Release Notes

Published
2026-01-08

Junos OS Evolved Release 23.4R1

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 23.4R1.

For more information on this release of Junos OS Evolved, see Introducing Junos OS Evolved.

Table of Contents

Junos OS Evolved Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.4R1 for ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509 devices. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New in 23.4R1-S1

IN THIS SECTION

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

Hardware

- **New ACX7332 router (ACX Series)**—Starting in Junos OS Evolved Release 23.4R1-S1, we introduce the Juniper Networks® ACX7332 Cloud Metro Router, an extended temperature-rated (E-Temp) platform from the ACX7300 series that supports a variety of deployment scenarios. With a compact 3-RU semi-modular form factor, it offers an aggregation solution that gives cloud providers and service providers the performance and scalability needed as networks grow.

  The ACX7332 router provides 1-Gigabit Ethernet (GbE) through 400GbE port flexibility and a throughput of 2.4 Tbps. The router has a fixed FPC with thirty-two 25GbE and eight 100GbE ports, dual Routing Engines, three bays for pluggable interface modules, redundant power supply modules (AC or DC), and four fan trays (two fans per tray).

  The ACX7332 router supports the following pluggable FPCs:

  - ACX7K3-FPC-2CD4C—Two 400GbE and four 100GbE ports
  - ACX7K3-FPC-16Y—Sixteen 50GbE ports

  The ACX7332 router runs Junos OS Evolved and provides several capabilities that include support for the latest protocol and traffic engineering technologies, enhanced security, and precision timing for mobile backhaul applications. These capabilities and features enable you to create converged, virtualized, and automated architectures to address the rapid growth of 5G, IoT, and the cloud.

  Table 1: ACX7332 Feature Support

Feature	Description
Authentication and Access Control	• Support for 802.1X authentication on Layer 2 and Layer 3 interfaces. [See 802.1X Authentication on Layer 2 Interfaces.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

Chassis	• Supports two Routing Engines, one Control Board, one Forwarding Engine Board (FEB), one fixed FPC, and three removable FPCs chassis supports:
	  • Platform FEB and FPC FRU presence and power-up.
	  • Infrastructure databases and services.
	  • Power management.
	  • Environment monitoring and cooling.
	  • System LED behavior.
	• Platform resiliency support for device chassis, RCB, PSM, fan tray, input, and output devices.

Class of service	• Support for classification and rewrite rules of all types (Inet-Prec/DSCP/DSCP-v6/IEEE-802.1p/IEEE-802.1ad) at the logical interface level. Supports logical interfaces classification and rewrite rules for MPLS, VPLS, Layer 3 VPN, Layer 2 circuit, CCC, IRB, and EVPN. [See Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview.] • Support for port shaping and scheduling with eight VoQ queues per port and two scheduling priority levels (strict-high and low). Supports multiple strict-high priority queues (RR scheduling), multiple low-priority queues (WFQ scheduling), low latency queues (LLQ), and default deep buffers. [See Schedulers Overview for ACX Series Routers and Shared and Dedicated Buffer Memory Pools on ACX Series Routers .] • Support for hierarchical class of service (CoS). Hierarchical CoS support for Layer 3 VPN, Layer 2 VPN, Layer 2 Circuit, VPLS, and EVPN services. [See Hierarchical Class of Service in ACX Series Routers.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

Dynamic Host Configuration Protocol	• DHCP server and DHCP relay configuration for IPv4 and IPv6 services.
	
	[See DHCP Overview.]
	
	• DHCP relay deployment of EVPN over MPLS, which includes:
	
	  • Edge-routed bridging (ERB)—Edge model where DHCP clients are connected and relayed in network leaf devices. The spine PEs do not perform DHCP relay functions, and the routers support transit spine functionality running protocols such as BGP for integrated routing and bridging (IRB).
	
	• The following functionalities:
	
	  • EVPN over MPLS Ethernet-LAN
	
	  • DHCPv4 and DHCPv6 relay options
	
	  • Stateless forward-only mode for DHCP relay over VPN
	
	  • Anycast IP address with IRB for a relay source
	
	  • Client VRFs only
	
	• DHCPv4 and DHCPv6 relay agent support for MC-LAG. DHCP relay agent support includes:
	
	  • DHCPv4 and DHCPv6 stateless forward-only option on Layer 3 static interfaces over MC-LAG.
	
	  • DHCPv4 and DHCPv6 stateless forward-only option on IRB interfaces over MC-LAG.
	
	  • DHCPv4 and DHCPv6 forward-snooped-clients on dual-stack configurations.
	
	[See DHCP Relay Agent and Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent.]
	
	[See DHCP Relay Agent in EVPN-MPLS Network.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

EVPN	• Support for the following EVPN-MPLS features on MAC-VRF instances:
	  • L2 flooding for broadcast, unknown unicast, and multicast (BUM) traffic
	  • Split-horizon between core interfaces
	  • Data plane and control plane MAC learning and aging, and static MAC
	  • MAC movement and MAC mobility on control plane only
	  • MAC limiting and MAC learning
	  • Input and output VLAN maps using normalization on user-to-network interfaces (UNIs)
	  • Aggregated Ethernet interfaces used for UNIs and network node interfaces (NNIs)
	  • Physical interfaces for VLAN tagging, stacked VLAN tagging, flexible VLAN tagging, and extended VLAN bridges using EVPN-MPLS as a service
	  • Ethernet bridge mode for logical UNIs
	  • VLAN ID lists, native VLAN ID supported logical UNIs, and priority-tagged logical interfaces
	  • Underlay networks with ECMP and Fast Reroute (FRR)
	  • Control-word support for EVPN
	  • EVPN Proxy Address Resolution Protocol (ARP) and ARP suppression
	  • EVPN-ELAN over segment routing
	  [See EVPN Feature Guide.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Virtual private wire service (VPWS) with EVPN signaling mechanisms and flexible cross-connect support.
	
	[See Overview of VPWS with EVPN Signaling Mechanisms.]
	
	• EVPN E-LAN active-active multihoming with EVPN aliasing support for ESI LAG.
	
	[See Example: Configuring EVPN Active-Active Multihoming.]
	
	• All-active multihoming redundancy in both Ethernet VPNvirtual private wire service (EVPN-VPWS) and EVPN-VPWS with flexible cross-connect.
	
	[See Overview of Flexible Cross-Connect Support on VPWS with EVPN.]
	
	• EVPN VPWS multihoming all-active forsegment routing over MPLS
	
	[See Overview of VPWS with EVPN Signaling Mechanisms.]
	
	• Entropy and flow label for EVPN-ELAN
	
	[See Configuring Entropy Labels.]
	
	• Support for the following EVPN-MPLS features:
	
	  • IRB with IPv4 and IPv6 addresses
	
	  • IRB virtual gateway
	
	  • IRB anycast gateway
	
	  • IRB with static mac
	
	  • EVPN asymmetric Type 2 and symmetric Type 5 routes
	
	  • EVPN E-LAN over BGP-LU
	
	  • EVPN proxy ARP and ARP suppression, and NDP and NDP suppression

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• EVPN routing policies • Ingress virtual machine traffic optimization (VMTO) [See EVPN with IRB Solution Overview, Anycast Gateways, Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes, Understanding EVPN Pure Type 5 Routes, EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression, Ingress Virtual Machine Traffic Optimization, and Routing policies for EVPN.] • Support for the following EVPN-VPWS features: • EVPN-VPWS FXC VLAN unaware service • EVPN-VPWS FXC VLAN aware service • EVPN-VPWS over segment routing • Single homing and all active multihoming support • Flow-aware transport (FAT) pseudowire labels • Entropy labels [See Overview of VPWS with EVPN Signaling Mechanisms.]
Firewall filters	• Support for firewall filters and policers. You can configure firewall filters with packet match conditions for the bridge domain, IPv4, IPv6, CCC, and MPLS families. In addition to packet match conditions, the count, discard, log, syslog, and policer actions are supported. [See Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview.] • Filter-based forwarding (FBF). [See Filter-Based Forwarding Overview.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• Firewall filter protocols: MPLS, CCC, virtual private LAN service (VPLS), and ANY. [See Firewall Filters Overview, Filter-Based Forwarding Overview, Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address, and Guidelines for gRPC and gNMI Services.]
High availability	• VRRP for IPv4 and IPv6. [See VRRP and VRRP for IPv6 Overview.] • BFD over label-switched paths (LSPs) or RSVP-based LSPs in a centralized mode. [See Bidirectional Forwarding Detection (BFD) for MPLS.] • High availability on these routers are supported at the hardware level. Graceful routing engine switchover is not supported in this release. • Support for loop-free alternate (LFA) routes for OSPF and IS-IS. LFA enables IP fast-reroute capability for OSPF and IS-IS. [See Loop-Free Alternate Routes for OSPF Overview and Understanding Loop-Free Alternate Routes for IS-IS.] • BFD-triggered fast reroute for unicast next hops. [See Bidirectional Forwarding Detection (BFD) for MPLS, session-id-change-limiter-indirect, and no-bfd-triggered-local-repair.]
Interfaces	• The ACX7332 router provides 1GbE through 400GbE port flexibility and a throughput of 2.4 Tbps. [See Port Speed.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Support for 6xQSFPDD and 16xSFP56 FPC line cards. The 6xQSFPDD FPC Line Card has two QSFPDD ports (Port 0 and 2) and four QSFP ports (Port 1, 3, 4 and 5). The 16xSFP56 FPC line card has 16 SFP56 ports (Port 0 to 15). Slot 1 and 2 supports 10-Gbps, 25-Gbps, and 50-Gbps speeds. Slot 3 supports 1-Gbps, 10-Gbps, and 25-Gbps speeds.
	
	[See Port Speed.]
	
	• Support for LACP link protection. We support 1:1 and N:N link protection.
	
	[See link-protection.]
	
	• Resiliency support for ASIC error and CM infra. Resiliency only supports logging and detection and not action.
	
	• Features supported for unnumbered interfaces:
	  • Bidirectional Forwarding Detection (BFD)
	  • BGP labeled unicast
	  • Ethernet VPN virtual private wire service (EVPN-VPWS)
	  • IS-IS protocol adjacency
	  • Label Distribution Protocol (LDP)
	  • Layer 2 VPN and Layer 2 circuit
	  • Layer 3 VPN
	  • Qualified next hop
	  • RSVP-TE
	  • Static subnet route
	  • Source Packet Routing in Networking (SPRING) over OSPFv2

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• SPRING-TE • Segment routing with MPLS • Static LSP • Source Packet Routing in Networking (SPRING) over OSPFv2 • SPRING-TE • Segment routing with MPLS • Static LSP [See Configure unnumbered Interfaces.]
Junos Telemetry Interface (JTI)	• Logical subinterface and Packet Forwarding Engine drop, pipe, and line-card counter sensor support for JTI. [See Junos YANG Data Model Explorer.] • Support for telemetry interfaces.
Layer 2 features	• Ethernet ring protection switching (ERPS) with G.8032 version 2. [See Understanding Ethernet Ring Protection Switching Functionality .] • Support for the following advanced Layer 2 (L2) features: • Bridge domain without a `vlan-id` *number* statement • Bridge domain with the `vlan-id` value set to None • Bridge domain with a single VLAN ID

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• MAC learning, ageing, and limiting
	• Single-learning domain per bridge domain
	• Ethernet service types:
	  • E-Line with these AC interface types: port, VLAN, Q-in-Q, VLAN list, and VLAN maps
	  • E-Line
	  • E-LAN
	  • E-Access
	  • E-Transit
	• LLDP
	• LACP
	• IRB interface
	• Link aggregation group (LAG) support with the following hashing algorithms:
	  • For family `multiservice`, destination and source MAC addresses
	  • For family `inet`, Layer 3 and Layer 4
	  • For family `inet6`, Layer 3 destination and source addresses
	  • For family `inet6`, Layer 4 destination and source ports

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Encapsulation types:
	  • `extended-vlan-bridge`
	  • `vlan-bridge`
	  • `ethernet-bridge`
	• Q-in-Q tunneling
	[See Understanding Layer 2 Bridge Domains and Q-in-Q Tunneling.]
	• Disable local switching in bridge domains.
	[See Configuring MAC Address Flooding and Learning for VPLS.]
	• Layer 2 protocol tunneling (L2PT) to send L2 protocol data units (PDUs) across the network and deliver them to devices that are not part of the local broadcast domain.
	• Storm control.
	[See Understanding Storm Control.]
	• Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP).
	[See Spanning-Tree Protocol Overview.]
	• MAC move limit and multiple trunk ports, virtual private LAN service (VPLS), and EVPN networks.
	[See Understanding MAC Limiting and MAC Move Limiting.]
	• Layer 2 Control Protocol (L2CP) BPDUs are transparently forwarded in hardware unless a specific protocol is configured on the incoming interface. This feature helps you to configure and enable L2PT.
	• VLAN sensor support.

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	[See Telemetry Sensor Explorer.]
	
	[See Understanding Layer 2 Bridge Domains on ACX Series and Q-in-Q Tunneling on ACX Series, Bridging and VLANs, and Configuring MAC Address Flooding and Learning for VPLS .]
	
	• Multichassis link aggregation groups (MC-LAGs). The following Layer 2 features are available on MC-LAGs:
	  • Layer 2 bridging for active-active and active-standby modes
	  • Layer 2 unicast with and without IGMP snooping
	  • Layer 3 unicast with and without IGMP snooping
	  • Layer 2 multicast with and without IGMP or MLD snooping
	  • Layer 3 multicast with and without IGMP or MLD snooping
	
	[See Understanding Multichassis Link Aggregation Groups.]

Layer 2 VPN	• Support for VPLS. The router supports a single VLAN for each virtual switch routing instance type. Junos OS Evolved does not support the `family vpls` option. To configure VPLS, configure the `instance-type virtual-switch` statement at the [`edit routing-instances` *`routing-instance-name`*] hierarchy level. If you configure normalized VLANs, either by not configuring VLAN IDs or by including the `vlan-id none` statement, then you must include the `service-type single` statement at the [`edit routing-instances` *`routing-instance-name`* `protocol vpls`] hierarchy level. [See Introduction to Configuring VPLS.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Support for control word and load-balancing capabilities using entropy and flow-aware transport of pseudowires (FAT) flow labels, across LDP-signaled pseudowires for virtual private LAN service (VPLS).
	
	[See control-word , Configuring Entropy Labels, and FAT Flow Labels Overview.]
	
	• Support for redundant pseudowires for virtual private LAN service (VPLS). The router supports VPLS with LDP hot-standby, cold-standby model, and without BFD or CFM trigger.
	
	[See Redundant Pseudowires for Layer 2 Circuits and VPLS.]
	
	• IRB support for VPLS.
	
	[See Configuring VPLS and Integrated Routing and Bridging.]
	
	• Layer 2 VPN and L2 circuit support:
	
	  • L2 circuit—Targeted LDP signaling pseudowires and interoperability between different types of supported attachment circuit for L2 circuit
	
	  • L2 VPN circuit—BGP signaling
	
	• MPLS fast reroute (FRR) on IGP, circuit attachment types (port, VLAN, and Q-in-Q tunneling), control word, pseudowire circuit on aggregated Ethernet interfaces, indirect next hops and composite next hops, pipe and uniform mode time-to-live (TTL), Tag Protocol Identifiers (TPIDs), and VLAN map on pop, push, or swap.
	
	[See Understanding Layer 2 VPNs and Understanding Layer 2 VPNs and Configuring Interfaces for Layer 2 Circuits.]
	
	• Flow-aware transport for pseudowires (FAT) label and entropy label support for Layer 2 circuit and Layer 2 VPN.
	
	[See Configuring Entropy Labels, and FAT Flow Labels Overview.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

Layer 3 features	• Support for the following Layer 3 features:
	  • IP forwarding and exception packet handling
	  • IEEE 802.1Q (VLAN trunk) on IRB interfaces
	  • Address Resolution Protocol (ARP), neighbor discovery, unicast reverse-path forwarding, and ECMP
	  • LPM and fragmentation handling, ICMP redirect handling, VLAN tagging modes, neighbor solicitation, and Interface-based routing
	  • Longest prefix match
	  • Exception packets handling
	  • VLAN tagging modes
	  • Integrated routing and bridging (IRB)
	  • IPv4 and IPv6
	The router also supports interior gateway protocols such as OSPF, IS-IS, RIP, and ECMP for IPv4 and IPv6. [See Configure ICMP Features, Enabling VLAN Tagging, Neighbor solicitation, Understanding Unicast RPF (Routers), OSPF Overview, IS-IS Overview, and RIP User Guide.]

Layer 3 VPN	Support for the following Layer 3 VPN features:
	NOTE: VT interface-based Layer 3 VPN is not supported. Layer 3 VPN ping is supported only with the `vrf-table-label` configuration.
	• IP-VPN services:
	  • Instance-type virtual routing and forwarding (VRF) and virtual-router
	  • All control plane configuration options

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• Per-prefix and per-table label signaling • Layer 3 VPN support with ECMP • BGP polices support for different Layer 3 VPN use cases (for example, full mesh VPN, hub-spoke VPN, management VPN, and leaking routes) • Layer 3 VPN with vrt-table-label mode • Layer 3 VPN with chained composite mode • Import or export of routes across local VRF and global VRF **NOTE**: Table next hop is not supported. • Inter-autonomous system (inter-AS) options A, B, and C **NOTE**: Inter-AS option B can be deployed in hierarchical network design within a single IGP AS. • PE to CE routing protocols—Static, eBGP, IS-IS, OSPF, and RIP • IPv6 Provider Edge (6PE)/IPv6 VPN routing over MPLS (6VPE) with PE-CE routing-static and PE-CE BGPv6 [See Layer 3 VPNs User Guide for Routing Devices.]
MACsec	Supports Media Access Control Security (MACsec). [See Understanding Media Access control Security (MACsec).]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

MPLS	• Support for the following MPLS features:
	  • IP/MPLS infrastructure feature set for the L3 VPN service
	  • Basic BGP control plane features such as LDP-DOD, CSPF, and single-area CSPF
	  • MPLS label stack
	  • MPLS protections:
	  • Fast reroute (FRR) and Make-before-break (MBB)
	  • Link protection
	  • Node protection
	  • Label-switching router (LSR)
	  • Shared Risk Link Group (SRLG) for MPLS
	  • RSVP label-switched path (LSP) over IPv4 includes refresh reduction
	  • Label Distribution Protocol (LDP) LSP over IPv4
	  • RSVP 1:1
	  • RSVP-Traffic Engineering (RSVP-TE)
	  • LDP over RSVP
	  • Inter-autonomous system LSP intra-area LSP
	  • [See MPLS Applications User Guide.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Support for MPLS LSP statistics and RSVP-TE auto-bandwidth features. Support includes:
	  • MPLS LSP statistics for the following LSP types:
	  • LDP-signaled LSPs
	  • RSVP-signaled LSPs
	  • Static LSPs
	  • Bypass LSPs
	  • Container LSPs
	  • RSVP-TE auto-bandwidth
	[See LSP Overview, LDP Overview, RSVP Overview, and Configuring Optimized Auto-bandwidth Adjustments for MPLS LSPs.]

Multicast	• Support for multicast snooping in a VPLS for the following protocols:
	  • IGMPv1, IGMPv2, and IGMPv3 snooping in VPLS
	  • MLDv1 and MLDv2 snooping in VPLS
	  • IGMP and MLD snooping in VPLS with integrated routing and bridging (IRB)
	  • Protocol Independent Multicast support over VPLS with IRB
	[See Multicast Snooping for VPLS.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Support for Layer 2 multicast-related features, including IGMP and MLD snooping. You can configure IGMP snooping with IGMPv1, IGMPv2, and IGMPv3, which includes support for:
	  • IGMP snooping in bridge domains
	  • IGMP snooping with integrated routing and bridging (IRB) configured in bridge domains
	  • MLD snooping in bridge domains
	  • MLD snooping with IRB configured in bridge domains
	  [See IGMP Snooping Overview and Understanding MLD Snooping.]
	• Support for IPv4 multicast for Layer 3. You can configure IGMP snooping with IGMPv2 and IGMPv3, which includes support for the following:
	  • Anycast RP
	  • IGMP filter
	  • IGMP querier
	  • Protocol Independent Multicast source-specific multicast (PIM SSM)
	  • PIM sparse mode (PIM SM)
	  [See IGMP Snooping Overview.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• Support for BGP MVPN. BGP over MPLS MVPN (also known as "next generation," or "NG," MVPN) running on multipoint LDP provider tunnels, where BGP MVPN is the intra-AS and PIM-SM and multipoint LDP point-to-multipoint (P2MP) tunnels is the data plane. [See Multiprotocol BGP MVPNs Overview.] • Multicast with IGMP or MLD snooping within VLANs for EVPN-MPLS. [See Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment.]
Network management and monitoring	• Support for port mirroring with analyzers and encapsulated remote Switch Port Analyzer (ERSPAN). [See Port Mirroring and Analyzers.] • Support for SNMP.

Operations, Administration, and Maintenance	• Support for OAM. You can configure connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM. You can also configure the following features of link-fault management (LFM):
	
	  • Discovery
	
	  • Link monitoring
	
	  • Remote fault detection
	
	[See ITU-T Y.1731 Ethernet Service OAM Overview and Introduction to OAM Link Fault Management (LFM).]
	
	• Support for IEEE 802.1ag OAM CFM.
	
	• Support for IEEE Standard 802.3ah and 802.1ag for OAM CFM down and up maintenance association end points (MEPs) over virtual private LAN service (VPLS).

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• Support for IEEE Standard 802.3ah and 802.1ag for OAM CFM up MEPs over EVPN. [See IEEE 802.3ah OAM Link-Fault Management Overview and IEEE 802.1ag OAM Connectivity Fault Management Overview.] • Support for CFM and performance monitoring (Y.1731) protocols over Ethernet interfaces for `bridge` and `inet` services. [See Ethernet OAM Connectivity Fault Management.] • Support for native Y.1731 operational state sensors to provide statistics such as frame loss ratio, frame delay, frame delay variation, and availability for Y.1731 performance monitoring.
Protection against DDoS attacks	• Support for control plane distributed denial of service (DDoS) protection. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]

Routing protocols	• Layer 3 and routing protocols IPv4, IPv6, BGP, IS-IS and ARP streaming sensor support using gRPC services.
	
	• Support for unicast reverse path forwarding (unicast RPF):
	
	  • Support for loose and strict mode
	
	  • Support for IPv4 and IPv6
	
	[See Understanding Unicast RPF (Routers).]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

	• Support for BGP flow specification (BGP flowspec).
	  • The following match conditions are not supported:
	  • Fragment for IPv6
	  • Packet length
	  • Port
	  • Source and destination prefix with offset
	• The following actions are not supported:
	  • Community
	  • Next-term
	  • Routing instance
	  • Sample
	  • Traffic marking
	  [See Understanding BGP Flow Routes for Traffic Filtering.]
	• Support for configuring interface groups in BGP flowspec filters.
	  [See Understanding BGP Flow Routes for Traffic Filtering and Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic.]
	• BGP PIC edge support for `inet` and MPLS VPNs. The following features are not supported:
	  • Session-based repair
	  • BGP PIC over LDP over RSVP tunnel
	  • BGP PIC over SR-MPLS

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• BGP PIC with RSVP • BGP-LU with PIC • BGP PIC edge protection for Layer 2 services • Protection with multilink failure [See Configuring BGP PIC Edge for MPLS Layer 3 VPNs and Use Case for BGP PIC for Inet.]
	• Support for entropy label for LDP, RSVP, L3VPN, and BGP-LU. [See Entropy label support for BGP Labeled Unicast (BGP-LU) and Configuring Entropy Labels.] • Support for BGP transport address family or BGP Classful Transport (BGP-CT) includes: • Service mapping over colored transport tunnels (RSVP, IS-IS flexible algorithm) to transport classes and map service routes over an intended transport class. The transport tunnels can span multiple domains (ASs or IGP areas). • Network slicing and interoperability between network domains. • IPv6 and segment routing–traffic engineered (SR-TE) color-only support. • IPv6 and BGP service routes with a color-only mapping community. • Enhanced transport-class configuration to provide precise resolution. [See use-transport-class.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description

Services Applications	• RFC 2544-based benchmarking tests. Support for Layer 2 reflection (bridge, L2CKT, L2VPN, EVPN-VPWS, EVPN-FXC, EVPN-MPLS, and VPLS), with family `ccc` or family `ethernet-switching` and for Layer 3 reflection (IPv4, L3VPN) with family `inet`. RFC 2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before activation of the service. The tests measure throughput, latency, frame loss rate, and back-to-back frames.
	
	[See RFC 2544-Based Benchmarking Tests for ACX Routers Overview.]
	
	• RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. You can configure the TWAMP monitoring service, which sends out probes to measure network performance. TWAMP is often used to check compliance with service-level agreements. The support for this service is limited to the following features:
	
	  • IPv4 and IPv6 source and target addresses for clients, control connections, and test sessions
	
	  • Probe statistics and history
	
	  • Control and test session status
	
	  • Test session probe generation and reception, as well as reflection
	
	  • Timestamps set by software (the Routing Engine or the Packet Forwarding Engine) or the hardware
	
	  • Error reporting through system log messages only
	
	  • Unauthenticated mode only
	
	[See Understand Two-Way Active Measurement Protocol.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• Support for sFlow monitoring (ingress). [See sFlow Monitoring Technology and Understanding How to Use sFlow Technology for Network Monitoring.] • Inline active flow monitoring support for IPFIX and v9 export formats. We support ingress and egress sampling of IPv4 and IPv6 traffic on aggregated Ethernet and IRB interfaces and interfaces mapped to non-default VRFs, for both the IPFIX and version 9 export formats. You can configure up to four IPv4 collectors for inline active flow monitoring. See Understand Inline Active Flow Monitoring.

Source Packet Routing in Networking (SPRING) or segment routing	• Support for the following segment routing features:
	
	  • Segment routing global block (SRGB) for OSPF, IS-IS, and fast reroute.
	
	  • Metro Ethernet services over segment routing infrastructure
	
	  • Segment routing services: L3VPN, IPv6 VPN Provider Edge (6VPE) , IPv6 Provider Edge (6PE), L2VPN, L2 circuit, and BGP-VPLS
	
	  • Static segment routing (node segment, prefix segment, adjacency, and anycast segments) for OSPF and IS-IS
	
	  • Topology-independent loop-free alternate (TI-LFA) with segment routing for OSPF and IS-IS

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	• Unnumbered interfaces support for segment routing with OSPF • Support for IPv6 L3VPN over IPv6 SR-TE and IPv6 underlay • Support for flexible algorithm in OSPF and IS-IS for segment routing traffic • Interoperability of segment routing with LDP • SPRING support for SR-TE • Support for BGP link-state distribution with SPRING extensions [See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS, Understanding Source Packet Routing in Networking (SPRING), Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING, Configure Unnumbered Interfaces, Understanding Static Segment Routing LSP in MPLS Networks, Link-State Distribution Using BGP Overview, Understanding OSPF Flexible Algorithm for Segment RoutingHow to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering, and Mapping Client and Server for Segment Routing to LDP Interoperability.] • Support for SRv6 network programming in BGP and IS-IS. [See Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP and How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for OAM ping and traceroute for Segment Routing for IPv6 (SRv6) network programming. [See ITU-T Y.1731 Ethernet Service OAM Overview and How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for SRv6 flexible algorithms in traffic engineering database (TED) and BGP Link State (BGP-LS)

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
	[See How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering and BGP Link-State Extensions for Source Packet Routing in Networking (SPRING).] • SRv6 support for static SR-TE policy. [See Understanding SR-TE Policy for SRv6 Tunnel.] • Support for SRv6 micro-SIDs in IS-IS transport. You can compress multiple SRv6 addresses into a single IPv6 address (micro-SID). For use cases that need to include more than six SRv6 SIDs, micro-SIDs can help in compressing multiple IPv6 addresses. [See How to Enable SRv6 Network Programming in IS-IS Networks.]
Software installation and upgrade	• Support for secure-boot implementation based on the UEFI 2.4 standard. [See Software Installation and Upgrade Guide.] • Zero-touch provisioning (ZTP) support for WAN interfaces and DHCPv6 options. [See Zero Touch Provisioning.] • Secure Zero Touch Provisioning. [See Secure Zero Touch Provisioning.]
System management	• Support for an alternate partition for device recovery. You can use an alternate partition called **/altconfig** to recover the device when the **/config** partition gets corrupted. In certain scenarios, the **/config** partition (which holds the last four committed configuration files along with the rescue configuration) gets corrupted during resets or power cycles. The **/altconfig** partition (which holds the **juniper.conf.gz** and **rescue.conf.gz** files) is used by the management daemon (mgd) to recover the device when the **/config** partition is corrupted. This is a boot-time feature and is enabled by default.

