

Release Notes

Published
2025-12-08

Junos OS Evolved Release 25.2R1

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 25.2R1.

For more information on this release of Junos OS Evolved, see [Introducing Junos OS Evolved](#).

HARDWARE HIGHLIGHTS

- ACX7020-AC
- ACX7020-DC

Table of Contents

Junos OS Evolved Release Notes for ACX Series

What's New | 1

Hardware | 2

Authentication and Access Control | 32

Chassis | 32

Class of Service | 33

EVPN | 33

High Availability | 34

Interfaces | 34

Junos Telemetry Interface | 35

Multicast | 40

Multichassis Link Aggregation (MC-LAG) | 41

Network Management and Monitoring | 41

Precision Time Protocol (PTP) | 41

Proxy ARP | 42

Routing Policy and Firewall Filters | 42

Source Packet Routing in Networking (SPRING) or Segment Routing | 42

VLANs | 44

Additional Features | 44

What's Changed | 47

Known Limitations | 50

Open Issues | 51

Resolved Issues | 52

Junos OS Evolved Release Notes for PTX Series

What's New | 55

Hardware	56
Authentication and Access Control	57
Class of Service	57
EVPN	57
Forwarding Options	58
High Availability	58
Interfaces	58
IPv6	59
Junos Telemetry Interface	60
Layer 2 VPN	62
Network Management and Monitoring	62
OpenConfig	64
Routing Options	64
Routing Policy and Firewall Filters	65
Routing Protocols	66
Serviceability	67
Services Applications	68
Software Installation and Upgrade	68
Source Packet Routing in Networking (SPRING) or Segment Routing	69
Additional Features	69

What's Changed | 72

Known Limitations | 76

Open Issues | 76

Resolved Issues | 78

Junos OS Evolved Release Notes for QFX Series

What's New | 83

- EVPN | 84
- Junos Telemetry Interface | 85
- Network Management and Monitoring | 86
- Multicast | 87
- OpenConfig | 87
- Precision Time Protocol (PTP) | 88
- Routing Policy and Firewall Filters | 88
- Routing Protocols | 89
- Serviceability | 90
- Additional Features | 90

What's Changed | 91

Known Limitations | 95

Open Issues | 96

Resolved Issues | 98

Upgrade Your Junos OS Evolved Software | 99

Licensing | 100

Finding More Information | 101

Requesting Technical Support | 101

Revision History | 103

Junos OS Evolved Release Notes for ACX Series

IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 47](#)
- [Known Limitations | 50](#)
- [Open Issues | 51](#)
- [Resolved Issues | 52](#)

These release notes accompany Junos OS Evolved Release 25.2R1 for ACX7020-AC, ACX7020-DC, 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 THIS SECTION

- [Hardware | 2](#)
- [Authentication and Access Control | 32](#)
- [Chassis | 32](#)
- [Class of Service | 33](#)
- [EVPN | 33](#)
- [High Availability | 34](#)
- [Interfaces | 34](#)
- [Junos Telemetry Interface | 35](#)
- [Multicast | 40](#)
- [Multichassis Link Aggregation \(MC-LAG\) | 41](#)
- [Network Management and Monitoring | 41](#)

- [Precision Time Protocol \(PTP\) | 41](#)
- [Proxy ARP | 42](#)
- [Routing Policy and Firewall Filters | 42](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 42](#)
- [VLANs | 44](#)
- [Additional Features | 44](#)

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 25.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [ACX7020-AC](#)
- [ACX7020-DC](#)
- [ACX7024](#)
- [ACX7024X](#)
- [ACX7100-32C](#)
- [ACX7100-48L](#)
- [ACX7332](#)
- [ACX7348](#)
- [ACX7509](#)

The following sections highlight the key features in this release.

Hardware

- **ACX7020 Cloud Metro Router (ACX Series)**—The Juniper Networks ACX7020 Cloud Metro Router is a high-density and temperature-hardened router with a compact 1-U form factor. With 100 Gbps of throughput and advanced timing features, the ACX7020 provides several capabilities that enable you to support mobile backhaul, residential, and metro Ethernet business services solutions.

The ACX7020 features four small form-factor pluggable (SFP28) ports (each configurable as 1GbE, 10GbE, or 25GbE) and sixteen SFP+ ports (each configurable as 1GbE or 10GbE). The router's

efficient thermal design allows the use of high-power SFP28 and SFP+ transceivers on all supporting ports.

We ship the ACX7020 routers with integrated fans for side-to-side airflow and built-in AC or DC power supply modules (PSMs).

To install the ACX7020 router and perform initial configuration, routine maintenance, and troubleshooting, see the [ACX7020 Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform.

Table 1: ACX7020 Feature Support

Feature	Description
Class of service	<ul style="list-style-type: none"> Support for classification and rewrite rules of all types (Inet-precedence classifier/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.] <p>NOTE: On IRB interfaces, only Layer 3 classification and rewrite are supported.</p> <ul style="list-style-type: none"> Support for port shaping and scheduling with eight Virtual Output Queues (VOQs) 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 1: ACX7020 Feature Support (Continued)

Feature	Description
Dynamic Host Configuration Protocol	<ul style="list-style-type: none"> • 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). ERB is an edge model where DHCP clients are connected and relayed in network leaf devices. The spine provider edge (PE) routers do not perform DHCP relay functions. The routers support transit spine functionality, running protocols such as BGP for integrated routing and bridging (IRB). • Support for the following functionalities: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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.

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
	<p>[See DHCP Relay Agent and Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent.]</p> <p>[See DHCP Relay Agent in EVPN-MPLS Network.]</p>
EVPN	<ul style="list-style-type: none"> • Support for the following EVPN-MPLS features on MAC-VRF instances: <ul style="list-style-type: none"> • 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

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
EVPN	<ul style="list-style-type: none"> • 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.] • NOTE: Data Center feature VXLAN and EVPN-VXLAN is not supported. • 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 VPN–virtual 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 for segment routing over MPLS [See Overview of VPWS with EVPN Signaling Mechanisms.] • Entropy and flow label for EVPN-ELAN

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
	<p data-bbox="898 359 1243 386">[See Configuring Entropy Labels.]</p> <ul data-bbox="862 422 1419 1146" style="list-style-type: none"> <li data-bbox="862 422 1419 449">• Support for the following EVPN-MPLS features: <ul data-bbox="898 485 1419 1146" style="list-style-type: none"> <li data-bbox="898 485 1419 512">• IRB with IPv4 and IPv6 addresses <li data-bbox="898 548 1419 575">• IRB virtual gateway <li data-bbox="898 611 1419 638">• IRB anycast gateway <li data-bbox="898 674 1419 701">• IRB with static MAC <li data-bbox="898 737 1419 806">• EVPN asymmetric Type 2 and symmetric Type 5 routes <li data-bbox="898 842 1419 869">• EVPN E-LAN over BGP-LU <li data-bbox="898 905 1419 974">• EVPN proxy ARP and ARP suppression, and NDP and NDP suppression <li data-bbox="898 1010 1419 1037">• EVPN routing policies <li data-bbox="898 1073 1419 1146">• Ingress virtual machine traffic optimization (VMTO) <p data-bbox="898 1182 1419 1461">[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.]</p> <ul data-bbox="862 1497 1419 1724" style="list-style-type: none"> <li data-bbox="862 1497 1419 1524">• Support for the following EVPN-VPWS features: <ul data-bbox="898 1560 1419 1724" style="list-style-type: none"> <li data-bbox="898 1560 1419 1587">• EVPN-VPWS FXC VLAN unaware service <li data-bbox="898 1623 1419 1650">• EVPN-VPWS FXC VLAN aware service <li data-bbox="898 1686 1419 1724">• EVPN-VPWS over segment routing

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • Single homing and all-active multihoming support • Flow-aware transport (FAT) pseudowire labels • Entropy labels <p>[See Overview of VPWS with EVPN Signaling Mechanisms.]</p>
Firewall filters	<ul style="list-style-type: none"> • Support for firewall filters and policers. You can configure firewall filters with packet match conditions for the any, ethernet switching, inet, inet6, ccc, and mpls families. In addition to packet match conditions, the ACX7020 supports the count, discard, log, syslog, and policer actions. <p>[See Firewall Filter Match Conditions and Actions in ACX Series Routers (Junos OS Evolved).]</p> <p>NOTE: On IRB interfaces, only Layer 3 classification and rewrite are supported.</p> <ul style="list-style-type: none"> • Filter-based forwarding (FBF). <p>[See Filter-Based Forwarding Overview.]</p> <ul style="list-style-type: none"> • Firewall filter protocols: MPLS, CCC, virtual private LAN service (VPLS), and Any <p>NOTE: On IRB interfaces, only Layer 3 classification and rewrite are supported.</p> <p>[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.]</p>

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
High availability	<ul style="list-style-type: none"> • 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.] • 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.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • Logical subinterface and Packet Forwarding Engine drop, pipe, and line-card counter sensor support for JTI [See Junos YANG Data Model Explorer.]

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Bridge domain without a <code>vlan-id number</code> statement • Bridge domain with the <code>vlan-id</code> value set to <code>None</code> • Bridge domain with a single VLAN ID • MAC learning, aging, and limiting • Single-learning domain per bridge domain • Ethernet service types: <ul style="list-style-type: none"> • 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

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • • Link aggregation group (LAG) support with the following hashing algorithms: <ul style="list-style-type: none"> • 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 • Encapsulation types: <ul style="list-style-type: none"> • 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.]

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • 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.] • Transparent forwarding of L2CP BPDUs. 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 configure and enable L2PT. • VLAN sensor support [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: <ul style="list-style-type: none"> • 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

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
	<ul style="list-style-type: none"> • Layer 3 multicast with and without IGMP or MLD snooping <p>[See Understanding Multichassis Link Aggregation Groups.]</p>
Layer 2 VPN	<ul style="list-style-type: none"> • 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 <i>routing-instance-name</i>] 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 <i>routing-instance-name</i> protocol vpls] hierarchy level. <p>[See Introduction to Configuring VPLS.]</p> <ul style="list-style-type: none"> • 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) <p>[See control-word , Configuring Entropy Labels, and FAT Flow Labels Overview.]</p> <ul style="list-style-type: none"> • 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. <p>[See Redundant Pseudowires for Layer 2 Circuits and VPLS.]</p>

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Layer 2 VPN	<ul style="list-style-type: none"> • IRB support for VPLS. [See Configuring VPLS and Integrated Routing and Bridging.] • Layer 2 VPN and L2 circuit support: <ul style="list-style-type: none"> • 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: ACX7020 Feature Support *(Continued)*

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • Support for the following Layer 3 features: <ul style="list-style-type: none"> • IP forwarding and exception packet handling • IEEE 802.1Q (VLAN trunk) on IRB interfaces • Address Resolution Protocol (ARP), neighbor discovery, unicast reverse-path forwarding (unicast RPF), 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 <p>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.]</p>

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
Layer 3 VPN	<p>Support for the following Layer 3 VPN features:</p> <p>NOTE: VT interface-based Layer 3 VPN is not supported. Layer 3 VPN ping is supported only with the vrf-table-label configuration.</p> <ul style="list-style-type: none"> • IP-VPN services: <ul style="list-style-type: none"> • Instance-type virtual routing and forwarding (VRF) and virtual router • All control plane configuration options • Per-prefix and per-table label signaling • Layer 3 VPN support with ECMP • BGP policies 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 <p>NOTE: Table next hop is not supported.</p> • Inter-autonomous system (inter-AS) options A, B, and C <p>NOTE: You can deploy inter-AS option B can in a hierarchical network design within a single IGP AS.</p>

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • 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 <p>[See Layer 3 VPNs User Guide for Routing Devices.]</p>
MPLS	<ul style="list-style-type: none"> • Support for the following MPLS features: <ul style="list-style-type: none"> • IP/MPLS infrastructure feature set for the L3VPN service • Basic BGP control plane features such as LDP-DOD, CSPF, and single-area CSPF • MPLS label stack • MPLS protections: <ul style="list-style-type: none"> • 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 including refresh reduction • Label Distribution Protocol (LDP) LSP over IPv4

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
MPLS	<ul style="list-style-type: none"> • RSVP 1:1 • RSVP-Traffic Engineering (RSVP-TE) • LDP over RSVP • Inter-autonomous system LSP intra-area LSP [See MPLS Applications User Guide.] • Support for MPLS LSP statistics and RSVP-TE auto-bandwidth features. Support includes: <ul style="list-style-type: none"> • MPLS LSP statistics for the following LSP types: <ul style="list-style-type: none"> • LDP-signaled LSPs • RSVP-signaled LSPs • Static LSPs • Bypass LSPs • Container LSPs • RSVP-TE auto-bandwidth <p>[See LSP Overview, LDP Overview, RSVP Overview, and Configuring Optimized Auto-bandwidth Adjustments for MPLS LSPs.]</p>

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Multicast	<ul style="list-style-type: none"> • Support for multicast snooping in a VPLS for the following protocols: <ul style="list-style-type: none"> • 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 <p>[See Multicast Snooping for VPLS.]</p> • 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: <ul style="list-style-type: none"> • 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 <p>[See IGMP Snooping Overview and Understanding MLD Snooping.]</p> • Support for IPv4 multicast for Layer 3. You can configure IGMP snooping with IGMPv2 and IGMPv3, which includes support for the following: <ul style="list-style-type: none"> • Anycast RP

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • IGMP filter • Protocol Independent Multicast source-specific multicast (PIM SSM) • IGMP querier • PIM sparse mode (PIM SM) <p>[See IGMP Snooping Overview.]</p> <ul style="list-style-type: none"> • Support for BGP MVPN. BGP over MPLS MVPN (also known as "next-generation" MVPN) runs 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. <p>[See Multiprotocol BGP MVPNs Overview.]</p> <ul style="list-style-type: none"> • Multicast with IGMP or MLD snooping within VLANs for EVPN-MPLS <p>[See Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment.]</p>

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> • Support for port mirroring with analyzers and remote port mirroring to IP address (GRE encapsulation) [See Port Mirroring and Analyzers.] • Support for SNMP [See SNMP Architecture and SNMP MIBs Overview.] • Support for Dying Gasp through SNMP in case of power failure. [See Dying Gasp Functionality and SNMP Traps and Informs.]
Operations, Administration, and Maintenance	<ul style="list-style-type: none"> • 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): <ul style="list-style-type: none"> • 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 802.3ah and 802.1ag for OAM CFM down and up maintenance association end points (MEPs) over virtual private LAN service (VPLS)

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
Operations, Administration, and Maintenance	<ul style="list-style-type: none"> Support for IEEE 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	<ul style="list-style-type: none"> Support for control plane distributed denial-of-service (DDoS) protection [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]
Routing protocols	<ul style="list-style-type: none"> 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): <ul style="list-style-type: none"> Support for loose and strict mode Support for IPv4 and IPv6 [See Understanding Unicast RPF (Routers).]

