. This is working on VMX (MX960) and physical PTX10008(Vale) This is being investigated and will be fixed in the next release (R2). [PR1695669](#)
- MX10008 and MX10004: G.8275.1: PTP over AE: PTP clock goes into FREERUN instead of HOLDOVER upon disabling phase locked slave aggregate Ethernet interface. [PR1696028](#)
- With JNP10K-LC2101 is backup/primary in a multi line card scenario, spikes can be seen in 2way time error during LAG switchover with primary and backup across the line cards.[PR1696527](#)
- Spikes seen in 2way time error with JNP10K-LC2101 is either ptp backup or primary and any switchover is done. [PR1697167](#)
- JNP10K-LC9600: G.8275.1: Noise generation performance fails post GRES. [PR1697602](#)
- On MX platforms, traffic egressing on the IRB (Integrated routing and bridging) interface with the underlying L2 (layer2) access port has VLAN tags imposed incorrectly.[PR1700321](#)
- When distributed multicast service is activated on several hundred subscribers, bbe-smg-upd process may crash. [PR1700571](#)
- JNP10K-LC9600: G.8275.1: Multiple GRES operation resulting in huge time error. [PR1701017](#)

High Availability (HA) and Resiliency

- When you perform GRES with the interface em0 (or fxp0) disabled on the primary Routing Engine, then enable the interface on the new backup Routing Engine, the network is unable to access. [PR1372087](#)

Interfaces and Chassis

- When the MX virtual-chassis was upgraded by using the Sequential Upgrade method, there is the possibility that pfe provisioning might start before link training completes and all PICs are online. In such scenario, the IFD provisioning is preserved and if the preserved state is applied to the Packet Forwarding Engine before fabric training has completed and all of the pics have been powered on, ifd missing errors will be seen. [PR1670345](#)
- MediaType value in SNMP/Jvision is not correct at the beginning after the switch comes up only for the DOWN interfaces where copper mediaType is connected till the link is not UP. This value is correct always in CLI output. You can bring the link up by connecting to the other side and restarting the dcd daemon to fix this issue. [PR1671706](#)

Layer 2 Ethernet Services

- If a client sends a DHCP request packet, and option 55 includes PAD option (0), a DHCP ACK will not be sent back to the client. [PR1201413](#)
- IPv4 ALQ not working with authentication and the following error message gets generated on the backup router: Message failed sanity test - the access-profile info is invalid. length:0. [PR1688272](#)

Layer 2 Features

- In case of the access-side interfaces used as SP-style interfaces, when a new logical interface is added and if there is already a logical interface on the physical interface, there is 20 to 50 milliseconds traffic drop on the existing logical interface. [PR1367488](#)

MPLS

- Ingress retries after lsp stay down for extended period of time or customer can clear lsp to speed up the retry. [PR1631774](#)
- On all Junos OS platforms, if CCC (Circuit Cross-Connect) is configured to use a label-switched-path that is IGP routed, i.e., no-cspf and no ERO (explicit route object) configuration, then restarting egress CCC node or restarting FPC on the egress CCC node containing receive-switch configuration multiple times may cause CCC to remain stuck in Remote-if-down state. Traffic loss will be there. [PR1694777](#)

Network Management and Monitoring

- When you configure maximum-password-length and try to configure password whose length exceeds configured maximum-password-length, an error message gets generated and 'ok' tag also gets emitted. The configuration does not get committed. [PR1585855](#)

Platform and Infrastructure

- With given multi dimensional scale, if configuration is removed and restored continuously for more than 24 times, MX Trio based FPC may crash and restart. During the reboot, there can be traffic impact if backup paths are not configured. [PR1636758](#)
- On all Junos OS platforms, a random IBGP (Interior Border Gateway Protocol) session flaps is observed immediately after committing unrelated configuration changes with the error **BGP_IO_ERROR_CLOSE_SESSION** and a connection reset. [PR1685113](#)
- PVSTP protocol packets is getting duplicated when it tunnelled through Layer2 tunnelling protocol. Other protocol data units PDUs(STP,VTP,CDP) are not impacted. [PR1686331](#)
- In EVPN-VxLAN, traffic drop can be seen for some local CEs which are multihomed to at least one MX devices. [PR1696106](#)

Routing Protocols

- When l2cpd (in the context of xSTP) clears the entries that it has programmed on pppmd, ie when you delete xSTP configs from the box, there can be a possibility of pppmd core. If pppmd is in distributed mode then there will be no service impact, else there can be service impact as packet transmission for various protocols will happen via if pppmd is in centralized mode. [PR1660299](#)
- BGP LU statistics does not report correct statistics when sharding is enabled. This is not specific to BGP CT feature of this RLI. [PR1684238](#)
- Junos OS release 22.3 and later, IS-IS yang gets uplifted to 1.0.0 version which has major change in existing OC path that was supported earlier. Since OC path has change, same need to reflected in translation script which is not done. As part of D27 release for cloud, translation script will be modified with newer OC path. Till then supported older OC config is broken. eventually D27 code will come back to DCB and things will work fine after that. [PR1686751](#)
- When Lsys is configured with family route-target, there is a certain corner case scenario where Lsys shutdown does not complete on a deactivate logical-system. This will be fix in the next maintenance release. Manually, restart routing logical-system <name> can be used to force shutdown of Lsys. [PR1695050](#)
- SR-TE secondary LSP should be only standby in forwarding table, however, it is also active and forwarding traffic due to the wrong metric calculation. [PR1696598](#)

Services Applications

- When a configured tunnel interface is changed to another one, flow-tap-lite functionality stops working that is packets don't get mirrored to content destination. But, this problem is not consistently seen.[PR1660588](#)

Resolved Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 97](#)
- [EVPN | 97](#)
- [Flow-based and Packet-based Processing | 98](#)
- [Forwarding and Sampling | 98](#)
- [General Routing | 98](#)
- [Infrastructure | 99](#)
- [Interfaces and Chassis | 99](#)
- [Juniper Extension Toolkit \(JET\) | 99](#)
- [Layer 2 Ethernet Services | 99](#)
- [MPLS | 100](#)
- [Network Management and Monitoring | 101](#)
- [Platform and Infrastructure | 101](#)
- [Routing Policy and Firewall Filters | 111](#)
- [Routing Protocols | 111](#)
- [Services Applications | 113](#)
- [Subscriber Access Management | 113](#)
- [User Interface and Configuration | 114](#)
- [VPN | 114](#)

Learn about the issues fixed in this release for MX Series routers.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- Packet drop might be seen on MX platforms with MPC10/MPC11/LC9600 line cards. [PR1671040](#)
- QoS may not work as expected on AE interfaces with explicit-null label. [PR1675781](#)
- The oid tree `jnxCosQstatEntry` returns nothing for some interfaces after restarting class-of-service. [PR1693977](#)
- The congestion details will be lost as ECN bits in DSCP are cleared after VXLAN decapsulation. [PR1683438](#)

EVPN

- EVPN next hops are not correctly synchronised between RE and PFE. [PR1633344](#)
- The `rpd` process crash might be seen due to memory allocation failure. [PR1636690](#)
- The kernel crash would be observed in an EVPN multi-homed scenario. [PR1649234](#)
- EVPN Proxy ARP doesn't work for the restricted feature on the IRB interface. [PR1650665](#)
- EVPN DF election will get stuck in the wrong state. [PR1662954](#)
- CRB EVPN MPLS is not working with control-word enabled. [PR1665130](#)
- Layer3 inter-subnet routing will fail if there is no reachability for the remote IP-host route. [PR1669585](#)
- Traffic drop might be observed in the EVPN-VPWS scenario. [PR1672749](#)
- The `rpd` crash would be observed when activating or deactivating the EVPN routing-instances. [PR1673157](#)
- RPD core upon receipt of a specific EVPN route by a BGP route reflector in an EVPN environment. [PR1675054](#)
- The ARP/ND entries are not learnt as expected on the spine with EVPN-VXLAN. [PR1677521](#)
- EVPN Proxy ARP does not work for the static VTEP interface. [PR1679115](#)

- EVPN MPLS traffic drop can be observed in a multi-vendor PE CE device setup with single-active LAG. [PR1680421](#)
- The `show ethernet switching table` command does not synchronize between two MH PE devices after GRES. [PR1686546](#)
- RPD (Routing Protocol Daemon) core is observed due to remote BGP routes being flashed as active routes. [PR1692249](#)

Flow-based and Packet-based Processing

- The hardware acceleration flag was not properly updated on `RT_FLOW_SESSION_CLOSE` logs. Additionally, the values for "Services-offload-sessions" for customers using SPC2's in their SRX5000-Series devices was incorrect. [PR1629216](#)
- In SD-WAN the association between VRF instance and VRF group fails for ISSU from 19.2~21.1 to 22.2R1. [PR1661935](#)

Forwarding and Sampling

- Traffic drop seen and filter not hitting as expected for match condition traffic class with FLT option configured. [PR1573350](#)
- The MPC/FPC crash is seen on specific LC's running BGP Flowspec. [PR1662955](#)
- MX platforms with specific line cards are affected due to increase in HEAP memory. [PR1668521](#)
- Traffic loss may be observed when changing firewall configuration. [PR1670622](#)
- The FPC crashes when the `show filter memory` command is used during a firewall filter configuration change. [PR1680849](#)

General Routing

- MAC address change is not taking effect in static route with qualified-next-hop. [PR1663484](#)
- Traffic loss might be seen when multicast route changes. [PR1669498](#)
- The `rpd-agent` process might restart post mastership switchover. [PR1669767](#)

Infrastructure

- Junos upgrade might fail due to file system corruption. [PR1668090](#)

Interfaces and Chassis

- VRRP flaps between MC-LAG peers. [PR1579016](#)
- 22.2TOT :SecPDT:Unified L4/L7 Use Case Sky ATP: reth1 interface down and DCD cores observed on node1 during test on 22.2TOT image. [PR1657021](#)
- The VRRP track might go down upon GRES. [PR1668280](#)
- When donor interface is admin down, borrower can't be pinged from local/remote. [PR1670362](#)
- VRRP master-master condition might occur when there are more than two devices in the VRRP group. [PR1680178](#)
- Unable to configure ae interfaces more than 256. [PR1681114](#)
- If vrrp authentication key is more than 16 characters instead of commit error it is ignoring remaining characters. [PR1683871](#)
- Traffic is getting impacted as interface hold-time is not working with wan-phy framing. [PR1684142](#)

Juniper Extension Toolkit (JET)

- Client's disconnect request using Stream.CloseSend() closing the connections. [PR1667855](#)
- Modify RPC Connection TC failure with disconnect(). [PR1677182](#)

Layer 2 Ethernet Services

- MX240:Verify VRRP stats is failed after Deactivate the Access interface. [PR1666943](#)
- The jdhcpd crash might be observed in a DHCP Relay Agent scenario. [PR1668015](#)
- The DHCP unicast acknowledge packet might be dropped. [PR1676573](#)

- DHCP packets sent to the client have the Option 82 Suboption length set to 0. [PR1684521](#)

MPLS

- RSVP refcount leak and the rpd crash observed post LSP churn. [PR1621771](#)
- The error severity of syslog message "ted_client reset" generated during commit is incorrect. [PR1649565](#)
- The rpd core is seen due to IGP database and BGP LS database out of synchronization. [PR1655031](#)
- Memory utilization keeps incrementing due to the path error message. [PR1657872](#)
- The rpd crash would be observed in a RSVP scenario. [PR1661526](#)
- A LSP might get stuck in the CSPF path computation stage. [PR1661954](#)
- Transit LSR might stop sending RESV msg if there is no RRO in the LSP's PATH message. [PR1667708](#)
- Traffic loss will be seen in an LDP->BGP-LU stitching scenario. [PR1670334](#)
- Premature RSVP Path Error BW-Unavailable originated by PLR. [PR1670638](#)
- VCCV BFD session will be down as the periodic ping will not work as expected in a seamless MPLS scenario. [PR1670711](#)
- LDP Traffic will be blackholed when the L-ISIS/L-OSPF route changes due to interface level configuration. [PR1671187](#)
- The rpd crash might be observed with Container LSPs. [PR1672804](#)
- CPU utilization of rpd process may reach 100% while reporting LSP states to pccd if the IS-IS update churn is high. [PR1673348](#)
- The rpd crashes very rarely when constructing LDP trace message irrespective of enable/disable LDP traceoptions. [PR1676503](#)
- LDP egress-policy for default route (0.0.0.0/0) with 'exact' option will make output label for the unrelated routes. [PR1676551](#)
- The traffic might drop when the Link State protocol with the least preference is set to active and fails the CSPF algorithm. [PR1677930](#)
- In an LDP -> BGP LU stitching scenario, Multiple LSPs will not be installed in the forwarding table, even if BGP Multipath and ECMP are enabled. [PR1680574](#)

- In the RSVP-TE scenario, with Entropy label capability is enabled during MBB issues handling Resv Messages. [PR1681403](#)
- The RE crashes when MPLS next-hop is created and deleted frequently. [PR1681892](#)
- RSVP PathTear is not encapsulated by MPLS header when Bypass is used. [PR1685182](#)
- On a controller based MPLS setup with container LSPs, rpd daemon crashes after LSP deletion occurs. [PR1690458](#)
- [MX]L2VPN ping is failing when UHP rsvp LSP is used. [PR1697982](#)
- Memory leak issue in TED. [PR1701800](#)

Network Management and Monitoring

- Observed memory leak in eventd leak during GRES. [PR1602536](#)
- The rpd crash will be observed post ephemeral database configuration commit synchronization. [PR1610713](#)
- The "snmpd" process might crash if SNMP timeout happens. [PR1666548](#)
- The snmpd core might be observed with filter-duplicates configuration. [PR1669510](#)
- While loading MIB file, saw error : "DESCRIPTION" is missing for "mib-jnx-chas-defines.txt". [PR1670858](#)
- AE (aggregated ethernet)interface beyond 1099 are allotted 0 snmp index. [PR1683264](#)

Platform and Infrastructure

- The core interface goes down. [PR1631217](#)
- The rmopd process crashes while deactivating all the TWAMP client control connections and executing show snmp mib walk ascii jnxTwampClientTestSessionsTable. [PR1650997](#)
- Multicast packet drop causes pixelization. [PR1655363](#)
- The CPU utilization might increase when a user login and logout to the device continuously. [PR1662172](#)
- The line cards MPC10/MPC11/LC-9600 might crash. [PR1667716](#)

- Traffic drop is observed with layer 2 circuit local switching with PS interface. [PR1669410](#)
- Traffic drop observed with SP style configuration for the logical tunnel in layer2 domain. [PR1669478](#)
- Layer 2 packets other than IPv4/IPv6 (e.g. CFM) will get forwarded as out of order via MPC10 and MPC11 in the egress direction. [PR1670316](#)
- DHCP bindings will fail for the client connected on an LT interface when DHCP snooping is enabled. [PR1677631](#)
- The line card gets crashed during node/interface statistics reporting with resource monitoring. [PR1681533](#)
- Traffic drop is seen after configuring fast-lookup-filter. [PR1682164](#)
- Probes received counter is not correct when set "moving-average-size" > "history-size" under TWAMP client configuration. [PR1685952](#)
- FPC might core if CFM flap trap monitor feature in use. [PR1536417](#)
- Error message seen in clksyncd logs with SyncE/PTP configs "ESYNC-Error:ferrari_zl30362_reg_write: Error, EEC(0) not yet initialized". [PR1583496](#)
- LTS19: MX2008: junos vmcore Dump failed. Partition too small. [PR1604755](#)
- DSLite might not work on MX platforms installed with MPC7E line card and SPC3 service PIC. [PR1632278](#)
- The `show chassis firmware` command does not show the revision for PIC FPGA. [PR1633187](#)
- The `ppman` process might crash and MPC cards will be stuck in the ISSU state when "Unified ISSU with Enhanced Mode" is performed. [PR1633286](#)
- `Pyrite_VC: em0` interface `pped` is reflecting as 10G instead of 1G. [PR1636668](#)
- Same VLAN cannot be used as data VLAN and VOIP VLAN together. [PR1637195](#)
- Error message will be seen during FPC boot up on MX10003. [PR1637756](#)
- The Layer 3 packets with the destination as IPv6 Link Local address will not work. [PR1638642](#)
- USB device not visible in Junos OS. [PR1639071](#)
- Enhanced subscriber management might not work when Junos is installed using USB image. [PR1641712](#)
- The `rpd` crash might be seen in backup RE. [PR1645457](#)
- The `qemu-img: command not found` error msg seen during PXE image install. [PR1648328](#)

- MX960 :: bbe-statsd core observed at vlogging,smid_reregister, sdb_db_check,Juniper: :SmidInterface:: isReady in bbe-smgd daemon restart test. [PR1648565](#)
- The MPC/FPC might crash or the traffic may be silently dropped/discarded. [PR1649499](#)
- BGP Sensor "/bgp-rib/afi-safis/afi-safi/ipv4-unicast/loc-rib/" not available as a 'periodic' sensor. [PR1649529](#)
- Traffic Loss will be observed with Virtual-Router. [PR1650335](#)
- VMcore gets triggered when control packets go over IRB and GRE. [PR1651273](#)
- PCS Errored blocks count will increment after Junos software upgrade to 20.2R1 or later. [PR1651526](#)
- After enabling inline-jflow(sampling), PPE traps/logs or traffic drops. [PR1652901](#)
- The bbe-smgd core is observed on MX platforms. [PR1653546](#)
- EVPN-VXLAN: With evpn type 5 and NSR configuration, an RE switchover may result in momentary traffic drop for about 2 to 3 seconds. [PR1655052](#)
- The 1G port always stays down while changing of 10G interface lane speed to 1G. [PR1655089](#)
- The protocol state sync on backup RE is affected. [PR1655249](#)
- Telemetry is reporting In-Error and CLI/SNMP does not report any counter in the In-Errors. [PR1655651](#)
- The lane 0 information might not be displayed when the SNMP poll is done using ifindex. [PR1656702](#)
- The mspmand will crash when service-set is configured with Syslog and SSL. [PR1657027](#)
- rpm-postinst fails on boot. [PR1657278](#)
- PTP passthrough packets are timestamped by certain line cards on MX platforms. [PR1657291](#)
- The **/interfaces/interface/state/counters/in-fcs-errors** is not streamed to jvision collector. [PR1657913](#)
- The jdhcpd process might be stuck at 99 percent if traceoptions is enabled in high DHCP traffic scenario. [PR1658087](#)
- TOS(DSCP+ECN) bits not getting copied from the Inner L3 header to Outer VXLAN header. [PR1658142](#)
- The packetio might core when router reboot or FPC reboot is triggered. [PR1658839](#)