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
Timing and synchronization	• Support for enhanced Ethernet equipment clock (eEEC). Enhanced EEC enables new clocks to operate with different quality levels defined in the Synchronous Ethernet chain. The ACX7332 router supports the following new clock quality levels for enhanced EEC: • Enhanced primary reference time clock (ePRTC) • Primary reference time clock (PRTC) • Enhanced primary reference clock (ePRC) • Enhanced Ethernet equipment clock [See enable-extended-ql-tlv, Ethernet Synchronization Message Channel Overview, and synchronization.] • Support for frequency synchronization using the Synchronous Ethernet protocol in accordance with the ITU-T G.8262 and G.8262.1 standards. [See Synchronous Ethernet Overview.] • Synchronous Ethernet over LAG with Ethernet Synchronization Message Channel (ESMC). [See Synchronous Ethernet and Ethernet Synchronization Message Channel (ESMC).] • SNMP MIB support for the Synchronous Ethernet timing feature. [See Configuring SNMP Trap Groups and Enterprise-Specific MIBs for Junos OS Evolved.] • Support for G.8275.1 telecom profile, Precision Time Protocol over Ethernet (PTPoE) encapsulation, and hybrid mode. [See Precision Time Protocol Overview and Understanding Hybrid Mode.] • PTP G.8275.1 support over Link Aggregation Group (LAG) [See G.8275.1 Telecom Profile.]

Table 1: ACX7332 Feature Support *(Continued)*

Feature	Description
Support for optics	• To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported on the ACX7332 router, see the Hardware Compatibility Tool.

Authentication and Access Control

- **Control device access privileges with exact match configuration (ACX7332, QFX5130-48C)**

  —Starting in Junos OS Evolved Release 23.4R1-S1, you can configure access privileges for login classes by allowing or denying full hierarchy strings with the `allow-configuration-exact-match` and `deny-configuration-exact-match` configuration options. The exact match configuration enables you to set separate permissions for `set`, `delete`, `activate`, or `deactivate` operators for any hierarchy.

  The `allow-configuration-exact-match` and `deny-configuration-exact-match` configuration options support full hierarchy strings as well as wildcard characters and regular expressions.

  [See Understanding Exact Match Access Privileges for Login Classes.]

Chassis

- **Software resiliency support (ACX7332)**—Starting in Junos OS Evolved release 23.4R1-S1, we provide resiliency support for Chassis, Routing and Control Board (RCB), Power Supply Module (PSM), fan tray, and input and output devices on the listed ACX Series routers. [See Chassis User Guide .]

IPv6

- **Operations, Administration, and Maintenance (OAM) ping and traceroute support for SRv6 uSID (ACX7332)**—Starting in Junos OS Evolved Release 23.4R1-S1, we support pinging an SRv6 micro-segment identifier (uSID) to verify that the uSID is reachable and is locally programmed at the target node. We also support tracerouting to an SRv6 uSID for hop-by-hop fault localization as well as path tracing to a uSID.

  As part of this feature, we support SRv6 uSID ping and traceroute for the following SIDs:

  - SRv6 IS-IS ping and traceroute for end behavior with NEXT-CSID (uN)/uN+End.X behavior with NEXT-CSID (uA)/uN+End.DT behavior with NEXT-CSID (uDT) SIDs
  - SRv6 IS-IS ping and traceroute for compressed SIDs (compressed SID to be provided by user) for uN/uA/uDT
  - SRv6 micro-SIDs-stack ping and traceroute for uN/uN+uA/nN+uDT SIDs

We've introduced the following commands for this feature:

- `ping srv6 spring-te micro-sids-stack nexthop-address <nh-addr> nexthop-interface <if-name> usids [usid1 usid2 …]`
- `traceroute srv6 spring-te micro-sids-stack nexthop-address <nh-addr> nexthop-interface <if-name> usids [usid1 usid2 …]`
- `traceroute srv6 spring-te micro-sids-stack nexthop-address <nh-addr> nexthop-interface <if-name> usids [usid1 usid2 …] probe-icmp`

Junos Telemetry Interface

- **Multicast telemetry support with IGMP and PIM operational state sensors (ACX7332)**—Junos OS Evolved Release 23.4R1-S1 supports IGMP and PIM sensors based on the OpenConfig data models **openconfig-igmp.yang** (version 0.3.0) and **openconfig-pim.yang** (version 0.4.2).

  [For sensors, see Junos YANG Data Model Explorer.]

- **SR-TE policy telemetry (ACX7332)**—In Junos OS Evolved Release 23.4R1-S1, we introduce support for telemetry streaming of operational state data for the segment routing–traffic engineering (SR-TE) policy. State sensors are based on the OpenConfig data model **openconfig-srte-policy.yang**. You can subscribe to SR-TE sensors using the resource path **/network-instances/network-instance/segment-routing/te-policies**.

  [See Junos YANG Data Model Explorer.]

- **IS-IS OpenConfig and operational state sensor support (ACX7332)**—Junos OS Evolved Release 23.4R1-S1 supports OpenConfig IS-IS configurations and sensors based on the OpenConfig data model **openconfig-isis.yang** (version 1.0.0). This feature closes some gaps in our OpenConfig configuration and sensor support in the IS-IS area.

  [For OpenConfig configuration, see Mapping OpenConfig ISIS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **MPLS OpenConfig and operational state sensor support (ACX7332)**—Junos OS Evolved Release 23.4R1-S1 supports OpenConfig MPLS configurations and sensors based on the OpenConfig data models **openconfig-mpls-ldp.yang** (version 3.2.0) and **openconfig-mpls-rsvp.yang** (version 4.0.0). This feature closes some gaps in our OpenConfig configuration and sensor support in the MPLS RSVP-TE and MPLS LDP areas.

  [For OpenConfig configuration, see Mapping MPLS OpenConfig MPLS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **OpenConfig interface configuration support (ACX7332)**—Junos OS Evolved Release 23.4R1-S1 introduces interface OpenConfig configuration support based on OpenConfig data model **openconfig-interfaces.yang**.

[For state sensors, see Junos YANG Data Model Explorer For OpenConfig configuration, see Mapping OpenConfig Interface Commands to Junos Configuration.}

- **802.1X configuration and operational state sensors using OpenConfig (ACX7332)**—Junos OS Evolved Release 23.4R1-S1 supports configuration and telemetry streaming of operational state data based on the OpenConfig data model **openconfig-if-8021x.yang**.

  [For state sensors, see Junos YANG Data Model Explorer. For OpenConfig configuration, see Mapping OpenConfig 802.1X Commands to Junos Configuration.]

- **STP OpenConfig and operational state sensor support (ACX7332)**—Junos OS Evolved Release 23.4R1-S1 supports OpenConfig STP configurations and sensors based on the OpenConfig data model **openconfig-spanning-tree** (version 1.0, revision 0.3.1).

  [For OpenConfig configuration, see Mapping OpenConfig STP Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

Layer 2 VPN

- **BFD Session support for VCCV (ACX7332)**—Starting in Junos OS Evolved 23.4R1-S1, you can configure Bidirectional Forwarding Detection (BFD) for virtual circuit connection verification (VCCV) with a minimum-interval equal to or greater than 1000ms. Distributed and Centralized sessions are both supported.

  [See Configuring BFD for VCCV for Layer 2 Circuits.]

Multicast

- **Recycle bandwidth management (ACX7332)**—Starting in Junos OS Evolved Release 23.4R1-S1, we support the following two modes of recycle bandwidth management:.

  - Default mode, where all supported recycle applications share the recycle bandwidth in a best-effort manner.
  - Configuration mode, where you can reserve a guaranteed bandwidth for each configured recycle application.

  To configure bandwidth recycling using the configuration mode, use the following CLI statements:

  - `set system packet-forwarding-options recycle-bandwidth-profiles` *`profile-name`*
  - `set system packet-forwarding-options recycle-bandwidth-profiles` *`profile-name`* `ingress-egress-mcast` *`percentage`* `srv6` *`percentage`* `srv6-ext` *`percentage`*

  [See Recycle Bandwidth Management.]

Routing Policy and Firewall Filters

- **Enhanced hierarchical policer (ACX7332)**-Starting in Junos OS Evolved Release 23.4R1-S1, you can use the enhanced hierarchical policer configuration to rate limit traffic based on packets classified on the traffic priority. You can configure traffic policing at four levels of hierarchies with respect to the traffic priority. Use `enhanced-hierarchical-policer` to configure these levels.

  [See enhanced-hierarchical-policer.]

Routing Protocols

- **Support for OSPFv2 HMAC SHA-1 keychain authentication and optimization for multi-active MD5 keys (ACX7332)**—Starting in Junos OS Evolved Release 23.4R1-S1, you can enable OSPFv2 HMAC-SHA1 authentication with keychain to authenticate packets reaching or originating from an OSPF interface. This feature ensures a smooth transition from one key to another for OSPFv2 with enhanced security.

  You can enable OSPFv2 to send packets authenticated with only the latest MD5 key after all the neighbors switch to the latest configured key. In Junos OS Evolved releases earlier than Release 23.4R1, we support advertising authenticated OSPF packets always with multiple active MD5 keys with a maximum limit of two keys per interface.

  To enable OSPFv2 HMAC-SHA1 authentication, configure the `authentication keychain <keychain name>` option at the `[edit protocols ospf area area-id interface interface_name` hierarchy level. To enable optimization of multiple active MD5 keys, configure the `delete-if-not-inuse` option at the `[edit protocols ospf area area-id interface interface_name authentication multi-active-md5]` hierarchy level.

  [See Understanding OSPFv2 Authentication.]

Software Installation and Upgrade

- **Optimize reboot times by disabling default initialization and startup of certain L2 applications (ACX7332)**—Starting in Junos OS Evolved Release 23.4R1-S1, when rebooting the device, the Layer 2 (L2) applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:

  - `[edit interface interface-name unit number family ethernet-switching]`
  - `[edit vlans]`
  - `[edit routing-instance instance-name instance-type virtual-switch]`
  - `[edit routing-instance instance-name instance-type mac-vrf]`
  - `[edit protocols l2-learning]`

Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

As a result of this change, if your configuration already contains these configuration statements and you then delete all of them, these L2 applications stop running.

[See request node reboot (re0 | re1) (Junos OS Evolved), request system reboot (Junos OS Evolved), and request system software add (Junos OS Evolved).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Dynamic tunnel support for SRv6 tunnels (ACX7332)**—Starting in Junos OS Evolved Release 23.4R1-S1, you can configure segment routing for IPv6 (SRv6) Layer 3 VPN dynamic tunnel over a traditional Layer 3 VPN network.

  The following functionalities are supported:

  - DT4, DT6, DT46, uDT4, uDT6, uDT46 SIDs.
  - Signal SRv6 locator based dynamic tunnel from BGP.
  - Resolve BGP route over dynamic tunnel route.
  - Resolve BGP route over dynamic tunnel and create transport tunnel composite next hop (TCNH) with BGP/IGP/Static as underlay. Have single or multiple router next hops.
  - Forwarding policies under dynamic tunnels.
  - DSCP propagation for dynamic tunnel at ingress.
  - Display dynamic tunnel (dyn-tunnel) flag information for SRv6 tunnel as part of `show route extensive command`.

Additional Features

We've extended support for the following features to these platforms.

- **Paragon Active Assurance test agent** (ACX7332). You can install a Paragon Active Assurance test agent on your router to help you monitor network quality, availability, and performance. Also, starting in this release for all ACX Series platforms that support this test agent, you now install the test agent by using the `test-agent` configuration statement at the `[edit services paa]` hierarchy level, instead of using the deprecated operational mode command `request services paa install`.

  [See Install the Paragon Active Assurance (PAA) Test Agent and test-agent.]

- **Support for EVPN-MPLS features** (ACX7332):

  - IRB with IPv4 and IPv6 addresses

- IRB virtual gateway
- IRB anycast gateway
- IRB with static mac
- EVPN asymmetric Type 2 and symmetric Type 5 routes
- EVPN E-LAN over BGP-LU
- EVPN proxy ARP and ARP suppression, and NDP and NDP suppression
- EVPN routing policies
- Ingress virtual machine traffic optimization (VMTO)

[See EVPN with IRB Solution Overview, Anycast Gateways, Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes, Understanding EVPN Pure Type 5 Routes, EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression, Ingress Virtual Machine Traffic Optimization, and Routing policies for EVPN.]

- **System, CPU, and memory UDP statistics support for Junos telemetry interface (JTI)** (ACX7332).

  Statistics support the health monitoring application.

  [For sensors, see Junos YANG Data Model Explorer.]

- **HTTP and TCP probe types for RPM** (ACX7332). You can now configure the `http-get`, `http-metadata-get`, and `tcp-ping` probe types for real-time performance monitoring (RPM) probes. You must configure the `offload-type none` statement to be able to commit the configuration.

  [See probe-server, probe-type, and rpm.]

- **Support for syslog over TCP and TLS** (ACX7332).

What's New in 23.4R1

IN THIS SECTION

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

To view features supported on the ACX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 23.4R1, click the Group by Release link. You can collapse and expand the list as needed.

- ACX7024
- ACX7024X
- ACX7100-32C
- ACX7100-48L
- ACX7348
- ACX7509

The following sections highlight the key features in this release.

Hardware

- **New ACX7348 router (ACX Series)**—Starting in Junos OS Evolved Release 23.4R1, we introduce the Juniper Networks® ACX7348 Cloud Metro Router, a multiservice platform from the ACX7300 series that addresses the growing demands of metro applications. With a compact 3-U semi-modular form factor, this industrial-rated router offers an aggregation solution that gives cloud providers and service providers the performance and scalability needed as networks grow.

  The ACX7348 router provides port flexibility of 1GbE through 400GbE and a throughput of 2.4 Tbps. The router has a fixed FPC with forty-eight 25GbE and eight 100GbE ports, dual Routing Engines, three bays for pluggable interface modules, redundant power supply modules (AC or DC), and four fan trays (two fans per tray). The ACX7348 router supports the following pluggable FPCs:

  - ACX7K3-FPC-2CD4C—Two 400GbE and four 100GbE ports
  - ACX7K3-FPC-16Y—Sixteen 50GbE ports

  Table 2: ACX7348 Feature Support

Feature	Description

Chassis	• Supports two Routing Engines, one Control Board, one Forwarding Engine Board (FEB), one fixed FPC, and three removable FPCs chassis supports:
	  • Platform FEB and FPC FRU presence and power-up.
	  • Infrastructure databases and services.
	  • Power management.
	  • Environment monitoring and cooling.
	  • System LED behavior.
	• Platform resiliency support for device chassis, RCB, FEB, and input/output card (IOC).
	• Support for telemetry interfaces.

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
Class of service	• Support for classification and rewrite rules of all types (Inet-Prec/DSCP/DSCP-v6/IEEE-802.1p/IEEE-802.1ad) at the logical interface level. Supports logical interfaces classification and rewrite rules for MPLS, VPLS, Layer 3 VPN, Layer 2 circuit, CCC, IRB, and EVPN. [See Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview.] • Support for port shaping and scheduling with eight VoQ queues per port and two scheduling priority levels (strict-high and low). Supports multiple strict-high priority queues (RR scheduling), multiple low-priority queues (WFQ scheduling), low latency queues (LLQ), and default deep buffers. [See Schedulers Overview for ACX Series Routers and Shared and Dedicated Buffer Memory Pools on ACX Series Routers .]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

Dynamic Host Configuration Protocol	• DHCP server and DHCP relay configuration for IPv4 and IPv6 services.
	
	[See DHCP Overview.]
	
	• DHCP relay deployment of EVPN over MPLS, which includes:
	
	  • Edge-routed bridging (ERB)—Edge model where DHCP clients are connected and relayed in network leaf devices. The spine PEs do not perform DHCP relay functions, and the routers support transit spine functionality running protocols such as BGP for integrated routing and bridging (IRB).
	
	• The following functionalities:
	
	  • EVPN over MPLS Ethernet-LAN
	
	  • DHCPv4 and DHCPv6 relay options
	
	  • Stateless forward-only mode for DHCP relay over VPN
	
	  • Anycast IP address with IRB for a relay source
	
	  • Client VRFs only
	
	• DHCPv4 and DHCPv6 relay agent support for MC-LAG. DHCP relay agent support includes:
	
	  • DHCPv4 and DHCPv6 stateless forward-only option on Layer 3 static interfaces over MC-LAG.
	
	  • DHCPv4 and DHCPv6 stateless forward-only option on IRB interfaces over MC-LAG.
	
	  • DHCPv4 and DHCPv6 forward-snooped-clients on dual-stack configurations.
	
	[See DHCP Relay Agent and Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent.]
	
	[See DHCP Relay Agent in EVPN-MPLS Network.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

EVPN	• Support for the following EVPN-MPLS features on MAC-VRF instances:
	  • L2 flooding for broadcast, unknown unicast, and multicast (BUM) traffic
	  • Split-horizon between core interfaces
	  • Data plane and control plane MAC learning and aging, and static MAC
	  • MAC movement and MAC mobility on control plane only
	  • MAC limiting and MAC learning
	  • Input and output VLAN maps using normalization on user-to-network interfaces (UNIs)
	  • Aggregated Ethernet interfaces used for UNIs and network node interfaces (NNIs)
	  • Physical interfaces for VLAN tagging, stacked VLAN tagging, flexible VLAN tagging, and extended VLAN bridges using EVPN-MPLS as a service
	  • Ethernet bridge mode for logical UNIs
	  • VLAN ID lists, native VLAN ID supported logical UNIs, and priority-tagged logical interfaces
	  • Underlay networks with ECMP and Fast Reroute (FRR)
	  • Control-word support for EVPN
	  • EVPN Proxy Address Resolution Protocol (ARP) and ARP suppression
	  • EVPN-ELAN over segment routing
	  [See EVPN Feature Guide.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Virtual private wire service (VPWS) with EVPN signaling mechanisms and flexible cross-connect support.
	
	[See Overview of VPWS with EVPN Signaling Mechanisms.]
	
	• EVPN E-LAN active-active multihoming with EVPN aliasing support for ESI LAG.
	
	[See Example: Configuring EVPN Active-Active Multihoming.]
	
	• All-active multihoming redundancy in both Ethernet VPNvirtual private wire service (EVPN-VPWS) and EVPN-VPWS with flexible cross-connect.
	
	[See Overview of Flexible Cross-Connect Support on VPWS with EVPN.]
	
	• EVPN VPWS multihoming all-active forsegment routing over MPLS
	
	[See Overview of VPWS with EVPN Signaling Mechanisms.]
	
	• Entropy and flow label for EVPN-ELAN
	
	[See Configuring Entropy Labels.]
	
	• Support for the following EVPN-MPLS features:
	
	  • IRB with IPv4 and IPv6 addresses
	
	  • IRB virtual gateway
	
	  • IRB anycast gateway
	
	  • IRB with static mac
	
	  • EVPN asymmetric Type 2 and symmetric Type 5 routes
	
	  • EVPN E-LAN over BGP-LU
	
	  • EVPN proxy ARP and ARP suppression, and NDP and NDP suppression

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• EVPN routing policies • Ingress virtual machine traffic optimization (VMTO) [See EVPN with IRB Solution Overview, Anycast Gateways, Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes, Understanding EVPN Pure Type 5 Routes, EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression, Ingress Virtual Machine Traffic Optimization, and Routing policies for EVPN.] • Support for the following EVPN-VPWS features: • EVPN-VPWS FXC VLAN unaware service • EVPN-VPWS FXC VLAN aware service • EVPN-VPWS over segment routing • Single homing and all active multihoming support • Flow-aware transport (FAT) pseudowire labels • Entropy labels [See Overview of VPWS with EVPN Signaling Mechanisms.]
Firewall filters	• Support for firewall filters and policers. You can configure firewall filters with packet match conditions for the bridge domain, IPv4, IPv6, CCC, and MPLS families. In addition to packet match conditions, the count, discard, log, syslog, and policer actions are supported. [See Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview.] • Filter-based forwarding (FBF). [See Filter-Based Forwarding Overview.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• Firewall filter protocols: MPLS, CCC, virtual private LAN service (VPLS), and ANY. [See Firewall Filters Overview, Filter-Based Forwarding Overview, Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address, and Guidelines for gRPC and gNMI Services.]
High availability	• VRRP for IPv4 and IPv6. [See VRRP and VRRP for IPv6 Overview.] • BFD over label-switched paths (LSPs) or RSVP-based LSPs in a centralized mode. [See Bidirectional Forwarding Detection (BFD) for MPLS.] • High availability on these routers are supported at the hardware level. Graceful routing engine switchover is not supported in this release. • Support for loop-free alternate (LFA) routes for OSPF and IS-IS. LFA enables IP fast-reroute capability for OSPF and IS-IS. [See Loop-Free Alternate Routes for OSPF Overview and Understanding Loop-Free Alternate Routes for IS-IS.] • BFD-triggered fast reroute for unicast next hops. [See Bidirectional Forwarding Detection (BFD) for MPLS, session-id-change-limiter-indirect, and no-bfd-triggered-local-repair.]
Interfaces	• The ACX7348 router provides 1GbE through 400GbE port flexibility and a throughput of 2.4 Tbps. [See Port Speed.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Support for 6xQSFPDD and 16xSFP56 FPC line cards. The 6xQSFPDD FPC Line Card has two QSFPDD ports (Port 0 and 2) and four QSFP ports (Port 1, 3, 4 and 5). The 16xSFP56 FPC line card has 16 SFP56 ports (Port 0 to 15). Slot 1 and 2 supports 10-Gbps, 25-Gbps, and 50-Gbps speeds. Slot 3 supports 1-Gbps, 10-Gbps, and 25-Gbps speeds.
	
	[See Port Speed.]
	
	• Support for 802.1X authentication on Layer 2 and Layer 3 interfaces.
	
	[See 802.1X Authentication on Layer 2 Interfaces.]
	
	• Support for LACP link protection. We support 1:1 and N:N link protection.
	
	[See link-protection.]
	
	• Resiliency support for ASIC error and CM infra. Resiliency only supports logging and detection and not action.
	
	• Features supported for unnumbered interfaces:
	
	  • Bidirectional Forwarding Detection (BFD)
	
	  • BGP labeled unicast
	
	  • Ethernet VPN virtual private wire service (EVPN-VPWS)
	
	  • IS-IS protocol adjacency
	
	  • Label Distribution Protocol (LDP)
	
	  • Layer 2 VPN and Layer 2 circuit
	
	  • Layer 3 VPN
	
	  • Qualified next hop
	
	  • RSVP-TE

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• Static subnet route • Source Packet Routing in Networking (SPRING) over OSPFv2 • SPRING-TE • Segment routing with MPLS • Static LSP • Source Packet Routing in Networking (SPRING) over OSPFv2 • SPRING-TE • Segment routing with MPLS • Static LSP [See Configure unnumbered Interfaces.]
Junos Telemetry Interface (JTI)	• Logical subinterface and Packet Forwarding Engine drop, pipe, and line-card counter sensor support for JTI. [See Junos YANG Data Model Explorer.]
Layer 2 features	• Ethernet ring protection switching (ERPS) with G.8032 version 2. [See Understanding Ethernet Ring Protection Switching Functionality .] • Support for the following advanced Layer 2 (L2) features: • Bridge domain without a `vlan-id` *number* statement • Bridge domain with the `vlan-id` value set to None • Bridge domain with a single VLAN ID

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• • MAC learning, ageing, and limiting • Single-learning domain per bridge domain • Ethernet service types: • E-Line with these AC interface types: port, VLAN, Q-in-Q, VLAN list, and VLAN maps • E-Line • E-LAN • E-Access • E-Transit • LLDP • LACP • IRB interface • Link aggregation group (LAG) support with the following hashing algorithms: • For family `multiservice`, destination and source MAC addresses • For family `inet`, Layer 3 and Layer 4 • For family `inet6`, Layer 3 destination and source addresses • For family `inet6`, Layer 4 destination and source ports

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Encapsulation types:
	  • `extended-vlan-bridge`
	  • `vlan-bridge`
	  • `ethernet-bridge`
	• Q-in-Q tunneling
	[See Understanding Layer 2 Bridge Domains and Q-in-Q Tunneling.]
	• Disable local switching in bridge domains.
	[See Configuring MAC Address Flooding and Learning for VPLS.]
	• Layer 2 protocol tunneling (L2PT) to send L2 protocol data units (PDUs) across the network and deliver them to devices that are not part of the local broadcast domain.
	• Storm control.
	[See Understanding Storm Control.]
	• Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP).
	[See Spanning-Tree Protocol Overview.]
	• MAC move limit and multiple trunk ports, virtual private LAN service (VPLS), and EVPN networks.
	[See Understanding MAC Limiting and MAC Move Limiting.]
	• Layer 2 Control Protocol (L2CP) BPDUs are transparently forwarded in hardware unless a specific protocol is configured on the incoming interface. This feature helps you to configure and enable L2PT.
	• VLAN sensor support.

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	[See Telemetry Sensor Explorer.]
	
	[See Understanding Layer 2 Bridge Domains on ACX Series and Q-in-Q Tunneling on ACX Series, Bridging and VLANs, and Configuring MAC Address Flooding and Learning for VPLS .]
	
	• Multichassis link aggregation groups (MC-LAGs). The following Layer 2 features are available on MC-LAGs:
	  • Layer 2 bridging for active-active and active-standby modes
	  • Layer 2 unicast with and without IGMP snooping
	  • Layer 3 unicast with and without IGMP snooping
	  • Layer 2 multicast with and without IGMP or MLD snooping
	  • Layer 3 multicast with and without IGMP or MLD snooping
	
	[See Understanding Multichassis Link Aggregation Groups.]

Layer 2 VPN	• Support for VPLS. The router supports a single VLAN for each virtual switch routing instance type. Junos OS Evolved does not support the `family vpls` option. To configure VPLS, configure the `instance-type virtual-switch` statement at the [`edit routing-instances` *`routing-instance-name`*] hierarchy level. If you configure normalized VLANs, either by not configuring VLAN IDs or by including the `vlan-id none` statement, then you must include the `service-type single` statement at the [`edit routing-instances` *`routing-instance-name`* `protocol vpls`] hierarchy level. [See Introduction to Configuring VPLS.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Support for control word and load-balancing capabilities using entropy and flow-aware transport of pseudowires (FAT) flow labels, across LDP-signaled pseudowires for virtual private LAN service (VPLS).
	
	[See control-word , Configuring Entropy Labels, and FAT Flow Labels Overview.]
	
	• Support for redundant pseudowires for virtual private LAN service (VPLS). The router supports VPLS with LDP hot-standby, cold-standby model, and without BFD or CFM trigger.
	
	[See Redundant Pseudowires for Layer 2 Circuits and VPLS.]
	
	• IRB support for VPLS.
	
	[See Configuring VPLS and Integrated Routing and Bridging.]
	
	• Layer 2 VPN and L2 circuit support:
	
	  • L2 circuit—Targeted LDP signaling pseudowires and interoperability between different types of supported attachment circuit for L2 circuit
	
	  • L2 VPN circuit—BGP signaling
	
	• MPLS fast reroute (FRR) on IGP, circuit attachment types (port, VLAN, and Q-in-Q tunneling), control word, pseudowire circuit on aggregated Ethernet interfaces, indirect next hops and composite next hops, pipe and uniform mode time-to-live (TTL), Tag Protocol Identifiers (TPIDs), and VLAN map on pop, push, or swap.
	
	[See Understanding Layer 2 VPNs and Understanding Layer 2 VPNs and Configuring Interfaces for Layer 2 Circuits.]
	
	• Flow-aware transport for pseudowires (FAT) label and entropy label support for Layer 2 circuit and Layer 2 VPN.
	
	[See Configuring Entropy Labels, and FAT Flow Labels Overview.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

Layer 3 features	• Support for the following Layer 3 features:
	  • IP forwarding and exception packet handling
	  • IEEE 802.1Q (VLAN trunk) on IRB interfaces
	  • Address Resolution Protocol (ARP), neighbor discovery, unicast reverse-path forwarding, and ECMP
	  • LPM and fragmentation handling, ICMP redirect handling, VLAN tagging modes, neighbor solicitation, and Interface-based routing
	  • Longest prefix match
	  • Exception packets handling
	  • VLAN tagging modes
	  • Integrated routing and bridging (IRB)
	  • IPv4 and IPv6
	The router also supports interior gateway protocols such as OSPF, IS-IS, RIP, and ECMP for IPv4 and IPv6. [See Configure ICMP Features, Enabling VLAN Tagging, Neighbor solicitation, Understanding Unicast RPF (Routers), OSPF Overview, IS-IS Overview, and RIP User Guide.]