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • Support for BGP flow specification (BGP flowspec). • The following match conditions are not supported: <ul style="list-style-type: none"> • Fragment for IPv6 • Packet length • Port • Source and destination prefix with offset • The following actions are not supported: <ul style="list-style-type: none"> • Community • Next-term • Routing instance • Sample • Traffic marking <p>[See Understanding BGP Flow Routes for Traffic Filtering.]</p> • Support for configuring interface groups in BGP flowspec filters • BGP Prefix-Independent Convergence (PIC) edge support for inet and MPLS VPNs. The following features are not supported: <ul style="list-style-type: none"> • Session-based repair <p>[See Understanding BGP Flow Routes for Traffic Filtering and Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic.]</p>

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • BGP PIC over LDP over RSVP tunnel • BGP PIC over SR-MPLS • BGP PIC with RSVP • BGP-LU with PIC • BGP PIC edge protection for Layer 2 services • Protection with multilink failure <p>[See Configuring BGP PIC Edge for MPLS Layer 3 VPNs and Use Case for BGP PIC for Inet.]</p> <ul style="list-style-type: none"> • Support for entropy label for LDP, RSVP, L3VPN, and BGP-LU <p>[See Entropy label support for BGP Labeled Unicast (BGP-LU) and Configuring Entropy Labels.]</p> <ul style="list-style-type: none"> • Support for BGP transport address family or BGP Classful Transport (BGP-CT) includes: <ul style="list-style-type: none"> • 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 engineering (SR-TE) color-only support • IPv6 and BGP service routes with a color-only mapping community

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
	<ul style="list-style-type: none"> Enhanced transport-class configuration to provide precise resolution <p>[See use-transport-class.]</p>
Services Applications	<ul style="list-style-type: none"> RFC 2544-based benchmarking tests. Support for Layer 2 reflection (bridge, Layer 2 circuit, 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. You perform RFC 2544 tests 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. <p>[See RFC 2544-Based Benchmarking Tests for ACX Routers Overview.]</p> <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Services Applications	<ul style="list-style-type: none"> • Timestamps set by software (the Routing Engine or the Packet Forwarding Engine) or hardware • Error reporting through system log messages only • Unauthenticated mode only • Support for sFlow monitoring (ingress) <p>[See sFlow Monitoring Technology and Understanding How to Use sFlow Technology for Network Monitoring.]</p> <ul style="list-style-type: none"> • 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 nondefault VRF instances, for both the IPFIX and version 9 export formats. You can configure up to four IPv4 collectors for inline active flow monitoring. <p>See Understand Inline Active Flow Monitoring.</p>
Source Packet Routing in Networking (SPRING) or segment routing	<ul style="list-style-type: none"> • Support for the following segment routing features: <ul style="list-style-type: none"> • Segment routing global block (SRGB) for OSPF, IS-IS, and fast reroute • Metro Ethernet services over segment routing infrastructure

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Source Packet Routing in Networking (SPRING) or segment routing	<ul style="list-style-type: none"> • 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 • 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 • 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. <p>[See How to Enable SRv6 Network Programming in IS-IS Networks.]</p>

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Source Packet Routing in Networking (SPRING) or segment routing	<ul style="list-style-type: none"> 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 Routing How 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 SRv6 flexible algorithms in traffic engineering database (TED) and BGP Link State (BGP-LS) [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.]

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> • 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.] • Support for Secure zero-touch provisioning (SZTP) [See Secure Zero Touch Provisioning.]
Timing and synchronization	<ul style="list-style-type: none"> • Support for Synchronous Ethernet over a Link Aggregation Group (LAG). You can define a LAG while configuring the Synchronous Ethernet clock source on these routers. [See Synchronous Ethernet and PTP Overview.] • Support for frequency synchronization using Synchronous Ethernet. [See Synchronous Ethernet Overview.]

Table 1: ACX7020 Feature Support *(Continued)*

Feature	Description
Timing and synchronization	<ul style="list-style-type: none"> • Support for SyncE MIB and PTP MIB. ACX7020 routers support: <ul style="list-style-type: none"> • Defect and event management capabilities for timing features. • SNMP get, get-next, and walk management capabilities for timing features. These capabilities are enabled through the Synchronous Ethernet MIB (SyncE MIB) and Precision Time Protocol MIB (PTP MIB) timing objects. <p>[See Timing Defects and Event Management on Routing Platforms and SNMP MIB for Timing on Routing Platforms.]</p> • Support for G.8275.1 telecom profile, Precision Time Protocol (PTP) over Ethernet encapsulation, and hybrid mode <p>[See Precision Time Protocol Overview and Understanding Hybrid Mode.]</p> • PTP G.8275.1 support over Link Aggregation Group (LAG) <p>[See G.8275.1 Telecom Profile.]</p> • Support for Precision Time Protocol (PTP) G.8275.1 enh and G.8275.2 enh profiles with PTP over IPv4 and IPv6 unicast. The G.8275.1 enh profile does not support unicast negotiation. <p>[See PTP profiles.]</p>

Table 1: ACX7020 Feature Support (Continued)

Feature	Description
Timing and synchronization	<ul style="list-style-type: none"> • Support for PTP G.8275.2 enhanced profile features compliant with the International Telecommunication Union Telecommunication Standardization (ITU-T) G.8273.4 standards. The following key features are supported: <ul style="list-style-type: none"> • Ordinary clocks and boundary clocks • The alternate best timeTransmitter clock algorithm • Full domain and packet-rate ranges • Primary and secondary asymmetry values • Manual mode (no unicast negotiation) and mixed mode • Primary (active) and secondary link configuration. <p>[See PTP Profiles.]</p> • Support for frequency and phase offset relaxation. <p>[See ptp.]</p> • Support for Precision Time Protocol (PTP) Transparent Clock- Regular Transparent Clock (without SyncE) and Syntonised Transparent Clock (with SyncE). <p>[See PTP Transparent Clocks and show ptp global-information.]</p>

- **Supported transceivers, optical interfaces, and DAC cables (ACX7020)**—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 the HCT and provide the first supported release information when the optic becomes available.

Authentication and Access Control

- **SSH enhancements for algorithm configuration (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)**—We've made the following updates to SSH algorithms:

- The CLI command `set system services ssh ca-signature-algorithms` should be used to configure the signature algorithms that are allowed for certificate authorities (CAs) to use when signing certificates.
- Under the `system services ssh hostkey-algorithm-list` hierarchy level, new options are introduced:
 - `set system service ssh hostkey-algorithm-list rsa-sha2-256`
 - `set system service ssh hostkey-algorithm-list rsa-sha2-512`

These options enable RSA hostkey signatures using the SHA-256 hash algorithm and SHA-512 hash algorithm.

- RSA signatures using the SHA-1 hash algorithm have been disabled by default. Consequently, the CLI command `set system services ssh hostkey-algorithm-list rsa` has been deprecated.
- SSH connections that require a subsystem (for example, `netconf`) need to explicitly use the `-s` option.

[See [hostkey-algorithm-list](#).]

Chassis

- **Optics EM policy support (ACX7509 and ACX7024)**—The Environment Monitoring (EM) policy now includes optics temperature sensors for ACX7509 and ACX7024 routers. The policy ensures efficient thermal management of high-power optical modules on your routers. Key functionalities include temperature monitoring integration and automatic shutdown procedures. You can use CLI commands to configure and manage EM policy on your router.

[See [Optics EM Policy](#).]

- **Chassis management support (ACX7020)**—The ACX7020 router supports these chassis management features:
 - Power Supply Module (PSM) redundancy on ACX7020 with AC power supply
 - Power feed redundancy on ACX7020 with DC power supply
 - Optics Environmental Monitoring (EM) policy for automatic fan speed management
 - Voltage and current monitoring for PSM management

In extremely cold temperatures—for example, at –40 °F (–40 °C)—the router turns off fans and disables the EM policy. The router restarts operations when the temperature increases to 14 °F (–10 °C).

We've also made the following CLI command updates:

- `show chassis power` displays combined output for PSM0 and PSM1.
- `show chassis hardware` displays PSMs as built-in.
- `show chassis environment psm` displays PSM state and DC output status.

[See [show chassis power](#), [show chassis hardware](#), [show chassis environment psm](#), and [Optics EM Policy Support](#).]

- **Resiliency support (ACX7020)**—We support resiliency for platform components on ACX7020 routers. Resiliency enables the system to monitor component health, alert you of errors, and take appropriate action to restore normal operation based on error severity.

[See [Resiliency](#).]

Class of Service

- **Automatic ingress port oversubscription management (ACX7000)**—The ACX7000 family of Cloud Metro Routers support automatic ingress port oversubscription management, also known as priority drop (PRD). When the overall data rate coming from the network exceeds the device's ingress capacity, PRD intelligently discards packets based on priority, determined through packet headers. PRD thus ensures that the device is less likely to drop high-priority and control traffic.

You don't need to configure the PRD feature. PRD is included in ACX7000 routers running Junos OS Evolved 25.1R1 or later.

[See [Class of Service User Guide for Routers](#).]

EVPN

- **SRv6 TE with fallback support (ACX7020-AC, ACX7020-DC, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, and ACX7024)**—You can implement SRv6 traffic engineering (TE) with fallback capabilities, using both Segment Routing Header (SRH) and micro-SID (uSID) formats. This feature allows for enhanced traffic management and reliability by providing alternative paths if the intended route fails.

[See [Configuring EVPN VPWS over SRv6 with Traffic Engineering](#).]

- **EVPN-VPWS FXC with SRH and micro-SID support (ACX7020-AC, ACX7020-DC, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, and ACX7024)**—You can configure Ethernet VPN–virtual private wire service (EVPN-VPWS) with flexible cross-connect (FXC) options in both SRH and micro-SID formats to improve network segmentation and connectivity. FXC is available in both VLAN-

aware and VLAN-unaware modes, enabling flexible and efficient Layer 2 service delivery over SRv6 transport.



NOTE: Egress protection is not supported for EVPN-VPWS FXC instances.

[See [Configuring EVPN VPWS over SRv6 with Traffic Engineering](#).]

- **Support for excluding MAC addresses from duplicate MAC detection (ACX7100-32C, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)**—You can configure an exclusion list for MAC addresses in EVPN networks to prevent legitimate MAC address movements from being marked as duplicates. Use `set protocols evpn mac-list list_name mac-address mac_address_with_prefix_len` to create the list and `set protocols evpn duplicate-mac-detection exclude-list list_name` to apply it. This feature helps maintain network stability by avoiding unnecessary duplicate MAC detection for specified addresses, particularly in scenarios involving virtual MAC configurations in redundant setups.

[See [EVPN Duplicate MAC Detection Exclusion Lists](#).]

High Availability

- **High availability enhancements for Packet Forwarding Engine sFlow, J-Flow, and port mirroring (ACX7348)**—The sFlow, J-Flow, and port mirroring processes in the Packet Forwarding Engine maintain high availability during a Routing Engine switchover. These processes restart on the new active node, causing session filters to persist in the hardware. The hardware continues sending packets, but session statistics are reset. Inline sessions remain unaffected, which ensures continuous monitoring and data collection with minimal disruption to service.

[See [Understand Graceful Routing Engine Switchover for Junos OS Evolved](#).]

- **BGP support for static tunnels (ACX7348)**—You can utilize static tunnels with BGP while running other protocols such as OSPF, IS-IS, RSVP, and LDP over these tunnels. Tunnel encapsulation and de-encapsulation entries synchronize from the primary Routing Engine to the backup Routing Engine for seamless switchover. However, after the switchover, statistics reset to zero on the new primary Routing Engine.

[See [BGP Overview](#).]

Interfaces

- **400ZR and 400G OpenZR+ support enhancements (ACX7100)**—We support 400ZR and 400G OpenZR+ optics enhancements on ACX7100 devices. The enhancements include application selection and configuration of target output power. You can view the advertised applications and can switch between the applications.

[See [Overview](#).]

- **Support for optical port speeds (ACX7020)**—We support port speeds on optical ports with a default speed set to 10 Gbps. Ensure that total configured port bandwidth does not exceed 100GbE to avoid packet loss. Interface names follow the *et-x/y/z* convention, where *x* and *y* are fixed at 0, and *z* represents the port index. Use the `show chassis pic fpc-slot 0 pic-slot 0` CLI command to check for bandwidth issues and adjust configurations accordingly. Specify unused ports to optimize bandwidth utilization.

[See [Port Speed](#).]

- **Support for CLI options for L2 VLANs and logical interface VLAN maps for asymmetric Bridge Domains (ACX7020, ACX7100, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—We've added support for the following CLI statements, which are configured at the interfaces hierarchy:
 - `inner-list` and `inner-range`: Use these CLI statements to configure the inner VLAN ID range for double-tagged interfaces. `inner-list` doesn't support Tag Protocol Identifier (TPID) values such as 0x8100, 0x88a8, 0x9100, and 0x9200, whereas `inner-range` supports all TPID values. Ensure that you configure the `number` option for both commands. You can configure a maximum of 32 unique ranges for VLAN IDs for each physical interface (IFD) and configure a maximum of 1024 VLAN ID lists at the system level.
 - `no-native-vlan-insert`: Use this statement if you don't want to add the native VLAN ID to untagged traffic at ingress. If you do not configure `no-native-vlan-insert`, the device adds the native VLAN ID to untagged traffic.
 - `vlan-id-range`: Use this statement to bind a range of VLAN IDs to a logical interface.

[See [no-native-vlan-insert](#), [vlan-id-range](#), [Configuring Dual VLAN Tags](#), and [Binding VLAN IDs to Logical Interfaces](#).]

Junos Telemetry Interface

- **Physical Ethernet interface sensor (ACX7509)**— Junos telemetry interface (JTI) supports ON_CHANGE and periodic streaming of physical Ethernet interface statistics by using RPC developed by Google (gRPC) or gRPC Network Management Interface (gNMI) from an ACX7509 device to an external collector. This feature supports OpenConfig model *openconfig-if-ethernet.yang* (physical interface level) version 2.6.2 (no configuration). Use the base sensor path `/interfaces/interface/ethernet/state/` in a gRPC or gNMI subscription to export statistics from the ACX7509 device to an external collector.

With this feature you can now subscribe physical ethernet interface sensors on ACX7509 platform.

[See [Junos YANG Data Model Explorer](#).]

- **MACsec statistics (ACX7100-32C, ACX7100-48L, ACX733, ACX7348, and ACX7509)**—Junos telemetry interface (JTI) provides ON_CHANGE and periodic streaming of Media Access Control

Security (MACsec) statistics by using RPC developed by Google (gRPC) or gRPC Network Management Interface (gNMI) from a device to an outside collector. This feature supports the data model *openconfig-macsec.yang*. Additional leaves outside the scope of the data model are augmented. To stream MACsec statistics, include the sensor path */macsec/* in a subscription.

[See [Junos YANG Data Model Explorer](#).]

- **Enhanced telemetry with multiple gRPC servers and multi-port gRPC services (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5700, and QFX5700E)**—You can configure multiple RPC developed by Google (gRPC) servers with distinct services, listening addresses, and ports by using the Junos telemetry interface (JTI). This feature enhances control over service management and telemetry data collection. You can also configure TLS certificates for secure communications. For example, you can configure a server to listen on a specific port and serve only designated gRPC services, enhancing flexibility and security in your telemetry setup.
- **Native YANG state model and telemetry support for network stack protocol statistics (ACX Series, QFX Series, and PTX Series)**—You can use a native YANG state model and telemetry to monitor network stack protocol statistics on EVO platforms. Telemetry provides real-time data streaming for protocols such as TTP, ICMP, MPLS, TCP, and more. These statistics, previously available through CLI commands, are now accessible through telemetry streaming, ensuring real-time updates. This feature provides a comprehensive and dynamic monitoring solution, configurable in either "on-change" or "periodic" mode, enhancing your network's observability and performance management.