- Speed change from 10G to 1G on MX causes all other lanes to flap. [PR1659087](#)
- JSD crash during cBNG container startup or restart from CLI from within a cBNG container. [PR1659175](#)
- Traffic loss might be seen when a VXLAN port is recovering from a failure. [PR1659533](#)
- Some of ports on MX platforms with MPC7E-10G do not come up with 1G speed. [PR1660154](#)
- Soft assertions in RPD will fail during GRES. [PR1660484](#)
- The port LEDs do not light up when 40G physical interfaces are up. [PR1660532](#)
- The I2circuit backup might not get reverted to primary in rare condition. [PR1661802](#)
- The fxpc crash might be observed with the RPF check enabled. [PR1662508](#)
- In the EVPN-VXLAN scenario, the DHCP packets will get dropped when the DHCP relay agent is configured. [PR1662524](#)
- Some packages name missing which is RMPD and MOBILE in the show version detail command output. [PR1662691](#)
- TCP MSS value might not get reflected to packets. [PR1662950](#)
- network-instance name for streaming telemetry to be changed from default to DEFAULT to align with CONFIG stanza. [PR1662999](#)
- The rpd core may be seen when there is a synchronization issue. [PR1663050](#)
- The offset value might be high on the downstream node while switching between line cards which impacts 5G services. [PR1663065](#)
- Primary and Backup NHG late binding is not supported, so the backup nhg should be created before the primary nhg and removed in the reverse order. [PR1663310](#)
- The forwarding plane is not updated properly in scaled MVPN scenario after receiving PIM leave messages. [PR1663568](#)
- Subscribers will be stuck in the Initializing or Terminating state. [PR1663689](#)
- Trinity-based line cards with VPLS and CFM configuration may crash when the indirect NH associated with LSI IFL is deleted. [PR1663725](#)
- Post switchover error message is seen during pccd initialization. [PR1664165](#)
- System shutdown might be observed for erroneous read for system temperature from the ASIC. [PR1664302](#)

- The show chassis fpc command displays inaccurate information about heap memory. [PR1664448](#)
- SPC3: Receive [Rx] queue of direct memory access might be stuck which may cause issues in packet processing. [PR1664517](#)
- The routing process on the device might crash when the IP address of local interface is changed to the IP address of BGP peer. [PR1664527](#)
- Switch Fabric Board information for supporting PTP on MX10k8 with MX10K-LC2101 LC(s). [PR1664569](#)
- Line card may crash after offline/online plane. [PR1664602](#)
- MAC addresses learned on the RTG interface are not aging out. [PR1664955](#)
- The link-degrade recovery will not work for a specific interface speed. [PR1664978](#)
- The rpf-check feature might not be working in a Junos Subscriber management scenario. [PR1665234](#)
- JDI-RCT:M/Mx: errors spew on junos upgrade @ rm: //opt/sdk/service-packages/...: is a directory. [PR1665411](#)
- The process pfe might crash upon using some CLI commands frequently. [PR1665515](#)
- Traffic loss will be seen when a high-priority class-of-service scheduler interface flaps. [PR1665783](#)
- MAC-IP bindings for IPv4 (ARP) and IPv6 (ND) may not be processed for IRB interfaces in an EVPN scenario. [PR1665828](#)
- In the SRTE scenario, sensors are wrongly populated for colored tunnel BSID routes when uncolored tunnels are enabled. [PR1665943](#)
- The traffic loss might be observed over the AF interface. [PR1666397](#)
- CmError: show system errors active detail fpc slot is returning empty output. [PR1666510](#)
- BGP-LU traffic might be dropped when "CCNH ingress labeled-bgp inet" is configured. [PR1666760](#)
- Traffic loss is observed when the VRRP is configured over the aggregated Ethernet interface. [PR1666853](#)
- Traffic drop might occur on AF interface when PFE gets in disabled state on GNF in NodeSlice platforms. [PR1666992](#)
- The hyper-mode might be set incorrectly after power cycle on MX platforms. [PR1667226](#)
- H-VPLS traffic blackhole when mesh group local-switching is disabled. [PR1667310](#)
- High numbers of PDs connected may result in high CPU utilization. [PR1667564](#)

- The FPC might fail to initialize on Junos platforms. [PR1667674](#)
- Shutting the CE interface and bringing back up causes traffic (going towards the core) drop. [PR1667724](#)
- The RSVP interfaces are not streamed when removing the interface configuration. [PR1667844](#)
- Shaping-rate is not taking 20 bytes of overhead into account. [PR1667879](#)
- The inline services softwires (like 6rd, map-e) will corrupt memory leading to FPC crash. [PR1667918](#)
- Periodic event generation doesn't work after Routing Engine reboot. [PR1668152](#)
- Timestamp uint_val isn't proper is streaming output. [PR1668265](#)
- Type-5 routes might not get programmed in PFE when the number of active interfaces and Vlan configured is high. [PR1668352](#)
- The BGP multipath might not install some of the available next-hops. [PR1668481](#)
- Kernel logs on cRPD containers running on the same host are incomplete. [PR1668794](#)
- EVPN PE router might respond traceroute with unexpected source IP address to remote CE devices. [PR1668837](#)
- Commit configuration check-out failed while configuring syslog stream host IP in specific range. [PR1668941](#)
- Traffic loss might be observed for the multicast traffic. [PR1668976](#)
- UDP: optics sensor data missing snmp_if_index under Optics_diag. [PR1669333](#)
- The rpd process restarts after generating core files. [PR1669346](#)
- jsd memory leak and may lead jsd restart. [PR1669426](#)
- The process fabspoked-pfe crash might be observed while executing CLI commands for fabric statistics. [PR1669435](#)
- Update Flush API response to current published proto definitions. [PR1669536](#)
- The error will be observed if eTree is used with EVPN-MPLS and the routing-instance is changed. [PR1669609](#)
- Periodic rebalancing of subscribers over AE interface might not work. [PR1669637](#)
- LLDP neighborhood might fail if the chassis-id format of the LLDP packet is xx:xx:xx:XX:XX:xx!. [PR1669677](#)
- Errors are seen on bringing SIB online. [PR1669713](#)

- Layer filters matching DMAC/Etype take no effect on L2 SP-style aggregated Ethernet interface. [PR1669718](#)
- Interoperability issue between legacy line cards and MPC10E/11E causes Layer 2 packet drop. [PR1669765](#)
- USB installation package loads with 32-bit smartd binary version. [PR1669892](#)
- EVPN multicast traffic might get impacted because of routes getting stuck in the kernel routing table (krt) queue. [PR1670435](#)
- Fabric Destination error and Fabric plane going in check state after changing the fabric redundancy mode. [PR1670507](#)
- Subscriber traffic drops are seen on all Junos MX platforms with reason of 'sw error' in PFE State Invalid after ISSU. [PR1670577](#)
- On MX platforms, back to back modification of the Interface profile results in ports not come up. [PR1670685](#)
- Routing Engine reboot can be seen when PPPoE subscribers login. [PR1671135](#)
- PTP server state stuck in acquiring state when configured on a port enabled with Ingress Queuing feature. [PR1671262](#)
- MX150 platform reports error for bandwidth license. [PR1671347](#)
- Traffic loss seen due to SPC3's packets getting stuck. [PR1671649](#)
- Traffic impact might be seen due to an unexpected reboot of SPC3 card. [PR1672819](#)
- Backup FEB1 links down after master FEB0 restart. [PR1673274](#)
- The new primary Routing Engine might self-reboot after the kernel crashes on an old primary Routing Engine. [PR1673306](#)
- Memory leaks by any change in IPv4 or IPv6 multicast prefixes. [PR1673341](#)
- Reporting-interval in the show jvision sensor info command gets stuck at 65000 when configured reporting rate is changed from 65000 to 68000. [PR1673476](#)
- Training failures reported on the MX2010/MX2020 Junos platforms post fabric plane offline-online. [PR1673806](#)
- OSPF state stuck in Init state in IGMP-snooping enabled scenario. [PR1674217](#)
- During the smooth upgrade from SFB1 to SFB2, SFB2 gets detected as "Unknown Fabric Board". [PR1674309](#)

- SNMP traps "Power Supply failed" and "Power Supply OK" are not generated. [PR1674322](#)
- The "nsd" may crash post NAT rule configuration change. [PR1674381](#)
- The 'kmd' process might crash due to SA re-negotiation failure during IKE phase-1. [PR1674585](#)
- The fragment-offset-except match condition will not work with some values. [PR1675482](#)
- In a rare case, 'pccd' will crash when the PCEP connection is down. [PR1675816](#)
- Traffic flow will be affected as interfaces will be removed from VLAN. [PR1675861](#)
- PFE core dump is seen when the CPCD service is modified. [PR1675985](#)
- MX-SPC3 PIC core dump is seen when a CPCD service is modified. [PR1675990](#)
- MPC stuck in present state with log " graceful offline in progress, returning false" flooding. [PR1676008](#)
- Minor memory leak in 'bbe-statsd' daemon may be seen when subscriber-management is enabled on MX devices. [PR1676049](#)
- Traffic would not go through on the management port at link speeds 10M and 100M. [PR1676433](#)
- While processing SNMP GetNext requests 'transportd' might reach 100% of CPU utilization. [PR1676593](#)
- The traffic does not re-route quickly causing traffic silently discarded. [PR1676740](#)
- Traffic drop can be seen on MX platforms with MPC10E-10C line card. [PR1676777](#)
- IS-IS packet drop observed for packets having GRE over FTI-VXLAN header. [PR1676912](#)
- Traffic drop can be seen for MPC7/8/9 during unified ISSU in a specific scenario. [PR1678130](#)
- Memory leak is observed after GRES. [PR1678217](#)
- The show interfaces diagnostics optics interface command displays all 0 on 100/400G port on MPC10E card. [PR1678716](#)
- The rpd process crashes when a delegated LSP with IPv6 install prefix is configured. [PR1678874](#)
- On linecards MPC10E, MPC11E and LC9600, no user configured MAC address on IRB IFL is used as source MAC in the transit path. [PR1678927](#)
- The pccd process might crash during MBB for an externally controlled LSP. [PR1678970](#)
- The l2ald is treating MAC as a duplicate causing traffic loss. [PR1680242](#)

- The process bbe-smgd on the router would stop processing new PPPoE subscribers session. [PR1680453](#)
- The dynamic tunnel route gets removed when a new tunnel is brought up for the same selector. [PR1680775](#)
- Traffic drop would be observed only when the backup link is up on link-protection LAG interface. [PR1680889](#)
- The PFE process crashes from 21.4R1 version onwards on VMhost platforms. [PR1681532](#)
- Fabric Plane check/error alarm would be seen due to the burst traffic in MS-MPC line cards. [PR1681624](#)
- The policy-multipath route inherits the attributes of active-route but does not undergo resolution, causing an incorrect metric value. [PR1683003](#)
- 'clear interfaces statistics all' taking more than 9 min due to invalid PIC configuration inside GNF. [PR1683312](#)
- The traffic drop would be observed with inter-VLAN configuration when deactivating and activating the EVPN routing instance. [PR1683321](#)
- Commit check error message is not thrown when DetNAT is configured with AMS load-balancing-options. [PR1683772](#)
- srv6-oam: more than one label stack is not supporting, gives as **Maximum number of sids supported is 0** error in srv6 ping in JNP10008-SF2/JNP10K-LC2301 [lc9600]. [PR1683883](#)
- The rpd crash when SRv6 service routes resolve over SRv6 SRTE policies using older resolution scheme. [PR1683993](#)
- Traffic would hit wrong queue post unified ISSU. [PR1684019](#)
- The l2cpd process crash might be observed when disabling RSTP on an interface. [PR1684072](#)
- An interface configured as 1G might flap on a port with the mixed speeds of 1G and 10G after a PIC restart. [PR1684728](#)
- Insufficient space for vmcores for JUNOS VM. [PR1684968](#)
- TI-LFA backup path is not computed which effects slow convergence in case of failures. [PR1685064](#)
- Multiple bbe-smgd cores might be observed resulting in subscribers being lost or failing to login in the Enhanced subscriber scenario. [PR1685070](#)
- When uncorrectable FEC/CRC errors above the threshold are injected the plane is not going to check state. [PR1685230](#)

- PICs on the GNF failed to come online after the chassisd restart. [PR1685453](#)
- Illinois: K8 CP: Telenor Norway CST: bbe-smd-cpd core (patricia_delete; - bbe_cos_drop_profile_remove_all .../bbe_cos_drop_profile.c:837) during commit after adding very large class-of-service stanza to CP configuration. [PR1685482](#)
- The l2ald core seen after zeroize. [PR1686097](#)
- The rpd crash would be observed when two separate next-hops in rpd map to the same next-hop-index in the kernel. [PR1686211](#)
- VPLS traffic loss might be seen when deleting and adding a routing-instance. [PR1686523](#)
- [bgp] [Cores] Scapa : Rpd core seen in __assert_fail_base, __GI___assert_fail, patricia_delete, spring_te_tunnel_delete_now, spring_te_tunnel_id_repl_entry_unlock, mirror_dequeue. [PR1687287](#)
- The FPC crash is observed with a **flexible-match-mask** condition. [PR1687862](#)
- On Junos and Junos Evolved platforms delegated LSP control will not be returned to the PCC in specific scenario. [PR1687885](#)
- The LLDP output packets are not transmitting on the em0 interface of Junos OS platforms. [PR1688023](#)
- A kernel crash can be seen with MIC-3D-8DS3-E3 installed. [PR1688315](#)
- The LACP would get stuck in a continuous update loop in the MC-LAG scenario. [PR1688958](#)
- DCSPF LSPs remain down indefinitely after router-id of the ingress router is changed. [PR1689067](#)
- The logical interface policer is not working as expected when applied to filter input-list and output-list. [PR1689199](#)
- "failed to get template var id" error messages are generated by FPC when BFD liveness detection is negotiated by DHCP subscriber which has lawful intercept enabled. [PR1689621](#)
- A 1G port on a QSFP-4x10G transceiver will be down sometimes after the FPC restart. [PR1689644](#)
- Use latest os-package when upgrading. [PR1691209](#)
- The firewall bridge filter policers (attached to AE interface) are not working on all Junos MX platform with MPC10 card upon deactivate-activate a term intended to limit overall traffic. [PR1692070](#)
- ALG child session will not be transported through the DS-Lite tunnel which might lead to traffic failures in absence of a direct route to the host. [PR1692525](#)
- JNP10K-LC9600: G.8275.1: 2way/cTE fails to meet class-B with asymmetric port combinations. [PR1692746](#)

- Traffic loss is observed when the ECMP path is IRB over AE (IPv4->MPLS). [PR1693424](#)
- Traffic loss will be seen when MACSEC is configured. [PR1693730](#)
- NDP cannot resolve neighbor after clearing IPv6 neighbor. [PR1694009](#)
- BMP EOR is sent with wrong peer address causing BMP failure. [PR1695320](#)
- MPC11E goes offline with fpc-slice configured. [PR1695510](#)
- JNP10K-LC9600: G.8275.1: 2way time error spikes seen upon switching between two upstream congruent primaries. [PR1696880](#)

Routing Policy and Firewall Filters

- The rpd process crashes whenever it is getting shut down with router reboot, rpd restart, RE switchover, and software upgrade. [PR1670998](#)
- lo0 egress filter with next-header option not supported. [PR1672315](#)
- The aftmand process crash might be observed. [PR1683361](#)

Routing Protocols

- An IGMPv2 snooping proxy will originate IGMPv3 membership for a new group join request towards a multicast router. [PR1637090](#)
- Ipv6 Inline BFD sessions are down when neighbor is not resolved. [PR1650677](#)
- The show security keychain detail CLI displays algorithm as hmac-* when ao authentication algorithm is configured. [PR1651195](#)
- Wrong next-hop weight might be observed with BGP PIC enabled. [PR1652666](#)
- Routing Process Daemon (rpd) crashes and restarts when a specific timing condition is hit with BGP configuration. [PR1659441](#)
- Incorrect inactive routes are being propagated to neighbors with add-path. [PR1660456](#)
- Damping policy not working as expected when import policy and damping policy are changed in one commit. [PR1660571](#)
- Traffic loss will be seen due to delay in BGP convergence time. [PR1663883](#)

- The v4 prefixes might not be advertised over the BGPv6 sessions. [PR1664168](#)
- The BSR information might not be flooded over NG-MVPN. [PR1664211](#)
- BGP labeled-unicast inactive routes might not be advertised when add-path is configured. [PR1665610](#)
- High CPU will be seen due to frequent triggering of SPF for IS-IS. [PR1667575](#)
- The rpd process might crash on removing the BGP-LU configuration. [PR1669514](#)
- RPD crash might be observed due to multiple sequences of flap events. [PR1669615](#)
- The rpd crash is observed while making configurational changes. [PR1669716](#)
- BGP inactive routes might not be advertised to peers in BGP-LU scenario. [PR1669930](#)
- The rpd crashes upon receiving BGP multi-nexthops inetflow route in the 21.4 software release and onward. [PR1670630](#)
- Source/Destination AS fields shows up as 0 in the flow record. [PR1670673](#)
- The rpd can crash while route exchange using BGP and LDP in a rare scenario. [PR1671081](#)
- The backup next hop computation might not be as expected for some random prefixes when there is a topology change. [PR1671672](#)
- MCSNOOPD will be restarted and will again learn the states after core. [PR1672488](#)
- Vrfs with color routes which needs resolution can trigger a crash in rpd when bgp peers are going down. [PR1673160](#)
- The IS-IS learnt routes might be downloaded to RIB again and again if the prefix attribute flags are different. [PR1673953](#)
- Segment-routing might incorrectly set a pop action for paths using a Strict SPF(1) Algorithm. [PR1674804](#)
- MX304: KRT queue shows deferred operation while creating IFL after FPC offline/online event. [PR1675212](#)
- The process rpd (route process daemon) crashes with BGP VPN (Border gateway protocol - Virtual Private Network) config, while ebgp (external bgp) routes exported into ibgp (internal bgp) core with vrf (virtual route forward) configured. [PR1675893](#)
- Micro BFD session state in Routing Engine remain UP even peer side session is down. [PR1675921](#)
- Label traffic will be dropped at the one-hop LSP stitching node if the packet has more than one label. [PR1677567](#)