Layer 3 VPN	Support for the following Layer 3 VPN features:
	NOTE: VT interface-based Layer 3 VPN is not supported. Layer 3 VPN ping is supported only with the `vrf-table-label` configuration.
	• IP-VPN services:
	  • Instance-type virtual routing and forwarding (VRF) and virtual-router
	  • All control plane configuration options

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• • Per-prefix and per-table label signaling • Layer 3 VPN support with ECMP • BGP polices support for different Layer 3 VPN use cases (for example, full mesh VPN, hub-spoke VPN, management VPN, and leaking routes) • Layer 3 VPN with vrt-table-label mode • Layer 3 VPN with chained composite mode • Import or export of routes across local VRF and global VRF **NOTE**: Table next hop is not supported. • Inter-autonomous system (inter-AS) options A, B, and C **NOTE**: Inter-AS option B can be deployed in hierarchical network design within a single IGP AS. • PE to CE routing protocols—Static, eBGP, IS-IS, OSPF, and RIP • IPv6 Provider Edge (6PE)/IPv6 VPN routing over MPLS (6VPE) with PE-CE routing-static and PE-CE BGPv6 [See Layer 3 VPNs User Guide for Routing Devices.]
MACsec	Supports Media Access Control Security (MACsec). [See Understanding Media Access control Security (MACsec).]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

MPLS	• Support for the following MPLS features:
	  • IP/MPLS infrastructure feature set for the L3 VPN service
	  • Basic BGP control plane features such as LDP-DOD, CSPF, and single-area CSPF
	  • MPLS label stack
	  • MPLS protections:
	  • Fast reroute (FRR) and Make-before-break (MBB)
	  • Link protection
	  • Node protection
	  • Label-switching router (LSR)
	  • Shared Risk Link Group (SRLG) for MPLS
	  • RSVP label-switched path (LSP) over IPv4 includes refresh reduction
	  • Label Distribution Protocol (LDP) LSP over IPv4
	  • RSVP 1:1
	  • RSVP-Traffic Engineering (RSVP-TE)
	  • LDP over RSVP
	  • Inter-autonomous system LSP intra-area LSP
	  • [See MPLS Applications User Guide.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Support for MPLS LSP statistics and RSVP-TE auto-bandwidth features. Support includes:
	  • MPLS LSP statistics for the following LSP types:
	  • LDP-signaled LSPs
	  • RSVP-signaled LSPs
	  • Static LSPs
	  • Bypass LSPs
	  • Container LSPs
	  • RSVP-TE auto-bandwidth
	[See LSP Overview, LDP Overview, RSVP Overview, and Configuring Optimized Auto-bandwidth Adjustments for MPLS LSPs.]

Multicast	• Support for multicast snooping in a VPLS for the following protocols:
	  • IGMPv1, IGMPv2, and IGMPv3 snooping in VPLS
	  • MLDv1 and MLDv2 snooping in VPLS
	  • IGMP and MLD snooping in VPLS with integrated routing and bridging (IRB)
	  • Protocol Independent Multicast support over VPLS with IRB
	[See Multicast Snooping for VPLS.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Support for Layer 2 multicast-related features, including IGMP and MLD snooping. You can configure IGMP snooping with IGMPv1, IGMPv2, and IGMPv3, which includes support for:
	  • IGMP snooping in bridge domains
	  • IGMP snooping with integrated routing and bridging (IRB) configured in bridge domains
	  • MLD snooping in bridge domains
	  • MLD snooping with IRB configured in bridge domains
	[See IGMP Snooping Overview and Understanding MLD Snooping.]
	• Support for IPv4 multicast for Layer 3. You can configure IGMP snooping with IGMPv2 and IGMPv3, which includes support for the following:
	  • Anycast RP
	  • IGMP filter
	  • IGMP querier
	  • Protocol Independent Multicast source-specific multicast (PIM SSM)
	  • PIM sparse mode (PIM SM)
	NOTE: In this Junos OS Evolved release, the ACX7348 router doesn't support IPv6 multicast or L3 multicast protocols (such as IGMP, MLD, or PIM) over IPv4 and IPv6 IRB interfaces.
	[See IGMP Snooping Overview.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• Support for BGP MVPN. BGP over MPLS MVPN (also known as "next generation," or "NG," MVPN) running on multipoint LDP provider tunnels, where BGP MVPN is the intra-AS and PIM-SM and multipoint LDP point-to-multipoint (P2MP) tunnels is the data plane. [See Multiprotocol BGP MVPNs Overview.] • Multicast with IGMP or MLD snooping within VLANs for EVPN-MPLS. [See Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment.]
Network management and monitoring	• Support for port mirroring with analyzers and encapsulated remote Switch Port Analyzer (ERSPAN). [See Port Mirroring and Analyzers.] • Support for SNMP.

Operations, Administration, and Maintenance	• Support for OAM. You can configure connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM. You can also configure the following features of link-fault management (LFM):
	
	  • Discovery
	
	  • Link monitoring
	
	  • Remote fault detection
	
	[See ITU-T Y.1731 Ethernet Service OAM Overview and Introduction to OAM Link Fault Management (LFM).]
	
	• Support for IEEE 802.1ag OAM CFM.
	
	• Support for IEEE Standard 802.3ah and 802.1ag for OAM CFM down and up maintenance association end points (MEPs) over virtual private LAN service (VPLS).

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• Support for IEEE Standard 802.3ah and 802.1ag for OAM CFM up MEPs over EVPN. [See IEEE 802.3ah OAM Link-Fault Management Overview and IEEE 802.1ag OAM Connectivity Fault Management Overview.] • Support for CFM and performance monitoring (Y.1731) protocols over Ethernet interfaces for `bridge` and `inet` services. [See Ethernet OAM Connectivity Fault Management.] • Support for native Y.1731 operational state sensors to provide statistics such as frame loss ratio, frame delay, frame delay variation, and availability for Y.1731 performance monitoring.
Protection against DDoS attacks	• Support for control plane distributed denial of service (DDoS) protection. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]

Routing protocols	• Layer 3 and routing protocols IPv4, IPv6, BGP, IS-IS and ARP streaming sensor support using gRPC services.
	
	• Support for unicast reverse path forwarding (unicast RPF):
	
	  • Support for loose and strict mode
	
	  • Support for IPv4 and IPv6
	
	[See Understanding Unicast RPF (Routers).]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

	• Support for BGP flow specification (BGP flowspec).
	  • The following match conditions are not supported:
	  • Fragment for IPv6
	  • Packet length
	  • Port
	  • Source and destination prefix with offset
	• The following actions are not supported:
	  • Community
	  • Next-term
	  • Routing instance
	  • Sample
	  • Traffic marking
	  [See Understanding BGP Flow Routes for Traffic Filtering.]
	• Support for configuring interface groups in BGP flowspec filters.
	  [See Understanding BGP Flow Routes for Traffic Filtering and Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic.]
	• BGP PIC edge support for `inet` and MPLS VPNs. The following features are not supported:
	  • Session-based repair
	  • BGP PIC over LDP over RSVP tunnel
	  • BGP PIC over SR-MPLS

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• BGP PIC with RSVP • BGP-LU with PIC • BGP PIC edge protection for Layer 2 services • Protection with multilink failure [See Configuring BGP PIC Edge for MPLS Layer 3 VPNs and Use Case for BGP PIC for Inet.]
	• Support for entropy label for LDP, RSVP, L3VPN, and BGP-LU. [See Entropy label support for BGP Labeled Unicast (BGP-LU) and Configuring Entropy Labels.] • Support for BGP transport address family or BGP Classful Transport (BGP-CT) includes: • Service mapping over colored transport tunnels (RSVP, IS-IS flexible algorithm) to transport classes and map service routes over an intended transport class. The transport tunnels can span multiple domains (ASs or IGP areas). • Network slicing and interoperability between network domains. • IPv6 and segment routing–traffic engineered (SR-TE) color-only support. • IPv6 and BGP service routes with a color-only mapping community. • Enhanced transport-class configuration to provide precise resolution. [See use-transport-class.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

Services Applications	• RFC 2544-based benchmarking tests. Support for Layer 2 reflection (bridge, L2CKT, L2VPN, EVPN-VPWS, EVPN-FXC, EVPN-MPLS, and VPLS), with family `ccc` or family `ethernet-switching` and for Layer 3 reflection (IPv4, L3VPN) with family `inet`. RFC 2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before activation of the service. The tests measure throughput, latency, frame loss rate, and back-to-back frames.
	
	[See RFC 2544-Based Benchmarking Tests for ACX Routers Overview.]
	
	• RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. You can configure the TWAMP monitoring service, which sends out probes to measure network performance. TWAMP is often used to check compliance with service-level agreements. The support for this service is limited to the following features:
	
	  • IPv4 and IPv6 source and target addresses for clients, control connections, and test sessions
	
	  • Probe statistics and history
	
	  • Control and test session status
	
	  • Test session probe generation and reception, as well as reflection
	
	  • Timestamps set by software (the Routing Engine or the Packet Forwarding Engine) or the hardware
	
	  • Error reporting through system log messages only
	
	  • Unauthenticated mode only
	
	[See Understand Two-Way Active Measurement Protocol.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• Support for sFlow monitoring (ingress). [See sFlow Monitoring Technology and Understanding How to Use sFlow Technology for Network Monitoring.] • Inline active flow monitoring support for IPFIX and v9 export formats. We support ingress and egress sampling of IPv4 and IPv6 traffic on aggregated Ethernet and IRB interfaces and interfaces mapped to non-default VRFs, for both the IPFIX and version 9 export formats. You can configure up to four IPv4 collectors for inline active flow monitoring. See Understand Inline Active Flow Monitoring.

Source Packet Routing in Networking (SPRING) or segment routing	• Support for the following segment routing features:
	
	  • Segment routing global block (SRGB) for OSPF, IS-IS, and fast reroute.
	
	  • Metro Ethernet services over segment routing infrastructure
	
	  • Segment routing services: L3VPN, IPv6 VPN Provider Edge (6VPE) , IPv6 Provider Edge (6PE), L2VPN, L2 circuit, and BGP-VPLS
	
	  • Static segment routing (node segment, prefix segment, adjacency, and anycast segments) for OSPF and IS-IS
	
	  • Topology-independent loop-free alternate (TI-LFA) with segment routing for OSPF and IS-IS

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	• Unnumbered interfaces support for segment routing with OSPF • Support for IPv6 L3VPN over IPv6 SR-TE and IPv6 underlay • Support for flexible algorithm in OSPF and IS-IS for segment routing traffic • Interoperability of segment routing with LDP • SPRING support for SR-TE • Support for BGP link-state distribution with SPRING extensions [See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS, Understanding Source Packet Routing in Networking (SPRING), Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING, Configure Unnumbered Interfaces, Understanding Static Segment Routing LSP in MPLS Networks, Link-State Distribution Using BGP Overview, Understanding OSPF Flexible Algorithm for Segment RoutingHow to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering, and Mapping Client and Server for Segment Routing to LDP Interoperability.] • Support for SRv6 network programming in BGP and IS-IS. [See Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP and How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for OAM ping and traceroute for Segment Routing for IPv6 (SRv6) network programming. [See ITU-T Y.1731 Ethernet Service OAM Overview and How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for SRv6 flexible algorithms in traffic engineering database (TED) and BGP Link State (BGP-LS)

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
	[See How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering and BGP Link-State Extensions for Source Packet Routing in Networking (SPRING).] • SRv6 support for static SR-TE policy. [See Understanding SR-TE Policy for SRv6 Tunnel.] • Support for SRv6 micro-SIDs in IS-IS transport. You can compress multiple SRv6 addresses into a single IPv6 address (micro-SID). For use cases that need to include more than six SRv6 SIDs, micro-SIDs can help in compressing multiple IPv6 addresses. [See How to Enable SRv6 Network Programming in IS-IS Networks.]
Software installation and upgrade	• Support for secure-boot implementation based on the UEFI 2.4 standard. [See .] • Zero-touch provisioning (ZTP) support for WAN interfaces and DHCPv6 options. [See Zero Touch Provisioning.] • Secure Zero Touch Provisioning. [See Secure Zero Touch Provisioning.]
System management	• Support for an alternate partition for device recovery. You can use an alternate partition called **/altconfig** to recover the device when the **/config** partition gets corrupted. In certain scenarios, the **/config** partition (which holds the last four committed configuration files along with the rescue configuration) gets corrupted during resets or power cycles. The **/altconfig** partition (which holds the **juniper.conf.gz** and **rescue.conf.gz** files) is used by the management daemon (mgd) to recover the device when the **/config** partition is corrupted. This is a boot-time feature and is enabled by default.

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description

Timing and synchronization	• Support for enhanced Ethernet equipment clock (eEEC). Enhanced EEC enables new clocks to operate with different quality levels defined in the Synchronous Ethernet chain.
	
	The ACX7348 router supports the following new clock quality levels for enhanced EEC:
	
	  • Enhanced primary reference time clock (ePRTC)
	
	  • Primary reference time clock (PRTC)
	
	  • Enhanced primary reference clock (ePRC)
	
	  • Enhanced Ethernet equipment clock
	
	[See enable-extended-ql-tlv, Ethernet Synchronization Message Channel Overview, and synchronization.]
	
	• Support for frequency synchronization using the Synchronous Ethernet protocol in accordance with the ITU-T G.8262 and G.8262.1 standards. [See Synchronous Ethernet Overview.]
	
	• Synchronous Ethernet over LAG with Ethernet Synchronization Message Channel (ESMC).
	
	[See Synchronous Ethernet and Ethernet Synchronization Message Channel (ESMC).]
	
	• SNMP MIB support for the Synchronous Ethernet timing feature.
	
	[See Configuring SNMP Trap Groups and Enterprise-Specific MIBs for Junos OS Evolved.]
	
	• Support for G.8275.1 telecom profile, Precision Time Protocol over Ethernet (PTPoE) encapsulation, and hybrid mode. [See Precision Time Protocol Overview and Understanding Hybrid Mode.]
	
	• PTP G.8275.1 support over Link Aggregation Group (LAG)
	
	[See G.8275.1 Telecom Profile.]

Table 2: ACX7348 Feature Support *(Continued)*

Feature	Description
Support for optics	• To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported on the ACX7348 router, see the Hardware Compatibility Tool.

The ACX7348 router runs Junos OS Evolved and provides several capabilities that include support for the latest protocol and traffic engineering technologies, enhanced security, and precision timing for mobile backhaul applications. These capabilities and features enable you to create converged, virtualized, and automated architectures to address the rapid growth of 5G, IoT, and the cloud.

- **New ACX7024X router (ACX Series)**—Starting in Junos OS Evolved Release 23.4R1, we introduce the Juniper Networks® ACX7024X Cloud Metro Router, a commercial-rated high-scale multiservice router that meets the growing demands of high-performance bandwidth requirements.

  With a compact 1-U fixed form factor and advanced timing capabilities, these routers are well suited to support Ethernet business services, residential access, and 5G mobile deployments.

  The ACX7024X router provides a system throughput of 360 Gbps and a port flexibility of 1GbE through 100GbE. With a powerful processor and 64 MB of RAM, these routers provide high-scale and low-latency capabilities in commercial-temperature applications.

  We ship the ACX7024X routers with front-to-back airflow (airflow out or AFO) and AC or DC power supply modules (PSMs).

Table 3: ACX7024X Feature Support

Feature	Description
Class of service	• Logical interfaces support classification and rewrite rules for MPLS, VPLS, L3VPN, L2CKT, CCC, IRB, and EVPN. [See Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview.] • Support for deep buffering of oversubscribed traffic and absorbs network bursts. [See Shared and Dedicated Buffer Memory Pools on ACX Series Routers .] • Hierarchical Class of Service (CoS) support for Layer 3 VPN, Layer 2 VPN, Layer 2 Circuit, VPLS, and EVPN services. [See Hierarchical Class of Service in ACX Series Routers.] • Support for multiple classifier rules on interfaces. [See Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces.]

EVPN	• Support for the following EVPN features:
	  • EVPN-ETREE
	  • VLAN-aware bundle service in EVPN-MPLS
	  • EVPN Active/Standby Multi-homing with LACP based functionality
	  • MC-LAG in active-standby mode in an EVPN configuration without configuration of an ICCP or ICL interface
	  • Configure `lacp-oos-on-ndf` statement at the [`edit interfaces` *`interface name`* `esi df-election-granularity per-esi`] hierarchy
	[See Understanding VLAN-Aware Bundle and VLAN-Based Service for EVPN and df-election-granularity.]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

Firewall filters	• Firewall filter protocols (MPLS, CCC, VPLS, and ANY). [See Firewall Filters Overview.]
	• Supports the following firewall filter actions to configure filter-based forwarding or policy based forwarding:
	  • next-ip firewall filter action for IPv4 address
	  • next-ip6 firewall filter action for IPv6 address
	  • next-interface firewall filter action for an interface
	  [See Filter-Based Forwarding Overview and Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address.]
	• IPv4 and IPv6 filter support for SRv6. [See Firewall Filter Match Conditions and Actions (ACX Series Routers).]

High availability	• BFD-triggered fast reroute for unicast next hops. [See Bidirectional Forwarding Detection (BFD) for MPLS, session-id-change-limiter-indirect, and no-bfd-triggered-local-repair.] • IPv6 link-local address support for TWAMP Light [`edit services monitoring twamp client control-connection` *`connection-name`*] hierarchy level. [See Configure TWAMP on ACX.] • Packet Forwarding Engine timestamping for Two-Way Active Measurement Protocol (TWAMP) IPv6 test probes. [See offload-type.] • VRRP for IPv4 and IPv6. [See VRRP and VRRP for IPv6 Overview.]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

	• Support for the following Bidirectional Forwarding Detection (BFD) features:
	  • BFD for IPv4 and IPv6 routes
	  • Single-hop BFD in inline mode with an interval range of 4 milliseconds to 1 second
	  • Single-hop BFD in distributed mode with an interval of 1 second or more
	  • Single-hop BFD in centralized mode with a minimum interval of 1 second to detect IRB failures
	  • Multihop BFD with an interval of 1 second or more
	  • Micro-BFD for LAG in centralized or distributed mode with an interval of 1 second or more
	  • Configure BFD over label-switched paths (LSPs) or RSVP-based LSPs in a centralized mode
	[See Bidirectional Forwarding Detection (BFD) for MPLS.]
	• Support for loop-free alternate (LFA) routes for OSPF and IS-IS. LFA enables IP fast-reroute capability for OSPF and IS-IS.
	[See Loop-Free Alternate Routes for OSPF Overview and Understanding Loop-Free Alternate Routes for IS-IS.]
	• Bidirectional Forwarding Detection (BFD) for virtual circuit connection verification (VCCV) for Layer 2 Circuits. [See Configuring BFD for VCCV for Layer 2 Circuits.]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

Interfaces	• Support for unnumbered interfaces:
	  • Bidirectional Forwarding Detection (BFD)
	  • BGP labeled unicast
	  • Ethernet VPN virtual private wire service (EVPN-VPWS)
	  • IS-IS protocol adjacency
	  • Label Distribution Protocol (LDP)
	  • Layer 2 VPN and Layer 2 circuit
	  • Layer 3 VPN
	  • Qualified next hop
	  • RSVP-TE
	  • Static subnet route
	  • Source Packet Routing in Networking (SPRING) over OSPFv2
	  • SPRING-TE
	  • Segment routing with MPLS
	  • Static LSP
	[See Configure unnumbered Interfaces.]

Junos Telemetry Interface (JTI)	• RPM and TWAMP statistics sensor support for Junos Telemetry Interface (JTI). [See Telemetry Sensor Explorer.]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description
	• Junos Telemetry Interface provides Packet Forwarding Engine sensor support for firewall filter stats. To generate periodic streaming of statistics, use the resource path /junos/system/linecard/firewall/ in a gRPC or gNMI subscription. JTI provides periodic streaming and ON_CHANGE support for interface queue statistics. The resource path is /junos/system/linecard/interface/queue/. [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).] • Layer 3 routing and protocols statistics for JTI. [See Telemetry Sensor Explorer.]

Layer 2 features	• Advanced Layer 2 features:
	  • Bridge domain without a `vlan-id` *number* statement
	  • Bridge domain with the `vlan-id` value set to None
	  • Bridge domain with a single VLAN ID
	  • MAC learning, ageing, and limiting
	  • Single-learning domain per bridge domain
	  • Ethernet service types (E-Line, E-LAN, E-Access, E-Transit, LLDP, LACP, IRB interfaces).
	• Link aggregation group (LAG) support with the following hashing algorithms:
	  • For family `multiservice`, destination and source MAC addresses
	  • For family `inet`, Layer 3 and Layer 4
	  • For family `inet6`, Layer 3 destination and source addresses
	  • For family `inet6`, Layer 4 destination and source ports

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description
	• Encapsulation types (`extended-vlan-bridge`, `vlan-bridge`, `ethernet-bridge`) • Q-in-Q tunneling • Disable local switching in bridge domains. [See Configuring MAC Address Flooding and Learning for VPLS.] • Layer 2 protocol tunneling (L2PT) to send L2 protocol data units (PDUs) across the network and deliver them to devices that are not part of the local broadcast domain. • Storm control. [See Understanding Storm Control.] • Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). [See Spanning-Tree Protocol Overview.] • MAC move limit, virtual private LAN service (VPLS), and EVPN networks. [See Understanding MAC Limiting and MAC Move Limiting.] • Layer 2 Control Protocol (L2CP) BPDUs are transparently forwarded in hardware unless a specific protocol is configured on the incoming interface. This feature helps you to configure and enable L2PT.

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

Layer 3 features	• Support for the following Layer 3 features:
	  • IP forwarding and exception packet handling
	  • IEEE 802.1Q (VLAN trunk) on IRB interfaces
	  • Address Resolution Protocol (ARP), neighbor discovery, unicast reverse-path forwarding, and ECMP
	  • LPM and fragmentation handling, ICMP redirect handling, VLAN tagging modes, neighbor solicitation, and Interface-based routing
	  • Longest prefix match
	  • Exception packets handling
	  • VLAN tagging modes
	  • Integrated routing and bridging (IRB)
	  • IPv4 and IPv6
	The router also supports interior gateway protocols such as OSPF, IS-IS, RIP, and ECMP for IPv4 and IPv6. [See Configure ICMP Features, Enabling VLAN Tagging, Neighbor solicitation, Understanding Unicast RPF (Routers), OSPF Overview, IS-IS Overview, and RIP User Guide.]

Operations, Administration, and Maintenance	• Support for OAM. You can configure connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM. You can also configure the following features of link-fault management (LFM):
	  • Discovery
	  • Link monitoring
	  • Remote fault detection
	[See ITU-T Y.1731 Ethernet Service OAM Overview and Introduction to OAM Link Fault Management (LFM).]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description
	• Support for IEEE 802.1ag OAM CFM. • Support for IEEE Standard 802.3ah and 802.1ag for OAM CFM down and up maintenance association end points (MEPs) over virtual private LAN service (VPLS). • Support for IEEE Standard 802.3ah and 802.1ag for OAM CFM up MEPs over EVPN. [See IEEE 802.3ah OAM Link-Fault Management Overview and IEEE 802.1ag OAM Connectivity Fault Management Overview.] • Support for CFM and performance monitoring (Y.1731) protocols over Ethernet interfaces for `bridge` and `inet` services. [See Ethernet OAM Connectivity Fault Management.] • Support for native Y.1731 operational state sensors to provide statistics such as frame loss ratio, frame delay, frame delay variation, and availability for Y.1731 performance monitoring.
Platform and Infrastructure	• Support for platform resiliency to handle failures and faults related to components such as CPU, fan trays, temperature sensors, power supply units, FPGA, and optics. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, communicating errors through LEDs, self-healing, and taking components out of service. [See show system errors active.]
Protection against DDoS attacks	• Support for distributed denial of service (DDoS) protection, which is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

Routing protocols	• Support for forwarding information base (FIB) compression. Using FIB compression, multiple routes with the same forwarding behavior are compressed into a single route. This helps to limit the number of forwarding entries stored in the hardware. FIB compression is enabled by default. You can store a larger number of IPv4 and IPv6 routes after FIB compression.
	
	Support for new MDB profiles to replace old MDB profiles, that replaces older profiles for lean-edge, cloud metro, carrier-ethernet and Border Network Gateways (BNGs). IPv6 multicast source specific multicast (SSM) is supported on all MDB profiles. The MDB resource for routes can now be used by either global or VRF routing tables and can be scaled to maximum per-profile capacity.
	
	Use the command `set system packet-forwarding-options hw-db-profile cloud-metro` to configure the cloud-metro profile.
	
	[See hw-db-profile.]
	
	• Layer 3 and routing protocols IPv4, IPv6, BGP, IS-IS and ARP streaming sensor support using gRPC services.
	
	• Support for unicast reverse path forwarding (unicast RPF):
	
	  • Support for loose and strict mode
	
	  • Support for IPv4 and IPv6
	
	[See Understanding Unicast RPF (Routers).]
	
	• Support for BGP flow specification (BGP flowspec).
	
	  • The following match conditions are not supported:
	
	    • Fragment for IPv6
	
	    • Packet length
	
	    • Port
	
	    • Source and destination prefix with offset

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

	• The following actions are not supported:
	  • Community
	  • Next-term
	  • Routing instance
	  • Sample
	  • Traffic marking
	  [See Understanding BGP Flow Routes for Traffic Filtering.]
	• Support for configuring interface groups in BGP flowspec filters.
	  [See Understanding BGP Flow Routes for Traffic Filtering and Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic.]
	• Unicast reverse path forwarding (unicast RPF) support for IPv4 and IPv6. [See Understanding Unicast RPF (Routers).]

Services Applications	• RFC 2544 benchmarking test reflector support for Layer 2 services (Layer 2 bridge, Layer 2 circuit, Layer 2 VPN, EVPN-VPWS, EVPN-FXC, EVPN-MPLS, and VPLS) and Layer 3 services. [See RFC 2544-Based Benchmarking Tests for ACX Routers Overview.]
	• RFC 2544-based benchmarking tests for the Layer 3 reflector function that supports family inet and IPv4 source and destination addresses.
	• RFC 2544 generation for the following services:
	  • Layer 3 VPN
	  • Basic layer 3 routing

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description

	• RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. In Junos OS Evolved, you configure TWAMP at the [`edit services monitoring twamp`] hierarchy level. The support for this service is limited to the following features:
	  • IPv4 and IPv6 source and target addresses for clients, control connections, and test sessions
	  • Probe statistics and history
	  • Control and test session status
	  • Test session probe generation and reception, as well as reflection
	  • Timestamps set by software (the Routing Engine or the Packet Forwarding Engine) or the hardware
	  • Error reporting through system log messages only
	  • Unauthenticated mode only
	Refer to the offload-type inline-timestamping option of the test-session statement. [See test-session (Junos OS Evolved).]
	[See Understand Two-Way Active Measurement Protocol.]

	• Inline active flow monitoring support for IPFIX and v9 export formats. We support ingress and egress sampling of IPv4 and IPv6 traffic on aggregated Ethernet and IRB interfaces and interfaces mapped to non-default VRFs, for both the IPFIX and version 9 export formats. You can configure up to four IPv4 collectors for inline active flow monitoring. See Understand Inline Active Flow Monitoring. • Real-time Performance Monitoring (RPM) IPv6 source and target address support. [See source-address (RPM) and target.]

Table 3: ACX7024X Feature Support *(Continued)*

Feature	Description
Support for optics	• To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported on the ACX7024X router, see the Hardware Compatibility Tool.

Authentication and Access Control

- **SSH Hostkey Algorithm Update (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, the `hostkey-algorithm` SSH configuration options has been replaced with `hostkey-algorithm-list`, and the `ecdsa-sha2-nistp384` and `ecdsa-sha2-nistp521` hostkey algorithms are now supported.

  You can find the `hostkey-algorithm-list` configuration option at the `[edit system services ssh]` hierachy level.

  [See hostkey-algorithm.]

- **Background File Transfer for SCP/SSH (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can transfer files in the background via SCP/SSH. To configure background file transfers, include the `archive-sites` configuration statement at the `[edit system archival configuration]` hierarchy level.

  [See Back Up Configurations to an Archive Site.]

- **Control device access privileges with exact match configuration (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**

  —Starting in Junos OS Evolved Release 23.4R1, you can configure access privileges for login classes by allowing or denying full hierarchy strings with the `allow-configuration-exact-match` and `deny-configuration-exact-match` configuration options. The exact match configuration enables you to set separate permissions for `set`, `delete`, `activate`, or `deactivate` operators for any hierarchy.

  The `allow-configuration-exact-match` and `deny-configuration-exact-match` configuration options support full hierarchy strings as well as wildcard characters and regular expressions.

  [See Understanding Exact Match Access Privileges for Login Classes.]

Class of Service

- **Support for low-latency queuing (ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, the listed ACX7000 routers support low-latency queuing (LLQ)

for both port-level and hierarchical scheduling. LLQ enables delay-sensitive data to have preferential treatment over other traffic. A queue configured as a low-latency queue has higher priority over any other priority queues, including `strict-high` queues.

To configure a queue as a low latency queue, set `priority low-latency` at the `[set class-of-service schedulers` *`scheduler-name`*`]` hierarchy level and apply the scheduler to the relevant forwarding class.

[See Schedulers Overview for ACX Series Routers.]

- **Support for six scheduler priority levels (ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509)**—In releases earlier than Junos OS Evolved Release 23.4R1, ACX7000 Series routers supported two scheduler priority levels: strict-high and low. Starting with Junos OS Evolved 23.4R1, the listed ACX7000 routers support six scheduler priority levels. The levels are, in order of priority:

  - `low-latency`
  - `strict-high`
  - `high`
  - `medium-high`
  - `medium-low`
  - `low`

  This support is for both port-level and hierarchical scheduling.

  [See Schedulers Overview for ACX Series Routers.]

EVPN

- **EVPN-VXLAN fabric with an IPv6 underlay (ACX7024, ACX7100-32C and ACX7100-48L)**—Starting in Junos OS Evolved Release 23.4R1, you can use either an IPv4 or an IPv6 underlay in an EVPN-VXLAN fabric. You cannot have IPv4 and IPv6 underlays running concurrently.

  To enable an IPv6 underlay, you need to configure the `vxlan-extended` system profile. After you enable this profile, the Packet Forwarding Engine restarts. Traffic might drop if any traffic is running.

  To enable the profile, configure the `vxlan-extended` statement statement at the `[edit system packet-forwarding-options system-profile]` hierarchy level.