[See [Junos YANG Data Model Explorer](#).]

- **Telemetry support through gRPC interface for real-time network monitoring (ACX7509)**—You can use telemetry to gain real-time insights into network performance and the current state of network elements. This feature supports new OpenConfig sensor paths that provide a high-frequency, asynchronous push model for exporting performance data, focusing on key metrics such as I/O, error counters, and queue statistics. Use the gRPC interface to integrate with client applications and model-driven telemetry to receive data in YANG models.

[See [Junos YANG Data Model Explorer](#).]

- **MPLS RSVP-TE sensors, OpenConfig for MPLS and LDP, and event-driven streaming (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—Junos telemetry interface (JTI) supports MPLS RSVP-TE sensors through the *openconfig-mpls-rsvp.yang* (version 4.0.0) data model. The feature supports LDP with the *openconfig-mpls-ldp.yang* data model. The feature also supports *openconfig-mpls.yang* (version 3.2.2), *openconfig-mpls-types.yang* (version 3.2.1), *openconfig-mpls-te.yang* (version 3.2.2), and *openconfig-mpls-static.yang* (version 3.2.2). The device supports these OpenConfig configurations *mpls global*, *mpls named-explicit-path*, and *mpls tunnels*. It supports the state groups *MPLS tunnels*, *MPLS named-*

explicit-path, *MPLS static label-switched-path*, *MPLS -TE interface attributes*, and *MPLS tunnel state counters*. Use ON_CHANGE notifications to stream MPLS LSP record route object statistics only when data changes. The resource path `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/sessions/session/record-route-objects/record-route-object/state/` is supported.

The feature supports the data models:

- **openconfig-mpls-ldp.yang**
- **openconfig-mpls.yang** (version 3.2.2)
- **openconfig-mpls-types.yang** (version 3.2.1)
- **openconfig-mpls-te.yang** (version 3.2.2)
- **openconfig-mpls-static.yang** (version 3.2.2)

The feature supports the OpenConfig configurations:

- `mpls global`
- `mpls named-explicit-path`
- `mpls tunnels`

The feature supports the state groups:

- MPLS tunnels
- MPLS named-explicit-path
- MPLS -TE interface attributes
- MPLS static label-switched-path
- MPLS tunnel state counters

The feature supports the sensor path `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/sessions/session/record-route-objects/record-route-object/state/`.

Use ON_CHANGE notifications to stream MPLS LSP record route object statistics only when data changes.

[See [Junos YANG Data Model Explorer](#).]

- **Support for OpenConfig network instance configuration and state (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7332, and ACX7348)**—Use Junos telemetry interface (JTI) to monitor new sensors for network instance statistics in the **openconfig-network-instance.yang**, **openconfig-local-routing.yang**, and **openconfig-routing-policy.yang** modules. OpenConfig network instance

commands map to Junos OS configurations. For mapping details, see [Mapping OpenConfig Network Instance Commands to Junos Operation | Junos OS | Juniper Networks](#).

[See [Junos YANG Data Model Explorer](#).]

- **Support for OSPF, IS-IS, BGP-RIB, ARP, and SR-TE OpenConfig configurations and state data streaming (ACX7020-AC, ACX7020-DC, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—Junos telemetry interface (JTI) supports the following features:
 - Support for OpenConfig OSPF data model **openconfig-ospfv2.yang** (*version 0.3.1*) includes OpenConfig configuration and streaming of operational state data under resource path `/network-instances/network-instance/protocols/protocol/ospfv2/`.
 - Support for OpenConfig IS-IS configurations and sensors based on the OpenConfig data model **openconfig-isis.yang** (*version 1.0.0*).
 - OpenConfig support for proxy Address Resolution Protocol (ARP) and IPv6 duplicate address detection (DAD) configurations based on the OpenConfig data model **openconfig-if-ip.yang** (*version 3.0.0*).
 - JTI support for operational state sensors based on the following latest OpenConfig BGP-RIB data models:

Table 2: BGP-RIB Data Models

File Name	Version
openconfig-rib-bgp-attributes.yang	<i>0.8.1</i>
openconfig-rib-bgp-ext.yang	<i>0.6.0</i>
openconfig-rib-bgp-shared-attributes.yang	<i>0.8.1</i>
openconfig-rib-bgp-table-attributes.yang	<i>0.8.1</i>
openconfig-rib-bgp-tables.yang	<i>0.8.1</i>
openconfig-rib-bgp-types.yang	<i>0.5.0</i>
openconfig-rib-bgp.yang	<i>0.8.1</i>

The following model versions are no longer supported:

Table 3: Unsupported BGP-RIB Data Models

File Name	Version
openconfig-rib-bgp-ext.yang	0.2.0
openconfig-rib-bgp-types.yang	0.2.0
openconfig-rib-bgp.yang	0.2.0

The following OpenConfig BGP models are upgraded to version 9.1.0:

Table 4: Upgraded BGP RIB Data Models

File Name	Version
openconfig-bgp-global.yang	9.1.0
openconfig-bgp-neighbor.yang	9.1.0
openconfig-bgp-peer-group.yang	9.1.0

The upgraded models introduce new leaves for operational state sensors and OpenConfig configurations.

- Telemetry streaming of operational state data for 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`.

New sensor paths are added for the above features.

[See [Junos YANG Data Model Explorer](#).]

- **TWAMP sensor and OpenConfig BFD (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**— Junos telemetry interface (JTI) supports:
 - Periodic streaming of Two-Way Active Measurement Protocol (TWAMP) statistics for IPv4 and IPv6 traffic in TWAMP-managed and TWAMP Light sessions

Proprietary RPC developed by Google (gRPC) and gRPC Network Management Interface (gNMI) support the sensor that streams this data.

- OpenConfig configuration and BFD state data streaming.

OpenConfig BFD configuration commands map to relevant Junos OS commands. See [Mapping OpenConfig BFD Commands to Junos Operation](#).

New sensor paths are added for these features.[See [Junos YANG Data Model Explorer](#).]

- **Health monitoring sensors (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**— Junos telemetry interface (JTI) provides native sensors to monitor device infrastructure health. Device streams health statistics that external collectors use to track performance.

Use the resource path `/state/system/infrastructure/junos-evolved/` to view the health statistics.

These sensors stream details such as cluster data, distributor statistics, Distributed Data Store (DDS) client information, common resources, and indexes for Identity Management and Device Management .

[See [Junos YANG Data Model Explorer](#).]

Multicast

- **RPT-SPT mode support for PE routers connected directly to sources and receivers (ACX7020-AC, ACX7020-DC, ACX7024, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—You can configure rendezvous-point tree–shortest-path tree (RPT-SPT) mode with sources and receivers directly connected to the PE router, without the need for a CE router. To enable this feature, configure the `sg-forwarding-only` statement on the PE routers under the `edit routing-instances <routing-instance-name> protocols mvpn mvpn-mode rpt-spt hierarchy`.

This configuration ensures that the (*, G) forwarding entries are not installed on the PE router.

It is recommended to add this configuration statement to all Junos OS Evolved-based ACX PE devices that are part of the deployment, instead of configuring only the PE device that is connected directly to a source or receiver.

[See [sg-forwarding-only](#), and [RPT- SPT Mode with Direct Sources and Receivers](#).]

- **MVPN bud node support with Looping back interface (LBI) (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**— You can terminate multicast virtual private network (MVPN) traffic to the correct routing instance in a bud node by using the Looping back interface (LBI). On detection of this interface, MVPN P2MP multipoint LDP modules install a newly labelled route, where the next hop is updated to include the LBI and an additional branch to the next egress PE router. The LBI is created automatically on boot up and needs no configuration.

[See [Bud Node Support with the Looping Back Interface \(LBI\)](#), and [show-interfaces-lbi](#).]

Multichassis Link Aggregation (MC-LAG)

- **Support for VPLS routing instances on MC-LAGs (ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024)**—The virtual private LAN service (VPLS) routing instances feature is available on multichassis link aggregation groups (MC-LAGs.)

[See [Understanding Multichassis Link Aggregation Groups](#).]

Network Management and Monitoring

- **Support for multiple gRPC servers hosting different service sets (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5700, and QFX5700E)**—You can configure multiple gRPC servers that host different sets of services on unique ports. Additionally, each server can support different certificates, listening addresses, and routing instances. You configure the gRPC servers at the [edit system services http servers] hierarchy level. Distributing gRPC services across different servers allows for better allocation of network resources, reducing the risk of port conflicts and optimizing server performance.

[See [Configure gRPC Services](#) and [server](#).]

- **SNMP MIB walk and traps support for coherent ZR optics (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Use SNMP MIB walk and traps to efficiently monitor and manage coherent ZR optics, including 100ZR, 400ZR, 400ZR-M, and 400ZR-M-HP. Use this feature to retrieve OID information sequentially and receive notifications when alarms are triggered or cleared. This enhancement extends SNMP MIB walks to include new digital optical monitoring (DOM) fields and implements SNMP traps for critical alarms.

[See [Enterprise-Specific MIBs for Junos OS Evolved](#) and [show snmp mib](#).]

Precision Time Protocol (PTP)

- **Assisted partial timing support over PTP (ACX7024)**—Assisted partial timing support (APTS) is a Global Navigation Satellite System (GNSS) backed by Precision Time Protocol (PTP), delivering accurate timing and synchronization in mobile backhaul networks.

APTS enhances network synchronization by combining local GNSS timing with PTP-based calibration from a core clock. This ensures accurate timekeeping even during GNSS outages, making it ideal for edge deployments in mobile and distributed networks.

[See [Assisted Partial Timing Support on Routing Platforms](#).]

- **Support for configurable hold over time interval (ACX7024)**—In an assisted partial timing support (APTS) setup, when both GNSS and the PTP time references are lost or inactive, the internal clock oscillator can provide synchronization. You can now set the required duration for synchronization with the newly introduced command set `protocols ptp holdover-in-spec-duration` on ACX7024. This functionality is enabled by default with a duration of 240 minutes.

[See [ptp](#).]

Proxy ARP

- **Support for proxy MAC addresses in an ARP request (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—Provider edge (PE) devices in an EVPN network that support Address Resolution Protocol (ARP) proxy can use a proxy MAC address in the ARP replies to a host. When a PE device receives an ARP or Neighbor Discovery Protocol (NDP) request, it searches for the requested IP address in the MAC-IP address binding database. If the device finds the MAC-IP address entry in its database, it responds to the request with the proxy MAC address. The proxy MAC address is derived from the integrated routing and bridging (IRB) interface or virtual gateway address (VGA) interface in an EVPN network with edge-routed bridging overlay. The address is derived from the manually configured MAC address in a centrally routed bridging overlay. If the PE device does not have an entry for the target IP address, the device resolves the ARP or NDP for that IP address and provides a proxy response for the following ARP or NS.

To enable this feature, configure the `proxy-mac` statement at the *[edit routing-instances routing-instance-name protocols evpn]* or *[edit routing-instances routing-instance-name bridge-domains domain_name]* hierarchy level.

[See [ARP and NDP Request with a proxy MAC address](#).]

Routing Policy and Firewall Filters

- **Automatic dropping for nonlocal packets (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, , PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—The device drops all packets that are not local to the Routing Engine, unless they are flagged as exception packets. This feature is automatically enabled on all supported platforms. The dropped packet count is available under the “non-local drops” counter in the `show system statistics` command.

[See [show system statistics](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **SRv6 unreachable prefix announcement (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**— To ensure scalability and prevent overwhelming all nodes with every prefix, route summarization at area

border routers (ABRs) conceal local domain details. Segment Routing for IPv6 (SRv6) further streamlines route summarization and condenses locators from remote domains and disseminates them into the core network, which can obscure local domain activities.

A provider edge router does not immediately detect the loss of reachability when a remote edge device becomes unreachable, resulting in a traffic drop until BGP sends a status update. The ABR assigns a maximum metric to prefixes from unreachable devices, ensuring they leak across domains as Unreachable Prefix Advertisements (UPAs).

To enable the UPA, include the `prefix-unreachable` statement at the `[edit protocols isis]` hierarchy level.

[See [prefix-unreachable](#).]

- **BGP-LS advertisements of PCE delegated and initiated SRv6 TE tunnels (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and PTX10002-36QDD)**—Report static Segment Routing for IPv6–Traffic Engineering (SRv6-TE) tunnels with static segment list with micro SID (uSID) configuration to Path Computation Element (PCE). When the PCE controller provisions an SRv6-TE tunnel with uSIDs, BGP-LS advertises the SRv6-TE tunnel with its uSID segment list. This feature supports the SID Structure TLV 1252 and the SRv6 endpoint behavior TLV 1250, which are now available in the PCE report. When the externally controlled and routed SRv6-TE receives a PCUpdate message with uSIDs from the controller, BGP-LS advertises the endpoint behavior of the uSIDs.

[See [Enable Segment Routing for the Path Computation Element Protocol](#) and [SRv6-TE Tunnels with micro-SIDs in PCEP](#) .]

- **SRv6 SR-TE binding SID support (ACX7020-AC, ACX7020-DC, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—Binding SID (BSID) segregates the network into multiple traffic engineering (TE) tunnels. You can allocate Segment Routing over IPv6 (SRv6) BSIDs under an SRv6-TE tunnel to reduce the stack size at the ingress node. The ingress node stores only a subset of TE database from other domains and transit nodes with the binding SID handles convergence. SRv6 binding supports both classic and micro SIDs (uSIDs).

[See [Binding SID \(BSID\) for Static SRv6 TE Tunnels](#)]

- **Support for SRv6-TE path computation (ACX7024, ACX7100-32C, ACX7348, and ACX7509)**—Segment Routing for IPv6–Traffic Engineering (SRv6-TE) path computation enhances your IPv6 network's routing efficiency by enabling the local computation of SRv6 TE paths using both classic SIDs and micro SIDs (uSIDs). You can embed explicit paths within IPv6 packets, optimizing routing paths and reducing overhead. By default, SRv6-TE path computation prefers uSIDs over classic SIDs, resulting in paths that may consist entirely of micro SIDs, classic SIDs, or a combination of both.

[See [Understanding SRv6 TE Tunnel Local Path Computation](#)].

VLANs

- **Support for uniform egress for L2 VLANs and logical interface VLAN maps for asymmetric Bridge Domains (ACX7020, ACX7100, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)–**

We've added support for VLAN maps in an asymmetric bridge domain (BD). Configure a uniform egress logical interface (IFL) VLAN map for an asymmetric bridge domain (BD) only if there is no VLAN ID under the BD. This enables you to configure a uniform VLAN map for the egress traffic of an asymmetric bridge domain and simplifies VLAN management by ensuring consistent VLAN tagging across the domain.

[See [Understanding VLAN Manipulation \(Normalization and VLAN Mapping\) on Ethernet Services](#) and [Optimized Intersubnet Multicast in EVPN Networks](#).]

Additional Features

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

- **EVPN E-LAN over SRv6 (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, and ACX7024)**

[See [EVPN E-LAN over SRv6](#).]