- High CPU is seen on the platforms running IPv6. [PR1677749](#)
- Inter-domain forwarding connectivity will be broken between different lo0s in the option-C network causing problems for the MPLS transit-route. [PR1677935](#)
- Traffic drops due to the generation of the FPC core, which makes the system unstable. [PR1678016](#)
- RV task replication will be stuck in the "NotStarted" state when routing-options validation is deactivated/activated. [PR1679495](#)
- The AGGREGATOR attribute will not be set correctly when the independent-domain is configured. [PR1679646](#)
- BGP auto-discovery sessions does not work any more after an interface flap. [PR1679950](#)
- InboundConvergencePending flag is set after the Routing Engine switchover. [PR1680360](#)
- Traffic loss is seen when the router in helper mode deletes the route for the router undergoing graceful restart. [PR1682506](#)
- The rpd process will crash and generate core post graceful restart. [PR1682778](#)
- On single PFE with Fusion satellite, LACP is not sending PDUs. [PR1687395](#)
- The rpd crash would be seen on a system running with IGP shortcuts. [PR1690231](#)
- RPD core is seen when using a BGP neighbor telemetry subscription in a sharding environment. [PR1692255](#)
- OSPF stuck in InitStrictBFD state for the neighbor which does not send LLS header. [PR1700966](#)

Services Applications

- L2TP session might not come up when L2TP access-line-information is not configured. [PR1667861](#)
- VMcore or Routing Engine crash might be triggered due to the memory corruption when the FPC is restarted for LNS subscribers. [PR1667950](#)

Subscriber Access Management

- New service profile provided by Radius during re-authentication triggered by DHCP Renew packet with changed actual data rate TLVs might not be applied. [PR1665947](#)

- CoA-NAK might not be sent for a coa-request-retry of the same service. [PR1667002](#)
- Errors are seen when the accounting server source address is IPv6. [PR1669284](#)
- The authd process crashes during GRES recovery phase. [PR1687998](#)

User Interface and Configuration

- Commit and commit check fails when you configure the interface-range command. [PR1656565](#)
- Configuration changes are not effective if special groups are applied using regex like, apply-groups "\$ {node}"; using explicit special groups names apply-groups [node0 node1]; solves the problem. [PR1660165](#)
- Commit failure when changing BGP well-known community attributes. [PR1669375](#)
- Configuration failure occurs after upgrade when modifying group for lsp auto-bandwidth. [PR1671038](#)
- Test Configuration might fail even though the config file is having valid configurations. [PR1671112](#)
- Upgrade/downgrade/rollback to 22.2R1 will fail if "system configuration-database extend-size" is configured. [PR1672348](#)
- "gethostbyname: Host name lookup failure" is displayed during commit. [PR1673176](#)
- Device is not entering in CLI mode; CLI core dumps are generated. [PR1673979](#)

VPN

- Traffic over IPSec tunnels may be dropped during ISSU. [PR1416334](#)
- 19.2TH:VPN:SRX5600: While verifying the show security ipsec next-hop-tunnels command output in device the IPsec SA and NHTB entry does not get cleared after configuring firewall filter. [PR1432925](#)
- Restore ATMVPN address family NLRI and use a new NLRI value for BGP Multicast NLRI. [PR1590331](#)
- Tunnel bringing up failed from strongswan when changing the configuration IKE in VR and the NO_PROPOSAL_CHOSEN notify error error message gets generated. [PR1627963](#)
- The multicast receiver receives no traffic in an extranet scenario having an SPT tree already established [PR1675099](#)

- Memory leak will be seen in rpd process. [PR1662239](#)
- Core files gets generated with multiple daemons restart. [PR1682573](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 116

This section contains the upgrade and downgrade support policy for Junos OS for MX Series routers. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

NOTE: Junos OS Release 22.4 is the last-supported release for the following SKUs:

- MS-MPC-128G-BB
- MS-MPC-128G-R
- MS-MPC-128G-SX
- MS-MIC-16G
- MS-MIC-16G-SX
- SCG-TM-BAS

We recommend upgrading to MX-SPC3 **only** for the following SKUS:

- MS-MPC-128G-BB
- MS-MPC-128G-R
- MS-MPC-128G-SX

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 7: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 117](#)
- [What's Changed | 117](#)
- [Known Limitations | 117](#)
- [Open Issues | 118](#)
- [Resolved Issues | 118](#)
- [Migration, Upgrade, and Downgrade Instructions | 119](#)

What's New

There are no new features or enhancements to existing features in Junos OS Release 22.4R1 for NFX.

What's Changed

There are no changes in behavior and syntax in this release for NFX Series devices.

Known Limitations

There are no known limitations in hardware or software in this release for NFX Series devices.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

IN THIS SECTION

- [High Availability | 118](#)

Learn about open issues in this release for NFX Series devices.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

High Availability

- On an NFX350 chassis cluster, when FPC0 (when node0 is primary) or FPC7 (when node1 is primary) is restarted by either using the request chassis fpc slot slot restart node local command or due to dcpfe core files on the primary, it restarts FPC1 or FPC8. This might break the preexisting TCP sessions and fail to restart the TCP sessions. The TCP sessions might require a manual restart. [PR1557607](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues | 119](#)

Learn about the issues fixed in this release for NFX Series.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

IN THIS SECTION

- [Routing Protocols | 119](#)
- [Virtual Network Functions \(VNFs\) | 119](#)

Routing Protocols

- The BSR information might not be flooded over NG-MVPN. [PR1664211](#)

Virtual Network Functions (VNFs)

- The NFX350 device stops responding after you configure VNF with SRIOV interfaces. Also, JDM becomes unreachable. [PR1664814](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 120](#)

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the NFX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

NOTE: For information about NFX product compatibility, see [NFX Product Compatibility](#).

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 8: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for PTX Series

IN THIS SECTION

- [What's New | 121](#)
- [What's Changed | 128](#)
- [Known Limitations | 130](#)
- [Open Issues | 131](#)
- [Resolved Issues | 132](#)
- [Migration, Upgrade, and Downgrade Instructions | 134](#)

What's New

IN THIS SECTION

- [Authentication and Access Control | 122](#)
- [EVPN | 122](#)
- [High Availability | 123](#)
- [Junos Telemetry Interface | 123](#)
- [MPLS | 124](#)
- [OpenConfig | 125](#)
- [Routing Policy and Firewall Filters | 125](#)
- [Routing Protocols | 125](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 128](#)
- [Additional Features | 128](#)

Learn about new features introduced in this release for the PTX Series.

Authentication and Access Control

- **OpenSSH certificate support (PTX1000, PTX5000)**—Starting in Junos OS Release 22.4R1, you can configure SSH certificate-based authentication for users and hosts. This lets you setup SSH access to a device with password-less login for users, and gives the capability to trust hosts without the need to verify key fingerprints.

The following new CLI configuration statements can be used to configure SSH certificate-based authentication:

- [system services ssh trusted-user-ca-key-file *filename*]—Configure the TrustedUserCAKey file at **/etc/ssh/sshd_config** which contains the public keys of an SSH certificate.
- [system services ssh host-certificate-file *filename*]—Configure the HostCertificatefile at **/etc/ssh/sshd_config** which contains the signed host certificate.
- [system services ssh authorized-principals-file *filename*]—Configure the AuthorizedPrincipalsfile at **/var/etc** which contains a list of names, one of which must appear in the certificate for it to be accepted for authentication.
- [system services ssh authorized-principals-command *program-path*]—Specify a program to be used for generating the list of allowed certificate principals found in the AuthorizedPrincipalsFile.

[See [Configure SSH Service for Remote Access to the Router or Switch.](#)]

EVPN

- **Support for Microsoft load-balancing node's static ARP entries with unicast MAC addresses (EX9208, MX-Series, and VMX)**—Starting in Junos OS Release 22.4R1, you can configure a Microsoft load-balancing node's static Address Resolution Protocol (ARP) entries for unicast MAC addresses on integrated routing and bridging (IRB) interfaces. On your provider edge (PE) device, you can create a static ARP entry for the Microsoft load-balancing node's virtual IP address and its unicast virtual MAC address. This static ARP configuration enables your PE devices to flood traffic for the Microsoft load-balancing node's virtual IP address to the virtual MAC address in an EVPN Layer 2 domain or any other Layer 2 domain.

To enable unicast MAC addresses on IRB interfaces, enable the `flood-as-unknown-unicas` option in the [edit interfaces irb unit <logical-interface-number> family inet address <local-ip-address>/<prefix-length> arp <MSLB-virtual IP address> mac <MSLB-unicast-VMAC>] hierarchy. The `flood-as-unknown-unicas` option enables flooding of virtual IP addresses and virtual MAC traffic flows from a Microsoft load-balancing cluster.

[See [EVPN User Guide.](#)]

- **VXLAN filter group prioritization (QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, and QFX5200)**—Starting in Junos OS Release 22.4R1, when your device boots up, the device initializes VXLAN dynamic filter groups before the CLI filter groups. If you configure a CLI

filter group before the VXLAN dynamic filter group, only the CLI filter group might be programmed in the TCAM space after you reboot your device. This situation might cause problems in your fabric.

[See [EVPN User Guide](#).]

High Availability

- **MVPN NSR with BGP Sharding enabled (MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we've enabled multicast virtual private network (MVPN) nonstop active routing (NSR) for border gateway function (BGP) sharding.

[See [Understanding BGP RIB sharding and BGP Update IO thread](#).]

Junos Telemetry Interface

- **Enhanced support for FIB telemetry streaming (MX240, MX960, MX2020, PTX5000, and PTX1000)**—Junos OS Release 22.4R1 introduces enhanced support for forwarding information base (FIB) telemetry streaming based on the OpenConfig Abstract Forwarding Table (AFT) model. We now support the following sensor paths:
 - `/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/origin-protocol`
 - `/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/origin-protocol`
 - `/network-instances/network-instance/afts/next-hops/next-hop/state/encapsulate-header`

[See [Telemetry Sensor Explorer](#).]

- **Event-driven streaming of sensor data for MPLS LSP record route objects (ACX5448, ACX7100, MX204, MX240, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX1000, and vMX)**—Junos OS Release 22.4R1 introduces ON_CHANGE notification for streaming MPLS label-switched path (LSP) record route object statistics. Using ON_CHANGE mode, data values are not streamed but sent only when data values change. Support includes leaf nodes under the resource path `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/sessions/session/record-route-objects/record-route-object/state/`.

[See [Telemetry Sensor Explorer](#).]

- **OpenConfig OSPF configuration and operational state sensors (ACX5448, ACX7100, MX150, MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, and PTX1000)**—Junos OS Release 22.4R1 introduces support for the OpenConfig OSPF data model `openconfig-ospfv2.yang (v.0.3.1)`. We now support configuration and streaming of operational state data under the resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.

To learn about OpenConfig configuration mappings, see [Mapping OpenConfig OSPF Commands to Junos Configuration](#). For state sensors, see [Telemetry Sensor Explorer](#).

- **System health reporting sensors support on gRPC (ACX5448 and ACX710 routers; MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10008, and MX10016 routers; and PTX10002 routers)**—Starting in Junos OS Release 22.4R1, Junos telemetry interface (JTI) Junos telemetry interface (JTI) supports data model `openconfig-system.yang` using gRPC remote procedure calls (gRPC) and provides new health-monitoring sensors.

[See [Telemetry Sensor Explorer](#).]

- **Support for gRPC tunnel sessions (MX204, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX-5000, PTX1000, PTX10002, VMX, and QFX5110)**—Starting with Junos OS Release 22.4R1, you can configure a gRPC tunnel session to establish a connection between an external TCP client and a TCP server. The gRPC tunnel session establishes a reverse connection when a TCP client can't reach the TCP server.

To establish a gRPC tunnel session, include the `grpc-tunnel` configuration statement in the `[edit system services]` hierarchy.

[See [gRPC Tunnels Overview](#).]

MPLS

- **Support for Software driven Wide Area Network (SWAN) ping and traceroute command for PRPD static routes (PTX10008 and PTX10016)**—Starting in Junos OS Release 22.4R1, SWAN ping and traceroute commands verify the PRPD static path by sending the MPLS echo request packet with the entire labelstack along the same data path that carries the data traffic. To forward this echo request, Junos OS follows a similar process it uses to forward any other packet.

[See [ping mpls segment routing static egress-ip](#) and [traceroute mpls segment-routing static egress-ip](#).]

- **PCEP multipath support for SR-TE (MX480, PTX10008, and QFX5200)**—Starting in Junos OS Release 22.4R1, you can configure the multipath feature (primary or secondary paths) for Path Computation Element Protocol (PCEP) segment routing-traffic engineering (SR-TE) as defined in *draft-ietf-pce-multipath-06*. We support the following multipath capabilities:

- When the PCEP multipath feature is enabled, you can configure multiple primary or secondary paths in a candidate path that you configure and control using Path Computation Client (PCC). Note that the PCEP multipath feature is enabled by default.
- When the PCEP multipath feature is disabled, you can configure only one primary path in a candidate path. Note that a secondary path configuration is not allowed.

The PCEP multipath feature removes the `compute-profile` restriction of 1 on the maximum number of segment lists (`maximum-computed-segment-lists`).

NOTE: When PCEP multipath is enabled, PCCD will not send constraints for PCC-controlled candidate paths.

[See [Configuring Multiple Paths for Path Computation Element Protocol SR-TE Overview](#).]

OpenConfig

- **OpenConfig OSPF configuration and operational state sensors (ACX5448, ACX7100, MX150, MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, and PTX1000)**—Starting in Junos OS Release 22.4R1, we support the OpenConfig OSPF data model `openconfig-ospfv2.yang` (version 0.3.1). We also support configuration and streaming of operational state data under the resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.

See [Mapping OpenConfig OSPF Commands to Junos Configuration](#) for OpenConfig configuration mappings. See [Telemetry Sensor Explorer](#) for state sensors.

Routing Policy and Firewall Filters

- **Support for transit traffic rates (PTX1000, PTX3000, PTX5000, and PTX10000)**—Starting in Junos OS Release 22.4R1, we support transit traffic rates in bits per second (bps) and packets per second (pps) for both IPv4 and IPv6 at the logical interface level.

Use the following commands to enable transit statistics accounting:

- For IPv4 traffic—set `forwarding-options family inet route-accounting`
- For IPv6 traffic—set `forwarding-options family inet6 route-accounting`

Use the `show interfaces interface statistics` command to display the traffic rates.

[See [Configuring IPv6 Accounting](#).]

Routing Protocols

- **BMP local RIB monitoring support for all RIBs with sharding (ACX Series, cRPD, PTX Series, QFX Series, and vRR)**—Starting in Junos OS Release 22.4R1, you can configure a policy to monitor routing information bases also known as routing table (RIBs) of virtual routers and virtual routing and forwarding instances (VRF). You can specify two separate sets of RIBs in the BGP Monitoring Protocol (BMP), one for monitoring and the other for reporting. With this feature, BMP can filter traffic based on the routes and routing instances.

[See [BGP Monitoring Protocol](#), `loc-rib`, and `rib-list`.]

- **Keep bypass tunnels operational during configuration changes (PTX1000)**—Starting in Junos OS Release 22.4R1, you can configure the OS to keep bypass tunnels operational until the tunnels no longer carry local repair traffic, even during configuration changes. Bypass tunnels that carry local repair traffic are in the BackupActive state. When you change the bypass-related configuration on software releases containing this feature, the OS keeps any bypass tunnels that are in BackupActive state up. When the bypass tunnels are no longer in BackupActive state, the operating system tears down the bypass tunnels. This feature ensures that all local repair traffic reaches its destination and prevents traffic loss on label-switched paths (LSPs).

Configure this feature at the [edit protocols rsvp interface all link-protection] hierarchy level. Use the `show rsvp session bypass` command to check whether the bypass routes protecting an interface remain operational in BackupActive state after the configuration changes.

[See [link-protection \(RSVP\)](#) and [Link Protection for MPLS LSPs](#).]

- **MD5 authentication key rotation with overlap for key transition for OSPF (MX204, MX480, MX10003, PTX1000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support advertising OSPF MD5 authentication with multiple active keys to send packets with a maximum limit of two keys per interface. Having multiple keys active at any one time at the interface enables the smooth transition from one key to another for OSPF. You can delete old keys without any impact on the OSPF session.

[See [Understanding OSPFv2 Authentication authentication](#).]

[See [show \(ospf | ospf3\) interface](#).]

- **OSPF FAPM and interarea support (ACX5448, MX204, MX240, MX480, MX960, MX10003, MX10008, MX2008, MX2010, MX2020, PTX1000, and QFX10002)**—Starting with Junos OS Release 22.4R1, the Flexible Algorithm Prefix Metric (FAPM) is defined to allow an optimal end-to-end path for an inter-area prefix. The Area Border Router (ABR) *must* include the FAPM when advertising the prefix between areas that areas reachable in that given Flex-Algorithm. When a prefix is unreachable, the ABR *must not* include that prefix in the Flex-Algorithm when advertising between areas. The defined FAPM provides inter-area support.

[See [Understanding OSPF Flexible Algorithm for Segment Routing](#).]

[See [show ospf database](#), [show route table](#), [show ted database](#)

- **Support for S-BFD over EPE SIDs (MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2010, MX2020, PTX5000, PTX1000, PTX10002, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.4R1, seamless BFD (S-BFD) running between ingress devices and autonomous system boundary routers (ASBRs) can track BGP egress peer engineering (EPE) segment identifiers (SIDs). With this feature, you can prevent null-route filtering if a BGP EPE SID goes down.

[See [sbfd](#).]

- **CLI support for BFD echo and echo-lite modes (MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2010, MX2020, PTX5000, PTX1000, PTX10002, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.4R1, you can configure BFD echo mode and echo-lite mode through the Junos OS CLI. When BFD echo mode is active, a neighbor device transmits and loops back BFD echo packets to ensure that a forwarding path is available. BFD echo mode requires both the local device and neighbor device to support the full BFD protocol. However, BFD echo-lite mode can function even if the neighbor device doesn't support BFD.

You can use the following new CLI configuration commands to configure BFD echo mode and echo-lite mode:

- **echo mode:** `set routing-options static route address bfd-liveness-detection echo minimum-interval interval`
- **echo-lite mode:** `set routing-options static route address bfd-liveness-detection echo-lite minimum-interval interval`

[See [bfd-liveness-detection](#).]

- **Support for bootstrapping route-validation database from a local file (cRPD, JRR200, MX204, PTX10008, and QFX10008)**—Starting in Junos OS Release 22.4R1, we support the ability to read validation records from a local binary file and install into the specified named route-validation databases within RPD. This feature implements syntactic and semantic checks on the content of the file to ensure that it is a well-specified set of validation records. If the syntactic and semantic checks fail, the entire file is rejected as a source of validation records. Use the `source-file` statement at the `[edit routing-options validation]` hierarchy level to source route-validation records from a local file source. You can use the `show validation source-file` command to display the properties of a local validation source file.