  To go back to using the default system profile, issue the `delete system packet-forwarding-options system-profile vxlan-extended` command. The PFE restarts after you revert to the default system profile. During this process, any traffic that's running might drop.

  [See vxlan-extended] .

- **EVPN-MPLS intersubnet multicast routing from sources inside the fabric to receivers inside the fabric (ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, we extend support for EVPN-MPLS multicast routing to include when the multicast source is inside the EVPN fabric and the receivers are also inside the EVPN fabric. In earlier releases, we supported intersubnet multicast routing to receivers in the fabric only when the multicast source was in an external Protocol Independent Multicast (PIM) domain.

  We support intersubnet multicast routing for sources inside the EVPN-MPLS fabric with:

  - Integrated routing and bridging (IRB) interfaces configured with PIM in `distributed-dr` mode
  - IGMP snooping with IGMPv2 and IGMPv3
  - Multicast Listener Discovery (MLD) snooping with MLDv1 and MLDv2
  - VLAN-based and VLAN-aware EVPN MAC-VRF instances
  - Active-active multihoming mode with EVPN Type 7 Join Sync routes and EVPN Type 8 Leave Sync routes (no single-active mode support)

  [See Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment.]

- **EVPN-VXLAN pure T5 host-route auto-generated community (ACX7024, ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, we've added support for the EVPN-VXLAN pure T5 host-route auto-generated community. This feature adds a community to MAC-IP ARP/NDP-based pure Type 5 host routes. Border leaf devices in edge-routed bridging (ERB) topologies with Type 5 connectivity to other leaf devices in the data center and Type 5 connections to external networks need to advertise aggregate routes to the external network instead of individual Type 5 routes. Border leaf devices can use this community to identify these routes and create an aggregate route to advertise to external EVPN networks.

  [See EVPN-VXLAN Pure T5 Host-Route Auto-Generated Community.]

- **OISM PEG DF election (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, you can customize the designated forwarder (DF) election method on border leaf devices that act as Protocol Independent Multicast (PIM) EVPN gateway (PEG) devices in an EVPN network running optimized intersubnet multicast (OISM). You can configure peer PEG devices to use one of the following DF election methods:

  - Mod-based DF election
  - Preference-based DF election with a specified preference value

  By default, a PEG device uses PIM-based DF election. When you configure this feature, the selected PEG DF election method replaces PIM-based DF election.

  [See peg-df-election and PEG DF Election.]

- **Support for EVPN E-LAN over SRv6 underlay (ACX7024, ACX7100, ACX7348, and ACX7509)**—EVPN E-LAN is a framework for delivering multipoint-to-multipoint VPN service with the EVPN signaling mechanisms. Service providers can use the E-LAN service to offer services that manage the Layer 2 (L2) learning very efficiently. Starting in Junos OS Evolved Release 23.4R1, you can configure a single-homed EVPN-ELAN service using Segment Routing for IPv6 (SRv6). To provide the SRv6 service, the egress provider edge (PE) device signals an SRv6 service segment identifier (SID) with the VPN route. The ingress PE device encapsulates the service SID in the VPN packet in an outer IPv6 header where the destination address is the SRv6 SID advertised by the egress PE device and is routable in the underlay. The nodes between the PE devices need to support only plain IPv6 forwarding. We support SRv6 micro-SID-based control planes and forwarding. Different endpoint behaviors are defined for SRv6 services on the egress node.

  [See Configuring EVPN E-LAN over SRv6.]

- **Support for EVPN-ELAN pure Type 5 routes over SRv6 (ACX7024, ACX7100, ACX7348, ACX7509)**—Starting in Junos OS Evolved 23.4R1, ACX Series routers running Junos OS Evolved support EVPN pure Type 5 routes over an SRv6 underlay.

  [See **Understanding EVPN Pure Type 5 Routes**.]

- **Static configuration of MAC-IP bindings with EVPN-VXLAN (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we've added the functionality to allow static configuration of MAC-IP bindings on an interface, similar to configuring static MACs on an interface. This feature enables the static configuration of IP and MAC entries for crucial services provided by management and infrastructure hosts. It proves particularly advantageous in Internet Exchange Point (IXP) networks where participant customer edge (CE) devices remain well-known and static, not transitioning to different provider edge (PE) devices.

  You can now utilize a new feature that establishes a static link between an IP address and a MAC address for a logical interface within a bridge domain or VLAN. When you provision a static MAC-IP entry on a PE device, the PE device initiates a probe following an exponential backoff pattern. The probe uses an all-zero sender IP address on the associated interface. If the entity owning the IP to MAC entry responds to the probe, the system will learn the IP to MAC binding as static. Subsequently, it will be propagated to remote PE devices through the BGP/EVPN Type 2 MAC advertisement route. The corresponding MAC address will be recognized as a dynamic entry. If you want to deactivate the probing mechanism for learning the IP to MAC binding, you can do so by configuring the new configuration option `arp-nd-probe-disable`. Without probing, both the MAC and IP to MAC binding will be acquired from network traffic and communicated using EVPN.

  We've introduced the following commands and configuration statements:

  - Configuration of static IP to MAC bindings

> **NOTE**: A maximum of eight MACs can be configured per static IP address.

The aforementioned commands provide an option to configure `router` and `override` bits for IPv6 entries. For example:

- Disable probing on configuration of static IP to MAC entries:

  To turn off the default probing on configuration of static IP to MAC entries, you can use the global configuration statement `arp-nd-probe-disable`.

  ```
  set protocols l2-learning arp-nd-probe-disable
  ```

- Enable logging for failed probing of static IP to MAC entries:

  To turn on the logging, configure the global configuration statement `arp-nd-probe-failed-log`.

  ```
  set protocols l2-learning arp-nd-probe-failed-log
  ```

- Enable GARP/unsolicited-NA for local and remote static entries

  If this feature is required, you must configure the global configuration statement `garp-na-enable`.

  ```
  set protocols l2-learning garp-na-enable
  ```

- Disable dynamic learning [all static provisioning]

  If dynamic learning of MAC-IP entries is not required, configure the statement `drop-unknown-macip` under BD/VLAN.

- Drop unicast ARP request

  To drop unicast address resolution requests (for instance, NUD NS messages), you can configure the statement `block-unicast-arp` at global level.

- [See EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression and interface-mac-ip-limit.]

- **EVPN-VPWS service over SRv6 underlay (ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, and ACX7024)**—Starting in Junos OS Evolved Release 23.4R1, you can configure single-active or all-active multi homed EVPN VPWS network using segment routing over IPv6 (SRv6).

  To enable EVPN-VPWS over SRv6, configure the following:

  1. Include the `end-dx2-sid` statement at the [`edit routing-instances` *`instance-name`* `protocols evpn source-packet-routing srv6 locator` *`name`*] hierarchy level or at the [`edit routing-instance` *`routing-instance-name`* `protocols evpn interface` *`interface-name`*] hierarchy level for the evpn-vpws instance type.

  2. Include the `enhanced-ip` statement at the [`edit chassis network-services`] hierarchy level.

3. Enable `advertise-srv6-service` and `accept-srv6-service` in the [`edit protocols bgp group` *`name`* `family evpn`] hierarchy level.

[See Configuring VPWS with EVPN Signaling Mechanisms and Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP.

High Availability

- **MC-LAG emulation in an EVPN deployment (ACX7024, ACX7100-32C, ACX7100-48C, ACX7348 and ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, you can emulate the function of an MC-LAG in active-standby mode in an EVPN configuration without having to configure an ICCP or ICL interface. In a standard EVPN configuration, logical interfaces configured on an aggregated Ethernet interface can have different designated forwarder (DF) election roles. To emulate an MC-LAG configuration, the designated forwarder takes on the role of the aggregated Ethernet interface. The provider edge (PE) device that is the non-DF device sends LACP unsynchronized packets to the CE device. This causes LACP to go down on the CE device, and the CE device does not use the links connected to the non-DF device for sending traffic.

  To emulate MC-LAG functionality, enable the `lacp-oos-on-ndf` statement at the [`edit interfaces` *`interface name`* `esi df-election-granularity per-esi`] hierarchy level.

  [See per-esi.]

Interfaces

- **Support port group - (ACX7100-32C)**—Starting in Junos OS Evolved Release 23.4R1, we've extended support for four port group configuration.

  [See Port Speed on PTX10001-36MR Router Overview.]

IPv6

- **SRv6-TE micro-SID support for transport and L3VPN (ACX7348 and ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, we extend the micro-segment identifier (micro-SID) support for SRv6 traffic engineering (TE). We support SR-TE micro-SID only with default block configurations across the whole network domain. However, if any block configurations are present, then those configurations must be the same throughout the network. The Packet Forwarding Engine supports the bit shifting operation for both <block>:<uN>:<uA> and <block>:<uA> routes. You must configure the full SID, the way it is advertised in the IS-IS IGP—that is, <block>:<uN> or <block>:<uN>:<uA>.

  We've introduced the following two configuration statements:

  - `micro-srv6-sid` statement under the `protocols source-packet-routing segment-list <name> <hop-name>` hierarchy to configure micro-SIDs in an SR-TE SRv6 segment list

  - `strict-adjacency` statement under the `protocols source-packet-routing segment-list <name> <hop-name>` hierarchy to strictly follow the micro-adjacency SID

You can configure the segment list containing micro-SIDs with the existing `srv6` configuration statement like the traditional SRv6 configuration. The only difference between the traditional and micro-SID configuration is that in the traditional SRv6-TE segment-list configuration, you must use the configuration statement `srv6-sid`. However, for micro-SID configuration, you must use the new configuration statement `micro-srv6-sid`.

- **Operations, Administration, and Maintenance (OAM) ping and traceroute support for SRv6 uSID (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, we support pinging an SRv6 micro-segment identifier (uSID) to verify that the uSID is reachable and is locally programmed at the target node. We also support tracerouting to an SRv6 uSID for hop-by-hop fault localization as well as path tracing to a uSID. As part of this feature, we support SRv6 uSID ping and traceroute for the following SIDs:

  - SRv6 IS-IS ping and traceroute for end behavior with NEXT-CSID (uN)/uN+End.X behavior with NEXT-CSID (uA)/uN+End.DT behavior with NEXT-CSID (uDT) SIDs
  - SRv6 IS-IS ping and traceroute for compressed SIDs (compressed SID to be provided by user) for uN/uA/uDT
  - SRv6 micro-SIDs-stack ping and traceroute for uN/uN+uA/nN+uDT SIDs

  We've introduced the following commands for this feature:

  - `ping srv6 spring-te micro-sids-stack nexthop-address <nh-addr> nexthop-interface <if-name> usids [usid1 usid2 …]`
  - `traceroute srv6 spring-te micro-sids-stack nexthop-address <nh-addr> nexthop-interface <if-name> usids [usid1 usid2 …]`
  - `traceroute srv6 spring-te micro-sids-stack nexthop-address <nh-addr> nexthop-interface <if-name> usids [usid1 usid2 …] probe-icmp`

- **Optimizing ARP, NDP and Default-Route handling in internal DB of DCD (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.4R1, DCD only deletes routing entries for addresses that are completely unlinked from all associated addresses. Additionally, we introduce checks to prevent configuring multiple static MAC addresses for a single ARP and NDP address, which helps improve system stability and avoid potential conflicts in network configurations.

Junos OS API and Scripting

- **Support for REST API functionality over nondefault VRF instances (ACX Series and PTX Series)**—

  Starting in Junos OS Evolved Release 23.4R1, we support REST API functionality with HTTPs over nondefault virtual routing and forwarding (VRF) instances. To hold all the configurations related to this feature, enable the `routing-instance` statement at the [`edit system services`] hierarchy level.

Junos Telemetry Interface

- **Multicast telemetry support with IGMP and PIM operational state sensors (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports IGMP and PIM sensors based on the OpenConfig data models **openconfig-igmp.yang** (version 0.3.0) and **openconfig-pim.yang** (version 0.4.2).

  [For sensors, see Junos YANG Data Model Explorer.]

- **SR-TE policy telemetry (ACX7348, PTX10003, PTX10004, PTX10008, and PTX10016)**—In Junos OS Evolved Release 23.4R1, we introduce support for telemetry streaming of operational state data for the segment routing–traffic engineering (SR-TE) policy. State sensors are based on the OpenConfig data model **openconfig-srte-policy.yang**. You can subscribe to SR-TE sensors using the resource path **/network-instances/network-instance/segment-routing/te-policies**.

  [See Junos YANG Data Model Explorer.]

- **IS-IS OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports OpenConfig IS-IS configurations and sensors based on the OpenConfig data model **openconfig-isis.yang** (version 1.0.0). This feature closes some gaps in our OpenConfig configuration and sensor support in the IS-IS area.

  [For OpenConfig configuration, see Mapping OpenConfig ISIS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **MPLS OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10003, PTX10004, and PTX10008)**—Junos OS Evolved Release 23.4R1 supports OpenConfig MPLS configurations and sensors based on the OpenConfig data models **openconfig-mpls.yang** (version 3.2.2), **openconfig-mpls-types.yang** (version 3.2.1), and**openconfig-mpls-te.yang** (version 3.2.2). The following OpenConfig configurations and state sensors are supported.

  - Configurations:

    - MPLS global TTL propagation (**/network-instances/network-instance/mpls/global/config/ttl-propagation**)

    - MPLS LSP PRI/SEC path-metric-bound-constraint (**/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/p2p-tunnel-attributes/p2p-primary-path/p2p-primary-path/path-metric-bound-constraints/path-metric-bound-constraint/config/**)

  - Sensors:

    - MPLS global (**/network-instances/network-instance/mpls/global/state/**)

- MPLS global interface attributes (**/network-instances/network-instance/mpls/global/interface-attributes/**)
- MPLS LSP autobandwidth (**/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/bandwidth/auto-bandwidth/state/**)
- MPLS LSP PRI/SEC path-metric-bound-constraint (**/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/p2p-tunnel-attributes/p2p-primary-path/p2p-primary-path/path-metric-bound-constraints/path-metric-bound-constraint/state/**)
- MPLS traffic engineering global attributes SRLG (**/network-instances/network-instance/mpls/te-global-attributes/srlgs/srlg/static-srlg-members/members-list/state/**)

[For OpenConfig configuration, see Mapping MPLS OpenConfig MPLS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **MPLS OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports OpenConfig MPLS configurations and sensors based on the OpenConfig data models **openconfig-mpls-ldp.yang** (version 3.2.0) and **openconfig-mpls-rsvp.yang** (version 4.0.0). This feature closes some gaps in our OpenConfig configuration and sensor support in the MPLS RSVP-TE and MPLS LDP areas.

  [For OpenConfig configuration, see Mapping MPLS OpenConfig MPLS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **OpenConfig interface configuration support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)**—Junos OS Evolved Release 23.4R1 introduces interface OpenConfig configuration support based on OpenConfig data model **openconfig-interfaces.yang**.

  [For state sensors, see Junos YANG Data Model Explorer For OpenConfig configuration, see Mapping OpenConfig Interface Commands to Junos Configuration.}

- **802.1X configuration and operational state sensors using OpenConfig (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)**—Junos OS Evolved Release 23.4R1 supports configuration and telemetry streaming of operational state data based on the OpenConfig data model **openconfig-if-8021x.yang**.

  [For state sensors, see Junos YANG Data Model Explorer. For OpenConfig configuration, see Mapping OpenConfig 802.1X Commands to Junos Configuration.]

- **STP OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L,ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports OpenConfig STP configurations and sensors based on the OpenConfig data model **openconfig-spanning-tree** (version 1.0, revision 0.3.1).

[For OpenConfig configuration, see Mapping OpenConfig STP Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **Upgrade of OpenConfig models for routing instances (ACX7100-32C, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports an upgrade for the following OpenConfig models:

  - **openconfig-local-routing.yang** to version 2.0.0.
  - **openconfig-routing-policy.yang** to version 3.3.0.

  The upgraded models introduce new leaves for operational state sensors and configuration in the following areas:

  - Inter-instance policies
  - Route limits
  - Router advertisement
  - Local aggregates
  - Static routes

  [For state sensors, see Junos YANG Data Model Explorer.

  For OpenConfig configuration, see Mapping OpenConfig Network Instance Commands to Junos Configuration.]

Layer 2 VPN

- **BFD Session support for VCCV (ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, and ACX7024, and ACX7024X)**—Starting in Junos OS Evolved 23.4R1, you can configure Bidirectional Forwarding Detection (BFD) for virtual circuit connection verification (VCCV) with a minimum-interval equal to or greater than 1000ms. Distributed and Centralized sessions are both supported.

  [See Configuring BFD for VCCV for Layer 2 Circuits.]

MPLS

- **Computation of unreserved bandwidth optimized RSVP dynamic bypass LSP (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, the Constrained Shortest Path First (CSPF) can optionally use a different approach to protect a link or a node by leveraging the computation based on unreserved bandwidths on traffic engineering (TE) links. To enable this feature, use the `optimize bandwidth` configuration statement at the `edit protocols rsvp interface` *`interface`* `link-protection` hierarchy level. While the default approach of RSVP bypass produces a bypass method that

optimizes traffic engineering (TE) metric, enabling the new configuration statement maximizes the end-to-end unreserved bandwidth.

[See Configuring Link Protection on Interfaces Used by LSPs.]

- **Map static IPv6 route to next-hop using service label (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024,PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can enable static IPv6 routes to be mapped to the next-hop over an IPv4 MPLS network. 6PE is a transitional IPv6 over IPv4 technology that uses MPLS tunnels to carry services.

  You can use the `explicit-null` configuration statement under the [`edit routing-options rib inet6.0 static route` *`ipv6-address`*] hierarchy level to push ingress service label as part of the static next hop configuration for static IPv6 routes. The `explicit-null` configuration statement only supports configuring IPv4 mapped IPv6 address.

  The `static` configuration statement under the [`edit routing-options forwarding-table chained-composite-next-hop ingress`] hierarchy provisions chained composite next-hop.

  > **NOTE**: The `static` configuration statement must be enabled before configuring the `explicit-null` configuration statement.

- **New CLI commands for MPLS LSPs (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.4R1, you can get more visibility into the current state of the MPLS LSPs on the router to debug suspected anomalies in high scale conditions with the following newly introduced CLI commands.

  - `show rsvp session bypass [bypass-name] [protected]` and `show rsvp session [unprotected]` provides visibility into LSPs protected by a specific bypass tunnel.
  - `show mpls lsp [make-before-break]` and `show rsvp session [multiple-lsp-sessions]` provides sisibility into LSPs undergoing make-before-break.
  - `show mpls tunnel-manager-statistics` provides statistics on all local repair and make-before-break events for LSPs.
  - `show rsvp session [fr-ingress]` provides visibility into LSPs on flood-reflector edge routers.

Multicast

- **IGMP and MLD snooping version configuration (ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can configure the version of IGMP or MLD snooping queries for VLANs or bridge domains associated with Layer 2 (L2) multicast. This configuration ensures that end hosts or CPE devices that are not compliant with RFC 4541 and can't normally respond to snooping queries of a later version are now able to process and respond to those snooping queries.

To configure the IGMP or MLD snooping version, use the following CLI statements:

- `set protocols igmp-snooping version` *version*
- `set protocols mld-snooping version` *version*

[See IGMP MLD Snooping Version Configuration.]

- **Recycle bandwidth management (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, we support the following two modes of recycle bandwidth management:.
  - Default mode, where all supported recycle applications share the recycle bandwidth in a best-effort manner.
  - Configuration mode, where you can reserve a guaranteed bandwidth for each configured recycle application.

  To configure bandwidth recycling using the configuration mode, use the following CLI statements:

  - `set system packet-forwarding-options recycle-bandwidth-profiles` *profile-name*
  - `set system packet-forwarding-options recycle-bandwidth-profiles` *profile-name* `ingress-egress-mcast` *percentage* `srv6` *percentage* `srv6-ext` *percentage*

  [See Recycle Bandwidth Management.]

Precision Time Protocol (PTP)

- **Support for Synchronous Ethernet over LAG support (ACX7024, ACX7100-32C, and ACX7100-48L)**—Starting in Junos OS Evolved Release 23.4R1, ACX7100-32C, ACX7100-48L, and ACX7024the listed ACX Series routers support Synchronous Ethernet over a Link Aggregation Group (LAG). You can now define a LAG while configuring the Synchronous Ethernet clock source on these routers.

  [See Synchronous Ethernet, PTP Overview]

- **Support for G.8275.1 over LAG and Synchronous Ethernet over LAG (ACX7509)**—Starting in Junos OS Evolved Release 23.4R1, the ACX7509 router supports G.8275.1 profile over link aggregation group (LAG) and Synchronous Ethernet over LAG with the following ITU recommendations:
  - G.8262/G.8262.1—Specifies timing characteristics of Synchronous Ethernet equipment clock (ECC).
  - G.8264—Describes the Ethernet Synchronization Message Channel (ESMC).

  [See Synchronous Ethernet, PTP Profiles]

Routing Protocols

- **Support for micro-SIDs in TI-LFA, microloop avoidance, flex algo, and IS-IS MT (ACX Series)**—Starting in Junos OS Evolved Release 23.4R1, we extend the support of compressing SRv6 addresses into a single IPv6 address (micro-SID) in topology-independent loop-free alternate (TI-LFA), microloop avoidance, and Flexible Algorithm (flex algo) path computations. From this release onward, you can also configure algorithms for micro-segment identifiers (micro-SIDs) to facilitate the new extended feature. We also support IPv6 unicast topology (part of IS-IS MT) in TI-LFA, microloop avoidance, and flex algo computations.

  To enable flex algo to install the ingress routes in transport class routing information bases (RIBs), configure the `use-transport-class` statement at the `[edit routing-options flex-algorithm` *`id`*`]` hierarchy level.

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

- **Support for OSPFv2 HMAC SHA-1 keychain authentication and optimization for multi-active MD5 keys (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can enable OSPFv2 HMAC-SHA1 authentication with keychain to authenticate packets reaching or originating from an OSPF interface. This feature ensures a smooth transition from one key to another for OSPFv2 with enhanced security.

  You can enable OSPFv2 to send packets authenticated with only the latest MD5 key after all the neighbors switch to the latest configured key. In Junos OS Evolved releases earlier than Release 23.4R1, we support advertising authenticated OSPF packets always with multiple active MD5 keys with a maximum limit of two keys per interface.

  To enable OSPFv2 HMAC-SHA1 authentication, configure the `authentication keychain` *`<keychain name>`* option at the `[edit protocols ospf area` *`area-id`* `interface` *`interface_name`* hierarchy level. To enable optimization of multiple active MD5 keys, configure the `delete-if-not-inuse` option at the `[edit protocols ospf area` *`area-id`* `interface` *`interface_name`* `authentication multi-active-md5]` hierarchy level.

  [See Understanding OSPFv2 Authentication.]

- **Support for Next-Hop Dependent Capability Attribute (ACX7100-32C and PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, we use the Entropy Label Capability (ELCv3) attribute defined within the IETF Next-Hop Dependent Capability Attribute for load balancing. This attribute replaces the existing ELCv2 attribute. To operate the ELCv2 attribute along with ELCv3, explicitly configure the `elc-v2-compatible` statement at the `labeled-unicast entropy-label` hierarchy level.

  [See Understanding Entropy Label for BGP Labeled Unicast LSP.]

- **Support for limiting the number of BGP sessions belonging to a subnet (ACX7100-32C and PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, we support limiting the number of BGP sessions belonging to a given subnet that is configured using the ‘`allow` statement. You can use this feature to configure wider subnets by limiting the number of BGP sessions over them. You can

set this limit using the `peer-limit value` statement at the `[edit protocols bgp group group-name dynamic-neighbor]` hierarchy level.

[See peer-limit.]

Services Applications

- **STAMP and changes to TWAMP support (ACX7024, ACX7024X, ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we support RFC 8762, *Simple Two-Way Active Measurement Protocol* (STAMP). RFC 8762 standardizes and expands upon the TWAMP Light operational mode, which was defined in Appendix I of RFC 5357, *Two-Way Active Measurement Protocol* (TWAMP). A STAMP-compliant reflector ensures symmetric payload size (in accordance with RFC 6038) and operates in either stateless or stateful mode, depending on whether the sequence number in the reflected payload is copied from the client frame or generated independently. A stateful reflector can detect in which direction drops have occurred. In previous releases, we supported symmetric payloads and stateless reflection. With this release, we support stateful reflection, full compliance with the STAMP standard, and unidirectional drop values for clients. We support unidirectional drop values not only for STAMP clients, but also for TWAMP Managed-mode clients.

  [See Understand Two-Way Active Measurement Protocol and twamp.]

Software Installation and Upgrade

- **Optimize reboot times by disabling default initialization and startup of certain Layer 2 applications (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.4R1, when rebooting the device, the Layer 2 (L2) applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:

  - `[edit interface interface-name unit number family ethernet-switching]`
  - `[edit vlans]`
  - `[edit routing-instance instance-name instance-type virtual-switch]`
  - `[edit routing-instance instance-name instance-type mac-vrf]`
  - `[edit protocols l2-learning]`

  Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

  As a result of this change, if your configuration already contains these configuration statements and you then delete all of them, these L2 applications stop running.

[See request node reboot (re0 | re1) (Junos OS Evolved), request system reboot (Junos OS Evolved), and request system software add (Junos OS Evolved).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for RFC 8814 (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we partially support RFC 8814, *Signaling Maximum SID Depth (MSD) Using the Border Gateway Protocol-Link State*. Currently, we support signaling the MSD using IS-IS for SRv6. For non-segment routing networks, the MSD value reflects the maximum label depth.

  A controller in a segment routing network learns the MSD of the participating router and computes the segment routing path. The controller ensures that the label stack is not greater than what the routers can support.

  [See Link-State Distribution Using BGP.]

- **SRv6 dynamic SID support for BGP and IS-IS protocols (ACX Series)**—Starting in Junos OS Evolved Release 23.4R1, we support dynamic segment IDs for BGP and IS-IS.

  To enable dynamic end SID, include the `dynamic-end-sid` configuration statement at the `[edit protocols isis source-packet-routing srv6 locator` *`locator-name`*`]` hierarchy level.

  To enable dynamic end x SID, include the `dynamic-end-x-sid` configuration statement at the `[edit protocols isis interface`*`int-name`* `level` *`level-number`* `srv6-adjacency-segment protected locator` *`locator-name`*`]` hierarchy level.

  [See.level and srv6.]

- **Dynamic tunnel support for SRv6 tunnels (ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, and ACX7024, and ACX7024X)**—Starting in Junos OS Evolved Release 23.4R1, you can configure sgment routing for IPv6 (SRv6) Layer 3 VPN dynamic tunnel over a traditional Layer 3 VPN network.

  The following functionalities are supported:

  - DT4, DT6, DT46, uDT4, uDT6, uDT46 SIDs.
  - Signal SRv6 locator based dynamic tunnel from BGP.
  - Resolve BGP route over dynamic tunnel route.
  - Resolve BGP route over dynamic tunnel and create transport tunnel composite next hop (TCNH) with BGP/IGP/Static as underlay. Have single or multiple router next hops.
  - Forwarding policies under dynamic tunnels.
  - DSCP propagation for dynamic tunnel at ingress.

- Display dynamic tunnel (dyn-tunnel) flag information for SRv6 tunnel as part of `show route extensive command`.

Subscriber Management and Services

- **Dynamic VLAN (DVLAN single tag and dual tag) support for subscriber services scaling, performance (ACX7100-48L)** – Starting in Junos OS Evolved Release 23.4R1, we support the following subscriber services scaling and performance features on the ACX7100-48L devices:
  - DHCP for IP-DEMUX lite and PPPoE subscribers (IPv4, IPv6, and dual stack) with CoS lawful intercept and filters.
  - Protocol groups DHCPv6, L2TP, PPP, and PPPoE.
  - Configuration of DVLAN (single and dual tag) with L2TP (LAC) DDoS policers for broadband edge (BBE) protocols. We've enabled individual BBE protocol and Distributed Denial of Service (DDoS) protocol group configurations . Only the aggregate DDoS protocol group `dhcpv4v6` is active.
  - Subscriber scale qualification with CoS for L2TP access controllers(LAC) subscriber interfaces for IPv4, IPv6, and dual stack.

    [See Understanding DVLAN (Single/Dual tag) for Subscriber Services Scaling (Junos Evolved for ACX7100-48L Devices, ddos-protection (DDOS), l2tp, protocols (DDoS) (ACX Series, PTX Series, and QFX Series), ppp, and pppoe.']

- • **Support for new modular database (MDB) profiles to replace old MDB profiles**—Starting in Junos OS Evolved Release 23.4R1, new MDB profiles replace older profiles for lean-edge, cloud metro, Carrier Ethernet, and Border Network Gateways (BNGs). Additionally, here are a few enhancements provided with this feature :
  - The new MDB profiles configuration automatically restart PFEs.
  - All MDB profiles now support IPv6 multicast source specific multicast (SSM)
  - We define the `lean-edge` profile as the default profile instead of the `balanced` MDB profile
  - ACX7024 platform supports only the `lean-edge` and `cloud-metro` profiles.
  - We've removed the earlier restriction for LPM distribution. We've deprecated the LPM distribution command option `hw-db-profile lpm-distribution (1 | 2 | 200 | 3 | 4)`.
  - You can use the MDB resource for either global routes or VRF routing tables and scale the resources to maximum per-profile capacity.