- **NIST purge method for media sanitization (ACX7024, ACX7024X, ACX7100-32C, ACX7332, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10008, QFX5700, and QFX5700E).** We've extended support for NIST media sanitization for SATA hard disk drives to include:

- Cryptographic scramble and block erase priorities for the purge method.
- Enhanced secure erase priority for the clear method.

[See [NIST Special Publication 800-88, Guidelines for Media Sanitization](#) and [request system zeroize](#).]

- **Support for asymmetric EVPN Type 5 routes (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)**

[See [EVPN Type 5 Route with VXLAN Encapsulation for EVPN-VXLAN](#) and [ip-prefix-routes](#).]

- **Support for asynchronous notification in E-LAN and E-Tree over EVPN-MPLS networks (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)**

[See [asynchronous-notification](#).]

- **Support for EVPN E-LAN and EVPN E-Tree over BGP-LU PIC Edge (ACX7100-32C, ACX7100-48L, ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)**

[See [BGP PIC for Layer 3 VPNs](#) and [labeled-unicast](#).]

- **Support for FEC 128 and FEC 129 VPLS with source packet routing (ACX7332)**

[See [site \(VPLS Multihoming for FEC 128\)](#), [site \(VPLS Multihoming for FEC 129\)](#), [Example: Configuring VPLS Multihoming \(FEC 129\)](#), and [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

- **Support for L2PT for VPLS** (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)

[See [Configuring L2PT for VPLS](#).]

- **Support for network-isolation in E-LAN and E-Tree over EVPN-MPLS networks** (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)

[See [network-isolation](#) and [network-isolation-profile](#).]

- **Support for no-local-switching in E-LAN over EVPN-MPLS networks** (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)

[See [no-local-switching](#) and [core-facing](#).]

- **Support for VPLS ping and disabling MAC learning** (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509). You can now disable MAC learning and use VPLS ping to send packets to a particular destination.

- To send ping packets to a destination for a specific VPLS instance, use the `ping vpls instance` command. [See [ping vpls instance](#).]

- To disable MAC learning for a particular VLAN, include the `no-mac-learning` statement at the `[edit routing-instances routing-instance-name vlans vlan-name switch-options]` hierarchy level. [See [no-mac-learning](#).]

- To disable MAC learning globally on a device, include the `global-no-mac-learning` at the `[edit protocols l2-learning]` hierarchy level. [See [global-no-mac-learning](#).]

- **QSFP-100G coherent ZR optics performance monitoring** (ACX7509). Monitor the performance of QSFP-100G coherent ZR optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Accumulate performance metrics into 15-minute and 1-day interval bins. Use the `show interfaces transport pm` command to view current and historical performance data.

[See [optics-options](#), and [show interfaces transport pm](#).]

- **Support for internet-options commands** (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR|PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5220). We have added additional traffic management command options to the `[set system internet-options]` hierarchy:

- `no-tcp-reset drop-all-tcp`

- `no-tcp-reset drop-tcp-with-syn-only`
- `ipv6-path-mtu-discovery-timeout` *minutes*

[See [internet-options](#).]

- **Paragon Active Assurance test agent** (ACX7020). Install and use the Paragon Active Assurance test agent on your router to monitor network quality, availability, and performance.

[See [Install the Paragon Active Assurance Test Agent](#).]

- **Support for IPv6 agent-ID in sFlow configuration** (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509). You can configure the IPv6 agent-ID statement at the `[edit protocols sflow]` hierarchy level. With this enhancement, you can utilize sFlow with improved agent identification, facilitating better network traffic monitoring and analysis.

[See [sflow](#) and [show sflow](#).]

- **Specify the install package name as a URL in the request system software add command** (ACX Series, PTX Series, and QFX Series)

[See [request system software add \(Junos OS Evolved\)](#).]

- **Support for file-system encryption with Trusted Platform Module (TPM 2.0)** (ACX7100-32C, ACX7100-48L, and PTX10002-36QDD)

[See [Encryption with Trusted Platform Module](#).]

- **Support for IRB in L3VPN VRF over SRv6** (ACX7020-AC, ACX7020-DC, ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)

[See [Layer 3 VPN Services over the SRv6 Core](#).]

- **Supported transceivers, optical interfaces, and DAC cables** (ACX Series). 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 the HCT and provide the first supported release information when the optical transceiver becomes available.

- **Support for packets-per-second (pps) policer** (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)

[See [Packets-Per-Second \(pps\)-Based Policer Overview](#).]

What's Changed

IN THIS SECTION

- [General Routing](#) | 47
- [Class of Service \(CoS\)](#) | 48
- [EVPN](#) | 48
- [Junos XML API and Scripting](#) | 48
- [Routing Protocols](#) | 49
- [System Management](#) | 49
- [User Interface and Configuration](#) | 49

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

General Routing

- **Changes to default behavior under forwarding-table**— `ecmp-fast-reroute` and `indirect-next-hop-change-acknowledgements` are enabled by default under the `edit routing-options forwarding-table` hierarchy. You can verify these defaults by running the `show configuration routing-options forwarding-table` in operational mode.

[See [ecmp-fast-reroute](#) and [indirect-next-hop-change-acknowledgements](#).]

- Support added for source and destination port optimization for port ranges for IPv6 input firewall filters.
- The `show subscribers extensive client-type dhcp | display xml validate` command has now been updated to display correct output instead of **Duplicate data element** error.
- **SFP Optics LOS alarms**— SFP Optics don't support Tx laser disabled alarm, Tx loss of signal functionality alarm, and Rx loss of signal alarm as diagnostics output.

See [show interfaces diagnostics optics](#).

- **Change in CLI output**—The CLI output for `show system license bandwidth`, `show system license bandwidth fpc`, and `show system license fpc` commands is updated.

See [Monitor Junos Licenses](#).]

- **Deprecated license trace (Junos OS Evolved)**—We've deprecated the CLI option `show system license liblicense-trace`.

Class of Service (CoS)

- In Junos OS Evolved, do not associate a default forwarding class name with a different queue number. Use a custom forwarding class name instead. Defining customer forwarding classes with factory default forwarding class names causes errors.

EVPN

- **Easy EVPN LAG (EZ-LAG) feature lightweight loop detection configuration change**— We have updated the configuration generated by the easy EZ-LAG features commit script to use logical interface names for the lightweight leaf to server loop detection feature. The EZ-LAG commit script previously generated the loop detection configuration with physical interface names, but the loop detection feature works only for logical interfaces. The lightweight loop detection configuration uses the `loop-detect` statement at the `edit protocols` hierarchy level.

[See [Easy EVPN LAG \(EZ-LAG\) Configuration](#).]

- **Duplicate MAC detection timeout**—The default setting for `auto-recovery-time` is 5 minutes on these platforms only.

[See [duplicate-mac-detection](#).]

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, PTX Series, and QFX Series)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [cert-file](#).]

Routing Protocols

- **Extension of traceoptions support for VLANs in IGMP/MLD snooping**— The traceoptions option is supported under the edit routing-instance protocols igmp-snooping vlan and edit routing-instance protocols mld-snooping vlan hierarchy. traceoptions can be enabled for both specific and all vlans.

See [vlan \(IGMP Snooping\)](#) and [vlan \(MLD Snooping\)](#).]

System Management

- **Core files generated for related processes are bundled into a compressed TAR file (ACX Series, PTX Series, and QFX Series)**—When a process generates a core file and the system generates a live core for its related process, the system bundles the core files into a compressed TAR file. In earlier releases, the system does not bundle the files.

The process pairs that support this feature are: bfdd and bfddagent; cfmd and cfmd-agent; dot1xd and dot1xd-agent; l2ald and l2ald-agent; l2cpd and l2cpd-agent; mcsnoopd and mcsnoopd-agent; ppmd and ppmdagent; routing and rpdagent.

User Interface and Configuration

- **Access privileges for request support information command (ACX Series, PTX Series, and QFX Series)** — The request support information command is designed to generate system information for troubleshooting and debugging purposes. Users with the specific access privileges maintenance, view, and view-configuration can execute request support information command.
- **Changes to the show system storage command output (ACX Series, PTX Series, and QFX Series)**—We've updated the show system storage command output to include only true (physical) storage and exclude any host/hypervisor level storage. In earlier releases, the output also includes a container/jail storage, which does not have a separate storage of its own.

[See [show system storage](#).]

- **Option to view combined disk space usage statistics for all configuration databases (ACX Series, PTX Series, and QFX Series)**—The show system configuration database usage command provides the merge option. When you include the merge option, the command output displays combined disk space usage

statistics for all configuration databases, including the static configuration database and all ephemeral configuration database instances.

[See [show system configuration database usage](#).]

- **Changes to `show system statistics` command output (ACX Series, PTX Series, and QFX Series)**—We've updated the output fields for the following commands:
 - `show system statistics mpls`—We've added a new field: Packets with IPv6 explicit NULL tag.
 - `show system statistics jtd`—We've renamed two fields to better reflect the underlying data representation. The MPLS IPv4 packet received field is renamed to MPLS IPv4 explicit NULL packets received, and the MPLS IPv6 packets received field is renamed to MPLS IPv6 explicit NULL packets received.

Known Limitations

IN THIS SECTION

- [General Routing | 50](#)

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

- When the original flow is egressing out through an aggregated Ethernet (AE) interface, the corresponding sampled sflow frame does not reflect the correct egress port number. This happens only when the flow is egressing out through an aggregated Ethernet (AE) interface. For non-AE egress interface, this works fine and the sflow frame reflects the correct egress port. [PR1647870](#)
- Only on ACX700 devices when the device is rebooted incorrect CT values are seen. Though incorrect CT values are seen, it does not impact any functionality. [PR1726761](#)

- On ACX7024 Junos OS Evolved platform, performance measurement impact is seen (measurement is incorrect since compensation removed) on deactivating both chassis and PTP configuration and power off USB. It is very unlikely sequence to be happened. [PR1868449](#)

Open Issues

IN THIS SECTION

- [General Routing | 51](#)
- [Infrastructure | 52](#)

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

- When ingress policer is configured, to drop ingress traffic, on an interface with upMep the CFM packets generated from the upMep is also dropped due to the policer. This leads to CFM session going down. [PR1754938](#)
- When DHCP trace options are enabled, there is a possibility that jdhcpd could generate a core file. In general, traceoptions should be enabled only for debugging. They should be disabled once debugging is done. [PR1771121](#)
- This is day-1 issue of Juniper Networks server. It exists on all ACX Series platforms where Juniper Networks server runs. Once the asymmetry configured at slave port, the error propagated to the down stream eventually causes performance issue of spike. [PR1793926](#)
- The application rpd-agent might restart with a core file after interface related event changes. This issue is likely to occur when the interfaces are activated/deactivated/created/deleted. This could be due to
 - Bulk change in interface configuration or

- Admin-initiated application restart through command `request system application node re0 restart app <application>`
- System triggered automatic application restart due to internal events.



NOTE: This issue occurs due to very rare combination of events and is not seen in most use cases. There is no service impact when the issue occurs, and router operations are unaffected.

[PR1885455](#)

- When a device is rebooted with PAA installed in 25.2R1 , PAA installation on reboot might fail due to **Not found default vrf** error. This can be resolved by deactivating and activating or freshly installing the PAA configuration. [PR1886928](#)

Infrastructure

- On all Junos OS Evolved ACX Series platforms, modifying an ECMP route triggers the deletion of the old Next Hop route, which can temporarily disrupt traffic flow until the new route is fully established. [PR1820482](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 53](#)
- [Interfaces and Chassis | 54](#)
- [Layer 2 Features | 54](#)
- [Subscriber Access Management | 55](#)

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

- [ACX7000 Series] DHCPv4/v6 packets might be dropped because DHCP packets are not routed to kernel after initial jdhcpd starts. [PR1816246](#)
- Traffic is dropped without entering the SRv6 dynamic tunnel. [PR1836457](#)
- Packet corruption in L3VPN traffic with preserve-nextthop-hierarchy. [PR1840722](#)
- ACX7509 : LACP interfaces stuck in Attached state. [PR1840790](#)
- High CPU utilization observed on all Junos OS Evolved ACK7000 platforms. [PR1841573](#)
- Junos OS Evolved-pfemamd process crashes on ACX7348/ACX7332/ACX7024/ACX7509 platforms. [PR1842389](#)
- Possible SSD **read disturbance** due to platform_monitor script accessing the drive every 5 seconds. [PR1846199](#)
- PTP packets which are incoming with padded bytes are dropped on certain Junos OS Evolved ACX platforms. [PR1848586](#)
- Default route installation fails on Junos OS Evolved ACX Series platforms. [PR1848599](#)
- Traffic statistics are missing for ipdemux lite subscribers for certain Junos OS Evolved ACX platforms. [PR1850651](#)
- The Devdb messages are seen continuously on ACX Junos OS Evolved platforms. [PR1851810](#)
- Firewall Policer causing unintended packet drops when VRF route leak with static route to next-table is enabled. [PR1853767](#)
- The IRB VGA is not shown after deleting one of the IRB interfaces with the same VGA address configured. [PR1853879](#)
- The evo-pfemamd process crashes on ACX Series platforms without specific trigger. [PR1854255](#)
- ACX platform configured as a data center gateway with EVPN-VXLAN failed to forward IP prefix routes for inter-data center communication. [PR1854710](#)
- Random evo-pfemamd.re.re0 crash is observed during boot up. [PR1856390](#)
- BFD single hop session flaps seen after a link down event on Junos OS Evolved ACX platform. [PR1857248](#)
- Jflow functionality stops working when configured with sflow. [PR1858954](#)

- The firewall filter using `next-ip` as an action does not forward the packets on Junos OS Evolved ACX platforms if the `next-ip` is resolved to a static route that points to itself. [PR1859053](#)
- The `rpd` crash due to overlapping flow route updates in a single transaction. [PR1860888](#)
- The `picd` process crashes on Junos OS Evolved ACX7509 platforms during Routing Engine switchover. [PR1863708](#)
- Log messages "OOPS: Crossing the LIMIT" are observed on dual Routing Engine platforms.
- The FEC leak is seen on all ACX Series Junos OS Evolved platforms. [PR1865858](#)
- Junos OS Evolved system might fail to response to system commands. [PR1866988](#)
- Packet drops are seen due to stale entries in multicast group when HCoS is enabled on ACX7100 platforms. [PR1867301](#)
- Packet drops are seen for default route when `rpf-check loose` mode is configured. [PR1869793](#)
- Interface down is observed while configuring 1G speed on ACX7332/ACX7348/ACX7509 platforms. [PR1870528](#)
- High memory and CPU usage due to unintended phone-home client activation. [PR1871802](#)
- Transient traffic loss during multicast route convergence scenarios. [PR1876781](#)
- Traffic disruption is observed when successive switchover is performed having single-hop BFD configured in the dual Routing Engine scenario. [PR1879925](#)

Interfaces and Chassis

- Interfaces with the same outer VLAN ID but different inner vlans or inner-lists/ranges report a commit error. [PR1859501](#)

Layer 2 Features

- A traffic outage is seen when disabling and enabling the LACP configuration. [PR1850803](#)

Subscriber Access Management

- Error message is observed after device is restarted. [PR1813456](#)

Junos OS Evolved Release Notes for PTX Series

IN THIS SECTION

- [What's New | 55](#)
- [What's Changed | 72](#)
- [Known Limitations | 76](#)
- [Open Issues | 76](#)
- [Resolved Issues | 78](#)

These release notes accompany Junos OS Evolved Release 25.2R1 for PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and PTX10002-36QDD 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

- [Hardware | 56](#)
- [Authentication and Access Control | 57](#)
- [Class of Service | 57](#)
- [EVPN | 57](#)
- [Forwarding Options | 58](#)
- [High Availability | 58](#)
- [Interfaces | 58](#)

- [IPv6 | 59](#)
- [Junos Telemetry Interface | 60](#)
- [Layer 2 VPN | 62](#)
- [Network Management and Monitoring | 62](#)
- [OpenConfig | 64](#)
- [Routing Options | 64](#)
- [Routing Policy and Firewall Filters | 65](#)
- [Routing Protocols | 66](#)
- [Serviceability | 67](#)
- [Services Applications | 68](#)
- [Software Installation and Upgrade | 68](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 69](#)
- [Additional Features | 69](#)

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 25.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [PTX10001-36MR](#)
- [PTX10002-36QDD](#)
- [PTX10003](#)
- [PTX10004](#)
- [PTX10008](#)
- [PTX10016](#)

The following sections highlight the key features in this release.