[See [validation](#).]

- **Support for BGP RIB sharding and update threading features (MX304, MX10003, MX10004, MX10008, MX10016, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.4R1, you can use a new CLI option to override the sharding and update-threading configuration that might be present either through platform defaults or through explicit configuration. To override the configuration, use the `no-rib-sharding` and `no-update-threading` options at the `[edit system processes routing bgp]` hierarchy level.

[See [bgp](#).]

- **MVPN feature support with sharding (cRPD, JRR200, MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support the following features:
 - Multicast virtual private network (MVPN) inactive route query from the main thread to shards
 - Extranet and auto-export support with sharding
 - Interact functions with RT-proxy client and server

- New data structure to store the inactive route data on the main thread
- Asynchronous route processing on the main thread

You can use `show mvpn c-multicast` to display the inactive route data stored on the main thread.

[See [rib-sharding](#) and [show mvpn c-multicast](#) .]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **BGP Classful Transport (CT) support for IPv6 and Segment Routing–Traffic Engineering (SR-TE) color-only support (MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2010, MX2020, PTX10008, PTX10016, VMX)**—Starting in Junos OS Release 22.4R1, we support BGP-CT with IPv6 and BGP service-routes with a color-only mapping community. We have also enhanced the transport-class configuration statement by enabling shards to provide strict resolution without falling back on best-effort tunnels.

[See [use-transport-class](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **Inline active flow monitoring support for configuring the options template ID, template ID, source ID, and observation domain ID (PTX1000, PTX10008, and PTX10016)**

[See [options-template-id](#), [template-id](#), [source-id](#), and [observation-domain-id](#).]

What's Changed

IN THIS SECTION

- [EVPN | 129](#)
- [Network Management and Monitoring | 129](#)
- [User Interface and Configuration | 129](#)

Learn about what changed in this release for the PTX Series.

EVPN

- **Flow-label configuration status for EVPN ELAN services**—The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

Network Management and Monitoring

- **Junos YANG modules for RPCs include the `junos:command` extension statement (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command` extension statement in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the YANG schemas on the local device, the YANG modules for RPCs include the `junos:command` extension statement.

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml` filter or the `| display json` filter to display command output in XML or JSON format in any release. We've removed these options from the CLI.

Known Limitations

IN THIS SECTION

- [General Routing | 130](#)
- [Routing Protocols | 130](#)

Learn about known limitations in this release for the PTX Series.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- The SDN-Telemetry process can crash during long running streaming telemetry collectors. There will be a loss of telemetry data streamed from the linecards till the process comes up. [PR1647568](#)
- On PTX10002 , all odd ports have a WAN re-timer connected to it , which might introduce more time during fault recovery , such as LocalFault clear . Hence sometimes even if the fault is of the order of milliseconds , the port might still get hold-time expired and flap when the configured "hold-time down" less than 3s. The behavior is confirmed as hardware limitation. [PR1687092](#)

Routing Protocols

- When routing-options transport-class fallback none is not configured - do not configure more than 10 transport-classes. Or advertise more than 10 distinct colors in SR-TE. [PR1648490](#)

Open Issues

IN THIS SECTION

- [General Routing | 131](#)
- [Multicast | 132](#)

Learn about open issues in this release for the PTX Series.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- On WRL8 based VMHost platforms (i.e., ACX6360/PTX10001/MX150/NFX150/NFX250/NFX350), there is no log rotation for resild log and temperature sensor info is incorrectly written into resild log which could result in continuous logs in resild log file. The disk usage might keep increasing due to this issue. The disk usage could be eventually full which could cause system to hang and reboot. [PR1480217](#)
- In the platform using indirect next-hop (INH), such as Unilist as route next hop type for multiple paths scenario (such as BGP PIC or ECMP), the session fast-reroute might be enabled in Packet Forwarding Engines. When the version-id of session-id of INH is above 256, the Packet Forwarding Engine might not respond to session update, which might cause the session-id permanently to be stuck with the weight of 65535 in Packet Forwarding Engine. It might lead Packet Forwarding Engine to have a different view of Unilist against load-balance selectors. Then either the BGP PIC or the ECMP-FRR might not work properly and traffic might be dropped or silently discarded. [PR1501817](#)
- On Junos OS platforms supporting GRES, the bandwidth flag on the backup Routing Engine for the aggregated Ethernet interfaces is set unconditionally. Therefore, the BW is struck at a value despite not being set by the user statically. When GRES is performed followed by link delete/add for the aggregated Ethernet interface, the interface might get stuck with the BW and gets synced with the backup RE. There can be seen partial service impact in the case of the RSVP (Resource Reservation Protocol) and the restoration to recover from this issue is to configure bandwidth manually and then remove the configuration. [PR1649958](#)
- In a scenario where static routes are added with varying weight values and verified, rpd core file is generated. rpd is recovered automatically and restarted by jlaunchd. [PR1699356](#)

Multicast

- On Junos OS PTX Series platforms traffic might silently drop because of the next-hop installation failure for multicast RSVP point to multipoint (P2MP) traffic. This issue might be encountered in a scaled RSVP P2MP environment after a network event which might cause reconvergence. [PR1653920](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 132](#)
- [Interfaces and Chassis | 133](#)
- [MPLS | 133](#)
- [Routing Protocols | 134](#)

Learn about the issues fixed in this release for PTX Series.

General Routing

- IS-IS adjacency does not come up through TCC I2circuit. [PR1590387](#)
- On Junos OS PTX Series platforms, traffic silently drops after interface flap. [PR1645488](#)
- BGP Sensor "/bgp-rib/afi-safis/afi-safi/ipv4-unicast/loc-rib/" not available as a 'periodic' sensor. [PR1649529](#)
- PCS errored blocks count will increment after upgrading to Junos OS release 20.2R1 and later. [PR1651526](#)
- IS-IS adjacency is not coming up through the Layer 2 domain. [PR1663134](#)
- The process pfe might crash upon using some CLI commands frequently. [PR1665515](#)
- PCS errored blocks count increments on PTX3000 and PTX5000 after Junos OS upgrade. [PR1669267](#)

- JSD memory might leak that restarts the JSD process. [PR1669426](#)
- Reporting-interval in show jvision sensor info is stuck at 65000 when configured reporting rate is changed from 65000 to 68000. [PR1673476](#)
- Issue with eth-lldp-stop.sh is seen after upgrading the Junos OS in PTX5000 (i40e-NVM). [PR1675177](#)
- The Packet Forwarding Engine process crashes from Junos OS release 21.4R1 and later on VMhost platforms. [PR1681532](#)
- Traffic drop is seen in MACsec enabled scenario on PTX Series platforms. [PR1682161](#)
- On PTX5000 platforms when a command is issued to power off an FPC, it gets stuck in the 'Announce Offline' state. [PR1683562](#)
- Logical components which do not have EEPROM, return empty-string for leaves hardware-version, part-no and ID. [PR1685968](#)
- You might observe the rpd crash when two separate next hops in rpd, map to the same next-hop-index in the kernel. [PR1686211](#)

Interfaces and Chassis

- reth1 interface down and DCD generates a core file on node1 during test. [PR1657021](#)

MPLS

- The error severity of syslog message ted_client reset generated during commit is incorrect. [PR1649565](#)
- Premature RSVP path error BW-Unavailable originated by PLR. [PR1670638](#)
- The rpd crash might be observed with container LSPs. [PR1672804](#)
- CPU utilization of rpd process might reach 100 percent while reporting LSP states to pccd if the IS-IS update churn is high. [PR1673348](#)
- The traffic might drop when the Link State protocol with the least preference is set to active and fails the CSPF algorithm [PR1677930](#)

- In an LDP -> BGP LU stitching scenario, Multiple LSPs will not be installed in the forwarding table, even if BGP Multipath and ECMP are enabled. [PR1680574](#)
- In the RSVP-TE scenario, with entropy label capability is enabled during MBB issues handling Resv messages. [PR1681403](#)

Routing Protocols

- A memory leak which will ultimately lead to an rpd crash will be observed when a peer interface flaps continuously in a Segment Routing scenario using OSPF (CVE-2023-22406). [PR1659366](#)
- Micro BFD session state in Routing Engine remain UP even peer side session is down. [PR1675921](#)
- The rpd crash might be seen on a system running with IGP shortcuts. [PR1690231](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 138](#)

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the PTX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

NOTE: Junos OS 22.4 is the last supported release on many PTX Series products. For more information on EOL dates, see: [PTX Series Hardware Dates & Milestones](#). See the [Junos OS Evolved release notes](#) for PTX Series products that run Junos OS Evolved.

Basic Procedure for Upgrading to Release 22.4

When upgrading or downgrading Junos OS, use the `jinstall` package. For information about the contents of the `jinstall` package and details of the installation process, see the [Installation and Upgrade Guide](#). Use other packages, such as the `jbundle` package, only when so instructed by a Juniper Networks support representative.

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

```
user@host>request system snapshot
```

NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [Installation and Upgrade Guide](#).

NOTE: We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

To download and install Junos OS Release 22.4R1:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.
4. Select the Software tab.
5. In the Install Package section of the Software tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new jinstall package on the router.

NOTE: We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/junos-install-ptx-
x86-64-22.4R1.9.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (limited encryption Junos OS package):

```
user@host> request system software add validate reboot source/junos-install-ptx-
x86-64-22.4R1.9-limited.tgz
```

Replace the source with one of the following values:

- ***/pathname***—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - ***ftp://hostname/pathname***
 - ***http://hostname/pathname***
 - ***scp://hostname/pathname***

The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Adding the `reboot` command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the `request vmhost software add` command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 22.4 `jinstall` package, you cannot return to the previously installed software by issuing the `request system software rollback` command. Instead, you must issue the `request system software add validate` command and specify the `jinstall` package that corresponds to the previously installed software.

NOTE: Most of the existing `request system` commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the [Installation and Upgrade Guide](#).

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 9: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for QFX Series

IN THIS SECTION

- [What's New | 139](#)
- [What's Changed | 147](#)
- [Known Limitations | 149](#)
- [Open Issues | 149](#)
- [Resolved Issues | 152](#)
- [Migration, Upgrade, and Downgrade Instructions | 158](#)

What's New

IN THIS SECTION

- [Authentication and Access Control | 140](#)
- [Ethernet Switching and Bridging | 140](#)
- [EVPN | 141](#)
- [High Availability | 143](#)
- [Interfaces | 143](#)
- [Juniper Extension Toolkit \(JET\) | 143](#)
- [Junos OS API and Scripting | 143](#)
- [Junos Telemetry Interface | 144](#)
- [MPLS | 144](#)
- [Routing Protocols | 145](#)
- [Additional Features | 146](#)

Learn about new features introduced in this release for QFX Series switches.

Authentication and Access Control

- **802.1X authentication with EVPN-VXLAN (EX4650, QFX5120)**—Starting in Junos OS Release 22.4R1, the EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM switches that act as access switches can use 802.1X authentication to protect an EVPN-VXLAN network from unauthorized end devices.

These switches support the following 802.1X authentication features on access and trunk ports:

- Access ports: single, single-secure, and multiple supplicant modes
- Trunk ports: single and single-secure supplicant modes
- Guest VLAN
- Server fail
- Server reject
- Dynamic VLAN
- Dynamic firewall filters
- RADIUS accounting
- Port bounce with Change of Authorization (CoA) requests
- MAC RADIUS client authentication
- Central Web Authentication (CWA) with redirect URL
- Captive portal client authentication
- Flexible authentication with fallback scenarios

[See [802.1X Authentication](#)

Ethernet Switching and Bridging

- **Support for multiple TPID configurations per VLAN on QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210**—Starting in Junos OS Release 22.4R1, the QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches support up to four Tag Protocol Identifier (TPID) configurations per VLAN. The system supports the following values for the TPID configuration: 0x8100, 0x9200, 0x88a8, and 0x9100. You have to define a TPID in the CLI for aggregated Ethernet interfaces and the Ethernet interfaces for the device, before you send or receive a TPID.

[See [Configuring Tag Protocol IDs \(TPIDs\) on QFX Series Switches](#) and [tag-protocol-id \(TPIDs Expected to Be Sent or Received\)](#).]

- **Support for transmitting maximum VLAN Name TLVs in LLDP**—Junos OS Release 22.4R1 supports transmission of more than 5 VLAN name type, length, and value (TLVs) in the Link Layer Discovery Protocol (LLDP) data units. The LLDP data units can carry as many TLVs as possible. The number of VLAN TLVs is decided dynamically based on the maximum transmission unit (MTU) of the interface and other necessary TLVs that the protocol data units need to carry. The command set `protocols lldp interface <name> transmit-max-vlan-tlv` enables the system to send the maximum VLAN TLVs in the LLDP data units.

The TLVs advertise information to peers. TLVs are used by LLDP neighbors to discover a device's capabilities.

[See [Configuring the Transmission of Maximum VLAN Name TLVs in LLDP](#), `show lldp`, and `show lldp neighbors-vlan-name-tlv-list interface`.]

EVPN

- **EVPN-VXLAN to EVPN-VXLAN seamless stitching for EVPN Type 5 routes (EX4100-24T, EX4400-24MP, EX4400-24P, EX4400-48F, EX4650, MX204, MX240, QFX10002-60C, QFX10002, QFX10008, QFX10016, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 22.4R1, you can configure Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) to EVPN-VXLAN seamless stitching with EVPN Type 5 (IP prefix) routes between two interconnected data centers or between two points of delivery (pods) in a data center.

In the EVPN-VXLAN fabric, border leaf or border spine devices act as interconnection gateways. You enable EVPN Type 5 routes in virtual routing and forwarding (VRF) instances on both sides of the interconnection. For each VRF instance, the server leaf devices in the first data center create VXLAN tunnels for Type 5 routes (with corresponding virtual network identifiers [VNIs]) toward their local gateway devices. The gateway devices map those VXLAN tunnels to an interconnection tunnel (with a new route distinguisher [RD], route target, and VNI) toward the second data center. The gateway devices in the second data center re-create the Type 5 VXLAN tunnels using their local RD.

We support one-to-one mapping of Type 5 VRF instances across the interconnection.

- **Support for VXLAN group-based policy with ingress and egress configuration (EX4100, EX4400, EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting in Junos OS Release 22.4R1, we've added enhancements to the group-based policy (GBP) feature and made some changes to the CLIs.

The enhancements are:

- You can enforce the policy on the ingress endpoint or the egress tunnel endpoint. Ingress enforcement optimizes the network bandwidth. To configure policy enforcement at the ingress endpoint, use the `set forwarding-options evpn-vxlan gbp ingress-enforcement` command.
- We support these match conditions for GBP tagging:
 - `interface <interface_name>`

- `mac-address <mac address>`
- `vlan-id <vlan id>`
- `ip-version ipv4 <ip address> or <prefix-list>`
- `ip-version ipv6 <ip address> or <prefix-list>`

[See [Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN.](#)]

- **Routing protocols on EVPN-VXLAN overlay IRB interfaces in the default routing instance (EX4400, EX4650, EX9200, EX9253, QFX5110, QFX5120, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.4R1, you can run routing protocols on Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) overlay integrated routing and bridging (IRB) interfaces in the IPv4 or IPv6 default routing instance associated with the underlay (default.inet.0 or default.inet6.0). To perform this task, you can set and export a policy with the `install-next-hop except overlay-vxlan-interfaces` policy qualifier option. The policy configuration avoids routing loops that can happen if the device uses overlay IRB routes for underlay VTEP reachability. To support this use case in releases prior to 22.4R1, you can configure the IRB interface in a routing instance of type `vrf` instead of in the default routing instance.

[See [install-nexthop.](#)]

- **Remote port mirroring for EVPN-VXLAN with pure IPv6 underlay (QFX10002, QFX10002-60C, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.4R1, you can configure remote port mirroring in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) environment with a pure IPv6 underlay. Remote port mirroring copies the packets of a traffic flow and delivers the mirrored packets to one or more destination hosts.

[See [MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.](#)]

- **Protect core support for EVPN-VXLAN (EX4300-MP, EX4400-48MP, EX4650, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 22.4R1, you can configure the protect core feature in an EVPN-VXLAN environment. You can use the protect core feature to install a route in the forwarding table for use as an alternative path when an existing route fails or if connectivity is lost.

[See [protect core.](#)]

- **Overlay and CE-IP ping and traceroute support for EVPN-VXLAN (EX4300-MP, EX4400, EX4650, QFX5110, QFX5120, QFX5200, QFX10002, QFX10002-60C, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.4R1, you can perform ping and traceroute operations within an EVPN-VXLAN overlay or to a specific customer edge (CE) device IP address (CE-IP) across an EVPN-VXLAN overlay. You can use ping, traceroute, CE-IP ping, and CE-IP traceroute utilities to detect and isolate faults in overlay networks.

[See [Understanding Overlay Ping and Traceroute Packet Support.](#)]

High Availability

- **MVPN NSR with BGP Sharding enabled (MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we've enabled multicast virtual private network (MVPN) nonstop active routing (NSR) for border gateway function (BGP) sharding.

[See [Understanding BGP RIB sharding and BGP Update IO thread.](#)]

Interfaces

- **Support for hold timer (QFX5120-32C)**—Starting in Junos OS Release 22.4R1, we support hold time on aggregated Ethernet (ae) interfaces.

[See [hold-time](#)

Juniper Extension Toolkit (JET)

- **Prevent script execution based on current system memory usage (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4300, EX4300-MP, EX4300-VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9208, EX9251, EX9253, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 22.4R1, you can configure the system memory usage threshold above which the device prevents the execution of Juniper Extension Toolkit (JET) scripts. You can configure the `start start-options mem-factor` statement for individual JET scripts or all JET scripts. The device doesn't execute the script if the system's memory usage exceeds the configured value at the time the script is invoked. This configuration ensures that a device executes only essential scripts when system resources are limited, thereby enabling the device to continue performing all critical network functions.

[See [Configure Script Start Options.](#)]

Junos OS API and Scripting

- **Prevent script execution based on current system memory usage (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4300, EX4300-MP, EX4300-VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9208, EX9251, EX9253, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 22.4R1, you can

configure the system memory usage threshold above which the device prevents the execution of certain op, event, or SNMP scripts. You can configure the `start start-options mem-factor` statement for individual scripts or all scripts of a given type. The device doesn't execute the script if the system's memory usage exceeds the configured value at the time the script is invoked. This configuration ensures that a device executes only essential scripts when system resources are limited, thereby enabling the device to continue performing all critical network functions.