  The following table lists new and old MDB profiles and indicates the default profiles :

New MDB Profile	Corresponding Old Profile

lean-edge (default)	l3-xl
cloud-metro	balanced (default)
carrier-ethernet	l2-xl
bng	balanced-exem

Use the `set system packet-forwarding-options hw-db-profile cloud-metro` command to configure the `cloud-metro` profile.

[See hw-db-profile.]

- **Hierarchical policer for subscriber services firewall (ACX7100-48L)** —Starting in Junos OS Evolved Release 23.4R1, we support the following hierarchical policers on DHCP and Point-to-Point Protcol over Ethernet (PPPoE) access models:
  - Single rate two-color marking policer
  - Single rate tri-color marking policer
  - Two rate tri-color marking policer
  - Hierarchical policer for DHCP and PPPoE
  - You can view the configuration by using the `show dynamic-profiles pppoe-client-policer-1-profile` command.

    [See Hierarchical Policer Overview and Example: Configuring a Hierarchical Policer for Subscriber Services Firewall (ACX7100-48L Devices).]

- **MAC source demultiplexing on control-plane interfaces (ACX7100-48L)**—Starting in Junos OS Evolved Release 23.4R1, we support MAC-based source demultiplexing on control-plane interfaces on ACX7100-48L devices. The feature includes DHCP server dual (IPv4-only or IPv6-only) support for stack subscribers over a 1:N customer VLAN (C-VLAN). Use the DHCP server command `dhcpv6 [dynamic-profile dhcp-prof]` under the `[edit system services dhcp-local-server]` hierarchy.

  [See dhcp-local-server (System Services) and dynamic-profile (Static Subscribers).]

Additional Features

We've extended support for the following features to these platforms.

- **HTTP and TCP probe types for RPM** (ACX7024, ACX7100, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016). You can now configure the `http-get`, `http-`

`metadata-get`, and `tcp-ping` probe types for real-time performance monitoring (RPM) probes. You must configure the `offload-type none` statement to be able to commit the configuration.

[See probe-server, probe-type, and rpm.]

- **Inline active flow monitoring support for IPFIX and v9 export formats** (ACX7024X). We support ingress and egress sampling of IPv4 and IPv6 traffic on aggregated Ethernet and IRB interfaces, for both the IPFIX and version 9 export formats. You can configure up to four IPv4 collectors for inline active flow monitoring.

  [See Understand Inline Active Flow Monitoring.]

- **Interconnecting EVPN-VXLAN data centers with EVPN-MPLS in WAN using gateway nodes** (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)

  [See Overview of EVPN-VXLAN Interconnect through EVPN MPLS WAN Using Gateways].

- **Layer 3 PIM support** (ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509):

  - PIMv4/v6 dense mode
  - PIMv4/v6 sparse-dense mode with local/static rendezvous point (RP)
  - Auto RP for PIMv4

  [See PIM Dense Mode, PIM Sparse-Dense Mode, and PIM Auto-RP.]

- **NETCONF event notifications support** (ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016):

  - NETCONF event notifications
  - Filtering capability
  - Interleave capability

  [See NETCONF Event Notifications.]

- **Paragon Active Assurance test agent** (ACX7348). You can install a Paragon Active Assurance test agent on your router to help you monitor network quality, availability, and performance. Also, starting in this release for all ACX Series platforms that support this test agent, you now install the test agent by using the `test-agent` configuration statement at the `[edit services paa]` hierarchy level, instead of using the deprecated operational mode command `request services paa install`.

  [See Install the Paragon Active Assurance (PAA) Test Agent and test-agent.]

- **QDD-400G-ZR-M-HP optics enhancement and channelization support** (ACX7509). QDD-400G-ZR-M-HP supports channelization of 1x400-Gbps, 4x100-Gbps, 3x100-Gbps, 2x100-Gbps, and 1x100-Gbps speeds.

- **QDD-400G-ZR-M-HP optics enhancement support** (ACX7100-32C, ACX7100-48L, and PTX10008 with PTX10K-LC1201-36CD line card)

- **Supported transceivers, optical interfaces, and DAC cables (ACX7024)**—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 HCT and provide the first supported release information when the optic becomes available.

- **Support for channelization, performance monitoring, and TCA** (ACX7509). We support channelization, performance monitoring, and threshold-crossing alert (TCA) information for the QDD-400G-ZR-M-HP optical transceiver modules. QDD-400G-ZR-M-HP supports channelization of 1x400-Gbps, 4x100-Gbps, 3x100-Gbps, 2x100-Gbps, and 1x100-Gbps speeds. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics by using the `show interfaces transport pm` command and can manage the optical transport link efficiently.

  [See show interfaces transport pm.]

- **Support for EVPN-ETREE** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)

  [See EVPN-ETREE Overview.]

- **Support for EVPN-MPLS features** (ACX7348):

  - IRB with IPv4 and IPv6 addresses
  - IRB virtual gateway
  - IRB anycast gateway
  - IRB with static mac
  - EVPN asymmetric Type 2 and symmetric Type 5 routes
  - EVPN E-LAN over BGP-LU
  - EVPN proxy ARP and ARP suppression, and NDP and NDP suppression
  - EVPN routing policies
  - Ingress virtual machine traffic optimization (VMTO)

  [See EVPN with IRB Solution Overview, Anycast Gateways, Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes, Understanding EVPN Pure Type 5 Routes, EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression, Ingress Virtual Machine Traffic Optimization, and Routing policies for EVPN.]

- **Support for EVPN-VXLAN OISM** (ACX7024, ACX7100-32C, and ACX7100-48L).

ACX Series devices support optimized intersubnet multicast (OISM) with IGMPv2 or IGMPv3 in MAC-VRF EVPN instances (`vlan-aware` and `vlan-based` service types). These devices can be OISM server leaf or border leaf devices, or can be lean spine devices in the OISM-enabled fabric. The border leaf devices support the classic L3 interface model to connect to an external multicast PIM domain.

[See Optimized Inter-Subnet Multicast in EVPN Networks, OISM, show evpn oism, multicast-replication, accept-join-always-from, show evpn multicast-snooping status.]

- **Support for FEC 128 and FEC 129 VPLS with source packet routing** (ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509). VPLS with Source Packet Routing in Networking (SPRING) with IS-IS, OSPF, and non-colored segment routing–traffic-engineering (SR-TE). Source packet routing or segment routing is applied in an MPLS network. You can use FEC 128 and FEC 129 VPLS with SPRING over MPLS as an alternative to LDP VPLS over MPLS.

  [See Example: Configuring a Multihomed VPLS (FEC 128), Example: Configuring VPLS Multihoming (FEC 129), and Understanding Source Packet Routing in Networking (SPRING).]

- **Support for FEC 128 and FEC 129 VPLS with source packet routing** (ACX7024, ACX7348, and ACX7509)

  [See Example: Configuring a Multihomed VPLS (FEC 128), Example: Configuring VPLS Multihoming (FEC 129), and Understanding Source Packet Routing in Networking (SPRING).]

- **Support for performance monitoring and TCA** (ACX7100-32C, ACX7100-48L, PTX10008, and PTX10008 with PTX10K-LC1201-36CD line card). We support performance monitoring and threshold-crossing alert (TCA) information for the QDD-400G-ZR-M-HP optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics by using the `show interfaces transport pm` command and can manage the optical transport link efficiently.

  [See show interfaces transport pm.]

- **Support for VLAN-aware bundle service in EVPN-MPLS** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)

  [See Understanding VLAN-Aware Bundle and VLAN-Based Service for EVPN.]

- **System, CPU, and memory UDP statistics support for Junos telemetry interface (JTI)** (ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, ACX7024, and ACX7024X).

  Statistics support the health monitoring application.

  [For sensors, see Junos YANG Data Model Explorer.]

- **Support for syslog over TCP and TLS** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016).

- **TWAMP and RPM statistics sensor support for Junos telemetry interface (JTI)** (ACX7348)

[See Junos YANG Data Model Explorer.]

What's Changed

IN THIS SECTION

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

General Routing

- **The active-user-count is defined as a numeric integer value in ODL request output** —The output for the get-system-uptime-information ODL request contains information for the active-user-count. The active-user-count is now defined as a numeric integer value and avoids an invalid value type error.

  [See show system uptime.]

- Two new alarms are added and can be seen with MPC11E when 400G-ZR optics are used. High Power Optics Too Warm: warning of the increase in chassis ambient temperature with no functional action taken on the optics Temperature too high for optics power on: New inserted optics when the chassis ambient temperature is elevated beyond the threshold will not be powered on and would need to be reinserted when the ambient temperature is within the acceptable range

- In TVP Platforms, When the FPC is configured to bring offline due to major errors, the alarm string **FPC x Offlined due to Major Errors** is shown in the `show chassis alarm` and `show system alarm` output instead of **FPC x Major Errors** alarm.

Interfaces and Chassis

- ACX7509: In the CLI using the command `request chassis feb slot slot-number offline` if you make the primary FEB offline, a traffic loss warning message is displayed and the FEB offline request is rejected. If offline or restart is still intended for primary FEB, use `force` option in addition to the command.

  WARNING message displayed in the CLI: warning: RCB and FEB work in the paired slot mode. FEB %s offline/restart will result in traffic loss and does not cause a switchover. Please re-try after initiating a mastership switchover using `request chassis routing-engine master switch` CLI.

Junos XML API and Scripting

- **Ability to commit `extension-service file` configuration when application file is unavailable**—When you set the `optional` option at the **edit system extension extension-service application file *file-name*** hierarchy level, the operating system can commit the configuration even if the file is not available at the **/var/db/scripts/jet** file path.

  [See file (JET).]

- **Ability to restart daemonized applications**—Use the request extension-service `restart-daemonize-app` *`application-name`* command to restart a daemonized application running on a Junos device. Restarting the application can assist you with debugging and troubleshooting.

  [See request extension-service restart-daemonize-app.]

- **XML output tags changed for `request-commit-server-pause` and `request-commit-server-start` (ACX Series and PTX Series)**—We've changed the XML output for the `request system commit server pause` command (`request-commit-server-pause` RPC) and the `request system commit server start` command (`request-commit-server-start` RPC). The root element is `<commit-server-operation>` instead of `<commit-server-information>`, and the `<output>` tag is renamed to `<message>`.

Licensing

- **Deprecated license revoke information**—Starting in Junos OS Evolved Release 23.4R1, we've deprecated the `show system license revoked-info` command. You can use the `show system license` and `show system license usage` commands to know the license information.

Network Management and Monitoring

- **Support for NETCONF event notifications filtering and interleave capabilities (ACX7100-32C and ACX7100-48L)**—NETCONF clients can specify top-level filters when they subscribe to NETCONF event notifications. Additionally NETCONF sessions support the interleave capability, which enables a NETCONF client to receive NETCONF event notifications and execute NETCONF requests in the same session.

  [See NETCONF Event Notifications.]

- **Changes to the RPC response for `<validate>` operations in RFC-compliant NETCONF sessions (ACX Series and PTX Series)**—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server emits only an `<ok/>` or `<rpc-error>` element in response to `<validate>` operations. In earlier releases, the RPC reply also includes the `<commit-results>` element.

- **NETCONF `<copy-config>` operations support a `file://` URI for copy to file operations (ACX Series and PTX Series)**—The NETCONF `<copy-config>` operation supports using a `file://` URI when `<url>` is the target and specifies the absolute path of a local file.

  [See <copy-config>.]

- Simplified gRIBI configuration—Starting in Junos OS Evolved Release 23.4R1, we have simplified the configuration to run gRPC Routing Information Base Interface (gRIBI) service remote procedure calls (RPCs). You no longer need to configure statements at the `[edit routing-options resolution]` hierarchy level.

  [See gRIBI.]

Routing Protocols

- In Junos OS Evolved platforms, `show route snooping` and `show route forwarding-table` does not show /56 routes in the VPLS address family table.

- Before this change most list were ordered by the sequence in which the user configured the list items, for example a series of static routes. After this change the list order is determined by the system with items displayed in numerical sequence rather than by the order in which the items were configured. There is no functional impact to this change.

Software Installation and Upgrade

- **`configuration` and `no-configuration` options for the `request system snapshot` command (ACX Series and PTX Series)**—When you omit or include the `configuration` option, the `request system snapshot` command copies the **/config** directory and the configuration stored for each installed software version to the alternate solid-state drive (SSD) as part of the snapshot. You can use the `no-configuration` option to exclude the **/config** directory and the configuration stored for each installed software version from the snapshot.

User Interface and Configuration

- **Information about users editing system configuration**— The `show system configuration database status` command displays information from the Junos OS configuration database that describes the users currently editing the system configuration.

  [See Display Users Currently Editing the Configuration.]

- **Viewing files with the `file compare files` command requires users to have `maintenance` permission**— The `file compare files` command in Junos OS and Junos OS Evolved requires a user to have a login class with `maintenance` permission.

  [SeeLogin Classes Overview.]

Known Limitations

IN THIS SECTION

Learn about limitations in this release for ACX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- Powering up and configuring data path of 400G-ZR and 400G-ZR-M optics is not instant and takes more time comparatively with other optics. So, the user has to wait enough time before switching to new speed configuration. If a user switches from one speed to another speed without waiting, it leads to interface down. PR1635443
- PSU's do not have fault simulation feature. So LEDs cannot be controlled by the software. PR1731084
- During first power cycle or reboot of router, primary Routing Engine (RE) takes 30 to 60 seconds more, to be displayed in CLI `show chassis routing-engine`. There is no functional impact due to this and it is limited to cosmetic display for a brief time only. PR1743012
- In Junos OS Evolved Release 23.4 on ACX7348 or ACX7332, syslog messages related to PEX configuration is seen at bootup with `sib-slot` which should be corrected as `cb-slot` in later releases.

  Example:

  ```
  Current log:
  syslog(LOG_INFO, "%s: sib%d: pexsw bringup successful", __func__, board_slot);
  Expected log:
  syslog(LOG_INFO, "%s: cb%d: pexsw bringup successful", __func__, board_slot);
  ```

  PR1761710
- CRC errors can be seen with SFP-T (Copper) optics. PR1771671

Open Issues

IN THIS SECTION

Learn about open issues in this release for ACX Series routers

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- IPv6 transit traffic statics output is missing. PR1662174
- With dhcp-snoop configuration, jdhcpd might crash during configuration commit on Junos OS Evolved based platforms.PR1688644
- On all Junos Evolved platforms, configuring Link Layer Discovery Protocol (LLDP) with **system services netconf notification** enabled triggers l2cpd crash. This crash causes the CPU to spike.PR1695057
- In ACX7509, time error spike is seen with PTP long. runPR1697093
- We observe core with rpd with BGP flowspec if secondary-independent-resolution is configured. PR1722715
- We might encounter jdhcpd core during initialization. This is rare, and there is no service impact because of this core (as the process recovers immediately). PR1730717
- ACX7348: 400G-AOC module support is not available in Junos OS Evolved Release 23.4. PR1732958
- On Junos OS Evolved ACX platforms with GRES (Graceful Routing Engine switchover), after performing GRES switchover jdhcpd does not start on the new primary Routing Engine due to which DHCPv4/v6 (Dynamic Host Configuration Protocol) session binding is lost resulting in traffic loss. PR1740530
- Servo sends some of DEL-REQ pakcets to 0:00:00:00:00:00. This is seen when moving the port from 400g to 25g/40g/100g.PR1740612
- `show pfe statistics traffic` command displays 0 in **Input packets:** and **Output packets:** after Packet Forwarding Engine app (evo-pfemand) restart for some time. Correct values are displayed after the delay. PR1745512
- During a continuous network churn, following errors related to pfe-spec filters are seen for Packet Forwarding Engine applicaiton (evo-pfemand). **[Error] compName = BrcmPlusMcast", tpName = "lrb", msg = " fn = installUnknownMcastLexemEntry, failed filter install for unit = 0, bdIndex = 3**.PR1748663
- In ACX7509, syncE clock failure is seen on 4x100G channelized interface. PR1756587

- Test chassis fpc slot <> power on or off command does not work on ACX7348 and ACX7332.PR1757288
- In scaled aggregated Ethernet cases during Routing Engine switchover when backup Routing Engine becomes primary, it takes upto 20 minutes for aggregated Ethernet interfaces statistics to be correctly reflected. The aggregated Ethernet traffic flows correctly and does not see any impact.PR1757394
- For heavily scaled VPLS scenarios (8K instances) , sometimes after a restart fpc command, a few MAC entries are not programmed in hardware. This issue is not seen in normal use case and with reduced scale.PR1759887
- On ACX7348 and ACX7332 Series running Junos OS Evolved multiple Routing Engine switchover can cause IDEEPROM failure on FPC and PSM, which results in PSM and FPC shown as unsupported. PR1760978
- In Junos OS Evolved 23.4 on ACX7348 or ACX7332 syslog messages related to PEX configuration are seen at bootup with **sib-slot** which should be corrected as **cb-slot** in later releases. Example: **Current log syslog(LOG_INFO, "%s: sib%d: pexsw bringup successful", __func__, board_slot); Expected log: syslog(LOG_INFO, "%s: cb%d: pexsw bringup successful", __func__, board_slot);**PR1761710
- On Junos OS Evolved ACX7024 and ACX7024X platforms, any MPLS encapsulated packet with a size less than or equal to 35 bytes on the 25G/10G/1G ports [Port No 4 to 27] gets dropped. The issue can also be observed in Point-to-Point Protocol over Ethernet (PPPoE) over Layer 2circuit scenario. PR1766889
- IPTV traffic is not forwarded to test agent. This is because route update comes before the filter add. Route update checks if netrounds filter is present. If so, punts the packet to netrounds queue. Otherwise it does not punt. Since filter add comes 10 minutes later, the traffic is not punted to netRounds Queue.PR1771527
- RC errors can be seen with SFP-T (Copper) optics due to vendor PHY limitation. PR1771671
- The following error message is seen with the delete vlan event trigger when configured with multiple VLANs with IRB and IGMP snooping/MLD snooping enabled. No functionality impact. **Error] compName = "BrcmPlusMcast", tpName = "Irb", msg = " fn = removeMemberFromL3FloodMcastGrp**.PR1771915
- On ACX7348 and ACX7332, during debugging of LMK alarm issue. There is a gap seen in the initial sequence of HW chips. There is no functional impact due to the issue. This is an internal fix to update the initial sequence as per discussion with HW team. PR1773212
- When the chassis is running headless (both the Routing Engines are removed from the system), peer interface remains up giving false indication that the link is up. PR1775785

- On ACX7509 platform having routing protocol with single-hop BFD (Bidirectional Forwarding Detection) configured over AE (Aggregated Ethernet) link and OAM (Operation, Administration, and Maintenance) CFM (Connectivity-Fault-Management) configured under unit 1 for the member interfaces, also if the member interfaces of the AE are in multiple line cards (one interface in 1-3 slots and other interfaces in 4-6 slots) might result in BFD session flaps and traffic impact due to programming issues. PR1776647
- 3% traffic drop is seen for some l2vpn instances after flapping the aggregated Ethernet interface which can be recovered by deactivating and activating that l2vpn instance in config. PR1777608
- RPD core can be see when running telemetry for protocols from top of tree and doing routing instance add delete operation. PR1778103
- We need to wait for the FPC and PIC to come online before the restart FPC command is executed. When the first restart command of the FPC is executed, it puts the FPC and the PIC offline and then online. Till both the FPC and PIC status is not shown online in `show chassis fpc pic-status` the next restart command for same FPC should not be attempted. PR1778163
- MPLS FRR with VPLS after multiple switchovers causes traffic to drop on a few VPLS instances. Deactivate/activate of the affected routing instances could recover the traffic. PR1779466
- On all Junos OS Evolved platforms, with Next-Generation Multicast Virtual Private Network (NG-MVPN) configured in the core and Ethernet Virtual Private Network (EVPN) in the access, multicast traffic from the EVPN sender is not forwarded to the NG-MVPN core as multicast next hop (NH) in Packet Forwarding Engine has some programming issues resulting in traffic loss. PR1781735
- During multiple Routing Engine switch-overs and router reboots, management port of one of the Routing Engines came up with 100 MBPS speed instead of 1 GBPS speed. There is no functional or stability impact due to this issue. Router monitoring might face a mild performance issue due to this lower speed of management port.PR1789156
- PICD core file can be seen during FPC offline/online, HA switchover, and system restart. PR1793824

Interfaces and Chassis

- On the ACX7024 platform, the tear-down rate is low. This is due to system CPU limitations.PR1659593

Network Management and Monitoring

- On all Junos Evolved platforms, a cosmetic change is made to an existing syslog message to print more information. The syslog error message that is logged when a user with wrong auth/privacy password sends a SNMP v3 request to router has been changed to add more information. **From: LIBJSNMP_NS_LOG_WARNING: WARNING: Authentication failed for <username> To: LIBJSNMP_NS_LOG_WARNING: WARNING: Authentication failed for <username> , SNMPv3 query from <NMS-ip> to <Router-IP>**PR1734549

  .

Routing Protocols

- BGP passing nh path's weight as 0 first and then passing nh path's weight as 1 later causes KRT to create 2 different next hops (as weight is now part of nh key) and results in installing same route in FIB as primary and backup. PR1745661
- When there are multiple events, the TI-LFA library seems to go into unnecessary multiple triggers. This can sometimes lead to pre-mature transition from MLA routes to global-convergence routes. However, we don't see this issue even when we do multiple link events manually. So, this might be some thing to do with so many multiple events intentionally/unintentionally caused by test-script. There is no service impact from this PR as it is only about pre-mature reversion to normal primary path. PR1794381

Resolved Issues

IN THIS SECTION

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

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- QSFP28-DD-2X100GBASE-LR4 link down on multiple FPC restarts/system reboots. PR1685520
- fibd core file is observed on Routing Engine switchover. PR1710227
- Multiple classification and rewrite support. PR1713158
- Continuous messages indicating duplicate IP address **L2ALM_DUPLICATE_IP_ADDR** is seen in MCLAG and VRRP scenario. PR1719868
- ACX7509: Major alarms getting reported for empty FPC slots. PR1721575
- Post changing MCAE mode from active-active to active-standby causes incomplete object state and system. PR1722626
- **/lib/systemd/system/docker.socket is marked executable** logs flood after system reboot. PR1727524
- IPv6 EBGP sessions might flap or be delayed in neighbour establishment due to hold timer expired error. PR1732443
- The evo-pfemand process crashes with simultaneous interface delete and statistics retrieval. PR1732077
- Timing protocols flap during a switchover event. PR1734457
- The inline IPv4 BFD session flaps on Junos OS Evolved platforms when hierarchical-scheduler is setup.PR1735836
- Control plane takes a long time to learn the multicast routes (scaled scenario). PR1736171
- After picd restart, traffic is not recovered on MACsec enabled ports. PR1738038
- JV DB misses leaf: **/interfaces/interface[name=\'ae0\']/state/counters/out-octets , out-pkts, out-unicast-pkts , out-broadcast-pkts , out-multicast-pkts , in-errors , out-errors, in-discards ,out-discards ,in-pause-pkts, out-pause-pkts** .PR1738395
- The PFE process crashes on Junos OS Evolved ACX platforms while deleting interfaces. PR1739175
- G.8273.2 noise generation tests fail for 400g ports (symmetric+asymmetric) .PR1739386

- DHCP daemon **jdhcpd** doesn't start in the new primary Routing Engine after GRES is performed. PR1740530
- IFD does not come up whenever optics is removed and inserted on all Junos OS Evolved platforms. PR1742772
- Transient multicast traffic drop on ACX Junos OS Evolved device. PR1742792
- Reachability for end host fails across Layer 2 circuit after enabling ECMP when the device acts as LSR. PR1743393
- Traffic loss observed on interface using ethernet-switching interface-mode trunk. PR1745163
- Bfd sessions flap after switchover. PR1746361
- Traffic loss observed in scenario where default route is received over multiple paths with link protection. PR1747512
- [timing] [ptp]: We observe the evo-pfemand error while changing the configuration from native vlan to dual vlan with PTPoLAG. PR1748089
- Traffic loss due to unknown multicast control packets getting dropped in non-default VRF. PR1748231
- In ACX7509, FEB 0 gets stuck in **Onlining Standby** after restarting it from CLI post jack out or in. PR1748450
- The authentication algorithm HMAC-SHA-256-128 for IPsec SA does not work and causes interoperability issues between Junos OS Evolved platforms and other devices PR1749779
- ARP dependency issue causes issue between IRB and the device. PR1751006
- GL: 16xSFP56: After ungraceful removal of 16xSFP56 in slot 2,3; packet drops are seen due to congestion. PR1752445
- ACX7100 : EvoPfemand core file at **PhyMarvelMacseclib::GetMacseclibPhyPortInfo** with evo-pfemand restart. PR1755883
- CPU host path not function after Routing Engine switchover. PR1756592
- Syslogs error will be seen while changing the port 32 to 4x1OG. PR1756606
- The evo-pfemand process is observed to crash on Junos OS Evolved platforms with LACP configuration. PR1756648
- ACX7509 : FEB-FPC-LINK error after admin disables on interfaces. PR1757125
- BFD sessions went down on BFDD restart with 1000 Scale BFD configuration. PR1757649

- After multiple reboot, we see rpdagent core on re0 primary. PR1761405
- Layer 2 loop can be seen on Junos OS Evolved ACX platforms after reboot. PR1765507
- Traffic impact is observed after device reboot or restart of picd on Junos OS Evolved ACX7100-48L. PR1766883
- xinted generates syslog messages **service ssh, accept: Invalid argument (errno = 22)** with high CPU usage. PR1767072
- ACX7100-32C - ifmon process 100% utilization. PR1768113
- IPv6 Neighbor Discovery not working resulting in traffic loss. PR1772838

Class of Service (CoS)

- Duplicate code points through code-point-aliases under a classifier results in cosd crash. PR1766873

EVPN

- Traffic drop is observed in EVPN-VXLAN CRB scenario. PR1734091
- BGP NH resolution must happen using locator and without extra policy at egress. PR1745991

Interfaces and Chassis

- All Junos OS Evolved ACX7024 stop transmitting LACP packet due to corruption of LACP packet PR1739254
- Multiple processes crash when more than 150 VLAN entries are configured in vlan-id-list under AE IFL. PR1774222

Routing Protocols

- The mscnoopd process crash is observed when snooping configuration is removed. PR1696374

- The RPD process will be stuck at a high CPU when OSPF areas are configured at a high scale and after starting the protocol. PR1728573
- The RPD process is stuck stuck at a high CPU when OSPF areas are configured at a high scale and after starting the protocol. PR1728573
- RPD crashes when multiple ISI-S processes are configured. PR1738222
- Traffic loss is seen in IPv6 only IS-IS topologies. PR1738901
- The rpd process crashes when routing is restarted. PR1751210

User Interface and Configuration

- The 'load replace' operation might result in mustd and mgd crash PR1740289
- Subsequent commits hang will be seen, when transfer-on-commit fails PR1752374

Junos OS Evolved Release Notes for PTX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.4R1 for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Packet Transport Routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

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

To view features supported on the PTX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 23.4R1, click the Group by Release link. You can collapse and expand the list as needed.

- PTX10001-36MR
- PTX10003

- PTX10004
- PTX10008
- PTX10016

The following sections highlight the key features in this release.

Hardware

- **New AC PSU and Active Blank for PTX Series Routers**—Starting in Junos OS Evolved Release 23.4R1, we introduce a new AC Power Supply Unit or PSU (JNP10K-PWR-AC3), and active blank (JNP10K-PWR-BLN3) for PTX10004 and PTX10008 routers.

  The new JNPR10K-PWR-AC3 power supply is a high capacity model that is designed to support AC systems in a 15-A and 20-A mode.

  The JNP10K-PWR-BLN3 active blank, as part of the power supply, helps in airflow and cooling in the PTX router.

  [See PTX10004 Power System and PTX10008 Power System.]

Authentication and Access Control

- **OpenSSH certificate support (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can set up SSH access to a device with password-less login for users. You can also trust hosts without the need to verify the key fingerprints.

  Use the following new CLI configuration statements 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 **HostCertificate** file at **/etc/ssh/sshd_config**, which contains the signed host certificate.
  - `system services ssh authorized-principals-file` *`filename`*—Configure the **AuthorizedPrincipals** file 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 **AuthorizedPrincipals** file.

  [See SSH Certificate-Based Authentication Overview.]

- **SSH Hostkey Algorithm Update (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, the `hostkey-`

algorithm SSH configuration options has been replaced with hostkey-algorithm-list, and the ecdsa-sha2-nistp384 and ecdsa-sha2-nistp521 hostkey algorithms are now supported.

You can find the hostkey-algorithm-list configuration option at the [edit system services ssh] hierachy level.

[See hostkey-algorithm.]

- **Background File Transfer for SCP/SSH (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can transfer files in the background via SCP/SSH. To configure background file transfers, include the archive-sites configuration statement at the [edit system archival configuration] hierarchy level.

  [See Back Up Configurations to an Archive Site.]

- **Control device access privileges with exact match configuration (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**

  —Starting in Junos OS Evolved Release 23.4R1, you can configure access privileges for login classes by allowing or denying full hierarchy strings with the allow-configuration-exact-match and deny-configuration-exact-match configuration options. The exact match configuration enables you to set separate permissions for set, delete, activate, or deactivate operators for any hierarchy.