Hardware

- **Supported transceivers, optical interfaces, and DAC cables (PTX10002-36QDD)**—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 the HCT and provide the first supported release information when the optic becomes available.

Authentication and Access Control

- **SSH enhancements for algorithm configuration (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)**—We've made the following updates to SSH algorithms:
 - The CLI command `set system services ssh ca-signature-algorithms` should be used to configure the signature algorithms that are allowed for certificate authorities (CAs) to use when signing certificates.
 - Under the `system services ssh hostkey-algorithm-list` hierarchy level, new options are introduced:
 - `set system service ssh hostkey-algorithm-list rsa-sha2-256`
 - `set system service ssh hostkey-algorithm-list rsa-sha2-512`

These options enable RSA hostkey signatures using the SHA-256 hash algorithm and SHA-512 hash algorithm.

- RSA signatures using the SHA-1 hash algorithm have been disabled by default. Consequently, the CLI command `set system services ssh hostkey-algorithm-list rsa` has been deprecated.
- SSH connections that require a subsystem (for example, `netconf`) need to explicitly use the `-s` option.

[See [hostkey-algorithm-list](#).]

Class of Service

- **On-chip buffer (OCB) with explicit congestion notification (ECN) (PTX10002-36QDD)**—On PTX10002-36QDD routers, virtual output queues (VOQs) with a buffer size equal to or less than 40 microseconds, at the configured buffer rate, are always on-chip. When you have a VOQ with an OCB and then configure an ECN profile for the queue, the queue remains on-chip after the system reaches 100% ECN marking probability at the defined queue occupancy. Keeping the queue on-chip eliminates the need to access slower external memory. This approach is ideal for low-latency, high-throughput applications.

[See [CoS Explicit Congestion Notification \(ECN\)](#).]

EVPN

- **Support for GRE over EVPN-VXLAN tunnels (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—You can handle GRE as a payload over EVPN-VXLAN Type 2 tunnels on the listed PTX Series routers that use the BT forwarding ASIC. This feature enables you to manage and encapsulate GRE payloads within EVPN-VXLAN tunnels. You can configure this functionality using the `tunnel-`

loopback option for IRB interfaces. This implementation supports IPv4 underlay only and is beneficial for scenarios requiring enhanced traffic management and billing accuracy in network architectures.

[See [GRE over EVPN-VXLAN](#).]

- **Support for excluding MAC addresses from duplicate MAC detection (ACX7100-32C, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)**—You can configure an exclusion list for MAC addresses in EVPN networks to prevent legitimate MAC address movements from being marked as duplicates. Use `set protocols evpn mac-list list_name mac-address mac_address_with_prefix_len` to create the list and `set protocols evpn duplicate-mac-detection exclude-list list_name` to apply it. This feature helps maintain network stability by avoiding unnecessary duplicate MAC detection for specified addresses, particularly in scenarios involving virtual MAC configurations in redundant setups.

[See [EVPN Duplicate MAC Detection Exclusion Lists](#).]

Forwarding Options

- **Override default unicast RPF and FBF behavior (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—When unicast reverse-path forwarding (unicast RPF) and filter-based forwarding (FBF) are configured on an interface, the source IP address lookup of an incoming packet is done by unicast RPF in the routing instance that the interface points to and destination IP address lookup is done in the routing instance specified by the FBF filter. You use `set forwarding-options no-rpf-fbf-handling` to override this default behavior and revert to the old behavior where the source IP lookup happens in the routing instance specified by FBF.

[See [no-rpf-fbf-handling](#).]

High Availability

- **Commit check for Routing Engine switchover (PTX10008 and PTX10016)**—If a commit is in progress, Routing Engine switchover is prohibited until the commit completes. You can ensure that active commit operations do not interfere with a Routing Engine switchover by using the `request chassis routing-engine master switch check` command. This feature prevents system crashes and undefined states by blocking the switchover until the commit process is complete.

[See [request chassis routing-engine master](#).]

Interfaces

- **LAG link protection with LACP and static configuration (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We've enhanced the reliability and availability of network links using the LAG Link-Protection feature, which supports both LACP and static LAG configurations. This feature ensures continuous network operation by seamlessly switching traffic to standby links in the event of primary links failure. You can manage failovers at individual link or subgroup levels. You can also configure primary and backup links explicitly. Additionally, various CLI

commands allow you to enable and manage link protection, including configuring fast failover and executing link switchover requests.

[See [Link Protection of Aggregated Ethernet Interfaces](#) .]

- **Support for appsel number configuration (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**—Use the `set interfaces interface name optics-options appselid id id` command to configure the appsel number.

[See [appselid](#), [optics-options](#) and [show interfaces diagnostics optics-applications](#).]

Support for display of Q-factor and Q-margin in 400ZR and 400G OpenZR+ optics (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)—Q-factor and Q-margin allows you to monitor performance, enhancing the reliability and efficiency of your network. Use the `show interfaces diagnostics optics interface name` command to display Q-factor and Q-margin.

[See [show interfaces diagnostics optics \(Routers\)](#).]

- **Support for 800G OpenZR+ Pluggable Modules (PTX10002-36QDD)**—You can enhance your data center and infrastructure connectivity with high-capacity 800G OpenZR+ pluggable modules. These modules support multiple optical modes and high-performance modulation formats, enabling transmission to reach up to 450 km at 800 Gbps. The optics modules support features such as application selection, wavelength configuration, optical loopback, and configuration of target output power.

[See [Overview](#).]

- **User-configurable performance monitoring intervals for more granular data collection (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can configure performance monitoring intervals from the default 15 minutes to shorter durations of 10 seconds, 30 seconds, 1 minute, or 5 minutes. Use the `set chassis optics pm interval-length (10s/30s/1min/5min/15min)` command to adjust these intervals according to your specific monitoring needs. This feature provides more granular performance data, helping you to better analyze network conditions. Note that changing the interval length affects current and historical performance monitoring bins and synchronizes with system time.

[See [User Configurable PM Interval Length](#), [Configure User Configurable PM Interval Length, interval-length \(Chassis\)](#), [optics \(Chassis\)](#), [Coherent Optics Performance Monitoring](#), and [show interfaces transport pm](#).]

IPv6

- **Support for NDP proxy and DAD proxy for multiple interfaces (PTX10002-36QDD)**—We support Neighbor Discovery Protocol (NDP) and duplicate address detection (DAD) for interface-restricted mode and interface-unrestricted mode. You can use existing CLI commands to enable and disable NDP proxy and DAD proxy for interface-restricted mode and interface-unrestricted mode.



NOTE: You cannot configure interface-unrestricted and interface-restricted options on the same interface simultaneously.

Junos Telemetry Interface

- **Support for controller card power management using gNMI and gNOI (PTX10001-36MR, PTX10008, and PTX10016)**—You can manage the power state of Routing Engines using the gRPC Network Management Interface (gNMI) and gRPC Network Operations Interface (gNOI) services. Use gNMI for configurations that persist across reboots, ensuring the controller card remains powered off. Use gNOI to temporarily power down a Routing Engine until the next reboot. This capability is particularly beneficial for troubleshooting and isolating faulty cards. To use the gNMI service, configure the path `/components/component/controller-card/config/power-admin-state` and set the value to "POWER_DISABLED" or "POWER_ENABLED" for the target Routing Engine. For the gNOI service, use the System service `Reboot()` RPC with the `POWERDOWN` or `POWERUP` option and specify the target Routing Engine.

[See [gNOI System Service](#) and [Sensor Power-State Management Support Using gNMI](#).]

[See [Junos YANG Data Model Explorer](#).]

- **Support for genstate YANG data models (PTX10003)**—You can subscribe to genstate YANG data models to access a subset of show command data. This feature enables a gNMI telemetry collector to subscribe to resource paths in the models, enabling you to query specific state data. This feature supports the `show snmp` command. The supported root resource paths are *genstate:genstate/snmp-statistics*, *genstate:genstate/rmon-alarm-information*, *genstate:genstate/rmon-event-information*, and *genstate:genstate/snmp-v3-information*. With this feature you can now subscribe to more specific state data on the device.

[See [Junos YANG Data Model Explorer](#) and <https://www.juniper.net/documentation/us/en/software/junos/interfaces-telemetry/topics/concept/genstate-gnmi-overview-telemetry.html>.]

- **Export timing data to collectors (ACX7100-32C, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**—Junos telemetry interface (JTI) supports export of timing attributes for Precision Time Protocol (PTP) and Synchronous Ethernet to a collector. Export of data is through native models (export of PTP data is through the YANG data model). This feature supports both periodic streaming and on-change notifications. The feature introduces the following subscription paths:
 - For Precision Time Protocol : `/state/protocols/ptp/instances/instance[instance-index]/`
 - For Synchronous Ethernet: `/state/protocols/synce/`

With this feature you can now view Precision Time Protocol and Synchronous Ethernet sensor information.

For a complete list of sensor paths supported by the device, see [Junos YANG Data Model Explorer](#).

- **Enhanced telemetry with multiple gRPC servers and multi-port gRPC services (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5700, and QFX5700E)**—You can configure multiple RPC developed by Google (gRPC) servers with distinct services, listening addresses, and ports by using the Junos telemetry interface (JTI). This feature enhances control over service management and telemetry data collection. You can also configure TLS certificates for secure communications. For example, you can configure a server to listen on a specific port and serve only designated gRPC services, enhancing flexibility and security in your telemetry setup.
- **Support for firewall filter monitoring under OpenConfig network-instance AFTs (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can configure firewall filters under the OpenConfig network instance policy-forwarding (OC-NI-PF) hierarchy, enabling state sensors capture. You can use the OpenConfig network-instance AFTs feature to monitor policy forwarding entries, such as IP prefixes, DiffServ code point (DSCP), IP protocol, and packet counters.

[See [Junos YANG Data Model Explorer](#).]

- **Native YANG state model and telemetry support for network stack protocol statistics (ACX Series, QFX Series, and PTX Series)**—You can use a native YANG state model and telemetry to monitor network stack protocol statistics on EVO platforms. Telemetry provides real-time data streaming for protocols such as TTP, ICMP, MPLS, TCP, and more. These statistics, previously available through CLI commands, are now accessible through telemetry streaming, ensuring real-time updates. This feature provides a comprehensive and dynamic monitoring solution, configurable in either "on-change" or "periodic" mode, enhancing your network's observability and performance management.

[See [Junos YANG Data Model Explorer](#).]

- **Support for genstate YANG data models (PTX10003)**—Use genstate YANG data models to access a subset of show command data. A gNMI telemetry collector can subscribe to resource paths in these models to query specific state data. The feature supports the show lacp command.

[See [Junos YANG Data Model Explorer](#), [Understanding Junos YANG Modules](#) and [gNMI Genstate Subscription](#).]

- **Health monitoring sensors (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**— Junos telemetry interface (JTI) provides native sensors to monitor device infrastructure health. Device streams health statistics that external collectors use to track performance.

Use the resource path `/state/system/infrastructure/junos-evolved/` to view the health statistics.

These sensors stream details such as cluster data, distributor statistics, Distributed Data Store (DDS) client information, common resources, and indexes for Identity Management and Device Management .

[See [Junos YANG Data Model Explorer](#).]

- **Per-segment list telemetry support for colored and uncolored SR-TE tunnels (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can configure per-segment-list sensors for segment routing–traffic engineering (SR-TE) tunnels to generate sensor IDs and collect traffic statistics from both ingress and transit points. The feature generates unique sensor IDs for each segment-list and provides the option to disable specific sensors. Additionally, updated SR-TE displays and route installations reflect per-path sensor information, ensuring comprehensive visibility and management of network telemetry.

Layer 2 VPN

- **VLAN ID range for Layer 2 circuit (PTX10001-36MR, PTX10002-36QDD, PTX10004, and PTX10008)**—You can configure VLAN ID ranges within the circuit cross-connect (CCC) encapsulation in both single-tagged and dual-tagged modes:
 - **Single VLAN Tagged**—The outer tag is set as a VLAN ID range. There is no inner VLAN tag.
 - **Dual VLAN Tagged**—The outer tag is a single VLAN ID, and the inner tag is configured as a VLAN ID range.

Configuring VLAN ID ranges allows you to bundle multiple VLANs within the same Layer 2 tunnel.

[See [Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface](#).]

Network Management and Monitoring

- **Support for multiple gRPC servers hosting different service sets (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5700, and QFX5700E)**—You can configure multiple gRPC servers that host different sets of services on unique ports. Additionally, each server can support different certificates, listening addresses, and routing instances. You configure the gRPC servers at the `[edit system services http servers]` hierarchy level. Distributing gRPC services across different servers allows for better allocation of network resources, reducing the risk of port conflicts and optimizing server performance.

[See [Configure gRPC Services](#) and [server](#).]

- **Support for controller card power management using gNMI and gNOI (PTX10001-36MR, PTX10008, and PTX10016)**—You can manage the power state of Routing Engines using the gRPC Network

Management Interface (gNMI) and gRPC Network Operations Interface (gNOI) services. Use gNMI for configurations that persist across reboots, ensuring the controller card remains powered off. Use gNOI to temporarily power down a Routing Engine until the next reboot. This capability is particularly beneficial for troubleshooting and isolating faulty cards. To use the gNMI service, configure the path `/components/component/controller-card/config/power-admin-state` and set the value to "POWER_DISABLED" or "POWER_ENABLED" for the target Routing Engine. For the gNOI service, use the `System service Reboot()` RPC with the `POWERDOWN` or `POWERUP` option and specify the target Routing Engine.

[See [gNOI System Service](#) and [Sensor Power-State Management Support Using gNMI](#).]

- **gNOI OS service RPC updates (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—We've updated the gRPC Network Operations Interface (gNOI) OS service remote procedure calls (RPCs) to more efficiently manage software upgrades. These changes enable you to install software without immediate activation, validate the current configuration against any installed software version, and activate any installed software version.

The `Install()` RPC copies the software installation package to the device, validates the configuration against the specified software version, and installs the software. The `Activate()` RPC sets the specified software version as the next boot version. If the specified software version is already installed on the device, `Install()` instead validates the current configuration against the existing software image and stores the validated current configuration as the running configuration associated with that image. In previous releases, the `Install()` RPC only copies the software installation package to the device. The `Activate()` RPC performs the rest of the actions—validates the configuration, installs the software, and sets the software version as the next boot version.