[See [Configure Script Start Options](#).]

Junos Telemetry Interface

- **Telemetry support for interfaces and chassis (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-MP, EX4100-F, EX4300-MP, EX4400-MP, EX4400, EX4650, QFX5110, QFX5120, QFX5200, and QFX5210)**—Junos OS Release 22.4R1 introduces support for streaming operational state statistics and counters for chassis and interfaces using OpenConfig sensor paths. We also support the following new Junos-specific sensor paths for statistics that are unsupported in OpenConfig:

- `/state/chassis/`
- `/state/interfaces/`

[See [Telemetry Sensor Explorer](#).]

- **Support for gRPC tunnel sessions (MX204, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX-5000, PTX1000, PTX10002, VMX, and QFX5110)**—Starting with Junos OS Release 22.4R1, you can configure a gRPC tunnel session to establish a connection between an external TCP client and a TCP server. The gRPC tunnel session establishes a reverse connection when a TCP client can't reach the TCP server.

To establish a gRPC tunnel session, include the `grpc-tunnel` configuration statement in the `[edit system services]` hierarchy.

[See [gRPC Tunnels Overview](#).]

MPLS

- **PCEP multipath support for SR-TE (MX480, PTX10008, and QFX5200)**—Starting in Junos OS Release 22.4R1, you can configure the multipath feature (primary or secondary paths) for Path Computation Element Protocol (PCEP) segment routing-traffic engineering (SR-TE) as defined in *draft-ietf-pce-multipath-06*. We support the following multipath capabilities:
 - When the PCEP multipath feature is enabled, you can configure multiple primary or secondary paths in a candidate path that you configure and control using Path Computation Client (PCC). Note that the PCEP multipath feature is enabled by default.

- When the PCEP multipath feature is disabled, you can configure only one primary path in a candidate path. Note that a secondary path configuration is not allowed.

The PCEP multipath feature removes the `compute-profile` restriction of 1 on the maximum number of segment lists (`maximum-computed-segment-lists`).

NOTE: When PCEP multipath is enabled, PCCD will not send constraints for PCC-controlled candidate paths.

[See [Configuring Multiple Paths for Path Computation Element Protocol SR-TE Overview](#).]

Routing Protocols

- **BMP local RIB monitoring support for all RIBs with sharding (ACX Series, cRPD, PTX Series, QFX Series, and vRR)**—Starting in Junos OS Release 22.4R1, you can configure a policy to monitor routing information bases also known as routing table (RIBs) of virtual routers and virtual routing and forwarding instances (VRF). You can specify two separate sets of RIBs in the BGP Monitoring Protocol (BMP), one for monitoring and the other for reporting. With this feature, BMP can filter traffic based on the routes and routing instances.

[See [BGP Monitoring Protocol](#), [loc-rib](#), and [rib-list](#).]

- **MD5 authentication key rotation with overlap for key transition for OSPF (MX204, MX480, MX10003, PTX1000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support advertising OSPF MD5 authentication with multiple active keys to send packets with a maximum limit of two keys per interface. Having multiple keys active at any one time at the interface enables the smooth transition from one key to another for OSPF. You can delete old keys without any impact on the OSPF session.

[See [Understanding OSPFv2 Authentication authentication](#).]

[See [show \(ospf | ospf3\) interface](#).]

- **OSPF FAPM and interarea support (ACX5448, MX204, MX240, MX480, MX960, MX10003, MX10008, MX2008, MX2010, MX2020, PTX1000, and QFX10002)**—Starting with Junos OS Release 22.4R1, the Flexible Algorithm Prefix Metric (FAPM) is defined to allow an optimal end-to-end path for an inter-area prefix. The Area Border Router (ABR) *must* include the FAPM when advertising the prefix between areas that areas reachable in that given Flex-Algorithm. When a prefix is unreachable, the ABR *must not* include that prefix in the Flex-Algorithm when advertising between areas. The defined FAPM provides inter-area support.

[See [Understanding OSPF Flexible Algorithm for Segment Routing](#).]

[See [show ospf database](#), [show route table](#), [show ted database](#)

- **Support for bootstrapping route-validation database from a local file (cRPD, JRR200, MX204, PTX10008, and QFX10008)**—Starting in Junos OS Release 22.4R1, we support the ability to read validation records from a local binary file and install into the specified named route-validation databases within RPD. This feature implements syntactic and semantic checks on the content of the file to ensure that it is a well-specified set of validation records. If the syntactic and semantic checks fail, the entire file is rejected as a source of validation records. Use the `source-file` statement at the `[edit routing-options validation]` hierarchy level to source route-validation records from a local file source. You can use the `show validation source-file` command to display the properties of a local validation source file.

[See [validation](#).]

- **MVPN feature support with sharding (cRPD, JRR200, MX2020, PTX5000, and QFX10002)**—Starting in Junos OS Release 22.4R1, we support the following features:
 - Multicast virtual private network (MVPN) inactive route query from the main thread to shards
 - Extranet and auto-export support with sharding
 - Interact functions with RT-proxy client and server
 - New data structure to store the inactive route data on the main thread
 - Asynchronous route processing on the main thread

You can use `show mvpn c-multicast` to display the inactive route data stored on the main thread.

[See [rib-sharding](#) and [show mvpn c-multicast](#) .]

Additional Features

Support for the following features has been extended to these platforms.

- **Lightweight leaf device to server loop detection in EVPN-VXLAN fabrics (QFX10002-60C, QFX10002, QFX10008, and QFX10016)**

We support this feature with both enterprise-style and service-provider-style interface configurations.

[See [EVPN-VXLAN Lightweight Leaf to Server Loop Detection](#) and [loop-detect](#).]

What's Changed

IN THIS SECTION

- [EVPN | 147](#)
- [General Routing | 147](#)
- [Network Management and Monitoring | 148](#)
- [Routing Protocols | 148](#)
- [User Interface and Configuration | 148](#)

Learn about what changed in this release for QFX Series Switches.

EVPN

- **Flow-label configuration status for EVPN ELAN services**—The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

General Routing

- **New ARP and NDP packet classification (QFX10002, QFX10008, and QFX10016)**—We've introduced two control plane classes for ARP and NDP packets received over VTEP interface. When your device identifies a packet as ARP or NDP, it performs an ingress port check which verifies whether the VTEP interface receives these packets. If VTEP interface receives the packet, datapath re-writes the control plane class to the newly defined values. Based on this new control plane class, the system performs the remaining packet processing and forwards the packets toward the host path. The system adds a separate DDoS policer to this ARP traffic, which ensures that the ARP traffic is not triggering underlay ARP DDoS violation.

Network Management and Monitoring

- **Junos YANG modules for RPCs include the `junos:command extension statement` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command extension statement` in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the YANG schemas on the local device, the YANG modules for RPCs include the `junos:command extension statement`.

Routing Protocols

- **AR replicators with OISM install multicast states only on the OISM SBD (QFX5130-32CD and QFX5700)**—In an EVPN-VXLAN ERB fabric with many VLANs, QFX5130-32CD and QFX5700 switches running as assisted replication (AR) replicators with optimized intersubnet multicast (OISM) might have scaling issues when they install multicast (*,G) states (with IGMPv2) or (S,G) states (with IGMPv3). As a result, these switches only install these multicast states on the OISM supplemental bridge domain (SBD) VLAN. They don't install these states on all OISM revenue bridge domain VLANs. On those devices, you see multicast group routes only on the SBD in `show multicast snooping route` command output.

[See [OISM and AR Scaling with Many VLANs](#)].

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit

output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml filter` or the `| display json filter` to display command output in XML or JSON format in any release. We've removed these options from the CLI.

Known Limitations

Learn about known limitations in this release for QFX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- QFX5000 platforms can support DF only at port level granularity (In other words, for all evpn instances hosted on an ESI, only one of the Multihomed QFX5000 nodes can be DF). The following configurations are recommended have `df-granularity` (with which QFX5000 platforms seem to have been qualified). Here, few bytes from esi value, instead of vlan-id are used for MOD-based DF. You can also use preference-based DF election. [PR1672383](#)
- On QFX10008, statistics for multicast packets is not as expected as the packets has L2 header stripped during replication in Packet Forwarding Engine because of which it is not forwarded to the next hop. [PR1678723](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 150](#)
- [General Routing | 150](#)
- [Interfaces and Chassis | 151](#)
- [Layer 2 Ethernet Services | 151](#)
- [Layer 2 Features | 152](#)
- [Virtual Chassis | 152](#)

Learn about open issues in this release for QFX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- On all Junos OS platforms, in a scaled scenario when some of the ge/xe/et interfaces are members of aggregated Ethernet (AE) and the Class of Service (CoS) forwarding-class-set configuration is applied with a wildcard for all the physical interfaces and aggregated Ethernet interfaces, it would trigger a Flexible PIC Concentrators (FPC) crash which leads to traffic loss. [PR1688455](#)

General Routing

- On the QFX5100 line of switches, inserting or removing optics on a port might cause a Packet Forwarding Engine Manager CPU spike and an eventual microcode failure. [PR1372041](#)
- VXLAN VNI (multicast learning) scaling on QFX5110 traffic issue is seen from VXLAN tunnel to Layer 2 interface. [PR1462548](#)
- In the platform using indirect next hop, such as Unilist as route next hop type for multiple paths scenario (such as BGP PIC or ECMP), the session fast-reroute might be enabled in Packet Forwarding Engines (PFEs). When the version-id of session-id of INH is above 256, the Packet Forwarding Engine might not respond to session update, which might cause the session-id permanently to be stuck with the weight of 65535 in Packet Forwarding Engine. It might lead Packet Forwarding Engine to have a different view of Unilist against load-balance selectors. Then either the BGP PIC or the ECMP-FRR might not work properly and traffic might be dropped or silently discarded. [PR1501817](#)
- On QFX5110 Virtual Chassis, FPC may disconnect with 24K DHCPv6 relay scaling, after the traffic is stopped. "pfe_listener_disconnect" error messages may be seen. [PR1594748](#)
- Pim VxLAN do not work on TD3 chipsets enabling VxLAN flexflow after Junos OS release 21.3R1. Customers Pim Vxlan or data plane VxLAN can use the version Junos OS Release 21.3R1. [PR1597276](#)
- On all devices running Junos OS or Junos OS Evolved, where this is a high BGP scale with flapping route and the BGP Monitoring Protocol (BMP) collector/station is slow, the rpd process might crash due to memory pressure. [PR1635143](#)
- When MACSEC and VRRP are enabled on QFX5120 VC, MACSEC sessions are flapping at random times. Without VRRP this issue is not seen. [PR1640031](#)

- On all QFX5100 Virtual Chassis platforms, after the reboot, Virtual Chassis port (VCP) ports may not establish a VCP connection and Cyclic Redundancy Check (CRC) errors are also observed.[PR1646561](#)
- On QFX Series platform, v6 logical interface statistics are being derived from the underlying physical interface statistics unlike on PTX Series where they are hardware assisted. Hence, they are not reliable and are at best, guesstimate. [PR1653671](#)
- When the remote end server/system reboots, QFX5100 platform ports with SFP-T 1G inserted may go into a hung state and remain in that state even after the reboot is complete. This may affect traffic after the remote end system comes online and resumes traffic transmission.[PR1665800](#)
- On QFX5100 platforms (both stand-alone and VC scenario) running Junos, occasionally during the normal operation of the device, PFE (Packet Forwarding Engine) can crash resulting in total loss of traffic. The PFE reboots itself following the crash.[PR1679919](#)
- On all Junos platforms, the maximum transmission unit (MTU) on the Integrated Routing and Bridging (IRB) interface is not getting reset when the MTU configuration on IRB or IFD is removed. When MTU configuration is removed from the IRB interface, the internal data structure for the MTU configuration is not getting reset which might affect traffic and it is a rare case.[PR1685406](#)
- Change from Enterprise (EP) to Service Provider (SP) style configuration result in reachability issue in pure L2 (Virtual Extensible LAN protocol) VXLAN setup.[PR1695058](#)

Interfaces and Chassis

- Release note needed [PR1649019](#)
- On QFX5100 Junos OS platforms configured with Virtual Chassis(VC), if a master member is unplugged or forced to power off, the unicast traffic is dropped due to mac-persistence-timer expiry there is a difference in mac addresses between logical aggregated parent interface and member aggregated ethernet(ae) interface. [PR1695663](#)

Layer 2 Ethernet Services

- On QFX5100 and QFX5110, vendor-id format might be incorrect for network ports. This does not impact the ZTP functionality or service. The DHCP client configuration is coming from two places, i.e AIU script and vsdk sandbox. The DHCP client configuration coming from AIU script has the serial Id in vendor id where as the default configuration from sandbox does not have.[PR1601504](#)

Layer 2 Features

- In case of the access-side interfaces used as SP-style interfaces, when a new logical interface is added and if there is already a logical interface on the physical interface, there is 20--50 ms traffic drop on the existing logical interface. [PR1367488](#)

Virtual Chassis

- On Junos EX4600 Virtual Chassis (VC), the master RE reboot and all-members reboot lead to the PFE Manager hogging logs when SFP-T pluggable is installed in. The PFE manager hogging logs has no functionality impact. [PR1685067](#)

Resolved Issues

IN THIS SECTION

- [Chassis Clustering | 153](#)
- [Class of Service \(CoS\) | 153](#)
- [EVPN | 153](#)
- [General Routing | 153](#)
- [Interfaces and Chassis | 157](#)
- [Layer 2 Ethernet Services | 157](#)
- [MPLS | 157](#)
- [Routing Protocols | 157](#)
- [User Interface and Configuration | 158](#)

Learn about the issues fixed in this release for QFX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Chassis Clustering

- GTP control packets might be incorrectly dropped/passed if there is more than one APN IMSI filter configured [PR1673879](#)

Class of Service (CoS)

- The congestion details will be lost as ECN bits in DSCP are cleared after VXLAN decapsulation. [PR1683438](#)

EVPN

- EVPN-VXLAN : after modifying ESI number, the old ESI number is still seen on peer devices. [PR1617068](#)
- The kernel crash would be observed in an EVPN multi-homed scenario [PR1649234](#)
- EVPN DF election will get stuck in the incorrect state. [PR1662954](#)
- BUM traffic might be blackholed for ESI configured CE interface flap [PR1669811](#)
- The ARP/ND entries are not relearnt as expected on the spine with EVPN-VxLAN [PR1677521](#)

General Routing

- On the QFX5120-48Y device, the filter action to change VLAN does not work. [PR1362609](#)
- Traffic loss might be observed on EX4650-48Y and QFX5120-48Y switches when there is a link flap. [PR1634495](#)
- em0 interface speed is reflecting as 10G instead of 1G. [PR1636668](#)
- Multicast traffic received on the INET interface might be dropped. [PR1636842](#)
- On Junos OS QFX Series platforms, traffic might silently get dropped after interface flaps. [PR1645488](#)
- The local-minimum-links feature not working as expected on QFX5100 Virtual Chassis platforms. [PR1649637](#)

- Traffic loss is seen with the virtual-router. [PR1650335](#)
- The MAC address from local CE may not be learned due to the VLAN programming issue. [PR1651827](#)
- The interface might not come up on EX Series platforms. [PR1656540](#)
- FEC link is down after disabling and then enabling interface. [PR1657534](#)
- TOS(DSCP+ECN) bits not getting copied from the Inner L3 header to outer VXLAN header. [PR1658142](#)
- BFD session session-state is showing DOWN while checking micro BFD sessions with authentication in Non-Distributed Mode. [PR1658317](#)
- Traffic loss might be seen when a VxLAN port is recovering from a failure [PR1659533](#)
- On QFX10K Junos platforms configuration of IGMP group range might result in traffic loss. [PR1659732](#)
- MACsec session configured over physical interface might be down when a logical interface configured on physical interface is disabled or deactivated. [PR1660070](#)
- OSPF flow check function violating RFC6864. [PR1660369](#)
- The port LEDs do not light up when 40G physical interfaces are up. [PR1660532](#)
- spmb0 Cell drops on sib 'x' pf 'x' errors seen in QFX10008/QFX10016 platforms without generating any alarms [PR1660699](#)
- High CPU utilization is observed with the Buffer Monitoring Feature on QFX5K platforms. [PR1660750](#)
- CoS might not get applied on VC ports [PR1660787](#)
- The dc-pfe process crash is observed with PTP Transparent clock on QFX platforms [PR1661602](#)
- BUM traffic might loop post adding/removing EVPN-VXLAN FRR configuration [PR1662515](#)
- IPv6 ND packets might be dropped in QFX5100 and QFX5110 platforms [PR1662707](#)
- L2 Multicast traffic loss observed on EX4400 Virtual Chassis platform. [PR1663102](#)
- IS-IS adjacency is not coming up through the Layer 2 domain [PR1663134](#)
- Verification of stats for BFD session is "UP" while checking BFD session [PR1663790](#)
- ALB stats not showing in CLI [PR1663881](#)

- The DHCP offer packets will not be sent to the clients when the DHCP relay agent is configured over Type-5 EVPN [PR1664656](#)
- On QFX5K series platforms, duplicate packets might be seen in the multihomed scenario in an EVPN-VxLAN fabric when unicast ARP packets are received [PR1665306](#)
- Static MACs are not programmed after reboot, resulting in floods of unicast traffic [PR1666399](#)
- PVLAN IGMP packet is forwarded between Isolated ports and also duplicated to primary vlan port (Promiscuous). [PR1667069](#)
- Multihop BFD sessions might remain down in inline mode [PR1667751](#)
- Shaping-rate is not taking 20bytes of overhead into account. [PR1667879](#)
- PFE crash upon receipt of specific genuine packets when sFlow is enabled (CVE-2023-22399). [PR1668330](#)
- Type-5 routes might not get programmed in PFE when the number of active interfaces and Vlan configured is high. [PR1668352](#)
- Route table and multicast add/change requests are got queued in the KRT queue post deleting EVPN enabled configuration followed by the rpd restart [PR1669161](#)
- On specific QFX5k platforms, member links may reduce their configured speed when the other side doesn't have auto-negotiation disabled. [PR1669436](#)
- FPC1 is getting disconnected after ISSU and before switchover while checking ISSU status [PR1669702](#)
- The dcpfe process might generate core-dumps and FPC might crash after line card reboot or switchover [PR1670240](#)
- EVPN multicast traffic may get impacted because of routes getting stuck in the kernel routing table (krt) queue [PR1670435](#)
- Packet drops are seen after flapping or changing a passive monitor interface [PR1671449](#)
- Flow sample packet is not sent to the collector when the destination is an ECMP path [PR1672121](#)
- QFX5120-48YM :: QFX-EVPN_VXLAN: ECN bits not getting copied to vxlan tunnel header at the encap node [PR1672308](#)
- The BFD packets will drop in an EVPN-VxLAN scenario due to incorrect layer3 offset being set in the host path [PR1674116](#)
- The traffic doesn't re-route quickly causing traffic blackholing [PR1676740](#)