  The allow-configuration-exact-match and deny-configuration-exact-match configuration options support full hierarchy strings as well as wildcard characters and regular expressions.

  [See Understanding Exact Match Access Privileges for Login Classes.]

Chassis

- **Packet Forwarding Engine-level recovery to support fabric resiliency (PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved Release 23.4R1, we've enhanced the fabric resiliency feature support by including a recovery configuration with Packet Forwarding Engine restart. Before you enable the enhanced fabric recovery at the Packet Forwarding Engine level, you must disable the default Packet Forwarding Engine restart action. To disable the default action, use the set chassis fabric event reachability-fault actions pfe-restart-disable command. To view the result of the configuration, run the show chassis fabric reachability detail command.
  [See reachability-fault, show chassis fabric reachability, and Fabric Resiliency.]

- **Support to monitor 100G, 400G optics temperature (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we support the following EM-related features:

  - 100G and400G optics firmware, which triggers optics shutdown when the high-temperature threshold is breached.

- Use of the `set chassis temperature-threshold` command to optionally support the threshold configuration for optics temperature sensors.
- Explicit disabling of the EM policy on WAN ports. By default, the EM policy is enabled on all 100Gand 400G. optics interfaces other than optics with DAC cables.Use the `set chassis fpc` *`fpc_slot`* `pic` *`pic_slot`* `port` *`port_no`* `no-temperature-monitoring` command to explicitly disable the EM policy on specific WAN ports.
- The following output is displayed in `show log messages` command output when you enable no-temperature-monitoring: `PICD_OPTICS_EMPOLICY_NOTIFICATION: et-1/0/29: Optics Temperature monitoring is disabled.`

  [See No Link Title and No Link Title.]

- **Platform and resiliency support (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)** – Starting in Junos OS Evolved Release 23.4R1, we support:
  - Major and minor alarms and repair during operational and diagnostic image bootup.
  - Lane remapping if a lane fails and hard repair information is preprogrammed.
  - Repair in bootup images when enabled.
- **Optics temperature monitoring and fan speed adjustment (PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, we support the EM policy to monitor 100G and 400G optics temperature sensors in PTX series devices. The support for the EM policy includes:
  - 100G, 400G optics firmware, which triggers optics shutdown when the high-temperature threshold is breached.
  - Use the `set chassis temperature-threshold` command to optionally support the fan speed adjustment for optics temperature management.
  - Use the `set chassis fpc` *`fpc_slot`* `pic` *`pic_slot`* `port` *`port_no`* `no-temperature-monitoring` command to explicitly disable the EM policy on specific WAN ports.
  - The following output is displayed in `show log messages` command output when you enable no-temperature-monitoring: `PICD_OPTICS_EMPOLICY_NOTIFICATION: et-1/0/29: Optics Temperature monitoring is disabled.`
  - By default, the EM policy is enabled on all 100G, 400G optics interfaces other than optics with DAC cables. [See No Link Title and No Link Title.]
- **IPv6 support for protocols (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting Junos OS Evolved Release 23.4R1, PTX series devices support IPv6 for SSH, Telnet, SCP, FTP, NETCONF over SSH, TLS, telemetry, and apps over gRPC.

  [See Junos Telemetry Interface User Guide.]

- **Source Redundancy and Feed Redundancy support on PTX10004 and PTX10008**–Starting in Junos OS Evolved Release 23.4R1, N+1 power redundancy is supported on PTX10004 and PTX10008 routers with JNP10K-PWR-AC3 power supply modules (PSMs). You can enable either source redundancy or feed redundancy for the PSM.

  [See Managing Power.]

- **Resiliency support (PTX10004 and PTX10008)** –Starting in Junos OS Evolved Release 23.4R1, the FRU resiliency support is provided on the PTX10004 and PTX10008 platforms with JNP10K-PWR-AC3 PSMs.

  [See Fabric Resiliency.]

EVPN

- **EVPN-VXLAN pure T5 host-route auto-generated community (ACX7024, ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, we've added support for the EVPN-VXLAN pure T5 host-route auto-generated community. This feature adds a community to MAC-IP ARP/NDP-based pure Type 5 host routes. Border leaf devices in edge-routed bridging (ERB) topologies with Type 5 connectivity to other leaf devices in the data center and Type 5 connections to external networks need to advertise aggregate routes to the external network instead of individual Type 5 routes. Border leaf devices can use this community to identify these routes and create an aggregate route to advertise to external EVPN networks.

  [See EVPN-VXLAN Pure T5 Host-Route Auto-Generated Community.]

- **OISM PEG DF election (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, you can customize the designated forwarder (DF) election method on border leaf devices that act as Protocol Independent Multicast (PIM) EVPN gateway (PEG) devices in an EVPN network running optimized intersubnet multicast (OISM). You can configure peer PEG devices to use one of the following DF election methods:

  - Mod-based DF election
  - Preference-based DF election with a specified preference value

  By default, a PEG device uses PIM-based DF election. When you configure this feature, the selected PEG DF election method replaces PIM-based DF election.

  [See peg-df-election and PEG DF Election.]

- **Static configuration of MAC-IP bindings with EVPN-VXLAN (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we've added the functionality to allow static configuration of MAC-IP bindings on an interface, similar to configuring static MACs on an interface. This feature enables the static configuration of IP and MAC entries for crucial services provided by management and infrastructure hosts. It proves particularly

advantageous in Internet Exchange Point (IXP) networks where participant customer edge (CE) devices remain well-known and static, not transitioning to different provider edge (PE) devices.

You can now utilize a new feature that establishes a static link between an IP address and a MAC address for a logical interface within a bridge domain or VLAN. When you provision a static MAC-IP entry on a PE device, the PE device initiates a probe following an exponential backoff pattern. The probe uses an all-zero sender IP address on the associated interface. If the entity owning the IP to MAC entry responds to the probe, the system will learn the IP to MAC binding as static. Subsequently, it will be propagated to remote PE devices through the BGP/EVPN Type 2 MAC advertisement route. The corresponding MAC address will be recognized as a dynamic entry. If you want to deactivate the probing mechanism for learning the IP to MAC binding, you can do so by configuring the new configuration option `arp-nd-probe-disable`. Without probing, both the MAC and IP to MAC binding will be acquired from network traffic and communicated using EVPN.

We've introduced the following commands and configuration statements:

- Configuration of static IP to MAC bindings

  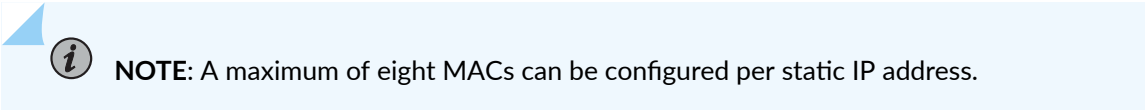

  NOTE: A maximum of eight MACs can be configured per static IP address.

  The aforementioned commands provide an option to configure `router` and `override` bits for IPv6 entries. For example:

- Disable probing on configuration of static IP to MAC entries:

  To turn off the default probing on configuration of static IP to MAC entries, you can use the global configuration statement `arp-nd-probe-disable`.

  ```
  set protocols l2-learning arp-nd-probe-disable
  ```

- Enable logging for failed probing of static IP to MAC entries:

  To turn on the logging, configure the global configuration statement `arp-nd-probe-failed-log`.

  ```
  set protocols l2-learning arp-nd-probe-failed-log
  ```

- Enable GARP/unsolicited-NA for local and remote static entries

  If this feature is required, you must configure the global configuration statement `garp-na-enable`.

  ```
  set protocols l2-learning garp-na-enable
  ```

- Disable dynamic learning [all static provisioning]

  If dynamic learning of MAC-IP entries is not required, configure the statement `drop-unknown-macip` under BD/VLAN.

- Drop unicast ARP request

  To drop unicast address resolution requests (for instance, NUD NS messages), you can configure the statement `block-unicast-arp` at global level.

- [See EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression and interface-mac-ip-limit.]

High Availability

- **NSR faster detection of connection slowness on replication layer (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, nonstop active routing (NSR) detects connection slowness on the replication layer faster. Faster detection helps reduce the negative impact that replication layer issues can have on NSR sessions, especially those with low hold timers.

  [See Nonstop Active Routing Concepts for Junos OS Evolved.]

- **NSR switchover support with low hold timers (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we support 4k BGP sessions with hold timers as low as 3 seconds on nonstop active routing (NSR) switchovers, depending on the number of BGP sessions. See the following table for timer details for the number of sessions.

BGP Hold-Timer Value (in seconds)	Number of BGP Sessions
3	80
5	120
10	200
15	400
30	600
45	800
90	1800

  [See hold-time and precision-timers.]

- **BFD session damping for LACP interfaces (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can use BFD session damping on LACP interfaces to suppress BFD session state change notifications for a configured time period when thresholds for session flapping are exceeded. Session damping helps reduce potential instability from excessive BFD notifications.

  Use the `set bfd-liveness-detection damping` configuration statement at the `[edit dynamic-profiles` *`name`* `interfaces` *`name`* `aggregated-ether-option]` hierarchy level to configure BFD session damping.

  [See BFD Session Damping Overview.]

Interfaces

- **Support for port bounce (PTX Series)**—Starting in Junos OS Evolved Release 23.4R1, you can shut down an interface for a given time by using the `request interface bounce` *`interface_name`* `interval` *`seconds`* command. The interface comes up at the end of the configured time.

  [See request interface bounce.]

- **400ZR and 400G Open ZR+ optical transceivers:**[See 400ZR and 400G OpenZR+.]

- **Signal power performance parameter (PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, the `show interface diagnostics optics et-<x/y/z>` command shows the Rx signal power in the mW/dBm format.[See 400ZR and 400G OpenZR+.]

IPv6

- **Optimizing ARP, NDP and Default-Route handling in internal DB of DCD (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.4R1, DCD only deletes routing entries for addresses that are completely unlinked from all associated addresses. Additionally, we introduce checks to prevent configuring multiple static MAC addresses for a single ARP and NDP address, which helps improve system stability and avoid potential conflicts in network configurations.

Junos OS API and Scripting

- **Support for REST API functionality over nondefault VRF instances (ACX Series and PTX Series)**—

  Starting in Junos OS Evolved Release 23.4R1, we support REST API functionality with HTTPs over nondefault virtual routing and forwarding (VRF) instances. To hold all the configurations related to this feature, enable the `routing-instance` statement at the `[edit system services]` hierarchy level.

Junos Telemetry Interface

- **New state data model for Juniper proprietary Remote Procedure (gRPC) service (PTX10001-36MR, PTX10003, PTX10008, PTX10016)**—Junos OS Evolved Release 23.4R1 contains a restructured

native state data model defining gRPC server instances. The new model includes common attributes and gRPC Network Management Interface (gNMI) service details.

The sensor **/state/system/services/http/servers/** and its leaves illustrate the new structure.

[For sensors, see Junos YANG Data Model Explorer.]

- **400G ZR+ optic enhancement for signal power statistics (PTX10001-36MR)**—Junos OS Evolved Release 23.4R1 supports enhanced signal power statistics for 400G ZR+ optics. Use the sensor **/components/component/optical-channel/state/input-power/** to export signal power statistics from a Juniper device to an outside collector. This feature supports the OpenConfig data model **openconfig-terminal-device.yang version** (version 1.9.0).

  [For sensors, see Junos YANG Data Model Explorer.]

- **Resource Public Key Infrastructure (RPKI) enhanced streaming telemetry support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports enhanced statistics for RPKI databases and RPKI sessions and validation-related statistics per route, per RIB and per BGP peer basis. Using these statistics, you can perform operational debugging on your network and take appropriate mitigating actions.

  These existing Junos operational mode commands contain new statistics:

  - `show route [extensive|detail]` displays origin validation information for each route entry
  - `show bgp neighbor validation statistics <peer>` displays BGP peer-RIB validation statistics
  - `show route validation-statistics` displays local routing information base (RIB) specific validation statistics
  - `show validation statistics` displays new counters for the VRP table

  These telemetry sensors (with leaves) are supported:

  - **/state/routing-instances/routing-instance/protocols/bgp/rib/afi-safis/afi-safi/[ipv4|ipv6]-unicast/loc-rib/routes/route/origin-validation-state**
  - **/state/routing-instances/routing-instance/protocols/bgp/rib/afi-safis/afi-safi/[ipv4|ipv6]-unicast/loc-rib/routes/route/origin-validation-invalid-reason**
  - **/state/routing-instances/routing-instance/protocols/bgp/groups/group/neighbors/neighbor/afi-safis/afi-safi[ipv4|ipv6]/validation-counters/**
  - **/state/routing-instances/routing-instance/protocols/bgp/groups/group/neighbors/neighbor/afi-safis/afi-safi[ipv4|ipv6]/validation-counters**
  - **/state/routing-instances/routing-instance/protocols/bgp/rib/afi-safis/afi-safi/[ipv4|ipv6]-unicast/loc-rib/validation-counters/**

- **/state/routing-instances/routing-instance/routing-options/route-validation/rpki-rtr/groups/group/sessions/session/rpki-session-counters/**
- **/state/routing-instances/routing-instance/routing-options/route-validation/route-validation-databases/route-validation-database/[ipv4|ipv6]/**
- **/state/routing-instances/routing-instance/routing-options/route-validation/rpki-rtr/groups/group/sessions/session/**

[For sensors, see Junos YANG Data Model Explorer.] For operational mode commands, see show route, show bgp neighbor validation statistics, show route validation-statistics, and show validation statistics.

- **Multicast telemetry support with IGMP and PIM operational state sensors (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports IGMP and PIM sensors based on the OpenConfig data models **openconfig-igmp.yang** (version 0.3.0) and **openconfig-pim.yang** (version 0.4.2).

  [For sensors, see Junos YANG Data Model Explorer.]

- **Telemetry streaming of operational state data for syslog messages (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports telemetry streaming of operational state data for syslog messages to an external gRPC Network Management Interface (gNMI) collector. Sensors are based on the native Junos data model under the **/state/system/syslog/messages** hierarchy. You can stream data using `ON_CHANGE` and `TARGET_DEFINED` modes.

  [See Junos YANG Data Model Explorer.]

- **SR-TE policy telemetry (ACX7348, PTX10003, PTX10004, PTX10008, and PTX10016)**—In Junos OS Evolved Release 23.4R1, we introduce support for telemetry streaming of operational state data for the segment routing–traffic engineering (SR-TE) policy. State sensors are based on the OpenConfig data model **openconfig-srte-policy.yang**. You can subscribe to SR-TE sensors using the resource path **/network-instances/network-instance/segment-routing/te-policies**.

  [See Junos YANG Data Model Explorer.]

- **FIB sensor support (PTX10008)**—Junos OS Evolved Release 23.4R1 supports the OpenConfig-Abstract Forwarding Table (oc-aft) model and the forwarding information base (FIB) to stream enhanced routing statistics. To deliver statistics to a collector, you add sensors to a subscription and also include the **set routing-options forwarding-table oc-tlv-support** statement at the [**edit**] hierarchy level to enable statistics collection. The Junos routing protocol process (rpd) sends the `origin-protocol` and `origin-network-instance` of a route, as well as the next hop using an opaque TLV to the collector. Include the following sensors in your subscription:

  - Next hops:
    - **/network-instances/network-instance/afts/next-hops/next-hop/state/pop-top-label**

    - **/network-instances/network-instance/afts/next-hops/next-hop/state/vni-label**
    - **/network-instances/network-instance/afts/next-hops/next-hop/state/vni-label**
    - **/network-instances/network-instance/afts/next-hops/next-hop/ip-in-ip/state/dest-ip**
  - State-synchronized:
    - **/network-instances/network-instance/afts/state-synced/state/ipv4-unicast**
    - **/network-instances/network-instance/afts/state-synced/state/ipv6-unicast**)
  - IPv4 and IPv6 unicast:
    - **/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/origin-network-instance**
    - **/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/origin-network-instance**)

  [For statement support, see forwarding-table. For state sensors, see Junos YANG Data Model Explorer.]

- **Hardware temperature alarm sensor support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports the data model **openconfig-platform.yang** (version 0.12.2) to stream granular hardware alarm temperature statistics. At present, temperature states are streamed assuming that there is one temperature sensor representing the whole physical component. However, a component can report alarms not only for the primary temperature sensor but also for multiple temperature sensors. We have created additional temperature sensors to determine alarm origins.

  For example, if you subscribe to **/components/component[name='CHASSIS0:FPC0:PEX_Temp_Sensor']/**, the data exported to the collector could look like this:

```
kv {
  key:__prefix__,
  str_value:/components/component[name='CHASSIS0:FPC0:PEX_Temp_Sensor']/
}
kv {
  key:state/name,
  str_value:CHASSIS0:FPC0:PEX_Temp_Sensor'
}
kv {
  key:state/type,
  str_value:SENSOR
```

```
}
kv {
  key:state/parent,
  str_value:CHASSIS0:FPC0
}
kv {
  key:state/removable,
  bool_value:False
}
kv {
  key:state/temperature/instant,
  double_value:78.0
}
kv {
  key:state/temperature/avg,
  double_value:77.0
}
kv {
  key:state/temperature/min,
  double_value:74.0
}
kv {
  key:state/temperature/max,
  double_value:80.0
}
kv {
  key:state/temperature/interval,
  uint_value:60000000000
}
kv {
  key:state/temperature/min-time,
  uint_value:1663653284020
}
kv {
  key:state/temperature/min-time,
  uint_value:1663663284020
}
kv {
  key:state/temperature/alarm-status,
  bool_value:False
}
kv {
  key:state/temperature/alarm-threshold,
```

```
  uint_value:107
}
kv {
  key:state/temperature/alarm-severity,
  str_value:WARNING
}
```

[For state sensors, see Junos YANG Data Model Explorer. For information about creating a subcription, see Configure a NETCONF Proxy Telemetry Sensor in Junos.]

- **IS-IS operational state sensors and configuration using OpenConfig (PTX10008 and PTX10016)**—Junos OS Evolved Release 23.4R1 introduces enhancements to IS-IS telemetry support based on the OpenConfig data model **openconfig-isis.yang** (version 1.0.0). Support includes new operational state paths and configuration paths.

  [For OpenConfig configuration, see Mapping OpenConfig ISIS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **IS-IS OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports OpenConfig IS-IS configurations and sensors based on the OpenConfig data model **openconfig-isis.yang** (version 1.0.0). This feature closes some gaps in our OpenConfig configuration and sensor support in the IS-IS area.

  [For OpenConfig configuration, see Mapping OpenConfig ISIS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **MPLS OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10003, PTX10004, and PTX10008)**—Junos OS Evolved Release 23.4R1 supports OpenConfig MPLS configurations and sensors based on the OpenConfig data models **openconfig-mpls.yang** (version 3.2.2), **openconfig-mpls-types.yang** (version 3.2.1), and**openconfig-mpls-te.yang** (version 3.2.2). The following OpenConfig configurations and state sensors are supported.

  - Configurations:

    - MPLS global TTL propagation (**/network-instances/network-instance/mpls/global/config/ttl-propagation**)

    - MPLS LSP PRI/SEC path-metric-bound-constraint (**/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/p2p-tunnel-attributes/p2p-primary-path/p2p-primary-path/path-metric-bound-constraints/path-metric-bound-constraint/config/**)

  - Sensors:

    - MPLS global (**/network-instances/network-instance/mpls/global/state/**)

- MPLS global interface attributes (**/network-instances/network-instance/mpls/global/interface-attributes/**)
- MPLS LSP autobandwidth (**/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/bandwidth/auto-bandwidth/state/**)
- MPLS LSP PRI/SEC path-metric-bound-constraint (**/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/p2p-tunnel-attributes/p2p-primary-path/p2p-primary-path/path-metric-bound-constraints/path-metric-bound-constraint/state/**)
- MPLS traffic engineering global attributes SRLG (**/network-instances/network-instance/mpls/te-global-attributes/srlgs/srlg/static-srlg-members/members-list/state/**)

[For OpenConfig configuration, see Mapping MPLS OpenConfig MPLS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **MPLS OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports OpenConfig MPLS configurations and sensors based on the OpenConfig data models **openconfig-mpls-ldp.yang** (version 3.2.0) and **openconfig-mpls-rsvp.yang** (version 4.0.0). This feature closes some gaps in our OpenConfig configuration and sensor support in the MPLS RSVP-TE and MPLS LDP areas.

  [For OpenConfig configuration, see Mapping MPLS OpenConfig MPLS Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **OpenConfig support for local station MAC address (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports configuration and streaming of the local station MAC address using OpenConfig.

  [For state sensor, see Junos YANG Data Model Explorer. For OpenConfig configuration, see Mapping OpenConfig System Management Model Commands to Junos Configuration.]

- **Platform reboot and switchover sensor support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports the data model **openconfig-platform.yang** (versions 0.15.0, 0.17.0, and 0.18.0) to stream state data and statistics related to component reboot and switchover activities.

  [For state sensors, see Junos YANG Data Model Explorer.]

- **QoS configuration and streaming with OpenConfig (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports additional leaves for the configuration and streaming of QoS statistics with OpenConfig for the following features:
  - MPLS classifiers
  - Classifier interface binding

- Rewrite rules
- Scheduler policies

[For state sensors, see Junos YANG Data Model Explorer.

For OpenConfig configuration, see Mapping OpenConfig QoS Commands to Junos Configuration.]

- **Streaming telemetry for hardware components using OpenConfig (PTX10001-36MR, PTX10004, PTX10008 and PTX10016)**—Junos OS Evolved Release 23.4R1 supports telemetry streaming for the following leaves based on the OpenConfig data model **openconfig-platform.yang** (version 0.19.0):

  - **/components/component/fabric/state/power-admin-state**
  - **/components/component/controller-card/state/power-admin-state**

  Support for the following leaves will be extended to all relevant hardware components:

  - **/components/component/state/mfg-name**
  - **/components/component/state/type**

  [See Junos YANG Data Model Explorer.]

- **Support for YANG data models in a single package and repository (PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, we provide all YANG data models for a given OS and release in a single download package and GitHub repository. The package and repository include the native configuration, state, and RPC data models and the OpenConfig and IETF data models that the OS supports. You can access the YANG data models from the Juniper/yang GitHub repository or the Juniper Networks download site.

  [See Understanding Junos YANG Modules and GitHub Resources.]

- **STP OpenConfig and operational state sensor support (ACX7024, ACX7100-32C, ACX7100-48L,ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports OpenConfig STP configurations and sensors based on the OpenConfig data model **openconfig-spanning-tree** (version 1.0, revision 0.3.1).

  [For OpenConfig configuration, see Mapping OpenConfig STP Commands to Junos Configuration. For state sensors, see Junos YANG Data Model Explorer.]

- **Telemetry streaming of operational state data for syslog messages (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**—Junos OS Evolved Release 23.4R1 supports telemetry streaming of operational state data for syslog messages to an external gNMI collector. Sensors are based on the native Junos data model under the **/state/system/syslog/messages/** hierarchy. You can stream data using `ON_CHANGE` and `TARGET_DEFINED` modes.

  [See Junos YANG Data Model Explorer.]

- **Upgrade of OpenConfig models for routing instances (ACX7100-32C, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.4R1 supports an upgrade for the following OpenConfig models:

  - **openconfig-local-routing.yang** to version 2.0.0.

  - **openconfig-routing-policy.yang** to version 3.3.0.

  The upgraded models introduce new leaves for operational state sensors and configuration in the following areas:

  - Inter-instance policies

  - Route limits

  - Router advertisement

  - Local aggregates

  - Static routes

  [For state sensors, see Junos YANG Data Model Explorer.

  For OpenConfig configuration, see Mapping OpenConfig Network Instance Commands to Junos Configuration.]

MPLS

- **Support for MLDP recursive forwarding equivalence class (FEC) (PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, we partially support RFC 6512. We've introduced the recursive opaque value type for the MLDP forwarding equivalence class (FEC) element. The recursive opaque value helps to form MLDP point-to-multipoint (P2MP) tunnels between two autonomous systems (ASs), where the intermediate nodes do not have the route to reach the root node.

  To enable the recursive opaque value, configure the `fec` statement at the `[edit protocols ldp p2mp recursive]` hierarchy level.

  [See Understanding Multipoint LDP Recursive FEC.]

- **Computation of unreserved bandwidth optimized RSVP dynamic bypass LSP (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, the Constrained Shortest Path First (CSPF) can optionally use a different approach to protect a link or a node by leveraging the computation based on unreserved bandwidths on traffic engineering (TE) links. To enable this feature, use the `optimize bandwidth` configuration statement at the `edit protocols rsvp interface` *`interface`* `link-protection` hierarchy level. While the default approach of RSVP bypass produces a bypass method that optimizes traffic engineering (TE) metric, enabling the new configuration statement maximizes the end-to-end unreserved bandwidth.

[See Configuring Link Protection on Interfaces Used by LSPs.]

- **Capability to compute diverse paths between a set of LSPs (PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, you can associate a group of LSPs (RSVP LSPs or SR MPLS LSPs) to the Path Computation Element Communication Protocol (PCEP) to compute diverse paths for the associated LSPs. The Junos PCC advertises to a Path Computation Element (PCE) that a particular LSP belongs to a diversity-association group. RFC 8800 defines PCEP protocol extensions to associate a set of LSPs that belong to the same association group. This enables a PCE to compute diverse paths for each of the LSPs in each diversity association group and then push the results to the PCC. A PCE can also associate set of LSPs across different PCCs.

  You can enable diversity-association capability in the open message by configuring the following statement:

  ```
  user@host# set protocol pcep diversity-association-capability
  ```

  After enabling diversity-association capability, you need to also configure the diversity-association group using the following statement:

  For RSVP LSPs:

  ```
  user@host# set protocols mpls label-switched-path lsp-name lsp-external-controller pccd
  diversity-association group group-name
  ```

  For SR LSPs:

  ```
  user@host# set protocols source-packet-routing source-routing-path lsp-name diversity-
  association group group-name
  ```

  You can provision and delegate the following LSPs:

  - Provision PCE initiated RSVP LSPs with diverse paths
  - Delegate RSVP LSPs with diverse association groups
  - Provision PCE initiated SR MPLS uncolored LSPs with diverse paths
  - Delegate SR MPLS uncolored LSPs with diverse association groups
  - Provision PCE initiated SR MPLS colored LSPs with diverse paths
  - Delegate SR MPLS colored LSPs with diverse association groups
  - Provision PCE initiated SRv6 colored LSPs with diverse paths

- Delegate SRv6 colored LSPs with diverse association groups
- Provision PCE initiated SRv6 uncolored LSPs with diverse paths
- Delegate SRv6 uncolored LSPs with diverse association groups

[See PCEP Configuration.]

- **Map static IPv6 route to next-hop using service label (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024,PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can enable static IPv6 routes to be mapped to the next-hop over an IPv4 MPLS network. 6PE is a transitional IPv6 over IPv4 technology that uses MPLS tunnels to carry services.

  You can use the `explicit-null` configuration statement under the [`edit routing-options rib inet6.0 static route` *`ipv6-address`*] hierarchy level to push ingress service label as part of the static next hop configuration for static IPv6 routes. The `explicit-null` configuration statement only supports configuring IPv4 mapped IPv6 address.

  The `static` configuration statement under the [`edit routing-options forwarding-table chained-composite-next-hop ingress`] hierarchy provisions chained composite next-hop.

  > **NOTE**: The `static` configuration statement must be enabled before configuring the `explicit-null` configuration statement.

- **PCE requests to allocate binding SIDs for SR-TE Colored LSPs (PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, a Path Computation Element (PCE) can request Path Computation Client (PCC) to allocate a binding SID from PCC's label space. PCE can request PCC to allocate a specific binding SID and can also allocate binding SID of PCC's choice.

  The following PCEP operations are now supported:

  - PCE requests PCC to allocate binding SID of PCC's choice for delegated LSPs
  - PCE requests PCC to allocate binding SID of PCC's choice for PCE-Initiated LSPs
  - PCE requests PCC to allocate a specific binding SID for delegated LSPs
  - PCE requests PCC to allocate a specific binding SID for PCE-Initiated LSPs
  - BGP LS for binding SID for colored SR LSP

  The following SRTE binding SID database and label show commands has been introduced to display all binding SIDs with brief and detail outputs:

  - `show spring-traffic-engineering binding-sid database brief`
  - `show spring-traffic-engineering binding-sid database detail`

- `show spring-traffic-engineering binding-sid database label label brief`
- `show spring-traffic-engineering binding-sid database label label detail`
- `show spring-traffic-engineering binding-sid label label brief`
- `show spring-traffic-engineering binding-sid label label detail`

[See PCEP Configuration.

- **New CLI commands for MPLS LSPs (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.4R1, you can get more visibility into the current state of the MPLS LSPs on the router to debug suspected anomalies in high scale conditions with the following newly introduced CLI commands.
  - `show rsvp session bypass [bypass-name] [protected]` and `show rsvp session [unprotected]` provides visibility into LSPs protected by a specific bypass tunnel.
  - `show mpls lsp [make-before-break]` and `show rsvp session [multiple-lsp-sessions]` provides sisibility into LSPs undergoing make-before-break.
  - `show mpls tunnel-manager-statistics` provides statistics on all local repair and make-before-break events for LSPs.
  - `show rsvp session [fr-ingress]` provides visibility into LSPs on flood-reflector edge routers.
- **PCC Policy Association with SR and RSVP LSP (PTX10008)**—Starting in Junos OS Evolved 23.4R1, PCC (Path Computation Clients) can link policies with a group of Label Switched Paths (LSPs). This enhancement allows Junos PCC to communicate with a Path Computation Element (PCE) using an extended communication protocol (PCEP). Through this extension, Junos PCC can tell the PCE that a specific LSP is part of a certain Policy Association Group.