[See [gNOI Operating System \(OS\) Service](#).]

- **SNMP MIB walk and traps support for coherent ZR optics (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Use SNMP MIB walk and traps to efficiently monitor and manage coherent ZR optics, including 100ZR, 400ZR, 400ZR-M, and 400ZR-M-HP. Use this feature to retrieve OID information sequentially and receive notifications when alarms are triggered or cleared. This enhancement extends SNMP MIB walks to include new digital optical monitoring (DOM) fields and implements SNMP traps for critical alarms.

[See [Enterprise-Specific MIBs for Junos OS Evolved](#) and [show snmp mib](#).]

- **SNMP support for coherent ZR optics performance monitoring and threshold alerts (PTX10001-36MR, PTX10002-36QDD, and PTX10003)**—You can monitor the performance of coherent ZR optics (100ZR, 400ZR, 400ZR-M, and 400ZR-M-HP) and receive threshold crossing alerts using SNMP. Retrieve real-time, historical, and statistical data for various performance parameters through SNMP Get requests. You also can receive trap notifications for threshold crossing alerts and clear events.

Use the updated enterprise MIB named Juniper-IFOPTICS-MIB to comprehensively monitor and manage coherent ZR series transceivers.

[See [show snmp mib](#), [Enterprise-Specific MIBs for Junos OS](#), and [SNMP MIB Explorer](#).]

- **Export of extended gateway structure in sFlow records (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can export an extended gateway structure in sFlow records. The extended gateway structure includes BGP next-hop, autonomous system (AS) numbers, and community data. You must configure the extended-gateway option at the [edit protocols sFlow family structure-list] hierarchy level to export an extended gateway structure in sFlow records.

The sFlow functionality is managed by the sflowapp application instead of the sflowd process. This enhancement improves the sFlow scaling and integrates BGP-related data, such as AS numbers and communities, into the sFlow records, facilitating more comprehensive network monitoring and analysis.

[See [sFlow Technology Overview](#) and [sFlow](#).]

- **sFlow configuration support for aggregated Ethernet bundles (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can configure sFlow over an aggregated Ethernet bundle and manage sFlow sampling for all interfaces within the bundle through a single configuration command. This feature improves efficiency with sFlow configured at the aggregated Ethernet interface level in the CLI, and individual interfaces configured at the physical interface level. Use the command `set protocols sFlow interfaces ae0 sample-rate ingress 1` to set the sampling rate for the aggregated Ethernet bundle.

[See [sFlow Technology Overview](#) and [sFlow](#).]

OpenConfig

- **Support for firewall filter monitoring under OpenConfig network-instance AFTs (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can configure firewall filters under the OpenConfig network instance policy-forwarding (OC-NI-PF) hierarchy, enabling state sensors capture. OpenConfig network instance AFTs feature allows you to monitor policy forwarding entries, such as IP prefixes, DSCP, IP protocol, and packet counters.

[See [Junos YANG Data Model Explorer](#).]

Routing Options

- **Route limiter for flow specification at global family level (PTX Series)**— Limit the number of flow specification routes at the global level across all routing instances for protection of flow specification filter resources. You can limit the flow routes at the IPv4 family level, IPv6 family level and at the global level. Use the new `global` option at the [edit routing-options flow] hierarchy level.

[See [Understanding BGP Flow Routes for Traffic Filtering](#).]

Routing Policy and Firewall Filters

- **Automatic dropping for nonlocal packets (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, , PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—The device drops all packets that are not local to the Routing Engine, unless they are flagged as exception packets. This feature is automatically enabled on all supported platforms. The dropped packet count is available under the “non-local drops” counter in the `show system statistics` command.

[See [show system statistics](#).]

- **Support for firewall filter match conditions for fast lookup filters (PTX10002-36QDD and PTX12008)**—New match conditions for fast lookup filters have been introduced in this release for these platforms.

[See [fast-lookup-filter \(PTX\)](#).]

- **View a CLI and non-CLI firewall filter's configured and compiled information (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—When firewall filters are configured, an optimization operation is performed on the configuration. The optimization process may merge or eliminate the terms of filters. This action can lead to differences between the configured filters and the filters programmed in the hardware. Two new `show` commands have been introduced to display a CLI or a non-CLI firewall filter's configured information or this same firewall filter's information after its compilation/optimization.

[See [show-firewall-configuration](#).]

- **Match source or destination ports in named list (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—You create a port-list to conveniently group multiple ports (source or destination ports) so that they can be referenced easily in firewall configurations as port-list, source-port-list and/or destination-port-list match conditions.

[See [port-list](#).]

- **Use policies to validate flow specification filters (PTX Series)**—Use policies to validate the flow specification filters at the edge routers signalling flow routes over external BGP (EBGP) session to the peers. By configuring the policies, you can prevent the flow routes from accidentally or maliciously blocking protocol sessions. You can also prevent the admission of malformed, unsupported, or undesired flow routes coming from the source.

Configure policies by specifying the match conditions and flow route actions at the `[edit policy-options flowspec-attribute]` hierarchy level.

[See [Configuring Policies for Flow Route Validation](#).]

- **Policy to enable per-route-accounting on selective flow routes (PTX Series)**—You can selectively enable individual counters for flow specification routes. Use the new policy action flow route accounting in the following statement format.

```
set policy-options policy-statement < term > then flow-route-accounting
```

[See [flowspec-attribute](#)].

- **New CLI option for flow family matching policy configuration (PTX Series)**—The following new CLI options are available for configuring policies to match against specific family routes. Use these options at the [edit policy-options policy-statement from family] hierarchy level:

inet-flow—IPv4 flow family

inet6-flow—IPv6 flow family

inet-vpn-flow—IPv4 VPN flow family

inet6-vpn-flow—IPv6 VPN flow family

[See [flowspec-attribute](#)].

Routing Protocols

- **BGP multipath prioritization with configurable priority queues (PTX10001-36MR and QFX5130-48C)**—Use BGP multipath prioritization to prioritize critical routes in BGP multipath computations with queues of low, medium, and high priorities. This feature reduces multipath-calculation latency for operator-prioritized routes in overall route convergence during high load. Additionally, you can manage delayed forwarding information base (FIB) convergence that results from frequent routing changes.

[See [BGP Route Prioritization](#) and [multipath \(Protocols BGP\)](#).]

- **BGP route leak prevention by using BGP roles and OTC attributes based on RFC 9234 (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5140-24CD8O, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, and QFX5250-64OE-DOT2L)**—You can prevent route leaks in BGP routing by using BGP roles and OTC attributes as defined in RFC 9234. The feature ensures that routes from providers or peers are only propagated to customers, reducing misconfigurations and errors. The BGP speaker automatically sets the OTC based on its configured role, and then advertises a prefix based on the OTC presence in the BGP update message, making the configuration straightforward and minimizing manual intervention. With this feature, you can maintain intended routing policies and prevent network delays and denial-of-service (DoS) attacks.

[See [BGP Route Leak Prevention and Detection](#), [otc-local-role](#), and [Supported Standards for BGP](#).]

- **Enhanced service route resolution for BGP multipath with list next hop (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD)**—This feature improves how service routes resolve over BGP multipath routes that use list-next hop structures. The resolver now tracks all contributing paths, not just the active one. When a service route resolves to BGP multipath helper route, RPD builds internal dependencies across each indirect next hop in the list. This ensures that your service route automatically re-resolves whenever any contributing path changes.

Previously, inactive paths could change without triggering a re-resolution. To fix this issue, RPD now links multipath routes to their full set of contributing routes using a patricia tree structure. The resolver can now detect changes across the entire path set and update service routes as needed.

You do not need to configure new CLI settings, but you must ensure that the BGP multipath list-nexthop command statement is enabled. You can use the updated `show route resolution list-nh` and `show krt indirect-next-hop` commands to inspect dependencies, contributing next hops, and the resulting forwarding decisions.

This enhancement improves resolution accuracy, supports re-evaluation during inactive path changes, and strengthens overall routing consistency in hybrid internal BGP (IBGP) and external BGP (EBGP) environments.

[See [Understanding BGP Path Selection](#).]

- **IGP-Metric-Based AIGP Path Selection for Flex- Algo Topologies (ACX Series and PTX Series)**—You can now use IGP metrics for AIGP path selection in SR-MPLS networks with flex-algorithm topologies, ensuring consistent low-latency routing across domains with different metric types. You can enable IGP metric computation for a flex-algorithm topology using the `compute-igp-metric` statement and assign multiple colors to a flex-algorithm with the `secondary-color` statement, both configured under `routing-options flex-algorithm <flex-algo-id>`. For finer control over route resolution, import policies can be managed within resolution schemes by using `import-policy-append` to add policies or `import-policy` to replace existing ones, configured under `routing-options resolution scheme <scheme-name>`. Additionally, policy statements now support the `actual-igp-cost` metric expression keyword, which allows precise adjustment of IGP metrics during route selection and is configured within `policy-options policy-statement <policy-name>`.

Serviceability

- **Command to show system alarm references (PTX10002-36QDD, PTX10008, and PTX10016)**—You can see a list of all active system-supported alarms on your router.

Use the `show system alarms reference operational-mode` command to see a list of the active system-supported alarms. Use the detail version of the command to see a list of remedies for the alarms in the list.

The output from the command includes:

- A list of active system-supported alarms for an individual FRU or all active system-supported alarms for your router
- A list of remedies if you include `detail` after `reference` in the command

[See [show system alarms reference](#).]

- **Sustained ingress Packet Forwarding Engine oversubscription log messages and alarms (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—The system will periodically monitor WAN ingress drop statistics, MAC pause frames, and flow control frames (depending on the flow control mechanism used), to report sustained WAN ingress Packet Forwarding Engine oversubscription. The system logs a major Cerror and raises a major alarm to indicate potential bursty traffic or unexpected packet drops. When the event is no longer present, the system clears the error and alarm and indicates this state in the INFO log. Use the `show chassis alarms operational mode` command to see the alarm.

Services Applications

- **New Information Element IDs 52 and 53 in some of the IPFIX and Version 9 data templates for inline active flow monitoring (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—The IPv4, IPv6, MPLS-IPv4, and MPLS-IPv6 templates within the IPFIX and Version 9 data templates now contain two elements that hold the minimum TTL value (ID 52) and the maximum TTL value (ID 53) for the sampled packet. Because the system does not maintain a flow cache, the minimum and maximum TTL values are the same. For the ingress direction of sampling, the TTL value received on the wire is the value reported in the export records. For the egress direction of sampling, the TTL value transmitted in the forwarded packet is the value reported in the export records.

[See [Understand Inline Active Flow Monitoring](#).]

Software Installation and Upgrade

- **Separate firmware installation packages for independent firmware updates (PTX10001-36MR, PTX10002-36QDD, and PTX10003)**—Use standalone firmware installation packages to manage firmware updates separately from full software upgrades. These packages address bug fixes, security updates, and hardware vendor updates. The packages are version-independent, so you can upgrade only the needed firmware without needing to upgrade all of the software on a system.

[See [Upgrade Firmware on Junos OS Evolved Devices](#), [Junos OS Evolved Installation Packages](#), [request system software add \(Junos OS Evolved\)](#), [request system firmware upgrade \(Junos OS Evolved\)](#), and [show system firmware \(Junos OS Evolved\)](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **SRv6 unreachable prefix announcement (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**— To ensure scalability and prevent overwhelming all nodes with every prefix, route summarization at area border routers (ABRs) conceal local domain details. Segment Routing for IPv6 (SRv6) further streamlines route summarization and condenses locators from remote domains and disseminates them into the core network, which can obscure local domain activities.

A provider edge router does not immediately detect the loss of reachability when a remote edge device becomes unreachable, resulting in a traffic drop until BGP sends a status update. The ABR assigns a maximum metric to prefixes from unreachable devices, ensuring they leak across domains as Unreachable Prefix Advertisements (UPAs).

To enable the UPA, include the `prefix-unreachable` statement at the `[edit protocols isis]` hierarchy level.

[See [prefix-unreachable](#).]

- **BGP-LS advertisements of PCE delegated and initiated SRv6 TE tunnels (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and PTX10002-36QDD)**—Report static Segment Routing for IPv6–Traffic Engineering (SRv6-TE) tunnels with static segment list with micro SID (uSID) configuration to Path Computation Element (PCE). When the PCE controller provisions an SRv6-TE tunnel with uSIDs, BGP-LS advertises the SRv6-TE tunnel with its uSID segment list. This feature supports the SID Structure TLV 1252 and the SRv6 endpoint behavior TLV 1250, which are now available in the PCE report. When the externally controlled and routed SRv6-TE receives a PCUpdate message with uSIDs from the controller, BGP-LS advertises the endpoint behavior of the uSIDs.

[See [Enable Segment Routing for the Path Computation Element Protocol](#) and [SRv6-TE Tunnels with micro-SIDs in PCEP](#) .]

Additional Features

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

- **NIST purge method for media sanitization (ACX7024, ACX7024X, ACX7100-32C, ACX7332, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10008, QFX5700, and QFX5700E)**. We've extended support for NIST media sanitization for SATA hard disk drives to include:
 - Cryptographic scramble and block erase priorities for the purge method.
 - Enhanced secure erase priority for the clear method.

[See [NIST Special Publication 800-88, Guidelines for Media Sanitization](#) and [request system zeroize](#).]

- **Support for EVPN-VPWS with transport class tunnels** (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)

[See [Configuring EVPN over Transport Class Tunnels](#) and [Example: Configuring EVPN-VPWS over Transport Class Tunnels](#).]

- **Support for EVPN-VXLAN Type 2 and Type 5 stitching** (PTX10002-36QDD)

[See [interconnect](#).]

- **Support for VPLS** (PTX10008). You can now configure VPLS on the PTX10008 router running Junos OS Evolved.

Keep in mind the following when you are configuring the PTX10008 router.

- To configure VPLS, configure the instance-type virtual-switch statement at the [edit routing-instances *routing-instance-name*] hierarchy level.
- You must enable control-word at the [edit routing-instances *routing-instance-name* protocols vpls] hierarchy level.
- We support single bridge domains in this release. You must configure the service-type single statement at the [edit routing-instances *routing-instance-name* vpls] hierarchy level.
- CE interfaces do not support encapsulation of ethernet-vpls and vlan-vpls.
- To display VPLS MAC address information, use the show ethernet-switching table command.
- On the PTX10008 router with the JNP10K-LC1301 line card installed, enable interoperability support for VPLS by using the set chassis interoperability express5-enhanced command.

We support the following VPLS features:

- Multihoming

[See [VPLS Multihoming Overview](#).]

- Hierarchical VPLS (H-VPLS)

[See [Example: Configuring H-VPLS With VLANs](#).]

- Flow Aware Transport (FAT) labels.

[See [FAT Flow Labels Overview](#).]

- Entropy labels

[See [entropy-label](#).]