- VLAN translation mapping gets deleted when one of the member interface removed from LAG [PR1676772](#)
- Interfaces with QFX-10000-30C and QFX10000-30C-M Line Cards will not work properly [PR1677325](#)
- Firewall functions will not work as expected when egress firewall filter is configured [PR1679574](#)
- ARP resolution will fail on QFX5120 VC [PR1679684](#)
- BFD sessions will remain down in the EVPN-VxLAN scenario [PR1680757](#)
- The PFE process crashes from 21.4R1 version onwards on VMhost platforms [PR1681532](#)
- LLDP neighborship fails to come up with a Private VLAN configuration [PR1681614](#)
- The dcpfe crash seen with PTP configuration on Junos platforms supporting boundary clock [PR1683308](#)
- Traffic loss is seen when MAC flaps between the MC-AE interface and the ICL interface. [PR1683771](#)
- Licenses on the device might become invalid when the device is upgraded from a legacy licensing-based release to an Agile licensing-based release [PR1684842](#)
- The protocol MTU for the IRB interface is not rolled back when the MTU of the IRB or IFD interfaces is modified or deleted [PR1685406](#)
- Traffic statistics verification fails as receiving packet count exceeds specified limit in evpn vxlan multicast scenario [PR1685467](#)
- Traffic via the ICL link to MC-AE peer box gets looped back to the VTEP tunnel on QFX5000 platforms. [PR1687024](#)
- QFX5120 will drop ingress traffic on an I2circuit configured interface on continuous flapping. [PR1687257](#)
- VxLAN configured on access port breaks L2 connectivity with "vxlan encapsulate-inner-vlan" knob [PR1687565](#)
- OVSDB certificate files are not copied from the Master to the Backup [PR1687847](#)
- ARP resolution to the CE port having EP style AE with multiple VLANs would get fail in the EVPN-VXLAN scenario [PR1687861](#)
- The LLDP output packets are not transmitting on the em0 interface of Junos and Junos OS Evolved platforms [PR1688023](#)
- On QFX10008/QFX10016 platforms fails to detect flaps even though the remote device connected has observed flaps [PR1688993](#)

- While verifying "show ethernet-switching global-mac-count | display xml" command "global-mac-count" is not as expected. [PR1689127](#)
- Packet Loss seen on the EVPN-VXLAN spine router router [PR1691029](#)
- Traffic loss is observed when the ECMP path is IRB over AE (IPv4->MPLS) [PR1693424](#)
- BMP EOR is sent with wrong peer address causing BMP failure [PR1695320](#)
- After upgrading to Junos OS Release 20.4R3-S5.3, the dcpfe core file generates and the device becomes unstable. [PR1695943](#)
- QFX 5130 as ARR has bfd sessions stuck in init state with BL(QFX 5120) and AR-LEAF (QFX 5120) [PR1696113](#)
- Traffic drop is observed after deleting or deactivating the logical interface. [PR1697827](#)

Interfaces and Chassis

- VRRP flaps between MC-LAG peers. [PR1579016](#)

Layer 2 Ethernet Services

- The DHCP unicast acknowledge packet might be dropped. [PR1676573](#)

MPLS

- Traffic loss will be seen in an LDP->BGP-LU stitching scenario. [PR1670334](#)
- The rpd crashes very rarely when constructing LDP trace message irrespective of enable/disable LDP traceoptions. [PR1676503](#)

Routing Protocols

- IPv6 Inline BFD sessions are down when neighbor is not resolved. [PR1650677](#)

- Routing Process Daemon (rpd) crashes and restarts when a specific timing condition is hit with BGP configuration. [PR1659441](#)
- Packets getting dropped on the Server leaf in EVPN-VXLAN with OISM. [PR1665791](#)
- High CPU will be seen due to frequent triggering of SPF for ISIS [PR1667575](#)
- MCSNOOPD will be restarted and will again learn the states after core [PR1672488](#)
- Traffic drops due to the generation of the FPC core, which makes the system unstable. [PR1678016](#)
- BGP auto-discovery sessions does not work any more after an interface flap [PR1679950](#)
- Traffic loss is seen when the router in helper mode deletes the route for the router undergoing graceful restart [PR1682506](#)

User Interface and Configuration

- "gethostbyname: Host name lookup failure" is displayed during commit [PR1673176](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 171](#)

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

Upgrading Software on QFX Series Switches

When upgrading or downgrading Junos OS, always use the jinstall package. Use other packages (such as the jbundle package) only when so instructed by a Juniper Networks support representative. For

information about the contents of the jinstall package and details of the installation process, see the [Installation and Upgrade Guide](#) and [Junos OS Basics](#) in the QFX Series documentation.

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **20.3** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.3 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.

NOTE: We recommend that you upgrade all software packages out of band using the console, because in-band connections are lost during the upgrade process.

Customers in the United States and Canada use the following command:

```
user@host> request system software add source/jinstall-host-qfx-5-x86-64-22.4-R1.n-secure-signed.tgz reboot
```

Replace *source* with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the switch.

- For software packages that are downloaded and installed from a remote location:
 - **ftp:// hostname/ pathname**
 - **http:// hostname/ pathname**
 - **scp:// hostname/ pathname** (available only for Canada and U.S. version)

Adding the `reboot` command reboots the switch after the upgrade is installed. When the reboot is complete, the switch displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 20.3 `jinstall` package, you can issue the `request system software rollback` command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a `junos-vmhost-install-x.tgz`.

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot. If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.

NOTE: If you have important files in directories other than `/config` and `/var`, copy the files to a secure location before upgrading. The files under `/config` and `/var` (except `/var/etc`) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the `request vmhost software add <pathname><source>` command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-22.4R1.9.tgz
```

If the Install Package resides remotely from the switch, execute the **request vmhost software add <pathname><source>** command.

For example:

```
user@switch> request vmhost software add ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-22.4R1.9.tgz
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the `show version` command.

```
user@switch> show version
```

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R1.

NOTE: On the switch, use the `force-host` option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the `force-host` option.

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the `show version` command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI `delete chassis redundancy` command when prompted. If GRES is enabled, it will be removed with the redundancy command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the `[edit routing-options]` hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the `request system software add` command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-x86-64-22.4R1.n-secure-signed.tgz
```

For more information about the `request system software add` command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the `request system reboot` command:

```
user@switch> request system reboot
```

NOTE: You must reboot the switch to load the new installation of Junos OS on the switch. To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the `request system software delete <package-name>` command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the `show version` command to verify the version of the software installed.

```
user@switch> show version
```

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the `request chassis routing-engine master` command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state           Backup
    Election priority       Master (default)

Routing Engine status:
  Slot 1:
    Current state           Master
    Election priority       Backup (default)
```

14. Install the new software package using the `request system software add` command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-
x86-64-22.4R1.n-secure-signed.tgz
```

For more information about the `request system software add` command, see the [CLI Explorer](#).

15. Reboot the Routing Engine using the `request system reboot` command:

```
user@switch> request system reboot
```

NOTE: You must reboot to load the new installation of Junos OS on the switch. To abort the installation, do not reboot your system. Instead, finish the installation and then issue the `request system software delete jinstall <package-name>` command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the `show version` command to verify the version of the software installed.
17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the `request chassis routing-engine master` command, see the [CLI Explorer](#).

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state           Master
    Election priority       Master (default)

Routing Engine status:
  Slot 1:
    Current state           Backup
    Election priority       Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- ["Preparing the Switch for Software Installation" on page 168](#)
- ["Upgrading the Software Using Unified ISSU" on page 169](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:

NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (Stateful Replication is Disabled), see [Configuring Nonstop Active Routing on Switches](#) for information about how to enable it.

- Enable nonstop bridging (NSB). See [Configuring Nonstop Bridging on EX Series Switches](#) for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the `request system snapshot` command.

Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in [Installing Software Packages on QFX Series Devices](#).
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where `package-name.tgz` is, for example, `jinstall-host-qfx-10-f-x86-64-22.4R1.n-secure-signed.tgz`.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get lost!  
ISSU: Validating Image  
ISSU: Preparing Backup RE  
Prepare for ISSU  
ISSU: Backup RE Prepare Done  
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...  
Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz completed  
Spawning the backup RE  
Spawn backup RE, index 0 successful  
GRES in progress  
GRES done in 0 seconds  
Waiting for backup RE switchover ready  
GRES operational
```

```

Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 10: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 172](#)
- [What's Changed | 176](#)
- [Known Limitations | 179](#)
- [Open Issues | 180](#)
- [Resolved Issues | 181](#)
- [Migration, Upgrade, and Downgrade Instructions | 186](#)

What's New

IN THIS SECTION

- [EVPN | 172](#)
- [High Availability | 173](#)
- [Interfaces | 174](#)
- [Juniper Extension Toolkit \(JET\) | 174](#)
- [Junos OS API and Scripting | 174](#)
- [Network Address Translation \(NAT\) | 175](#)
- [Content Security \(UTM\) | 175](#)
- [VPNs | 175](#)

Learn about new features introduced in this release for SRX Series devices.

EVPN

- **Pure EVPN Type 5 routes with EVPN-VXLAN (SRX Series and vSRX)**—Starting in Junos OS Release 22.4R1, you can configure pure Type 5 routes in an Ethernet VPN–Virtual Extensible LAN (EVPN-

VXLAN) environment. These devices use EVPN Type 5 routes to advertise IP prefixes for intersubnet connectivity within and across data centers.

[See [Understanding EVPN Pure Type 5 Routes](#) and [ip-prefix-routes](#).]

High Availability

- **Active-active support in Multinode High Availability (SRX5400, SRX5600, and SRX5800 with SPC3, IOC3, SCB3, SCB4, and RE3)**—Starting in Junos OS Release 22.4R1, you can operate Multinode High Availability in the active-active mode with the support of multiple services redundancy groups (SRGs). In this mode, some SRGs remain active on one node and some SRGs remain active on another node. Based on the SRG activeness, you can utilize the bandwidth of both the devices.

With this enhancement, we introduce the following changes:

- Establish multiple active tunnels based on SRG activeness on both nodes.
- Support a failover domain for each SRG.

[See [Multinode High Availability](#).]

- **APBR support in Multinode High Availability (SRX5400, SRX5600, and SRX5800 with SPC3, IOC3, SCB3, SCB4, and RE3, SRX4600, SRX4200, SRX4100, SRX1500, and vSRX)**—Starting in Junos OS Release 22.4R1, Multinode High Availability supports advanced policy-based routing (APBR). APBR classifies a session based on applications, and applies the configured rules to reroute the traffic.

[See [Multinode High Availability](#).]

- **Associating an IPsec VPN configuration with a particular SRG (SRX5400, SRX5600, and SRX5800 with SPC3, IOC3, SCB3, SCB4, and RE3)**—Starting in Junos OS Release 22.4R1, you can selectively and flexibly associate IPsec VPN services to one of the multiple service redundancy groups (SRGs) configured on SRX Series firewalls in Multinode High Availability.

Releases before 22.4R1 supported only SRG0 and SRG1, and SRG1 was associated to IPsec VPN by default. In 22.4R1, SRG1 is not associated to the IPsec VPN service by default. You must associate the IPsec VPN service to any of the SRGs by specifying the following statement:

```
[edit]
user@host# set chassis high-availability services-redundancy-group srg-number managed-
services ipsec
```

[See [IPsec VPN Support in Multinode High Availability](#).]

Interfaces

- **Support for 1-Gbps speed on SRX5K-IOC4-10G card (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 22.4R1, you can change the speed configuration of a 10-Gbps port to operate at 1-Gbps. You can make this change by configuring the speed value as 1-Gbps in the set interfaces <interface-name> gigether-options speed 1g command. After you commit the configuration, the operating speed of the 10-Gbps port changes to 1 Gbps.

[See [Port Speed on SRX5K-IOC4-MRATE](#).]

Juniper Extension Toolkit (JET)

- **Prevent script execution based on current system memory usage (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4300, EX4300-MP, EX4300-VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9208, EX9251, EX9253, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 22.4R1, you can configure the system memory usage threshold above which the device prevents the execution of Juniper Extension Toolkit (JET) scripts. You can configure the start start-options mem-factor statement for individual JET scripts or all JET scripts. The device doesn't execute the script if the system's memory usage exceeds the configured value at the time the script is invoked. This configuration ensures that a device executes only essential scripts when system resources are limited, thereby enabling the device to continue performing all critical network functions.

[See [Configure Script Start Options](#).]

Junos OS API and Scripting

- **Prevent script execution based on current system memory usage (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4300, EX4300-MP, EX4300-VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9208, EX9251, EX9253, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 22.4R1, you can configure the system memory usage threshold above which the device prevents the execution of certain op, event, or SNMP scripts. You can configure the start start-options mem-factor statement for individual scripts or all scripts of a given type. The device doesn't execute the script if the system's memory usage exceeds the configured value at the time the script is invoked. This configuration

ensures that a device executes only essential scripts when system resources are limited, thereby enabling the device to continue performing all critical network functions.

[See [Configure Script Start Options](#).]

Network Address Translation (NAT)

- **Support to retain existing NAT session with destination NAT (SRX Series)**—Starting in Junos OS release 22.4R3, with FQDN based Destination NAT, we support to retain existing NAT sessions even when the DNS resolved IP address changes for the Destination NAT Pool. To retain the existing NAT sessions, you can enable session-retain at [security nat destination pool pool-name] hierarchy. When session-retain is enabled, FQDN based destination NAT sessions remain in the session table, and cleared only upon connection termination from clients, or timeout due to sessions being inactive, or when sessions are explicitly cleared by the user through CLI.

Common DNS cache for NAT and Policy: The NAT and configured policy FQDNs use a single cache. The use of the single cache helps avoid packet drops if you've configured the same FQDN in the policy and NAT.

[See [pool \(Security Destination NAT\)](#).]

Content Security (UTM)

- **Integration of Content Filtering module with JDPI parser**—

Starting in Junos OS Release 22.4R1, Content Filtering (CF) module is integrated with the JDPI parser and the JDPI contexts are used to invoke the CF functionalities.

Content Security (UTM) CF packet and stream plug-ins are added to handle plain traffic.

The Content Filtering (CF) module does not support the notify-mail-sender CLI configuration for mail protocols.

[See [Content Filtering](#).]

VPNs

- **Multiple certificate types support on IKEv2 (MX240, MX480, and MX960 in USF mode, SRX1500, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0 running IKED process)**—Starting in Junos OS Release 22.4R1, you can establish the IKEv2 and IPsec SA tunnels irrespective of the type of certificate used on an initiator and a responder.

To support the multiple certificate types, configure the authentication method as certificates using the certificates option at the [security ike proposal proposal-name authentication-method] hierarchy.

[See [proposal \(Security IKE\)](#).]

- **ACME protocol (SRX Series and vSRX)**—Starting in Junos OS release 22.4R1, we support Automated Certificate Management Environment (ACME) protocol. The ACME protocol allows the enrollment of certificates from Let's Encrypt server or PKI servers.

The SRX Series devices allows usage of certificates issued by Let's Encrypt server or PKI server using ACME.

[See [Understanding Certificate Enrollment with CMPv2](#), [Enroll a CA Certificate](#), [Certificate-Based Validation Using EAP-TLS Authentication](#), and [ACME Protocol](#).]

- **Post-quantum Pre-shared Key (SRX1500, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Starting in Junos OS Release 22.4R1, we support Post-quantum Pre-shared Key (PPK), as defined in the RFC 8784.

The RFC 8784 defines Mixing Pre-shared Keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for post quantum security support in the IKED process through the **junos-ike package** to negotiate quantum secured IKE and IPsec SAs.

The Junos Key Manager (JKM) is introduced to manage different types of quantum keys or PPKs for client applications to make respective infrastructure quantum secured. The IKED process uses the JKM to provide support for quantum secured SAs.

Two out-of-band key retrieval mechanisms are supported to get PPKs:

- **Pre shared key:** You can configure static keys on concerned gateways and do not need share static keys over the Internet.
- **Quantum Key Distribution:** A secure key distribution method based on Quantum Key Distribution (QKD) to generate and distribute keys that are quantum safe. These keys are dynamic.

[See [IPsec VPN Overview](#).]

What's Changed

IN THIS SECTION

- [EVPN | 177](#)
- [High Availability | 177](#)
- [Network Management and Monitoring | 178](#)
- [Platform and Infrastructure | 178](#)
- [User Interface and Configuration | 178](#)

Learn about what changed in this release for SRX Series.

EVPN

- Flow-label configuration status for EVPN ELAN services. The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

High Availability

- In Junos OS releases before 22.4R1, when an SRG changes into Ineligible state due to control-plane failure, a system reboot was required to recover the SRG. Starting in Junos OS Release 22.4R1, the system reboot is not required to recover the SRG, you can restart the control plane process by using the `restart ike-key-management` command.
- Starting in Junos OS Release 22.4R1, you can associate IPsec VPN services to one of the multiple service redundancy groups (SRGs) configured on SRX Series firewalls in Multinode High Availability.

Releases before 22.4R1 supported only SRG0 and SRG1, and SRG1 was associated to IPsec VPN by default. In 22.4R1, SRG1 is not associated to the IPsec VPN service by default. You must associate the IPsec VPN service to any of the SRGs by specifying the following statement:

```
[edit]
user@host# set chassis high-availability services-redundancy-group srg-number managed-
services ipsec
```

[See [Multinode High Availability](#).]

Network Management and Monitoring

- **Junos YANG modules for RPCs include the `junos:command` extension statement (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command` extension statement in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the YANG schemas on the local device, the YANG modules for RPCs include the `junos:command` extension statement.

Platform and Infrastructure

- **from-zone and to-zone are optional when policy match is done for global policies (SRX Series)**—When you use match criteria to troubleshoot traffic problems for global policies, from-zone and to-zone need not be provided while performing the policy match.

[See [show security match-policies](#).]
- **Time zone support for local certificate verification (SRX1500 and SRX5600)**—Starting in this release, when the local certificate verification fails, you can see the time zone for the failed local certificate in the command output and system log messages.

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml` filter or the `| display json` filter

to display command output in XML or JSON format in any release. We've removed these options from the CLI.

VPNs

- **Removal of power mode IPsec Intel QAT option in IPsec VPN (SRX Series)**—We have removed the option `power-mode-ipsec-qat` at `[edit security flow]` hierarchy level from Junos CLI for display. This option is now hidden as it is not recommended to be configured with multiple IPsec VPN tunnels. We continue to use AES-NI in PMI mode for better performance than QAT.

[See [Improving IPsec Performance with PowerMode IPsec.](#)]

Known Limitations

IN THIS SECTION

- [High Availability | 179](#)
- [Chassis Cluster | 180](#)
- [Network Address Translation \(NAT\) | 180](#)

Learn about known limitations in this release for SRX Series devices.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

High Availability

- In Multinode High Availability for vSRX instances in VMWare ESXi environment, configuration of virtual MAC address is not supported in the following statement:

```
[set chassis high-availability services-redundancy-group <number> virtual-ip <id> use-virtual-mac
```

Chassis Cluster

- On SRX Series devices, when you install junos-ike package and configure VPN, the ISSU upgrade from any Junos OS Release prior to 22.4R1 to the Junos OS release 22.4R1 or later is not supported. As a workaround, you can perform the following:
 - Use the minimal downtime procedure to upgrade. See https://supportportal.juniper.net/s/article/SRX-How-to-upgrade-an-SRX-cluster-with-minimal-down-time?language=en_US
 - Deactivate the security ike and VPN configuration before proceeding with ISSU.

[PR1722689](#)

Network Address Translation (NAT)

- While port ranges are configured as part of NAT source pool, port affinity allocation might fail as when the affinity allocation is failed for a flow then the port random allocation is set. The random allocation can allocate any port and the allocation failure can grow.[PR1678563](#)

Open Issues

Learn about open issues in this release for SRX Series devices.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Layer 2 Ethernet Services

- If a client sends a DHCP request packet, and option 55 includes PAD option (0), a DHCP ACK will not be sent back to the client. [PR1201413](#)

Platform and Infrastructure

- IPsec rekey fails when SRX Series devices are configured with kilobyte based lifetime in remote access solution. [PR1527384](#)
- With SSL proxy configured along with Web proxy, the client session might not closed on the device even though proxy session ends gracefully.[PR1580526](#)

- HA active/passive mode on-box logging in logical systems and tenant systems, Intermittently Security log contents of binary log file in LSYS are not as expected. [PR1587360](#)
- The IMAP or IMAPS email permitted counter is not incremented in AAMW email statistics. [PR1646661](#)
- On SRX Series devices with chassis cluster enabled, the redundant Ethernet interface might not go up due to speed mismatch when the redundant Ethernet interface speed is changed after RGO failover. [PR1658276](#)
- The show security firewall-authentication users identifier 1 and show security firewall-authentication users address 10.1.1.1 commands does not display user group information. [PR1659115](#)
- The show services user-identification authentication-table ip-address src_ip command is failing when auth entry boundary testing with auth entry containing maximum length group-name and resource-group-name is used. [PR1665691](#)
- The dns-proxy service is not working when enabled. [PR1688481](#)

VPNs

- Sometimes after manual failover, IKE-SA rekey does not succeed. In order to recover from this scenario, enable DPD with always-send. [PR1690921](#)

Resolved Issues

Learn about the issues fixed in this release for SRX Series devices.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Application Layer Gateways (ALGs)

- The flowd process might stop on SRX5000 line of devices where Central Point is present. [PR1658370](#)
- SIP 200 OK (INVITE) response packets are dropped leading to SIP call failure. [PR1677554](#)
- SIP calls are getting dropped due to NAT failure and SIP ALG is enabled. [PR1686613](#)

Chassis Clustering

- In the MNHA SRG scenario on the IPv6 switching mode, not using Virtual MAC as the source MAC address for G-NDP. [PR1670309](#)
- GTP control packets might be incorrectly dropped or passed if there is more than one APN IMSI filter configured. [PR1673879](#)
- Chassis cluster IP monitoring on the secondary node failed after the system reboot on the SRX Series devices. [PR1691071](#)

Class of Service (CoS)

- The show interfaces queue command output not correctly displaying bps values for throughput higher than 4.25Gbps. [PR1596172](#)

Flow-Based and Packet-Based Processing

- The hardware acceleration flag was not properly updated on RT_FLOW_SESSION_CLOSE logs. Additionally, the values for "Services-offload-sessions" for customers using SPC2's in their SRX5000-Series devices was incorrect. [PR1629216](#)
- The GRE performance acceleration might cause VPLS traffic drop. [PR1661409](#)
- The Routing Engine and Packet Forwarding Engine sync issue with NAT configuration and closed scan session counter issues. [PR1661796](#)
- In SD-WAN the association between VRF instance and VRF group fails for ISSU from Junos OS Release 19.2, 19.3, 19.4, and 21.1 to Junos OS Release 22.2R1. [PR1661935](#)
- vSRX not processing fragmented packets. [PR1668898](#)
- The non-fragmented packets might get dropped on the SPC3 card. [PR1683835](#)
- The flow sessions traversing the IOC2 card would time out early when Express Path is enabled. [PR1688658](#)
- SOF was incorrectly offloading short lived flows leading to early exhaustion of NP memory, reducing overall device performance. [PR1692100](#)

Interfaces and Chassis

- The reth1 interface down and DCD cores files are seen on node1. [PR1657021](#)

Intrusion Detection and Prevention (IDP)

- Execute RSI on SRX5000 line of devices might generate flowd process core files and trigger data plane failover. [PR1665442](#)

J-Web

- All the security policies on Junos SRX Series devices can get deleted while trying to delete any particular policy through J-Web. [PR1681549](#)

Network Management and Monitoring

- High logging rate might cause eventd to increase Routing Engine CPU utilization. [PR1661323](#)

Platform and Infrastructure

- A major alarm DPDK Tx stuck issue of SRX4100 and SRX4200 devices. [PR1626562](#)
- SMS channel down alarm on primary node of HA pair after upgrade. [PR1629972](#)
- Packet loss might be seen on SRX4100 and SRX4200 devices from Junos OS Release 20.2R2. [PR1650112](#)
- Split tunneling feature might not work. [PR1655202](#)
- Archived file which created by non-root user might not include some files under /var/log/ directory. [PR1657958](#)
- After ISSU upgrade completed, RG1 nodes priority remains in CS state and fab interfaces are down. [PR1658148](#)
- The CPU utilization might increase when a user login and logout to the device continuously. [PR1662172](#)
- Cache miss counter increments twice instead of one. [PR1663678](#)
- SRX alarming SMS control channel down without SMS feature configured. [PR1666420](#)
- NG custom APPID fails. [PR1667221](#)
- IPv6 feature not working on SRX5000 line of devices. [PR1668473](#)
- The monitored IP addresses for a redundancy group are reachable despite removing the redundant Ethernet interface from a zone. [PR1668532](#)
- Traffic loss seen due to SPC3 packets getting stuck. [PR1671649](#)

- The forwarding plane stops during HA failover. [PR1672378](#)
- Information about users groups is not displayed completely. [PR1673125](#)
- VPN tunnel will not be established in exclusive client scenario. [PR1674522](#)
- A flowd process stops might occur when AAMW encounters a memory leak. [PR1675722](#)
- NetBIOS traffic is getting dropped post upgrade on the SRX Series devices. [PR1675853](#)
- PKID process stops when validating the certificate chain of a certificate. [PR1679067](#)
- DOD mode on DL interface not working as expected. [PR1680405](#)
- The NSD_CLEAR_POLICY_DNS_CACHE_ENTRY_IP log is not found on the device after keying DNS cache entry unchanged. [PR1684268](#)
- The cluster fabric link will be down post reboot of node or power cycle. [PR1684756](#)
- The unexpected default event-rate value for event mode logging. [PR1687244](#)
- The system might stops when Jflow inactive timeout is configured to be less than previous flow-inactive-timeout + 180 seconds. [PR1688627](#)
- SNMP MIB walk for jnxBoxDescr OID returns incorrect value. [PR1689705](#)
- SRX cluster might fail in a rare scenario when node status changes to disabled state without going through the ineligible state. [PR1692611](#)

Routing Policy and Firewall Filters

- The utility monitor security packet-drop now correctly reports policy-related drops for unified policy. [PR1576150](#)
- Junos OS: SRX Series: Cache poisoning vulnerability in BIND used by DNS Proxy (CVE-2021-25220). [PR1656324](#)
- Security policy state might be invalid on SRX Series devices. [PR1669386](#)
- The rpd process stops whenever it is getting shut down with router reboot, rpd restart, Routing Engine switchover, and software upgrade. [PR1670998](#)
- SRX Series devices stop refreshing the FQDNs used in the security policies and NAT. [PR1680749](#)

Routing Protocols

- The BSR information might not be flooded over NG-MVPN. [PR1664211](#)

- High CPU is seen on the platforms running IPv6. [PR1677749](#)

User Interface and Configuration

- IPsec tunnel will flap post MNHA configuration commit. [PR1669104](#)
- The gethostbyname: Host name lookup failure is displayed during commit. [PR1673176](#)

VLAN Infrastructure

- Traffic stops when the MAC address of a node changes in Layer 2 secure wire SOF. [PR1597681](#)
- The OSPF neighbor does not establish under transparent mode when neighborship across different zone. [PR1599891](#)

VPNs

- Traffic over IPsec tunnels might be dropped during ISSU. [PR1416334](#)
- While verifying show security ipsec next-hop-tunnels command output in device the IPsec SA and NHTB entry is not getting cleared after configuring firewall filter. [PR1432925](#)
- Tunnel bringing up failed from strongSwan when changing the configuration IKE in VR and observed the NO_PROPOSAL_CHOSEN notify error message. [PR1627963](#)
- Severity is unknown at some IPsec SYSLOG messages. [PR1629793](#)
- Packets traversing through a policy based VPN get dropped when PowerMode is enabled. [PR1663364](#)
- IPsec tunnels might flap on SRX Series devices. [PR1665332](#)
- Master encryption password is not accessible when system is in FIPS mode. [PR1665506](#)
- High control plane CPU utilization while the kmd process is stuck after the core file. [PR1673391](#)
- With active/active Multi SRGs, the address pools used by SRGs in the access profile must not overlap. [PR1687654](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 186](#)

This section contains the upgrade and downgrade support policy for Junos OS for SRX Series devices. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 11: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for vMX

IN THIS SECTION

- [What's New | 188](#)
- [What's Changed | 190](#)
- [Known Limitations | 191](#)
- [Open Issues | 191](#)
- [Resolved Issues | 192](#)
- [Upgrade Instructions | 192](#)

What's New

IN THIS SECTION

- [EVPN | 188](#)
- [Junos Telemetry Interface | 189](#)

Learn about new features introduced in this release for vMX.

EVPN

- **EVPN-MPLS E-LAN flow-aware transport (FAT) label load balancing (MX Series, EX9200, vMX)** – Starting in Junos OS Release 22.4R1, you can configure provider edge (PE) devices to use FAT labels in an Ethernet VPN-MPLS (EVPN-MPLS) routing instance, according to Request for Comments (RFC) 6391. PE devices use these labels to load-balance EVPN-MPLS unicast packets across equal-cost multipaths (ECMPs) without performing deep packet inspection of the MPLS payload. This feature supports emulated LAN (ELAN) with single-homing and multi-homing active/standby and active/active topologies and supports the VLAN-based, VLAN-bundle, and VLAN-aware bundle EVPN-MPLS variants.

NOTE: This feature does not support MX Series devices with Advanced Forwarding Toolkit (AFT) cards.

NOTE: On MX Series devices, a configuration where the local PE has a static-flow-label and the remote PE does not have a static-flow-label, the remote PE can process packets without dropping any traffic.

Enabling Load Balancing Using Fat Labels for EVPN Routing Instances:



WARNING: Configuring a flow label or deleting a flow label with the following CLI commands causes a catastrophic event for the routing instance. As a best practice, perform these CLI commands during a maintenance period to avoid network disruptions.

- Configure the `flow-label-static` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level on PE devices to insert FAT flow labels into pseudowire packets sent to remote PE devices.
- Configure the `flow-label` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level on PE devices to signal flow-label capability in the EVPN Layer 2 Attributes Extended Community by setting the flow-label (F) bit in the EVPN Type 3 route.

[See [flow-label](#) and [flow-label-static](#).]

- **Support for Microsoft load-balancing node's static ARP entries with unicast MAC addresses (EX9208, MX-Series, and VMX)**—Starting in Junos OS Release 22.4R1, you can configure a Microsoft load-balancing node's static Address Resolution Protocol (ARP) entries for unicast MAC addresses on integrated routing and bridging (IRB) interfaces. On your provider edge (PE) device, you can create a static ARP entry for the Microsoft load-balancing node's virtual IP address and its unicast virtual MAC address. This static ARP configuration enables your PE devices to flood traffic for the Microsoft load-balancing node's virtual IP address to the virtual MAC address in an EVPN Layer 2 domain or any other Layer 2 domain.

To enable unicast MAC addresses on IRB interfaces, enable the `flood-as-unknown-unicast` option in the `[edit interfaces irb unit <logical-interface-number> family inet address <local-ip-address>/<prefix-length> arp <MSLB-virtual IP address> mac <MSLB-unicast-VMAC>]` hierarchy. The `flood-as-unknown-unicast` option enables flooding of virtual IP addresses and virtual MAC traffic flows from a Microsoft load-balancing cluster.

[See [EVPN User Guide](#).]

Junos Telemetry Interface

- **Event-driven streaming of sensor data for MPLS LSP record route objects (ACX5448, ACX7100, MX204, MX240, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX1000, and vMX)**—Junos OS Release 22.4R1 introduces `ON_CHANGE` notification for streaming MPLS label-switched path (LSP) record route object statistics. Using `ON_CHANGE` mode, data values are not streamed but sent only when data values change. Support includes leaf nodes under the resource path `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/sessions/session/record-route-objects/record-route-object/state/`.

[See [Telemetry Sensor Explorer](#).]

- **Support for gRPC tunnel sessions (MX204, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, PTX-5000, PTX1000, PTX10002, VMX, and QFX5110)**—Starting with Junos OS Release 22.4R1, you can configure a gRPC tunnel session to establish a connection between an external TCP client and a TCP server. The gRPC tunnel session establishes a reverse connection when a TCP client can't reach the TCP server.

To establish a gRPC tunnel session, include the `grpc-tunnel` configuration statement in the `[edit system services]` hierarchy.

[See [gRPC Tunnels Overview](#).]

What's Changed

IN THIS SECTION

- [EVPN | 190](#)
- [Network Management and Monitoring | 190](#)
- [User Interface and Configuration | 191](#)

Learn about what changed in this release for vMX.

EVPN

- Flow-label configuration status for EVPN ELAN services. The output for the `show evpn instance extensive` command now displays the `flow-label` and `flow-label-static` operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

Network Management and Monitoring

- Junos YANG modules for RPCs include the `junos:command extension statement` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command extension statement` in schemas emitted with extensions. The statement defines the CLI command for the corresponding RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions statement` at the `[edit system services netconf yang-modules]` hierarchy level and generate the

YANG schemas on the local device, the YANG modules for RPCs include the `junos:command` extension statement.

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml` filter or the `| display json` filter to display command output in XML or JSON format in any release. We've removed these options from the CLI.

Known Limitations

There are no known limitations in hardware or software in this release for vMX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Open Issues

IN THIS SECTION

- [Platform and Infrastructure | 192](#)

Learn about open issues in this release for vMX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Platform and Infrastructure

- We have seen traffic drop for some streams with IPv6 tunneling on vMX10008 or vMX304 devices. [PR1695669](#)

Resolved Issues

IN THIS SECTION

- [Platform and Infrastructure | 192](#)

Learn about the issues fixed in this release for vMX.

Platform and Infrastructure

- The vMX generate core files. [PR1691459](#)

Upgrade Instructions

You cannot upgrade Junos OS for the vMX router from earlier releases using the `request system software add` command.

You must deploy a new vMX instance using the downloaded software package.

Remember to prepare for upgrades with new license keys and/or deploying Agile License Manager.

Junos OS Release Notes for vRR

IN THIS SECTION

- [What's New | 193](#)
- [What's Changed | 194](#)
- [Known Limitations | 194](#)
- [Open Issues | 194](#)
- [Resolved Issues | 194](#)

What's New

IN THIS SECTION

- [Routing Protocols | 193](#)

Learn about new features introduced in this release for vRR.

Routing Protocols

- **BMP local RIB monitoring support for all RIBs with sharding (ACX Series, cRPD, PTX Series, QFX Series, and vRR)**—Starting in Junos OS Release 22.4R1, you can configure a policy to monitor routing information bases also known as routing table (RIBs) of virtual routers and virtual routing and forwarding instances (VRF). You can specify two separate sets of RIBs in the BGP Monitoring Protocol (BMP), one for monitoring and the other for reporting. With this feature, BMP can filter traffic based on the routes and routing instances.

[See [BGP Monitoring Protocol](#), [loc-rib](#), and [rib-list](#).]

What's Changed

There are no changes in behavior and syntax in this release for vRR.

Known Limitations

There are no known limitations in hardware or software in this release for vRR.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

To learn more about common BGP or routing known limitations in Junos OS 22.4R1, see "[Known Limitations](#)" on page 85 for MX Series routers.

Open Issues

There are no known issues in hardware or software in this release for vRR.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues

Learn about the issues fixed in this release for vRR.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Platform and Infrastructure

- A 802.1Q tagged Ethernet traffic with an expected VLAN ID and with a non-zero 802.1P value ingressing a JRR200 VLAN enabled interface is dropped. [PR1691694](#)

Junos OS Release Notes for vSRX

IN THIS SECTION

- [What's New | 195](#)
- [What's Changed | 197](#)
- [Known Limitations | 198](#)
- [Open Issues | 199](#)
- [Resolved Issues | 199](#)
- [Migration, Upgrade, and Downgrade Instructions | 201](#)

What's New

IN THIS SECTION

- [EVPN | 195](#)
- [Interfaces | 196](#)
- [Network Address Translation \(NAT\) | 196](#)
- [VPNs | 196](#)

Learn about new features introduced in this release for vSRX.

EVPN

- **Pure EVPN Type 5 routes with EVPN-VXLAN (SRX Series and vSRX)**—Starting in Junos OS Release 22.4R1, you can configure pure Type 5 routes in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) environment. These devices use EVPN Type 5 routes to advertise IP prefixes for intersubnet connectivity within and across data centers.

[See [Understanding EVPN Pure Type 5 Routes](#) and [ip-prefix-routes](#).]

Interfaces

- **Increased interface support for AWS instances (vSRX 3.0)**—Starting with Junos OS Release 22.4R1, vSRX 3.0 supports c5n.18XL Amazon Web Services (AWS) instance types for AWS deployments. With this support, you can run vSRX on up to 36 virtual CPUs (vCPUs) and support 16 interfaces.

[See [Requirements for vSRX on AWS.](#)]

Network Address Translation (NAT)

VPNs

- **Multiple certificate types support on IKEv2 (MX240, MX480, and MX960 in USF mode, SRX1500, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0 running IKED process)**—Starting in Junos OS Release 22.4R1, you can establish the IKEv2 and IPsec SA tunnels irrespective of the type of certificate used on an initiator and a responder.

To support the multiple certificate types, configure the authentication method as certificates using the certificates option at the [security ike proposal proposal-name authentication-method] hierarchy.

[See [proposal \(Security IKE\).](#)]

- **ACME protocol (SRX Series and vSRX)**—Starting in Junos OS release 22.4R1, we support Automated Certificate Management Environment (ACME) protocol. The ACME protocol allows the enrollment of certificates from Let's Encrypt server or PKI servers.

The SRX Series devices allows usage of certificates issued by Let's Encrypt server or PKI server using ACME.

[See [Understanding Certificate Enrollment with CMPv2, Enroll a CA Certificate, Certificate-Based Validation Using EAP-TLS Authentication, and ACME Protocol.](#)]

- **Post-quantum Pre-shared Key (SRX1500, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Starting in Junos OS Release 22.4R1, we support Post-quantum Pre-shared Key (PPK), as defined in the RFC 8784.

The RFC 8784 defines Mixing Pre-shared Keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for post quantum security support in the IKED process through the **junos-ike package** to negotiate quantum secured IKE and IPsec SAs.

The Junos Key Manager (JKM) is introduced to manage different types of quantum keys or PPKs for client applications to make respective infrastructure quantum secured. The IKED process uses the JKM to provide support for quantum secured SAs.

Two out-of-band key retrieval mechanisms are supported to get PPKs:

- Pre shared key: You can configure static keys on concerned gateways and do not need share static keys over the Internet.
- Quantum Key Distribution: A secure key distribution method based on Quantum Key Distribution (QKD) to generate and distribute keys that are quantum safe. These keys are dynamic.

[See [IPsec VPN Overview](#).]

What's Changed

IN THIS SECTION

- [EVPN | 197](#)
- [Network Management and Monitoring | 197](#)
- [User Interface and Configuration | 198](#)
- [VPNs | 198](#)

Learn about what changed in this release for vSRX.

EVPN

- Flow-label configuration status for EVPN ELAN services. The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

Network Management and Monitoring

- Junos YANG modules for RPCs include the `junos:command extension statement` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)—The Junos YANG modules that define RPCs for operational mode commands include the `junos:command extension statement` in schemas emitted with extensions. The statement defines the CLI command for the corresponding

RPC. The Juniper [yang](#) GitHub repository stores the RPC schemas with extensions in the `rpc-with-extensions` directory for the given release and device family. Additionally, when you configure the `emit-extensions` statement at the `[edit system services netconf yang-modules]` hierarchy level and generate the YANG schemas on the local device, the YANG modules for RPCs include the `junos:command` extension statement.

User Interface and Configuration

- **Changes to the JSON encoding of configuration data for YANG leaf nodes of type `identityref` (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—If a YANG leaf node is type `identityref`, Junos devices emit the namespace-qualified form of the identity in the JSON encoding of that node. In addition, Junos devices accept both the simple (no namespace) and the namespace-qualified form of an identity in JSON configuration data. In earlier releases, Junos devices only emit and accept the simple form of an identity. Emitting and accepting the namespace-qualified identity ensures that the device can properly resolve the value in the event that the YANG data model defines an identity and a leaf node containing the `identityref` value in different modules.
- **The `file copy` command supports only text-formatted output in the CLI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The `file copy` command does not emit output when the operation is successful and supports only text-formatted output when an error occurs. The `file copy` command does not support using the `| display xml` filter or the `| display json` filter to display command output in XML or JSON format in any release. We've removed these options from the CLI.

VPNs

- **IKEv1 Tunnel establishment not allowed with HSM enabled (vSRX3.0)**—On vSRX 3.0, you can safeguard the private keys used by `pkid` and `iked` processes using Microsoft Azure Key Vault hardware security module (HSM) service. But, you cannot configure Internet Key Exchange version 1 (IKEv1) after enabling the HSM service. If you still try to configure IKEv1 when HSM is enabled, a warning message is displayed.

Known Limitations

Learn about known limitations in this release for vSRX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Chassis Clustering

- The BFD detection interval is 16 seconds. If the detection interval is too large, no BFD down event gets notified by BFDD process to jsrpd process. The jsrpd process is not aware that ICL once goes down since BFD is the single source of MNHA ICL link failure detection. [PR1671622](#)

VPNs

- In case of IKEv2, if the IKE and IPsec security association setup fails in the IKE-SA-AUTH exchange at the initiator end due to authentication failure, it will lead to a situation where in the responder would have already started the IKE and IPsec security association and there would be no delete notification sent from initiator to the responder. To avoid such a scenario, it is recommended to enable dead-peer-detection (DPD) on the responder end which will ensure that the IKE and IPsec SAs gets deleted on the responder. [PR1680885](#)

Open Issues

Learn about open issues in this release for vSRX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Flow-Based and Packet-Based Processing

- With SSL proxy configured along with Web proxy, the client session might not closed on the device even though proxy session ends gracefully. [PR1580526](#)
- When APBR profile is configured as a policy and not attached to a security zone, if there is a failover occurs in between a long-lived ALG (FTP-DATA) session, then the APBR information does not populated in the AppTrack session close log from the backup node. This issue will be seen only when the (FTP) control session and the ALG(FTP-DATA) session are not "Active" on the same node. [PR1688021](#)

Resolved Issues

Learn about the issues fixed in this release for vSRX.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Flow-Based and Packet-Based Processing

- The traffic in the power mode still passthrough when the ingress logic interface is manually disabled. [PR1604144](#)
- Expected TCP sequences not found in ICMP6 core files. [PR1611202](#)
- vSRX not processing fragmented packets. [PR1668898](#)
- Packet loss on GRE tunnel due to improper route look up for tunnel destination. [PR1683334](#)

Platform and Infrastructure

- AMR first session traffic is not copying over multiple paths for IPv6 traffic over IPv6 IPsec tunnel mode. [PR1643570](#)
- Split tunneling feature might not work. [PR1655202](#)
- Cache miss counter increments twice instead of one. [PR1663678](#)
- The evaluation license reappears after deletion and reboot. [PR1664434](#)
- SRX alarming SMS control channel down without SMS feature configured. [PR1666420](#)
- NG custom APPID fails on Junos SRX Series devices. [PR1667221](#)
- vSRX instance in GCP gets stuck in halt state randomly when trying to reboot multiple times. [PR1680874](#)
- ARP will not get learned if reth interface is configured with VLAN. [PR1681042](#)
- The NSD_CLEAR_POLICY_DNS_CACHE_ENTRY_IP log not found on the device after keying DNS cache entry unchanged. [PR1684268](#)

Routing Policy and Firewall Filters

- The utility monitor security packet-drop now correctly reports policy-related drops for unified policy (includes the exact policy that dropped the packet). [PR1576150](#)

User Interface and Configuration

- Configuration changes are not effective if special groups are applied using regex like, apply-groups "\$ {node}"; using explicit special groups names apply-groups [node0 node1]; solves the problem. [PR1660165](#)
- Device is not entering in CLI mode; CLI core files are generated. [PR1673979](#)

VPNs

- The VPN monitoring might not work as expected when PMI enable. [PR1669110](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 207](#)

This section contains information about how to upgrade Junos OS for vSRX using the CLI. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

You also can upgrade to Junos OS Release 22.4R1 for vSRX using J-Web (see [J-Web](#)) or the Junos Space Network Management Platform (see [Junos Space](#)).

Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Releases 17.4, 18.1, 18.2, 18.3,18.4, 19.1, 19.2 and 19.4 is supported.

The following limitations apply:

- Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Release 19.3 and higher is not supported. For upgrade between other combinations of Junos OS Releases in vSRX and vSRX 3.0, the general Junos OS upgrade policy applies.
- The file system mounted on /var usage must be below 14% of capacity.

Check this using the following command:

```
show system storage | match " /var$" /dev/vtbd1s1f
2.7G      82M      2.4G      3% /var
```

Using the request system storage cleanup command might help reach that percentage.

- The Junos OS upgrade image must be placed in the directory `/var/host-mnt/var/tmp/`. Use the `request system software add /var/host-mnt/var/tmp/<upgrade_image>`
- We recommend that you deploy a new vSRX virtual machine (VM) instead of performing a Junos OS upgrade. That also gives you the option to move from vSRX to the newer and more recommended vSRX 3.0.
- Ensure to back up valuable items such as configurations, license-keys, certificates, and other files that you would like to keep.

NOTE: For ESXi deployments, the firmware upgrade from Junos OS Release 15.1X49-Dxx to Junos OS releases 17.x, 18.x, or 19.x is not recommended if there are more than three network adapters on the 15.1X49-Dxx vSRX instance. If there are more than three network adapters and you want to upgrade, then we recommend that you either delete all the additional network adapters and add the network adapters after the upgrade or deploy a new vSRX instance on the targeted OS version.

Upgrading Software Packages

To upgrade the software using the CLI:

1. Download the **Junos OS Release 22.4R1 for vSRX .tgz** file from the [Juniper Networks website](#). Note the size of the software image.
2. Verify that you have enough free disk space on the vSRX instance to upload the new software image.

```
root@vsrx> show system storage
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/vtbd0s1a   694M      433M      206M    68%      /
devfs           1.0K      1.0K        0B    100%     /dev
/dev/md0        1.3G      1.3G        0B    100%     /junos
/cf             694M      433M      206M    68%     /junos/cf
```

```

devfs          1.0K      1.0K      0B      100% /junos/dev/
procfs         4.0K      4.0K      0B      100% /proc
/dev/vtbd1s1e 302M      22K      278M     0% /config
/dev/vtbd1s1f 2.7G      69M      2.4G     3% /var
/dev/vtbd3s2   91M      782K     91M      1% /var/host
/dev/md1       302M      1.9M     276M     1% /mfs
/var/jail      2.7G      69M      2.4G     3% /jail/var
/var/jails/rest-api 2.7G      69M      2.4G     3% /web-api/var
/var/log       2.7G      69M      2.4G     3% /jail/var/log
devfs          1.0K      1.0K      0B      100% /jail/dev
192.168.1.1:/var/tmp/corefiles 4.5G      125M     4.1G     3% /var/crash/
corefiles
192.168.1.1:/var/volatile 1.9G      4.0K     1.9G     0% /var/log/host
192.168.1.1:/var/log 4.5G      125M     4.1G     3% /var/log/hostlogs
192.168.1.1:/var/traffic-log 4.5G      125M     4.1G     3% /var/traffic-log
192.168.1.1:/var/local 4.5G      125M     4.1G     3% /var/db/host
192.168.1.1:/var/db/aamwd 4.5G      125M     4.1G     3% /var/db/aamwd
192.168.1.1:/var/db/secinteld 4.5G      125M     4.1G     3% /var/db/secinteld

```

3. Optionally, free up more disk space, if needed, to upload the image.

```

root@vsrx> request system storage cleanup
List of files to delete:
Size Date      Name
11B Sep 25 14:15 /var/jail/tmp/alarmd.ts
259.7K Sep 25 14:11 /var/log/hostlogs/vjunos0.log.1.gz
494B Sep 25 14:15 /var/log/interactive-commands.0.gz
20.4K Sep 25 14:15 /var/log/messages.0.gz
27B Sep 25 14:15 /var/log/wtmp.0.gz
27B Sep 25 14:14 /var/log/wtmp.1.gz
3027B Sep 25 14:13 /var/tmp/BSD.var.dist
0B Sep 25 14:14 /var/tmp/LOCK_FILE
666B Sep 25 14:14 /var/tmp/appidd_trace_debug
0B Sep 25 14:14 /var/tmp/eedebg_bin_file
34B Sep 25 14:14 /var/tmp/gksdchk.log
46B Sep 25 14:14 /var/tmp/kmdchk.log
57B Sep 25 14:14 /var/tmp/krt_rpf_filter.txt
42B Sep 25 14:13 /var/tmp/pfe_debug_commands
0B Sep 25 14:14 /var/tmp/pkg_cleanup.log.err
30B Sep 25 14:14 /var/tmp/policy_status
0B Sep 25 14:14 /var/tmp/rtsdb/if-rtsdb
Delete these files ? [yes,no] (no) yes

```

```
<
output omitted>
```

NOTE: If this command does not free up enough disk space, see [\[SRX\] Common and safe files to remove in order to increase available system storage](#) for details on safe files you can manually remove from vSRX to free up disk space.

4. Use FTP, SCP, or a similar utility to upload the Junos OS Release 22.4R1 for vSRX .tgz file to `/var/crash/corefiles/` on the local file system of your vSRX VM. For example:

```
root@vsrx> file copy ftp://username:prompt@ftp.hostname.net/pathname/
junos-vsrx-x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz /var/crash/corefiles/
```

5. From operational mode, install the software upgrade package.

```
root@vsrx> request system software add /var/crash/corefiles/junos-vsrx-
x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz no-copy no-validate reboot
Verified junos-vsrx-x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING: This package will load JUNOS 20.4 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.
Saving the config files ...
Pushing Junos image package to the host...
Installing /var/tmp/install-media-srx-mr-vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31 junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31 junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz
Setting up Junos host applications for installation ...
=====
Host OS upgrade is FORCED
```

```

Current Host OS version: 3.0.4
New Host OS version: 3.0.4
Min host OS version required for applications: 0.2.4
=====
Installing Host OS ...
upgrade_platform: -----
upgrade_platform: Parameters passed:
upgrade_platform: silent=0
upgrade_platform: package=/var/tmp/junos-srx-mr-vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-
linux.tgz
upgrade_platform: clean install=0
upgrade_platform: clean upgrade=0
upgrade_platform: Need reboot after staging=0
upgrade_platform: -----
upgrade_platform:
upgrade_platform: Checking input /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz ...
upgrade_platform: Input package /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz is valid.
upgrade_platform: Backing up boot assets..
cp: omitting directory '.'
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK
initrd.cpio.gz: OK
upgrade_platform: Checksum verified and OK...
/boot
upgrade_platform: Backup completed
upgrade_platform: Staging the upgrade package - /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz..
./
./bzImage-intel-x86-64.bin
./initramfs.cpio.gz
./upgrade_platform
./HOST_COMPAT_VERSION
./version.txt
./initrd.cpio.gz
./linux.checksum
./host-version
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK
upgrade_platform: Checksum verified and OK...

```

```

upgrade_platform: Staging of /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz completed
upgrade_platform: System need *REBOOT* to complete the upgrade
upgrade_platform: Run upgrade_platform with option -r | --rollback to rollback the upgrade
Host OS upgrade staged. Reboot the system to complete installation!
WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software rollback'
WARNING:      command as soon as this operation completes.
NOTICE: 'pending' set will be activated at next reboot...
Rebooting. Please wait ...
shutdown: [pid 13050]
Shutdown NOW!
*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY
Shutdown NOW!
System shutdown time has arrived\x07\x07

```

If no errors occur, Junos OS reboots automatically to complete the upgrade process. You have successfully upgraded to Junos OS Release 22.4R1 for vSRX.

NOTE: Starting in Junos OS Release 17.4R1, upon completion of the vSRX image upgrade, the original image is removed by default as part of the upgrade process.

6. Log in and use the show version command to verify the upgrade.

```

--- JUNOS 20.4-2020-10-12.0_RELEASE_20.4_THROTTLE Kernel 64-bit
JNPR-11.0-20171012.170745_fbsd-
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@:~ # cli
root> show version
Model: vsrx
Junos: 20.4-2020-10-12.0_RELEASE_20.4_THROTTLE
JUNOS OS Kernel 64-bit [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs [20171012.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20171012.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20171012.170745_fbsd-builder_stable_11]

```



```

JUNOS py extensions [20171017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20171017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20171012.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20171012.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx Data Plane Crypto Support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20171017.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20171017.110007_ssd-builder_release_174_throttle]

```

Validating the OVA Image

If you have downloaded a vSRX .ova image and need to validate it, see [Validating the vSRX .ova File for VMware](#).

Note that only .ova (VMware platform) vSRX images can be validated. The .qcow2 vSRX images for use with KVM cannot be validated the same way. File checksums for all software images are, however, available on the download page.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 12: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that Juniper Networks has developed over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see [Juniper Flex Program](#).

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.

<https://apps.juniper.net/feature-explorer/>

- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.

<https://prsearch.juniper.net/InfoCenter/index?page=prsearch>

- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.

<https://apps.juniper.net/hct/home>

NOTE: To obtain information about the components that are supported on the devices and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#).

<https://pathfinder.juniper.net/compliance/>

Requesting Technical Support

IN THIS SECTION

- Self-Help Online Tools and Resources | 210
- Creating a Service Request with JTAC | 211

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>

- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net/>
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

Revision History

8 August 2024—Revision 16, Junos OS Release 22.4R1.

4 June 2024—Revision 15, Junos OS Release 22.4R1.

23 May 2024—Revision 14, Junos OS Release 22.4R1.

22 February 2024—Revision 13, Junos OS Release 22.4R1.

15 February 2024—Revision 12, Junos OS Release 22.4R1.

31 August 2023—Revision 11, Junos OS Release 22.4R1.

20 July 2023—Revision 10, Junos OS Release 22.4R1.

1 June 2023—Revision 9, Junos OS Release 22.4R1.

4 May 2023—Revision 8, Junos OS Release 22.4R1.

19 April 2023—Revision 7, Junos OS Release 22.4R1.

17 March 2023—Revision 6, Junos OS Release 22.4R1.

16 March 2023—Revision 5, Junos OS Release 22.4R1.

23 February 2023—Revision 4, Junos OS Release 22.4R1.

12 January 2023—Revision 3, Junos OS Release 22.4R1.

23 December 2022—Revision 2, Junos OS Release 22.4R1.

16 December 2022—Revision 1, Junos OS Release 22.4R1.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2024 Juniper Networks, Inc. All rights reserved.