Multicast

- **IGMP and MLD snooping version configuration (ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can configure the version of IGMP or MLD snooping queries for VLANs or bridge domains associated with Layer 2 (L2) multicast. This configuration ensures that end hosts or CPE devices that are not compliant with RFC 4541 and can't normally respond to snooping queries of a later version are now able to process and respond to those snooping queries.

  To configure the IGMP or MLD snooping version, use the following CLI statements:

  - `set protocols igmp-snooping version version`
  - `set protocols mld-snooping version version`

  [See IGMP MLD Snooping Version Configuration.]

Network Management and Monitoring

- **Junos OS Evolved YANG modules published (PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, we publish the native YANG data models for Junos OS Evolved. You can access the YANG data models from the Juniper/yang GitHub repository or the Juniper Networks download site.

  [See Understanding Junos YANG Modules.]

- **Support for YANG data models in a single package and repository (PTX10008)**—Starting in Junos OS Evolved Release 23.4R1, we provide all YANG data models for a given OS and release in a single download package and GitHub repository. The package and repository include the native configuration, state, and RPC data models and the OpenConfig and IETF data models that the OS supports. You can access the YANG data models from the Juniper/yang GitHub repository or the Juniper Networks download site.

  [See Understanding Junos YANG Modules and GitHub Resources.]

OpenConfig

- **OpenConfig translation and telemetry support for sFlow (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we introduce OpenConfig translation and telemetry support for sFlow based on the following data models:

  - `openconfig-sampling.yang`
  - `openconfig-sampling-sflow.yang`.

  These models define top-level configuration and operational state data related to traffic sampling. For modularity purposes, the top-level sampling container provides a natural attachment point for implementations such as sFlow.

  [See OpenConfig User Guide.]

Routing Options

- **Support for configuring route priority for BGP static routes and route prioritization during reconfiguration (PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, you can configure a route priority for static routes. Include the `priority` statement at the `[edit routing-options static route destination next-hop]` hierarchy level. In addition, when you perform a route reconfiguration, a new routing table policy mechanism ensures that routes are processed based on the configured priority.

  [See **BGP Route Prioritization**.]

Routing Policy and Firewall Filters

- **Support for EVPN-VXLAN filtering and port mirroring based on VNI match conditions (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release

23.4R1, you can construct a firewall filter to filter EVPN-VXLAN traffic by using the VXLAN network identifier (VNI) values in the match condition on ingress interfaces. This feature supports redirecting traffic to a global port-mirroring instance.

To filter traffic based on the VNI, use the following commands:

- `set firewall filter filter-name term term-name from vxlan vni vni-value`
- `set firewall filter filter-name term term-name from vxlan vni-except vni-value`

  vni-value can be a numeric value or range of numeric values.

[See Firewall Filter Match Conditions and Actions (PTX Series Routers).]

Routing Protocols

- **BGP triggered dynamic IP encapsulation (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can configure the new attributes for GRE and UDP tunnel types:

  - IPv4 and IPv6 payload support for UDP tunnels
  - GRE tunnels

  We've added the following configuration statements at the `[edit routing-options tunnel-attributes]` hierarchy level:

  - `dynamic-tunnel-ttl` statement to set the set the tunnel TTL value (the default value is 0)
  - `dynamic-tunnel-payload-port-profile` statement to set the global UDP port profile
  - `dynamic-tunnel-forwarding-class-name` statement to set the forwarding class name

  We've also added the `gre` configuration statement at the `[edit routing-options dynamic-tunnels]` hierarchy level to set the encapsulation as GRE.

- **Support for EIBGP multipath ECMP for defined prefixes (PTX Series)**—Junos OS Evolved Release 23.4R1 supports EBGP and IBGP (EIBGP) multipath. In the existing BGP multipath, EBGP routes take priority over IBGP routes because both have different metrics. After you enable EIBGP multipath and there is equal load sharing between the EBGP and IBGP routes, Junos OS Evolved initiates ECMP using a blend of both EBGP and IBGP.

  Feature-specific policies specify prefixes that support EIBGP multipath. You can configure the policy to choose the prefixes based on any match condition.

  To enable EIBGP multipath, configure the `allow-external-internal` option at the `[edit protocols bgp multipath]` or `[edit logical-systems logical-system-name protocols bgp multipath]` hierarchy level.

  [See multipath (Protocols BGP).]

- **Support for OSPFv2 HMAC SHA-1 keychain authentication and optimization for multi-active MD5 keys (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, you can enable OSPFv2 HMAC-SHA1 authentication with keychain to authenticate packets reaching or originating from an OSPF interface. This feature ensures a smooth transition from one key to another for OSPFv2 with enhanced security.

  You can enable OSPFv2 to send packets authenticated with only the latest MD5 key after all the neighbors switch to the latest configured key. In Junos OS Evolved releases earlier than Release 23.4R1, we support advertising authenticated OSPF packets always with multiple active MD5 keys with a maximum limit of two keys per interface.

  To enable OSPFv2 HMAC-SHA1 authentication, configure the `authentication keychain <keychain name>` option at the `[edit protocols ospf area area-id interface interface_name` hierarchy level. To enable optimization of multiple active MD5 keys, configure the `delete-if-not-inuse` option at the `[edit protocols ospf area area-id interface interface_name authentication multi-active-md5]` hierarchy level.

  [See **Understanding OSPFv2 Authentication**.]

- **Support for Next-Hop Dependent Capability Attribute (ACX7100-32C and PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, we use the Entropy Label Capability (ELCv3) attribute defined within the IETF Next-Hop Dependent Capability Attribute for load balancing. This attribute replaces the existing ELCv2 attribute. To operate the ELCv2 attribute along with ELCv3, explicitly configure the `elc-v2-compatible` statement at the `labeled-unicast entropy-label` hierarchy level.

  [See Understanding Entropy Label for BGP Labeled Unicast LSP.]

- **Support for EIBGP multipath (PTX Series)**—

  Starting in Junos OS Release 23.4R1, we support EIBGP multipath. To enable this feature, use the `allow-external-internal prefix-policy policy-name` statement under global BGP multipath configuration hierarchies.

  [See multipath (Protocols BGP).

- **Support for limiting the number of BGP sessions belonging to a subnet (ACX7100-32C and PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.4R1, we support limiting the number of BGP sessions belonging to a given subnet that is configured using the ‘`allow` statement. You can use this feature to configure wider subnets by limiting the number of BGP sessions over them. You can set this limit using the `peer-limit value` statement at the `[edit protocols bgp group group-name dynamic-neighbor]` hierarchy level.

  [See **peer-limit**.]

- **Deterministic ordering of qualified next-hops (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.4R1, a new CLI configuration option `order` is introduced to ensure that multiple paths in the static routes maintain the order and do not get re-

ordered on RPD restart or RPD switchover. The order of the path is maintained even if one of the links goes down and comes back again.

To configure deterministic ordering of static route next-hops, include the `order` *`number`* configuration statement at the [`edit routing-options static route prefix qualified-next-hop nh-address`] hierarchy level.

[See qualified-next-hop (Static Routes).]

Services Applications

- **STAMP and changes to TWAMP support (ACX7024, ACX7024X, ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we support RFC 8762, *Simple Two-Way Active Measurement Protocol* (STAMP). RFC 8762 standardizes and expands upon the TWAMP Light operational mode, which was defined in Appendix I of RFC 5357, *Two-Way Active Measurement Protocol* (TWAMP). A STAMP-compliant reflector ensures symmetric payload size (in accordance with RFC 6038) and operates in either stateless or stateful mode, depending on whether the sequence number in the reflected payload is copied from the client frame or generated independently. A stateful reflector can detect in which direction drops have occurred. In previous releases, we supported symmetric payloads and stateless reflection. With this release, we support stateful reflection, full compliance with the STAMP standard, and unidirectional drop values for clients. We support unidirectional drop values not only for STAMP clients, but also for TWAMP Managed-mode clients.

  [See Understand Two-Way Active Measurement Protocol and twamp.]

Software Installation and Upgrade

- **Optimize reboot times by disabling default initialization and startup of certain Layer 2 applications (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.4R1, when rebooting the device, the Layer 2 (L2) applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:

  - `[edit interface` *`interface-name`* `unit` *`number`* `family ethernet-switching]`
  - `[edit vlans]`
  - `[edit routing-instance` *`instance-name`* `instance-type virtual-switch]`
  - `[edit routing-instance` *`instance-name`* `instance-type mac-vrf]`
  - `[edit protocols l2-learning]`

  Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

As a result of this change, if your configuration already contains these configuration statements and you then delete all of them, these L2 applications stop running.

[See request node reboot (re0 | re1) (Junos OS Evolved), request system reboot (Junos OS Evolved), and request system software add (Junos OS Evolved).]

- **Enhanced USB, console, AUX access management (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, users can now enable or disable USB, console, and AUX access. Users can mount USB mass storage devices as read-only, allowing only file viewing while blocking write attempts with a **Read-only file system** error. Port status remains persistent. Admins can remotely deactivate USB ports, forcibly un-mounting and disabling devices. Ongoing transactions on mounted devices will be gracefully aborted for system stability.

  [See show chassis usb storage.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for RFC 8814 (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we partially support RFC 8814, *Signaling Maximum SID Depth (MSD) Using the Border Gateway Protocol-Link State*. Currently, we support signaling the MSD using IS-IS for SRv6. For non-segment routing networks, the MSD value reflects the maximum label depth.

  A controller in a segment routing network learns the MSD of the participating router and computes the segment routing path. The controller ensures that the label stack is not greater than what the routers can support.

  [See Link-State Distribution Using BGP.]

- **BGP classful transport support for IPv4 DTM segment routing traffic engineered (SR-TE) tunnels (PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.4R1, we support transport-rib model for V4 DTM SR-TE tunnels by configuring the `use-transport-class` statement at the [`edit dynamic-tunnels` *`tunnel-name`* `spring-te`] hierarchy level.

  If the `use-transport-class` statement is not configured then catch all route and application route is created in the inetcolor.0 table. If the `use-transport-class` statement is configured then catch all route and application route is created in color.inet.3 table. This behavior is irrespective of including the `use-transport-class` statement at the [`edit protocols source-packet-routing`] hierarchy. For dynamic tunnels, SR-TE honors the `use-transport-class` statement under the dynamic-tunnel configuration rather than source-packet-routing configuration.

  The following IPv4 endpoint for DTM SR-TE tunnels with transport-rib model is supported:

  - DCSPF support (using compute-profile)
  - Dynamic segment list support. Configured segment list must not have any IPv6 address and MPLS SID based of IPv6.

- Delegation to PCEP controller
- sBFD support
- SPRING-TE route is added only into color.inet.3 table

For IPv4 endpoint for DTM SR-TE tunnels with inetcolor.0 model, if the `use-transport-class` statement is configured under SR-TE, then dynamically triggered SR-TE tunnel routes is created in both inetcolor.0 table and color.inet.3 table. The `use-transport-class` statement under dynamic-tunnels hierarchy decides if the SR-TE tunnels need to be placed in color.inet.3 table. SPRING-TE route is added only into inetcolor.0 table for DTM SRTE tunnels for IPv4 endpoints and inetcolor.0 model.

Traffic steering based on extended color community is supported. For transport-rib model for DTM SR-TE tunnels (IPv4 destinations only), enable the computation and setup of interdomain segment routing paths using express-segments with SR-Policy underlay.

- **Mitigate traffic congestions using tactical traffic engineered (TTE) tunnels (PTX10008 and PTX10016)**—Starting with Junos OS Evolved Release 23.4R1, you can avoid congestions on oversubscribed links or domains using the dynamic tactical traffic engineered (TTE) tunnel solution. The dynamic TTE tunnel solution allows you to define congestion for a link by configuring high and low bandwidth thresholds. If the traffic load on the link exceeds the high threshold, then load-sharing is increased. If the traffic load falls below the low threshold, then load-sharing is decreased.

  The TTE solution helps you to:

  - Load-balance traffic towards destination prefixes using the congested outgoing interface or through a dynamically installed Tactical TE (TTE) tunnel..
  - Monitor the cumulative load and subsequent deactivation of the TTE tunnel(s) when congestion is no longer detected.

  To enable congestion protection, include the `congestion-protection` statement at the [`edit routing-options`] hierarchy level. Define high and low bandwidth thresholds by including the `high-threshold` and `low-threshold` statements at the [`edit routing-options congestion-protection template` *`template-name`*] hierarchy level. You also need to include the `export isis-export` statement at the [`edit protocols isis`] hierarchy level.

  The TTE tunnel solution supports ISIS and uses TI-LFA backup routes for congestion mitigation.

Subscriber Management and Services

- **IPv6 support for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016**—Starting in Junos OS Evolved Release 23.4R1, we support SSH, Telnet, SCP, FTP, NetConf over SSH and TLS, Telemetry Apps over gRPC.

  [See Broadband Subscriber Access Protocols User Guide.]

Additional Features

We've extended support for the following features to these platforms.

- **Firewall filter support for bitwise logical operations for TCP flag match** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See Firewall Filter Match Conditions Based on Bit-Field Values.]

- **HTTP and TCP probe types for RPM** (ACX7024, ACX7100, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016). You can now configure the `http-get`, `http-metadata-get`, and `tcp-ping` probe types for real-time performance monitoring (RPM) probes. You must configure the `offload-type none` statement to be able to commit the configuration.

  [See probe-server, probe-type, and rpm.]

- **Interconnecting EVPN-VXLAN data centers with EVPN-MPLS in WAN using gateway nodes** (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)

  [See Overview of EVPN-VXLAN Interconnect through EVPN MPLS WAN Using Gateways].

- **Juniper Resiliency Interface** (PTX10003). You can use the Juniper Resiliency Interface (JRI) to detect, correlate, and mitigate exceptions.

  [See Juniper Resiliency Interface.]

- **NETCONF event notifications support** (ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016):

  - NETCONF event notifications
  - Filtering capability
  - Interleave capability

  [See NETCONF Event Notifications.]

- **QDD-400G-ZR-M-HP optics enhancement support** (ACX7100-32C, ACX7100-48L, and PTX10008 with PTX10K-LC1201-36CD line card)

- **Support for EVPN-ETREE** (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)

  [See EVPN-ETREE Overview.]

- **Support for firewall filter-based de-encapsulation** (PTX10003)

  [See decapsulate (Firewall filter)].

- **Support for performance monitoring and TCA** (ACX7100-32C, ACX7100-48L, PTX10008, and PTX10008 with PTX10K-LC1201-36CD line card). We support performance monitoring and

threshold-crossing alert (TCA) information for the QDD-400G-ZR-M-HP optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics by using the `show interfaces transport pm` command and can manage the optical transport link efficiently.

[See show interfaces transport pm.]

- **Support for retrieving NETCONF state information** (PTX10008). NETCONF clients can retrieve NETCONF state information for the following `netconf-state` subtrees:
  - `capabilities`—Supported NETCONF operations
  - `datastores`—Supported configuration datastores
  - `sessions`—Active NETCONF sessions
  - `statistics`—NETCONF server performance data

  [See NETCONF Monitoring.]

- **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 HCT and provide the first supported release information when the optic becomes available.

- **Support for filter-based forwarding** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address.]

- **Support for syslog over TCP and TLS** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016).

What's Changed

IN THIS SECTION

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

EVPN

- **EVPN-VXLAN tracing configuration**— The `set services trace evpn-vxlan` configuration invokes a built-in commit script to generate tracing configurations for troubleshooting EVPN-VXLAN in multiple modules and hierarchies.

  See [trace (EVPN-VXLAN).]

- **New commit check for MAC-VRF routing instances with the encapsulate-inner-vlan statement configured**— We introduced a new commit check that prevents you from configuring an IRB interface and the `encapsulate-inner-vlan` statement together in a MAC-VRF routing instance. Please correct or remove these configurations prior to upgrading to 23.2R2 or newer to avoid a configuration validation failure during the upgrade.

  See [encapsulate-inner-vlan.]

- **Default behavior changes and new options for the easy EVPN LAG configuration (EZ-LAG) feature**— The easy EVPN LAG configuration feature now uses some new default or derived values, as follows:

  - Peer PE device `peer-id` value can only be 1 or 2.

  - You are required to configure the loopback subnet addresses for each peer PE device using the new `loopback peer1-subnet` and `loopback peer2-subnet` options at the **edit services evpn device-attribute** hierarchy level. The commit script uses these values for each peer PE device's loopback subnet instead of deriving those values on each PE device. These replace the `loopback-subnet` option at the **edit services evpn device-attribute** hierarchy level, which has been deprecated.

- If you configure the `no-policy-and-routing-options-config` option, you must configure a policy statement called EXPORT-LO0 that the default underlay configuration requires, or configure the new no-underlay-config option and include your own underlay configuration.
- The commit script generates "notice" messages instead of "error" messages for configuration errors so you can better handle **edit services evpn** configuration issues.
- The commit script includes the element names you configure (such as IRB instance names and server names) in description statements in the generated configuration.
- This feature also now includes a few new options so you have more flexibility to customize the generated configuration:
- `no-underlay-config at` the **edit services evpn** hierarchy level—To provide your own underlay peering configuration.
- `mtu overlay-mtu` and `mtu underlay-mtu` options at the **edit services evpn global-parameters** hierarchy level—To change the default assigned MTU size for underlay or overlay packets.

See [Easy EVPN LAG Configuration.]

.

- **Change in options and generated configuration for the EZ-LAG configuration IRB subnet-address statement**With the EZ-LAG `subnet-address inet` or `subnet-address inet6` options at the **edit services evpn evpn-vxlan irb *irb-instance*** hierarchy, you can now specify multiple IRB subnet addresses in a single statement using the list syntax **addr1 addr2 ...** . Also, in the generated configuration for IRB interfaces, the commit script now includes default `router-advertisement` statements at the **edit protocols** hierarchy level for that IRB interface.

  See [subnet-address (Easy EVPN LAG Configuration).]

Forwarding and Sampling

- For firewall filter action and match on slice ID, family-name can be inet, inet6, mpls, or any.

General Routing

- **The active-user-count is defined as a numeric integer value in ODL request output** —The output for the get-system-uptime-information ODL request contains information for the active-user-count. The active-user-count is now defined as a numeric integer value and avoids an invalid value type error.

[See show system uptime.]

- Two new alarms are added and can be seen with MPC11E when 400G-ZR optics are used. High Power Optics Too Warm: warning of the increase in chassis ambient temperature with no functional action taken on the optics Temperature too high for optics power on: New inserted optics when the chassis ambient temperature is elevated beyond the threshold will not be powered on and would need to be reinserted when the ambient temperature is within the acceptable range

- The packet rate and byte rate fields for LSP sensors on AFT (with the legacy path) have been renamed as jnx-packet-rate and jnx-byte-rate and is in parity with the UKERN behavior. Previously, these rate fields were named as packetRate and byteRate.

- You can specify the minimum and maximum value for the hold-time down and hold-time up interval between 0 through 3600000 milliseconds at the `edit protocols network-isolation group group-name detection` hierarchy level.PR1726039

- **Single source of data for operational state sensor leaves (PTX10008)**— You can use the suppress-interface-leaf CLI statement to suppress telemetry streaming of the following sensors from the packet forwarding engine (PFE): /interfaces/interface/state/high-speed /interfaces/interface/state/oper-status This might be required for collectors that require a single source of data for each sensor.

  [See suppress-interface-leaf.]

- LSPs scale license-Starting in 22.4R3 and 23.4R1 Junos OS Evolved releases, the license counts only the Resource Reservation Protocol-Traffic Engineering (RSVP-TE) and Segment Routing Traffic Engineering (SR-TE) LSPs. Prior to the 22.4R3 and 23.4R1 Junos OS Evolved releases, license counts all Traffic Engineering (TE) such as RSVP-TE and SR-TE and non-TE such as LDP, BGP-LU, L-ISIS, L-OSPF LSPs.

- The backup routing instance then starts index (hrStorageIndex) at 1, incrementing each time after assigning index to an entry in this list. If a mount is unmounted, the other mounts indices will not shift. hrstoragetable indices persist during the lifetime of mib2d. When a new mount is detected, it is assigned the next free index in the system. If there is no free index, it is assigned the last index +1.

- **Enhanced DDoS status operational command (PTX Series)**—We've enhanced the aggregate DDoS status output field to display the aggregate count of all sub packet types.

  Earlier to this release, the aggregate DDoS status output displayed only the packet type level output information.

  [See show ddos-protection protocols.]

- In TVP Platforms, When the FPC is configured to bring offline due to major errors, the alarm string **FPC x Offlined due to Major Errors** will be shown in the `show chassis alarm` and `show system alarm` output instead of **FPC x Major Errors** alarm.

- Before this change most list were ordered by the sequence in which the user configured the list items, for example a series of static routes. After this change the list order is determined by the system with items displayed in numerical sequence rather than by the order in which the items were configured. There is no functional impact to this change.
- While running request system snapshot recovery command on all VMHost based Routing Engines, disable or stop reporting any warning message.

Interfaces and Chassis

- Starting in Junos OS Evolved release 23.2R1-EVO, the output of `show chassis power` command displays the state of the power supply in PTX10003 platform.

  [See **show chassis power**.]

- When all the members of the AE have the same speed (x) and no mixed speed configured. If you change the speed value of any member of the AE to a value other than x, the commit succeeded in earlier releases. From this release, the commit fails. When there are et interfaces with different speeds and you want them to be part of an AE interface. If you change the speed of all the members of the interfaces to be the same speed (x), configure the AE interface, and commit, the commit failed in earlier releases. From this release, such commits succeed.

Junos XML API and Scripting

- **Ability to commit `extension-service file` configuration when application file is unavailable**—When you set the `optional` option at the **edit system extension extension-service application file *file-name*** hierarchy level, the operating system can commit the configuration even if the file is not available at the **/var/db/scripts/jet** file path.

  [See file (JET).]

- **Ability to restart daemonized applications**—Use the request extension-service `restart-daemonize-app` *`application-name`* command to restart a daemonized application running on a Junos device. Restarting the application can assist you with debugging and troubleshooting.

  [See request extension-service restart-daemonize-app.]

- **XML output tags changed for `request-commit-server-pause` and `request-commit-server-start` (ACX Series and PTX Series)**—We've changed the XML output for the `request system commit server pause` command (`request-commit-server-pause` RPC) and the `request system commit server start` command (`request-commit-`

server-start RPC). The root element is <commit-server-operation> instead of <commit-server-information>, and the <output> tag is renamed to <message>.

Licensing

- **Deprecated license revoke information**—Starting in Junos OS Evolved Release 23.4R1, we've deprecated the show system license revoked-info command. You can use the show system license and show system license usage commands to know the license information.

- **LSPs scale license**—Starting in 22.4R3 and 23.4R1 Junos OS Evolved releases, the license counts only the Resource Reservation Protocol-Traffic Engineering (RSVP-TE) and Segment Routing Traffic Engineering (SR-TE) LSPs. Prior to the 22.4R3 and 23.4R1 Junos OS Evolved releases, license counts all Traffic Engineering (TE) such as RSVP-TE and SR-TE and non-TE such as LDP, BGP-LU, L-ISIS, L-OSPF LSPs.

  [See Software Licenses for PTX Series Routers.]

Network Management and Monitoring

- **Changes to the RPC response for <validate> operations in RFC-compliant NETCONF sessions (ACX Series and PTX Series)**—When you configure the rfc-compliant statement at the [edit system services netconf] hierarchy level, the NETCONF server emits only an <ok/> or <rpc-error> element in response to <validate> operations. In earlier releases, the RPC reply also includes the <commit-results> element.

- **NETCONF <copy-config> operations support a file:// URI for copy to file operations (ACX Series and PTX Series)**—The NETCONF <copy-config> operation supports using a file:// URI when <url> is the target and specifies the absolute path of a local file.

  [See <copy-config>.]

- **gNOI OS RPCs use the software version string instead of the package filename (PTX Series)**—The version field in the gnoi.os.OS Activate(), Install(), and Verify() RPCs uses the software version string (as displayed in /system/state/software-version) instead of the package name.

- Simplified gRIBI configuration—Starting in Junos OS Evolved Release 23.4R1, we have simplified the configuration to run gRPC Routing Information Base Interface (gRIBI) service remote procedure calls (RPCs). You no longer need to configure statements at the [edit routing-options resolution] hierarchy level.

  [See gRIBI.]

Routing Protocols

- Prior to this change the output of the `show isis spring flex-algorithm | display xml` command was invalidly formatted when multiple flex algorithm instances were configured. With the change, the XML output is properly structured showing flex algorithm information for each instance. A new XML tag **isis-spring-flex-algorithm** is added to bundle information for each instance.
- After this change ISIS export policies support setting the down bit configuring the **set-down-bit** action in an export policy term.
- Starting in Junos OS Evolved 23.4R1, we have enabled the `process-non-null-as-null-register` configuration statement under `edit protocols pim rp local` by default. For earlier releases, you must configure this statement explicitly.
- In Junos OS Evolved platforms, `show route snooping` and `show route forwarding-table` does not show /56 routes in the VPLS address family table.
- Before this change most list were ordered by the sequence in which the user configured the list items, for example a series of static routes. After this change the list order is determined by the system with items displayed in numerical sequence rather than by the order in which the items were configured. There is no functional impact to this change.
- Starting in Junos OS Evolved 23.4R1, we have enabled the `process-non-null-as-null-register` configuration statement under `edit protocols pim rp local` by default. For earlier releases, you must configure this statement explicitly.

Software Installation and Upgrade

- **`configuration` and `no-configuration` options for the `request system snapshot` command (ACX Series and PTX Series)**—When you omit or include the `configuration` option, the `request system snapshot` command copies the **/config** directory and the configuration stored for each installed software version to the alternate solid-state drive (SSD) as part of the snapshot. You can use the `no-configuration` option to exclude the **/config** directory and the configuration stored for each installed software version from the snapshot.

User Interface and Configuration

- **Information about users editing system configuration**— The `show system configuration database status` command displays information from the Junos OS configuration database that describes the users currently editing the system configuration.

  [See Display Users Currently Editing the Configuration.]

- **Output for request system software status | display xml validate has a tag mismatch error (all platforms)**—The output contains the error message `CRITICAL ERROR: Root tag 'package-status' is either not defined in ODL or does not have 'flag root' set. Please check. Rendering may not work properly`. We have removed the improperly defined `package-status` tag and replaced it with a new child tag `package-status-message`.

- Starting in Junos OS Evolved 23.4R1, you need not set up passwordless communication explicitly by issuing `request security ssh password-less-authentication` command for `transfer-on-commit` and syslog archival operation over scp.

- **Viewing files with the `file compare files` command requires users to have `maintenance` permission**— The `file compare files` command in Junos OS and Junos OS Evolved requires a user to have a login class with `maintenance` permission.

  [SeeLogin Classes Overview.]

- Starting in Junos OS Evolved 23.4R1, you need not set up passwordless communication explicitly by issuing `request security ssh password-less-authentication` command for `transfer-on-commit` and syslog archival operation over scp.