- Flood policer
[See [Configuring Firewall Filters and Policers for VPLS](#).]
- MAC address limits and actions
[See [mac-move-limit](#).]
- CFM support for VPLS topology
[See [Configure a MEP to Generate and Respond to CFM Protocol Messages](#).]
- **Support for internet-options commands (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR|PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5220).** We have added additional traffic management command options to the `[set system internet-options]` hierarchy:
 - `no-tcp-reset drop-all-tcp`
 - `no-tcp-reset drop-tcp-with-syn-only`
 - `ipv6-path-mtu-discovery-timeout minutes`
 [See [internet-options](#).]
- **Two-Way Active Measurement Protocol (TWAMP) monitoring service (RFC5357) hardware timestamp support (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**
[See the `offload-type inline-timestamping` option of the [test-session](#) statement.]
- **Specify the install package name as a URL in the `request system software add` command** (ACX Series, PTX Series, and QFX Series)
[See [request system software add \(Junos OS Evolved\)](#).]
- **Support for displaying load balancing decision result for Layer 3 (L3) unicast traffic (PTX10002-36QDD)**
[See [show forwarding-options load-balance](#).]
- **Support for file-system encryption with Trusted Platform Module (TPM 2.0)** (ACX7100-32C, ACX7100-48L, and PTX10002-36QDD)
[See [Encryption with Trusted Platform Module](#).]
- **Support for performance monitoring and TCA (PTX10002-36QDD).** We support performance monitoring for the QDD-400G-ZR 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 optical transport link efficiently.

[See [show interfaces transport pm](#).]

- **Support for displaying load balancing decision result for L3 unicast traffic (PTX10002-36QDD)**

[See [show forwarding-options load-balance](#).]

- **Support for firmware upgrade on optics (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**

[See [request system firmware upgrade \(Junos OS Evolved\)](#), [show system firmware \(Junos OS Evolved\)](#), [Upgrading Firmware on the CMIS Optics Module](#), and [request system firmware switch optics \(Junos OS Evolved\)](#).]

- **Enhanced Address Detection for Reliable Connectivity (PTX Series)**—We've improved our network address detection process to deliver more reliable connectivity and uninterrupted performance. This update prevents disruptions caused by duplicate address detection (DAD) failures under rare network conditions. By integrating advanced algorithms and unique identifiers, we reduce false detections and ensure smooth data flow, keeping your network running seamlessly.

What's Changed

IN THIS SECTION

- [General Routing | 73](#)
- [Class of Service \(CoS\) | 74](#)
- [EVPN | 74](#)
- [Interfaces and Chassis | 74](#)
- [Junos XML API and Scripting | 74](#)
- [Routing Protocols | 75](#)
- [System Management | 75](#)
- [User Interface and Configuration | 75](#)

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

General Routing

- Support added for source and destination port optimization for port ranges for ipv6 input firewall filters.
- Control Board offline delay for system stability (PTX10008)—After initiating a node halt, you must wait 1 minute before doing Control Board (CB) offline. Attempting to offline the CB within this period will result in an error message. This delay helps maintain the stability and proper functioning of the system.

[See [request chassis cb.](#)]

- **Log file location during system downgrade (PTX10002-36QDD)** — When performing a system downgrade, if the downgrade fails, the validation failure log message now points to the log file location at /var/log/validation_config.log.
- The show subscribers extensive client-type dhcp | display xml validate command has now been updated to display correct output instead of **Duplicate data element** error.
- Starting in Junos OS Evolved Release 24.4R2, you can set the set protocols ptp delay-comp interface ifd-name transmit tx-comp-in-nsec receive rx-comp-in-nsec command to compensate for external latencies for a given physical interface (IFD) on PTX10008 routers. You can use it to compensate for transceiver latencies, which are significantly higher in PAM4 optics, and other delays such as optics latency and cable asymmetry. The range of latency that can be compensated is from -1000ns to +1000ns. See "ptp", ?delay-comp?, and "Precision Time Protocol Delay Compensation"
- **Change in CLI output (PTX Series)**—The CLI output for show system license bandwidth, show system license bandwidth fpc, and show system license fpc commands is updated.

[See [Monitor Junos Licenses.](#)]

- **Deprecated license trace (Junos OS Evolved)**—We've deprecated the CLI option show system license liblicense-trace.
- **SMAC accounting mismatch (PTX10002-36QDD, and PTX10008)** — Source MAC (SMAC) accounting over accounts the byte counter by including the Layer 2 overhead in an IP packet. Both ingress and egress accounting for a SMAC learnt on any interface is affected. The packet accounting and the number of SMAC addresses learnt is correct.

[See [MAC address accounting for L3 interfaces and aggregated Ethernet interfaces.](#)]

Class of Service (CoS)

- In Junos OS Evolved, do not associate a default forwarding class name with a different queue number. Use a custom forwarding class name instead. Defining customer forwarding classes with factory default forwarding class names causes errors.

EVPN

- **Easy EVPN LAG (EZ-LAG) feature lightweight loop detection configuration change**—We have updated the configuration generated by the easy EZ-LAG feature's commit script to use logical interface names for the lightweight leaf to server loop detection feature. The EZ-LAG commit script previously generated the loop detection configuration with physical interface names, but the loop detection feature works only for logical interfaces. The lightweight loop detection configuration uses the `loop-detect` statement at the `edit protocols` hierarchy level.

[See [Easy EVPN LAG \(EZ-LAG\) Configuration](#).]

Interfaces and Chassis

- **Vlan Tagging** 1. For PTX EVO platforms, if you have configured an Interface Device (IFD) with the family ethernet switching vlan members configuration, you cannot use both VLAN tagging and flexible VLAN tagging CLI commands on the IFD at the same time. This configuration is not supported, and a warning is issued if you try to commit this configuration. 2. For EVO platforms, if you have configured any Logical Interface (IFL) on an Interface Device (IFD) with the family ethernet-switching configuration, you cannot configure any other families on a different IFL unless you configure the IFD with the flexible-ethernet-services encapsulation type. This configuration is not supported, and a warning is issued if you try to commit this configuration.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, PTX Series, and QFX Series)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first

configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [cert-file](#).]

Routing Protocols

- **Extension of traceoptions support for VLANs in IGMP/MLD snooping**— The `traceoptions` option is supported under the `edit routing-instance protocols igmp-snooping vlan` and `edit routing-instance protocols mld-snooping vlan` hierarchy. `traceoptions` can be enabled for both specific and all VLANs.

[See [vlan \(IGMP Snooping\)](#) .]

System Management

- **Core files generated for related processes are bundled into a compressed TAR file (ACX Series, PTX Series, and QFX Series)**—When a process generates a core file and the system generates a live core for its related process, the system bundles the core files into a compressed TAR file. In earlier releases, the system does not bundle the files.

The process pairs that support this feature are: `bfdd` and `bfddagent`; `cfmd` and `cfmd-agent`; `dot1xd` and `dot1xd-agent`; `l2ald` and `l2ald-agent`; `l2cpd` and `l2cpd-agent`; `mcsnoopd` and `mcsnoopd-agent`; `ppmd` and `ppmdagent`; `routing` and `rpdagent`.

User Interface and Configuration

- **Access privileges for request support information command (PTX Series Firewalls)**— The `request support information` command is designed to generate system information for troubleshooting and debugging purposes. Users with the specific access privileges `maintenance`, `view`, and `view-configuration` can execute `request support information` command.
- **Changes to the show system storage command output (ACX Series, PTX Series, and QFX Series)**—We've updated the `show system storage` command output to include only true (physical) storage and exclude any host/hypervisor level storage. In earlier releases, the output also includes a container/jail storage which does not have a separate storage of its own.

[See [show system storage](#).]

- **Option to view combined disk space usage statistics for all configuration databases (ACX Series, PTX Series, and QFX Series)**—The `show system configuration database usage` command provides the `merge` option. When you include the `merge` option, the command output displays combined disk space usage statistics for all configuration databases, including the static configuration database and all ephemeral configuration database instances.

[See [show system configuration database usage](#).]

- **Changes to `show system statistics` command output (ACX Series, PTX Series, and QFX Series)**—We've updated the output fields for the following commands:
 - `show system statistics mpls`—We've added a new field: Packets with IPv6 explicit NULL tag.
 - `show system statistics jtd`—We've renamed two fields to better reflect the underlying data representation. The MPLS IPv4 packet received field is renamed to MPLS IPv4 explicit NULL packets received, and the MPLS IPv6 packets received field is renamed to MPLS IPv6 explicit NULL packets received.

Known Limitations

There are no known limitations in hardware or software 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.

Open Issues

IN THIS SECTION

- [General Routing | 77](#)
- [Interfaces and Chassis | 78](#)
- [Routing Policy and Firewall Filters | 78](#)

Learn about open issues 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

- When DHCP trace options are enabled, there is a possibility that jdhcpd could core. In general, traceoptions should be enabled only for debugging. They should be disabled once debuffing is done. [PR1771121](#)
- On PTX10004, PTX10008, and PTX10016 Junos OS Evolved platforms, !Minor alarm LED on FPM(Front Panel Module) is glowing Yellow, although there is no active alarms/errors shown by show system alarms CLI. [PR1782498](#)
- The output display of show ptp global-information has been updated to align the columns and to include only those items that are relevant to the configured profile. [PR1791957](#)
- On Junos Evolved PTX platforms, the log message "CFMMAN: Parse Error: CFM CCM pkt TLV parsing error" might appear when an FPC or PFE is restarted and comes online. This message indicates that a CFM Continuity Check Message (CCM) was truncated, leading to an error while parsing the TLV (Type-Length-Value) fields in the packet. The error is transient and typically affects only a small number (usually 1 or 2) of CCM packets received immediately after the FPC or PFE becomes operational. It has no functional impact on CFM sessions or overall network operations. CFM continues to function normally, and no persistent degradation is observed. This issue does not consistently affect any specific CFM session and occurs only during the brief initialisation period of the hardware component. [PR1810549](#)
- On Junos OS Evolved based platforms, while using ping, traceroute, or other utility that requires host name resolution, an error is raised indicating hostname resolution has failed. [PR1822994](#)
- This is a transient log which sometimes is seen when LSI is recreated. This has no functional impact. It's generated because of additional dependency of LSI with BD. System takes care of cleaning the token for which this error is generated. This can be confirmed through VTY command show sandbox token. [PR1834443](#)
- DNS resolution for traceroute does not work for a router using the mgmt_vrf with 23.4R2-EVO. [PR1858650](#)
- On all Junos OS Evolved based platforms, when the RE (Routing Engine) node experiences switchover, offline/online transitions, or rebooting, a 'Sysman.re' crash file might appear in rare cases and could cause traffic impact. [PR1859095](#)
- On Junos Evolved PTX, AE (Aggregated Ethernet) re-anchoring is not working correctly when CFM (Connectivity Fault Management) is configured on untagged AE interfaces. Deleting an anchor child interface results in the associated CFM session going down. This can cause loss of CFM session continuity during AE re-anchoring operations involving untagged interfaces. [PR1864491](#)

- **INLINEKA:CFM-SLM-RESPONDER** No entry error might be seen with CFM configuration having more than one remote peer MEP in inline mode. There is no functional impact due to this error. [PR1865524](#)
- On Junos Evolved PTX platforms, a resource allocation error occurs when configuring inline CFM sessions at scales beyond the supported limits. This might cause session setup failures or unexpected behavior. [PR1865698](#)
- If the evo-aftmand-bt process restarts, the connection with the timingd process will get disconnected. The issue is that this does not get re-established resulting in the phy synchronization logic to fail in timingd after a few retries from its end. This might result in time error as the system is then no longer keeping the timestamp in the PHY/MAC in sync with the PTP FPGA. [PR1878029](#)

Interfaces and Chassis

- On PTX10003 systems, it is not allowed to configure ZR optics (400G) through CLI. [PR1851078](#)

Routing Policy and Firewall Filters

- On all Junos OS Evolved platforms, the Open Shortest Path First (OSPF) neighbourship configured with Internet Protocol Security (IPsec) authentication will go down during the Routing Engine (RE) switchover. [PR1807830](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 79](#)
- [Flow-based and Packet-based Processing | 81](#)
- [Forwarding and Sampling | 81](#)
- [Interfaces and Chassis | 81](#)
- [Network Management and Monitoring | 82](#)
- [Platform and Infrastructure | 82](#)

- [Routing Policy and Firewall Filters | 82](#)
- [Routing Protocols | 82](#)
- [VPNs | 82](#)

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.

General Routing

- The cfmman process crashes with high scaled CFM configuration. [PR1780805](#)
- USB Media installation shows up minor alarm **Host 0 Voltage Threshold Crossed**. [PR1799443](#)
- The Netconf output for xml format shows ping-success tag under ping failure situation. [PR1835388](#)
- **bt_mtip_chpcs_hw_clear_stats: failed** error will be seen post rebooting the device. [PR1838394](#)
- Traffic is still forwarded after disabling ingress logical interface. [PR1843667](#)
- Traffic blackhole is observed for IPv4 /32 LDP prefixes advertised over BGP-LU when BGP sharding is configured. [PR1845425](#)
- Retimer/gearbox ports on LC1202 line card might take long duration to bring an interface back up. [PR1846379](#)
- Traffic going over LDP or LDPoRSVP route matching the CBF is getting impacted when PTX Junos OS Evolved platform working as PHP router with CBF configuration. [PR1847169](#)
- Traffic loss observed over load-balanced ESI LAG towards Provide Edge router in an EVPN-MPLS path on Junos OS Evolved PTX Series platforms. [PR1849188](#)
- Extend the interfaces optics-options configuration of DP init max duration for 400G ZR-M-HP optical transceiver. [PR1849814](#)
- FLT alpha algorithm gracefully rollback support. [PR1853326](#)
- MACSec statistics will not show up in the show interfaces *intf-name* statistics detail command output. [PR1853676](#)

- Router flag is not getting set in Neighbor Advertisement message. [PR1853868](#)
- [DCF14] [EVPN-VxLAN] After running a series of negative trigger tests on scale setup for 16 hours, traffic loss is seen in one IPv6 stream on PTX10001-36mr in CRB Spine and Border Router Role. [PR1854283](#)
- The rpd gets struck with 100% CPU usage after enabling BGP RIB-Sharding. [PR1854481](#)
- The evo-aftmand process crashes on PTX Junos OS Evolved platforms. [PR1855307](#)
- Chassis synchronization clock enters holdover state upon secondary BITS input failure. [PR1855958](#)
- Missing traffic statistics through telemetry on PTX Junos OS Evolved platforms with MPC10E, MPC11E, and LC9600 line cards. [PR1856528](#)
- Extra link flap will be seen when the link fails for a short time on PTX10000 with LC1202. [PR1857490](#)
- Route change is not synced after rpd restart due to rib-fib inconsistency. [PR1858750](#)
- The IPv6 firewall filter applied on loopback 0 is not processing NS and NA packets. [PR1859044](#)
- EVPN next-hop installation for ESI over VTEP fails after remote device restart on PTX platforms. [PR1859302](#)
- PFH Interface Down After PFE Offline and Re-anchor. [PR1859387](#)
- SNMP walk failure due to missing OID after FCO/TPA delete compression. [PR1859621](#)
- The rpd crash due to overlapping flow route updates in a single transaction. [PR1860888](#)
- PPPoE over EVPN-MPLS tunnel does not work. [PR1861208](#)
- Control plane is not properly managing the label 1048575 for MPLS traffic, leading to service impact. [PR1861803](#)
- NPU's (Network Processing Unit) KHT (Kernal Hash Table) utilization shows 100% utilization post configuration commit. [PR1862048](#)
- Traffic drops will be observed when multiple interfaces are part of the same ASIC and configured with the same MAC address. [PR1863060](#)
- Wedge detected after multiple, rapid PFE restarts. [PR1866487](#)
- IS-IS adjacency between CE devices is not established over Layer 2 circuit. [PR1867003](#)
- Applying a Layer2 filter on interface on results in protocol flaps and traffic disruption on other interfaces on the same PFE. [PR1869773](#)