Known Limitations

IN THIS SECTION

Learn about limitations in this release for PTX Series routers.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- Powering up and configuring data path of 400G-ZR and 400G-ZR-M optics is not instant and takes more time comparatively with other optics. So, user has to wait enough time before switching to new speed configuration. If user switches from one speed to other speed without waiting will lead to interface down. PR1635443
- Packet Forward Engine restart trigger followed by graceful switchover is an unsupported feature. The target PFE could get stuck in TRANSITION_OFFLINE or READY state. The user must wait for the target PFE state to be ONLINE before issuing switchover request. PR1740391
- On native CPTX product, we assign the lowest of all family MTUs on Junos OS Evolved interface and assign to the kernel interface as we do not have explicit IFF representation in Linux. This can result in unexpected kernel interface state in some scenarios. For example, If the lowest family MTU is on IPv4 on evo interface and it is < 1280, the kernel interface mtu is no longer IPv6-compatible as IPv6 has a minimum mtu requirement of 1280. So, kernel gets rid of all IPv6 addresses on the interface. Similarly, Linux has minimum MTU requirement of 68 for IPv4. There is no recovery from this situation even if the IPv4 MTU configuration is undone as there are no resulting IPv6 specific DDS events from this undo operation. The only way to recover from this situation is to delete interface and re-configure the interface with the valid MTU. This is a product limitation and hence whenever family MTU configurations are performed via CLI on native CPTX, the user is expected to understand the impact of the configuration on all the families present on the interface and configure MTU that suits all families. PR1747980

Interfaces and Chassis

- On native CPTX product, we assign the lowest of all family MTUs on evo interface and assign to the kernel interface as we do not have explicit IFF representation in Linux. This can resulted unexpected kernel interface state in some scenarios. For example, If the lowest family MTU is on IPv4 on Junos OS Evolved interface and it is < 1280, the kernel interface mtu is no longer IPv6-compatible as IPv6 has a minimum mtu requirement of 1280. So, kernel gets rid of all IPv6 addresses on the interface. There is no recovery from this situation even if the IPv4 MTU configuration is undone as there are no resulting IPv6 specific DDS events from this undo operation. The only way to recover from this situation is to delete interface and re-configure the interface with the valid MTU. This is a product limitation and hence whenever family MTU configurations are performed through CLI on native

CPTX, the user is expected to understand the impact of the configuration on all the families present on the interface and configure MTU that suits all families.PR1746443

- On native CPTX product, we assign the lowest of all family MTUs on evo interface and assign to the kernel interface as we do not have explicit IFF representation in Linux. This can resulted unexpected kernel interface state in some scenarios. For example, If the lowest family MTU is on IPv4 on Junos OS Evolved interface and it is < 1280, the kernel interface mtu is no longer IPv6-compatible as IPv6 has a minimum mtu requirement of 1280. So, kernel gets rid of all IPv6 addresses on the interface. There is no recovery from this situation even if the IPv4 MTU configuration is undone as there are no resulting IPv6 specific DDS events from this undo operation. Similarly Linux has minimum MTU requirement of 68 for IPv4. The only way to recover from this situation is to delete interface and re-configure the interface with the valid MTU. This is a product limitation and hence whenever family MTU configurations are performed through CLI on native CPTX, the user is expected to understand the impact of the configuration on all the families present on the interface and configure MTU that suits all families.PR1747705

Open Issues

IN THIS SECTION

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- G.8273.2 Synce to PTP and SyncE to 1PPS Transient Response test fails. PR1681527
- Class B performance as per G.8273.2 fails for SyncE to PTP and SyncE to 1pps Noise transfer for lower frequencies. PR1681884

- Class B performance as per G.8273.2 shall be supported only when FEC is enabled on both the primary and secondary ports of the T-BC(default option). PR1683579
- With sharding enabled, BGP flags like the following are not displayed on active route in `show route extensive` output: **Accepted Multipath MultipathContrib MultiNexthop** per shard view, using `show route extensive prefix rib-sharding shard-name` will show these flags. PR1693207
- On all Junos OS Evolved platforms, configuring Link Layer Discovery Protocol (LLDP) with **system services netconf notification** enabled will trigger the l2cpd crash. This crash causes the CPU to spike.PR1695057
- In high scaled (beyond 14000 s,g routes - 7000 ipv4 and 7000 ipv6)NGMVPN SPMSI scenarios, core might be seen on PTX10003 platforms due to memory getting exhausted. PR1708454
- On all Junos OS Evolved platforms, when a client sends a neighbor solicitation, the VRRPv3(Virtual Router Redundancy Protocol) node replies with neighbor advertisement and virtual MAC of the wrong interface. There can be some traffic drop because of incorrect virtual MAC entry. PR1708712
- We might encounter jdhcpd core during initialization. The core in rare, and there is no service impact because of this core (as the process recovers immediately). PR1730717
- On all Junos OS Evolved platforms, VMcores are seen when MACsec (Media Access Control Security) key-chains and BGP(Border Gateway Protocol) configurations are applied through Netconf. PR1732611
- Many issues were observed in Junos OS Evolved library regarding the DNS resolution and these are fixed. PR1733616
- The PFE applications ppman-aft-bt, packetio-bt and aft-healthmonitor-bt on PTX10002-36QDD are now renamed to ppman-aft-bx, packetio-bx and aft-healthmonitor-bx. PR1743457
- Ubuntu 20.04 has a bug realted to cgroups and docker that causes communication issues in a VRF. Alternative is to ether use Ubuntu 22.04 as the host to launch ncPTX or add the following cgroup params at /etc/default/grub as below, update-grub and reboot when running ncPTX on Ubuntu 22.04. GRUB_CMDLINE_LINUX_DEFAULT="maybe-ubiquity cgroup_no_v1=net_prio,net_cls". PR1744602
- In PTX Junos OS Evolved devices, power is available from the base power allocated. The power allocation for ABPM is small and the ABPM shares the base power. The power is not allocated explicitly. PR1750747
- When using a Junos OS Evolved device in FIPS mode, key generation using ssh-keygen might fail with the following error: Saving key "/path/to/key/file" failed: error in libcrypto To mitigate this error, use the -o option in ssh-keygen: For example: ssh-keygen -o other options or arguments. PR1756930

- On all Junos OS Evolved PTX platforms, RIB (Routing Information Base) and FIB (Forwarding Information Base) tables are not synchronized properly, causing the P2MP (Point-to-Multipoint) LSP (label-switched-path) traffic outage when executing the CLI command `clear rsvp session`.PR1757635
- The `set chassis ambient-temperature` command is not supported in PTX10004, PTX10008, and PTX10016 platforms. It will be removed in future releases. PR1767840
- EEPROM read failures for PSM are not shown on the PSM LED since the LED of the PSM is internally controlled by the PSM firmware and not by the system software. PR1770991
- On PTX10004, PTX10008, and PTX10016 Junos OS Evolved platforms, after GRES, the backup Routing Engine BITS left over alarms is still in CM alarms even backup RE BITS does not function. The backup RE BITS alarms should be cleared. PR1777209
- On all Junos OS Evolved platforms, committed configuration files are not preserved post software version rollback operation. The actual configuration is not affected. PR1779593
- On PTX10008 Junos OS Evolved, there is no fan name in `jnxFruName` instead it has only fan tray name and FTC. PR1754833
- On PTX10008 Junos OS Evolved, the FTC and SIB LED status on fan tray does not match with SNMP MIB information. SNMP information always is marked as 2 for FTC and SIB value. LED on FTC module is off. PR1765184
- During the programming of the firmware, all Pri/Sec/Led/Comm firmware shows as programming even if any one single firmware gets programmed. However, the firmware upgrade will be done only on the firmware that must be programmed based on the current and available versions. PR1774769
- 1. To avoid a PSU ending up on "Unsupported" state, when a PSU is inserted into a live system, it **must be** pushed completely in (a little bit beyond the point when LEDs light up), and the thump screw must be tightened completely for proper operation. 2. If a PSU gets into the **unsupported** state, it can be slightly pulled out, and re-inserted completely after 30 seconds. This will fix the issue. 3. Alternatively, after all the PSUs are completely inserted into the system with their thumb-screws completely tightened, a chassis-level power cycle will result in detection of all PSUs correctly.PR1784345

Infrastructure

- On all Junos OS Evolved platforms, Transmission Control Protocol (TCP) packets whose size is larger than its interface Maximum Transmission Unit (MTU) will be dropped impacting TCP connections. PR1718999

MPLS

- On all Junos OS Evolved platforms, when MPLS (Multiprotocol Label Switching) statistics is configured without LSP (Label-Switched Path) configuration, the rpd process will crash and impact the routing protocols. This leads to traffic disruption due to the loss of routing information. PR1698889
- On all Junos OS Evolved platforms, a memory leak is observed in the traffic engineering database (TED) when the inet table is not cleaned up after routing instance deactivation. PR1701800

Network Management and Monitoring

- On all Junos Evolved platforms, a cosmetic change is made to an existing syslog message to print more information. The syslog error message that is logged when a user with wrong auth/privacy password sends a SNMP v3 request to router has been changed to add more information. **From: LIBJSNMP_NS_LOG_WARNING: WARNING: Authentication failed for <username> To: LIBJSNMP_NS_LOG_WARNING: WARNING: Authentication failed for <username> , SNMPv3 query from <NMS-ip> to <Router-IP>**PR1734549

  .

Routing Protocols

- When rib (Routing Information Base) contains IPv4 routes with IPv6 next-hops, these routes do not get re-advertised by IPv4 EBGP sessions unless export policy is configured to change it to IPv4 next-hop. PR1712406

Resolved Issues

IN THIS SECTION

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

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

Class of Service (CoS)

- Duplicate code points through code-point-aliases under a classifier results in cosd crash. PR1766873

EVPN

- Traffic drop is observed in EVPN-VXLAN CRB scenario. PR1734091
- EVPN-VXLAN comp nh is not installed in PFE after peer reboot. PR1739686
- ARP/FIB are added even if IRB in EVPN is disabled. PR1743529

General Routing

- Small memory leak is seen on whenever anomalies command `show platform object-info anomalies` is executed all Junos OS Evolved platforms. PR1706565

- The fibd core is observed on Routing Engine switchover. PR1710227
- The dcpfe process will crash due to memory fragmentation. PR1711860
- EVO_REG: VSCAPA: FPC stuck in Onlining state after FPC restart when 100 lsp are configured. PR1712252
- BMP station will not receive the RIBs as expected. PR1715886
- The evo-aftmand-bt process might restart when an application exit. PR1719739
- The **/lib/systemd/system/docker.socket is marked executable** logs flood after system reboot. PR1727524
- Auto-sw-sync does not trigger upgrade or restart of Routing Engine. PR1731877
- In TCP flow, the initial SYN+ACK packet will not be marked with specified CoS related action on Junos OS Evolved platforms. PR1733509
- PTP will get stuck in acquiring state which leads to improper time synchronization after system reboot. PR1734235
- Incorrect queue-counters-trans-bytes-rate ('0')observed while verifying traffic rate for aggregated ethernet as the AE member is not configured with the shaping rate. PR1735087
- JDI_REG:EVO:Brackla:Observing evo-aftmand-zx.re cored at jexpr_fdb_ht_ent_add (hndl=0x7fb46c431740, params=params@entry=0x7fb40cdc4e00, pfe_inst=pfe_inst@entry=16) at ../../../jfdb/common/jexpr_fdb.c:765. PR1735288
- The interface comes up with a mismatch FEC after setting FEC91 on one end and will be stuck in down state after deleting FEC91. PR1736206
- BGP session flaps due to hold time expiration. PR1736428
- Junos OS Evolved is using old MAC address for forwarding leading to traffic drops. PR1736699
- PTX10000 Junos OS Evolved - BITS port LED color of Physical/CLI/MIB do not match. PR1738022
- EVO_REG: PTX10003: JV DB is missing leaf: /interfaces/interface[name=\'ae0\']/state/counters/out-octets , out-pkts, out-unicast-pkts , out-broadcast-pkts , out-multicast-pkts , in-errors , out-errors, in-discards ,out-discards ,in-pause-pkts, out-pause-pkts. PR1738395
- The interface goes down and the error message floods due to the FD leak in the picd process. PR1738854
- FTC X FTC FPGA minimum supported firmware version mismatch alarm raised by OIR FTC. PR1739842
- Error/warning message is missing for unsupported speed in syslog file. PR1740145

- Fans might stop working after removal and insertion of fan tray. PR1742174
- IFD does not not come up whenever optics is removed and inserted on all Junos OS Evolved platforms. PR1742772
- In P2MP-MPLS-LSP set-up traffic drop/traffic blackhaul/label swap/ttl being set to 0 seen due to ARP timeout. PR1743034
- PTP disruption is seen as the back up goes re-acquiring on GM CC change from 7->6. PR1744746
- Error observed when configuring AE interface. PR1745528
- The hwdre application restarted due to memory leak. PR1745749
- The rpd crashes when BGP sharding, multipath and dynamic tunnel are configured. PR1746012
- PTx10001-36mr: no debug logs created post boot of the DUT. PR1746103
- Control board is stuck in present state. PR1747567
- The memory consumption increases due to memory leak. PR1747992
- Child interfaces deleted from aggregate Ethernet interfaces are still shown as part of AE. PR1748236
- The picd crash can be seen on all Junos OS Evolved platforms. PR1748505
- PTX10000 EVO - rpdagent might core after FMBB configuration statement is enabled and deleted then system is rebooted. PR1749431
- The default lo0 firewall filter is not used for traffic sent from a Routing Instance to a host when there is no lo0 configured under the RI. PR1751076
- License is missing on device on performing upgrade. PR1751384
- PTX10000 Junos OS Evolved - CMerror not raised post LAH(link auto-heal) fails during training failure. PR1751581
- TTL value of the explicit null label is ignored on certain PTX platforms. PR1752262
- Traffic null routes due to next-hops are stuck in the pending-delete in evo-aftmand. PR1752267
- Automatic software synchronization mechanism does not function as expected. PR1755616
- FPC unreachable due to running out of Guid space. PR1756452
- License-service crash is seen on Junos OS Evolved platforms. PR1759618
- Traffic loss is observed in a scaled scenario of node-link-protected LSP. PR1759664
- Interface queue statistics are not displayed on `show interfaces queue` CLI command. PR1760134

- Sflow functionality will report packets source subnet mask as 0 for the sampled copy. PR1761350
- CFM DMM with aggregate Ethernet might not work with PTX10001-36mr. PR1763629
- Idmdsensor process might spontaneously crash if the standby Routing Engine has been frequently crashing. PR1764408
- Transit traffic loss during P2MP LSP change. PR1764775
- CCM interval change and rollback resulting in possible traffic loss. PR1766560
- The xinted generates syslog messages **service ssh, accept: Invalid argument (errno = 22)** with high CPU usage. PR1767072
- PCS errors on Ethernet interfaces on certain PTX platforms running Junos OS Evolved 21.2R1-S1, 21.2R2, 21.3R1 or later releases. PR1768453
- JDI-REG : [EVO] : [VPTX10003] : Observing evo-aftmand-zx @ jprds_nh_seq_free,JexprHandleReplicate::unbindPfe,AftWorkRequest::HandleRequest,net::juniper::workqueue::WorkQueuePeerImpl::HandleOneRequest,std::execute_native_thread_routine. PR1768610
- L2TPV3 load-balancing not working properly and create out of order packet flow. PR1769545
- Unexpected evo-aftmand-bt error logs: Jexpr: Invalid HwScId 0, Jexpr: Invalid pfeId 32 while route updated. PR1770432
- The `show interfaces extensive | no-more` command is taking a longer time to display the output. PR1773428
- The orchestratord core dump during JSU. PR1776669

High Availability (HA) and Resiliency

- The traffic drop is observed during the graceful restart on Junos OS Evolved platforms. PR1727957

Infrastructure

- Upon kernel panic on a dual Routing Engine system primary-role relinquishment will take longer than 5 seconds causing FPC restart. PR1759541

Interfaces and Chassis

- Changing speed and adding to AE in the same commit fails. PR1743461

Junos XML API and Scripting

- OpenConfig data obtained with gNMI GetRequest in json format displays module prefix. PR1736286

MPLS

- MPLS LSP stats will not increment post the rpd restart. PR1719162
- LSP with auto bandwidth enabled is not updating its Max AvgBW value, preventing the LSP from being resized. PR1740226

Network Management and Monitoring

- Syslog filter not functioning with generating /etc/syslog.conf+ file after syslog config is deactivated and re-activated. PR1726925
- The snmpd crash is observed after FPC restart. PR1737682
- Custom scripts may fail in Junos OS Evolved single Routing Engine platforms. PR1753283
- The snmpd-subagent cored. PR1760937

Platform and Infrastructure

- [gnoi-pblq] few pblq packets are sent out from port which was under test , after completion of pblq create request process in generator mode. PR1753228

Routing Policy and Firewall Filters

- On all Junos OS Evolved platforms support for a firewall policer rate needs to be increased. PR1743233
- Commit error seen **Fast-lookup-filter <filter name> cannot be configured with next-header match condition**. PR1753514
- Firewall modification will cause the AFTman crash. PR1760210

Routing Protocols

- The mscnoopd process crash will be observed when snooping configuration is removed. PR1696374
- Junos OS Evolved: A crafted BGP UPDATE message allows a remote attacker to de-peer (reset) BGP sessions (CVE-2023-4481). PR1709837
- The PE advertises incorrect next-hop towards CE although BGP export policy configured with next-hop under policy-statement. PR1712527
- The rpd process will crash in a scaled BGP setup with traceoptions configured. PR1732087
- Enabling BGP traceoptions flags will log frequently to the trace file. PR1735189
- The rpd crash files are seen due to a use-after free of objects. PR1737679
- Partial application of BGP import policy with BMP configuration and after back-to-back commits changes BGP import policy. PR1742222
- RPD scheduler slip is observed when the BGP session flaps and subsequent configuration changes for the same peer. PR1742416
- Traffic loss observed in SR-LDP stitch scenario when ECMP is enabled on PTX platforms. PR1746349

User Interface and Configuration

- The **load replace** operation might result in mustd and mgd crash. PR1740289
- hasGlobalIP: Attribute GLOBALIPOWNER does not exist is reported on primary Routing Engine when commit synce to backup Routing Engine. PR1741284

- Commit confirm and commit race condition crashes the firewall functionality. PR1743038

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.4R1-S1 for QFX5130-48C switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New in 23.4R1-S1

IN THIS SECTION

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

Hardware

- **New QFX5130-48C switch (QFX Series)**—Starting in Junos OS Evolved Release 23.4R1-S1, we introduce the Juniper Networks® QFX5130-48C Switch. QFX5130-48C is our first 1-U fixed form factor switch that is completely optimized for 100GbE server connections. The QFX5130-48C

switch offers high-density 100GbE access ports in a SFP-DD form factor optimized for servers, along with high-density 400GbE ports in a QSFP-56 form factor optimized for easy uplinks to data centers. The QFX5130-48C provides a throughput of 8 Tbps by means of:

- Forty-eight high-density 100GbE access ports that support SFP-DD transceivers optimized for servers.
- Eight high-density 400GbE ports that support QSFP56 transceivers optimized for easy uplinks to the spine layer in data centers.

The QFX5130-48C runs Junos OS Evolved. We've designed it to meet the needs of demanding data center environments such as high-performance computing and research networks and cloud and service provider data centers.

Table 4: QFX5130-48C Feature Support

Feature	Description
Class of service	CoS support on EVPN-VXLAN. [See CoS Support on EVPN VXLANs.]
Ethernet switching and bridging	• Support for Q-in-Q tunneling with a service-provider-style configuration. [See Configuring Q-in-Q Tunneling.] • LLDP support. [See Device Discovery Using LLDP.] • Support for MAC move limit with EVPN-VXLAN. [See Understanding MAC Move Limiting.]
Forwarding options	• Support for port mirroring in EVPN-VXLAN environments. [See How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics.]
High availability	• VRRP support on Packet Forwarding Engine. [See VRRP Overview].

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
Interfaces	• Support for BGP flowspec. [See BGP.]
Junos Telemetry Interface (JTI)	• JTI streaming support for hardware Routing Engine-based sensors. Subscribe to the /components/sensor to stream hardware operational stages. Statistics include Routing Engine, power supply units (PSUs), Control Boards, FPCs, and PICs states. [See Junos YANG Data Model Explorer.]
Multicast	• MLD snooping and IRB stitching support . [See Understanding MLD Snooping.] • Support for multicast forwarding. [See Multicast Overview.] • IGMP snooping support. [See PIM Overview.] • IGMP, MLD multicast snooping, and IRB elaboration with MBB. [See IGMP Snooping Overview.] [See Understanding MLD Snooping.]
Network management and monitoring	• Support for sFlow. [See Overview of sFlow Technology.] • Support for analyzers and port mirroring. [See Understanding Port Mirroring and Analyzers.]

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Inband Flow Analyzer (IFA) 2.0 transit node support. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.] • IPsec support for OSPFv2 and OSPFv3. [See Overview of IPsec.] [See Configuring OSPF Authentication.] [See Configuring IPsec Security Associations.] • DHCP stateless relay MIB support. [See Enterprise-Specific MIBs for Junos OS Evolved.]
Protection against DDoS attacks	• Supports DDoS protection, which is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.] and protocols (DDoS) (ACX Series, PTX Series, and QFX Series).]
Platform and infrastructure	• Platform resiliency support for hardware components of each FRU. If a failure is detected on a hardware component, Junos OS Evolved: • Logs the message to give clear indication of failure details, including time stamp, module name, component name & failure details. • Raises or clears alarms if applicable. • Performs local action, such as self-healing and taking the component out of service.

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description

Precision Time Protocol	• Transparent Clock support
	  • With or without VLAN encapsulation
	  • With PTP over IPv4
	  • With PTP unicast or multicast
	  • On LAG and MC-LAG
	  • On all physical, IRB, and AE interfaces

Routing options	• Support for Unified Forwarding Table (UFT). [See Understanding the Unified Forwarding Table.]

Routing protocols	• Support for redistribution of IPv4 routes with IPv6 next hop into BGP.
	  [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]
	• Support for collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding.
	  [See Telemetry Sensor Explorer.]
	• Support for maximum reference bandwidth increased to 4 TB for IGP protocols.
	  [See reference-bandwidth (Protocols IS-IS).]
	  [See reference-bandwidth (Protocols OSPF).]
	• Support for check for AS matches in BGP policy AS paths without regular expressions.
	  [See Improve the Performance of AS Path Lookup in BGP Policy.]

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for stripping or replacing BGP private AS. [See Autonomous Systems for BGP Sessions.] • BMP local RIB monitoring support for all RIBs with sharding. [See BGP Monitoring Protocol]. [See loc-rib.] [See rib-list.] • Support for bootstrapping route-validation database from a local file. [See validation (Origin Validation for BGP).]
Routing policy and firewall filters	• Sharding support for conditional route manager. [See Routing Policy Match Conditions]. [See rib-sharding.] [See show policy conditions.] • Support for fast lookup of origin and neighbor autonomous systems (ASs). [See policy-options.] [See policy-statement.]

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Firewall filter support on Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions.] • Support for profiles to improve the firewall filter scale. [See Planning the Number of Firewall Filters to Create.] • EVPN-VXLAN firewall filtering and policing. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]
System management	• Secure boot and secure BIOS support. [See Secure Boot.] • CLI-based hash and ECMP resilient hashing support. [See enhanced-hash-key.] [See ecmp-resilient-hash].
	• Support for dynamic load balancing (DLB). [See enhanced-hash-key.] • Support for configuring firewall filters and interfaces programmatically using JET APIs. [See Overview of JET APIs.]
Software installation and upgrade	• ZTP support. [See Zero Touch Provisioning.]

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
Services applications	• Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agent.]
Support for optics supported on QSFP-DD ports, SFP+ ports, SFP56-DD optics, and DAC cables	• To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported , see the Hardware Compatibility Tool.

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
VPNs	• Support for EVPN Type 5 routes. [See Understanding EVPN Pure Type-5 Routes.] • Assisted replication (AR) integrated with optimized intersubnet multicast (OISM) in an EVPN-VXLAN edge-routed bridging (ERB) fabric support. [See Assisted Replication Multicast Optimization in EVPN Networks.] [See Optimized Inter-Subnet Multicast in EVPN Networks.] • EVPN-VXLAN support with MAC-VRF routing instances. [See EVPN User Guide.] • Support for EVPN-VXLAN fabric with an IPv6 underlay. [See EVPN-VXLAN with an IPv6 Underlay.] [See Example: Configure an IPv6 Underlay for Layer 2 VXLAN Gateway Leaf Devices.] • Support for symmetric IRB with EVPN Type 2 routes. [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics.] [See irb-symmetric-routing.] • Support for MLDv1, MLDv2, and MLD snooping with OISM and AR in EVPN-VXLAN fabrics. [See Optimized Intersubnet Multicast in EVPN Networks.]

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for determining IRB interface state changes based on local and remote connectivity states in EVPN fabrics. [See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics.] [See interface-state.] [See network-isolation.] • Overlay and CE-IP ping and traceroute support for EVPN-VXLAN. [See Understanding Overlay ping and traceroute Packet Support.]

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for blocking asymmetric EVPN Type 5 routes. [See EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN.] [See ip-prefix-routes.] • Support for DHCP relay in EVPN-VXLAN. [See DHCP Relay Agent over EVPN-VXLAN.] • Support for coexistence of EVPN Type 2 and Type 5 routes . [See EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN.] • Support for Interconnecting EVPN-VXLAN in a data center to an EVPN-VXLAN control plane in a WAN using a gateway model. [See Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network.] • Support for OISM in an EVPN-VXLAN fabric. [See Optimized Inter-Subnet Multicast in EVPN Networks.] • Support for service-provider-style interface configuration on EVPN-VXLAN Layer 3 gateways. [See Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network.] • Overlapping VLAN support in EVPN-VXLAN fabrics on edge-routed bridging (ERB) overlay leaf devices.

Table 4: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	[See Overlapping VLAN Support Using VLAN Translation in EVPN-VXLAN Networks.] [See vlan-rewrite.]

Software Installation and Upgrade

- **Firmware upgrade support (QFX5130-48C)**—Starting in Junos OS Evolved Release 23.4R1-S1, QFX5130-48C devices support the following commands:
  - `request system firmware upgrade re bios`
  - `request system firmware upgrade re fpga`
  - `request system firmware upgrade re ssd`
  - `request system firmware upgrade re xmcfpga`
  - `request system firmware upgrade re fancpld`
  - `request system firmware upgrade fpc opticscpld<0/1/2/3>`
  - `request system firmware upgrade psm`

  [See request system firmware upgrade.]

What's Changed

There are no changes in behavior and syntax in this release for QFX Series switches.

Known Limitations

IN THIS SECTION

Learn about limitations in this release for QFX Series switches.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

System Management

- On QFX5130-48C, two alarms are raised when power-supply feed is disconnected.

```
# run show system alarms
..
2 alarms currently active
Alarm time               Class  Description
2024-01-01 11:17:41 PDT  Major  PSM 0 Unit Offline
2024-01-01 11:17:41 PDT  Major  PSM 0 Input Under Voltage Failure
```

PR1725674

Open Issues

IN THIS SECTION

Learn about open issues in this release for QFX5130-48C switches.

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- Whenever user issues `request system power-cycle` CLI command to perform remote power cycle of the device, an incomplete message string might be displayed on the console. PR1775828
- In an EVPN-VxLAN multicast scenario with BDs-not-being (BD - Bridge Domain) everywhere, PIM (Protocol Independent Multicast) Register Messages are not sent from the Border Leaf to the PIM-RP (Protocol Independent Multicast Rendezvous point), when the source of the multicast traffic is behind a Server Leaf and the Source-VLAN of the multicast traffic is not hosted by the Border Leaf device. PR1777493
- Whenever a QFX5130-48C device is reset using rear panel reset button by pressing the button for short/long duration, the port LEDs on ports 56/57 might glow in amber for a few seconds during boot up. The port LEDs amber status turns off and reflects correct port state once the device is up.PR1792619
- On QFX5130-48C, when a USB is plugged into the chassis, the USB disk does not show up in `show chassis hardware detail`CLI command output until a subsequent reboot. Similarly when the USB stick is plugged out of the chassis, the `show chassis hardware detail` CLI command continues to show up the USB stick as plugged into the chassis until subsequent reboot of the device. PR1793934
- On QFX5130-48C, when image upgrade is done using CLI, the log messages have kernel trace reported. There is no functionality impact due to the messages. PR1798213
- QFX5130-48C: TX laser disable alarm for 400G DR4 optics is seen when multiple optics are jacked out and jacked in in quick succession. PR1798463

Infrastructure

- QFX5130-48C running Junos OS Evolved 23.4R1-S1 image might report a vmcore during bootup. The device reboots itself to recover from the issue state. PR1768516

Resolved Issues

IN THIS SECTION

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

For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- PFE crash seen on Junos OS Evolved platforms due to upstream interface change. PR1771209
- The system storage disk layout size might be different across QFX5130-48C devices. PR1781237

Upgrade Your Junos OS Evolved Software

For products impacted, see Feature Explorer.

Follow these steps to upgrade your Junos OS Evolved software:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: https://www.juniper.net/support/downloads/
2. In the Find a Product box, enter the Junos OS platform for the software that you want to download.
3. Select Junos OS Evolved from the OS drop-down list.
4. Select the relevant release number from the Version drop-down list.
5. In the **Install Package** section, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.

9. Copy the software to the device or to your internal software distribution site.
10. Install the new 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.

For more information about software installation and upgrade, see Software Installation and Upgrade Overview (Junos OS Evolved). For more information about EOL releases and to review a list of EOL releases, see https://support.juniper.net/support/eol/software/junosevo/.

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

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/

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 Juniper Support Portal: Case Management, Product Support & More.
- 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

08 January 2026—Revision 16, Junos OS Evolved Release 23.4R1

13 August 2025—Revision 15, Junos OS Evolved Release 23.4R1

11 August 2025—Revision 14, Junos OS Evolved Release 23.4R1

12 March 2025—Revision 13, Junos OS Evolved Release 23.4R1

28 February 2025—Revision 12, Junos OS Evolved Release 23.4R1

19 September 2024—Revision 11, Junos OS Evolved Release 23.4R1

23 May 2024—Revision 10, Junos OS Evolved Release 23.4R1

16 May 2024—Revision 9, Junos OS Evolved Release 23.4R1

07 May 2024—Revision 8, Junos OS Evolved Release 23.4R1-S1

06 May 2024—Revision 7, Junos OS Evolved Release 23.4R1-S1

27 March 2024—Revision 6, Junos OS Evolved Release 23.4R1

29 February 2024—Revision 5, Junos OS Evolved Release 23.4R1

8 February 2024—Revision 4, Junos OS Evolved Release 23.4R1

9 January 2024—Revision 3, Junos OS Evolved Release 23.4R1

4 January 2024—Revision 2, Junos OS Evolved Release 23.4R1

1 January 2024—Revision 1, Junos OS Evolved Release 23.4R1