- PTX10002-36QDD showing incorrect power status. [PR1870153](#)
- Route installation for MPLS labels fails when an RSVP LSP goes down. [PR1874004](#)
- Rare assert with an Junos OS Evolved process. [PR1874151](#)
- Transient traffic loss during multicast route convergence scenarios. [PR1876781](#)
- Log messages getting disappeared after upgrade. [PR1878365](#)
- Display issue is observed for command output show forwarding-options port-mirroring self-mirror. [PR1878786](#)
- MACsec traffic impact due to incorrect port mapping. [PR1879375](#)
- Rare evo-cda-bt core in PTX10008 Junos OS Evolved with multiple FPCs, failure of one FPC triggers get-state, causing rapid offline of its PFEs. [PR1879439](#)

Flow-based and Packet-based Processing

- PTX10000 Junos OS Evolved: JFLOW exported flow data InputInt is set to zero when redirecting to routing-instance and redirecting back to default instance. [PR1858721](#)

Forwarding and Sampling

- The update of the parcel_dump script to support the new Aegon linecard. [PR1851396](#)
- JNP10K-LC1301 running on ASAN image will go into fault state after restart. [PR1853920](#)

Interfaces and Chassis

- Mixing EP and SP style configurations on AE bundle causes traffic loss and protocol session failure on Junos OS Evolved PTX Series platforms. [PR1856651](#)
- LACP session destroy time is not getting updated in PFE. [PR1859633](#)
- Transit multicast traffic is getting blackhole due to peer interface flap. [PR1867231](#)
- Coredump during LAG disable. [PR1867603](#)

- Disabling a single AE IFL in a MAC-VRF can bring the entire AE interface down. [PR1875917](#)
- Interface remains administratively down state, fails to recover after execution of CLI command `request interface bounce`. [PR1880635](#)

Network Management and Monitoring

- The eventd memory leak on Syslog over TLS with unconfigured PKI certificate [PR1845058](#)
- [PTX10000 Junos OS Evolved] Cold start SNMP trap is unexpectedly sent when doing primaryship switchover from RE1 to RE0, or backup RE reboot followed by primaryship switchover. [PR1875093](#)

Platform and Infrastructure

- Process `evo-aftmand-bt` crashes on PTX10000 Junos OS Evolved platforms when command `request chassis fpc slot pfe-instance restart` is executed. [PR1844523](#)

Routing Policy and Firewall Filters

- The `firewalld` process crashes post specific set of filter related modifications performed in single commit. [PR1875725](#)

Routing Protocols

- The changes from instance-type `DEFAULT_INSTANCE` to others and vice versa will not be allowed. [PR1663776](#)
- The `advertise-inactive` configuration does not work as expected when `add-path multipath` is configured and negotiated with the neighbor. [PR1861799](#)

VPNs

- The MVPN traffic forwarding is affected when BGP PIC is enabled. [PR1861726](#)

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

- [What's New | 83](#)
- [What's Changed | 91](#)
- [Known Limitations | 95](#)
- [Open Issues | 96](#)
- [Resolved Issues | 98](#)

These release notes accompany Junos OS Evolved Release 25.2R1 for QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-64OD, QFX5240-QD, QFX5700, and QFX5700E switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- [EVPN | 84](#)
- [Junos Telemetry Interface | 85](#)
- [Network Management and Monitoring | 86](#)
- [Multicast | 87](#)
- [OpenConfig | 87](#)
- [Precision Time Protocol \(PTP\) | 88](#)
- [Routing Policy and Firewall Filters | 88](#)
- [Routing Protocols | 89](#)
- [Serviceability | 90](#)
- [Additional Features | 90](#)

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

To view features supported on the QFX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 25.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [QFX5130-32CD](#)
- [QFX5130-48C](#)
- [QFX5130-48CM](#)
- [QFX5130E-32CD](#)
- [QFX5220-32CD](#)
- [QFX5220-128C](#)
- [QFX5230-64CD](#)
- [QFX5240-64OD](#)
- [QFX5240-QD](#)
- [QFX5700](#)
- [QFX5700E](#)

EVPN

- **Exception policy for enhanced OISM to avoid multicast traffic loss on packets with TTL=1 (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)**—Enhanced optimized intersubnet multicast (OISM) routes most multicast traffic on the OISM supplemental bridge domain (SBD) rather than on the source VLAN, even if the destination OISM device hosts the source VLAN. This extra routing decrements a packet's time-to-live (TTL) more than once, so packets with TTL=1 don't reach the receivers. To avoid this problem on enhanced OISM devices, use the following steps to configure the devices to use the source VLAN instead of the SBD to forward multicast data to remote receivers:
 1. Configure a routing policy *policy-name* at the [edit policy-options policy-statement] hierarchy level to match the multicast groups (or sources and groups) for which to forward multicast traffic on the source VLAN.
 2. Set the forward-policy *policy-name* option at the [edit routing-instances *VRF-instance-name* protocols evpn oism enhanced forward-on-source-bridge-domain] hierarchy level to enable forwarding on the source VLAN instead of on the SBD for the multicast groups (or sources and groups) that match the policy.

You can configure and apply multiple policies with the forward-policy option.

[See [forward-on-source-bridge-domain](#) and [Enhanced OISM Exception Policy to Forward on Source VLAN Instead of SBD for Packets with TTL=1.](#)]

- **Optimized EVPN-VXLAN DCI multicast with enhanced OISM and an IPv6 underlay (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—**

You can configure enhanced optimized intersubnet multicast (OISM) and seamless Data Center Interconnect (DCI) with EVPN-VXLAN instances on an IPv6 underlay. In EVPN-VXLAN DCI fabrics with enhanced OISM and an IPv6 underlay, DCI gateway (iGW) devices send EVPN Type 6 Selective Multicast Ethernet Tag (SMET) routes to remote iGW devices when hosts subscribe to multicast groups. iGW devices in the source data center selectively forward multicast traffic for a group across the DCI only if the remote data center has receivers subscribed to that group. Previously, the iGW devices always flooded multicast traffic across the interconnection even when the remote data center had no subscribed receivers.

[See [EVPN-VXLAN DCI Multicast with Enhanced OISM.](#)]

- **Support for excluding MAC addresses from duplicate MAC detection (ACX7100-32C, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—**You can configure an exclusion list for MAC addresses in EVPN networks to prevent legitimate MAC address movements from being marked as duplicates. Use `set protocols evpn mac-list list_name mac-address mac_address_with_prefix_len` to create the list and `set protocols evpn duplicate-mac-detection exclude-list list_name` to apply it. This feature helps maintain network stability by avoiding unnecessary duplicate MAC detection for specified addresses, particularly in scenarios involving virtual MAC configurations in redundant setups.

[See [EVPN Duplicate MAC Detection Exclusion Lists.](#)]

Junos Telemetry Interface

- **Support for tail-drop sensor in OpenConfig model (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—**This feature supports the OpenConfig model `openconfig-if-ethernet.yang` (physical interface level) version 2.6.2 (no configuration). You can now monitor tail-drop metrics using the OpenConfig model, enhancing your network management capabilities. The sensor supports tracking tail-drop packets and bytes for both ingress and egress queues. This integration enables you to use standardized data modeling for better interoperability and streamlined network performance analysis.

[See [Junos YANG Data Model Explorer.](#)]

- **Enhanced telemetry with multiple gRPC servers and multi-port gRPC services (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5700, and QFX5700E)—**You can configure multiple RPC developed by Google (gRPC) servers with distinct

services, listening addresses, and ports by using the Junos telemetry interface (JTI). This feature enhances control over service management and telemetry data collection. You can also configure TLS certificates for secure communications. For example, you can configure a server to listen on a specific port and serve only designated gRPC services, enhancing flexibility and security in your telemetry setup.

- **Support for tail-drop sensor in OpenConfig mode (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—This feature supports OpenConfig model `openconfig-if-ethernet.yang` (physical interface level) version 2.6.2 (no configuration). You can now monitor tail-drop metrics using the OpenConfig model, enhancing your network management capabilities. The sensor supports tracking tail-drop packets and bytes for both ingress and egress queues. This integration allows you to utilize standardized data modeling for better interoperability and streamlined network performance analysis.

[See [Junos YANG Data Model Explorer](#).]

- **Native YANG state model and telemetry support for network stack protocol statistics (ACX Series, QFX Series, and PTX Series)**—You can use a native YANG state model and telemetry to monitor network stack protocol statistics on EVO platforms. Telemetry provides real-time data streaming for protocols such as TTP, ICMP, MPLS, TCP, and more. These statistics, previously available through CLI commands, are now accessible through telemetry streaming, ensuring real-time updates. This feature provides a comprehensive and dynamic monitoring solution, configurable in either "on-change" or "periodic" mode, enhancing your network's observability and performance management.

[See [Junos YANG Data Model Explorer](#).]

- **Health monitoring sensors (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**— Junos telemetry interface (JTI) provides native sensors to monitor device infrastructure health. Device streams health statistics that external collectors use to track performance.

Use the resource path `/state/system/infrastructure/junos-evolved/` to view the health statistics.

These sensors stream details such as cluster data, distributor statistics, Distributed Data Store (DDS) client information, common resources, and indexes for Identity Management and Device Management .

[See [Junos YANG Data Model Explorer](#).]

Network Management and Monitoring

- **Support for multiple gRPC servers hosting different service sets (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR,**

PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5700, and QFX5700E)—You can configure multiple gRPC servers that host different sets of services on unique ports. Additionally, each server can support different certificates, listening addresses, and routing instances. You configure the gRPC servers at the [edit system services http servers] hierarchy level. Distributing gRPC services across different servers allows for better allocation of network resources, reducing the risk of port conflicts and optimizing server performance.

[See [Configure gRPC Services](#) and [server](#).]

- **Egress Inline sFlow with EP Recirculation and Flex Sampler (QFX5240-64OD and QFX5240-64QD)**—You can use the Egress Inline sFlow feature to achieve true egress mirroring through EP Recirculation using an Egress Flex Sampler. When a packet passes through the egress processing pipeline, a copy is sent to a loopback port for a second pass, ensuring all egress changes are mirrored accurately. This feature supports Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) environments, enhancing your network monitoring and diagnostics capabilities.

[See [sFlow Technology Overview](#) and [sFlow](#).]

- **TAP aggregation (QFX5130-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Test access point (TAP) aggregation, similar to port mirroring, is a network monitoring and troubleshooting tool. Unlike port mirroring, TAP aggregation provides many-to-many packet replication, enabling you to capture different types of data in real time so that you quickly see what is happening in your network. You configure the TAP aggregation feature at the [edit forwarding-options tap-aggregation] hierarchy level.

[See [TAP Aggregation for Network Monitoring](#).]

Multicast

- **Support for RFC 6395 (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)**—Junos OS Evolved supports the exchange of local interface-id and router-id PIM hello options on PIM-enabled routers by default. You can disable this behavior by configuring the disable-interface-id-tlv statement under the edit protocols pim hello-options hierarchy. You can also configure cluster-id as part of PIM hello options by configuring the cluster-id statement under the edit protocols pim hierarchy.

[See [cluster-id](#) and [disable-interface-id-tlv](#).]

OpenConfig

- **Support for tail-drop sensor in OpenConfig mode (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—This feature supports OpenConfig model openconfig-if-ethernet.yang (physical interface level) version 2.6.2 (no configuration). You can now monitor tail-drop metrics using the OpenConfig model, enhancing your network management capabilities. The sensor supports tracking tail-drop packets and bytes for both ingress and egress queues. This integration allows you to utilize

standardized data modeling for better interoperability and streamlined network performance analysis.

[See [Junos YANG Data Model Explorer](#).]

Precision Time Protocol (PTP)

- **Support for enterprise and media profiles with PTP ordinary clocks and boundary clocks (QFX5130-48CM)**—The enterprise and media profiles functionality with ordinary and boundary clocks is introduced for QFX5130-48CM devices, which includes support for:
 - Enterprise profile using PTP ordinary clock and PTP boundary clock applications.
 - Enterprise profile using PTP over IPv4 multicast transport.
 - SMPTE, AES67, and AES67+SMPTE profiles.
 - Media profiles using PTP ordinary clock and PTP boundary clock applications.
 - Media profiles using PTP over IPv4 multicast transport.

The enterprise profile is tailored for enterprise networks, enabling precise time synchronization using both multicast and unicast communication. It supports scalable, dynamic clock discovery and selection, making it ideal for large, modern IT infrastructures.

The media profile is designed specifically for media and broadcast networks, enabling precise synchronization of audio, video, and data streams. It ensures low-latency, phase-aligned timing across devices to support seamless media production and playout.

[See [PTP Enterprise Profile](#) and [PTP Media Profile](#).]

- **Support for PTP over IRB (QFX5130-48CM)**— Precision Time Protocol over Integrated Routing and Bridging (PTP over IRB) enables multiple PTP streams to share a common local IP address while forwarding packets through Layer 2 switching. This setup is particularly useful for broadcast media applications, where synchronization accuracy is crucial.

PTP over IRB ensures precise timing for applications like SMPTE and AES67 and reduces the need for multiple physical interfaces. It supports multicast PTP over IP, allowing seamless integration into existing network infrastructures.

[See [PTP over IRB for Broadcast Profiles](#).]

Routing Policy and Firewall Filters

- **Automatic dropping for nonlocal packets (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7020-FLTR, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, , PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220,**

QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)—The device drops all packets that are not local to the Routing Engine, unless they are flagged as exception packets. This feature is automatically enabled on all supported platforms. The dropped packet count is available under the “non-local drops” counter in the `show system statistics` command.

[See [show system statistics](#).]

Routing Protocols

- **BGP multipath prioritization with configurable priority queues (PTX10001-36MR and QFX5130-48C)**—Use BGP multipath prioritization to prioritize critical routes in BGP multipath computations with queues of low, medium, and high priorities. This feature reduces multipath-calculation latency for operator-prioritized routes in overall route convergence during high load. Additionally, you can manage delayed forwarding information base (FIB) convergence that results from frequent routing changes.

[See [BGP Route Prioritization](#) and [multipath \(Protocols BGP\)](#).]

- **BGP route leak prevention by using BGP roles and OTC attributes based on RFC 9234 (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5140-24CD8O, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, and QFX5250-64OE-DOT2L)**—You can prevent route leaks in BGP routing by using BGP roles and OTC attributes as defined in RFC 9234. The feature ensures that routes from providers or peers are only propagated to customers, reducing misconfigurations and errors. The BGP speaker automatically sets the OTC based on its configured role, and then advertises a prefix based on the OTC presence in the BGP update message, making the configuration straightforward and minimizing manual intervention. With this feature, you can maintain intended routing policies and prevent network delays and denial-of-service (DoS) attacks.

[See [BGP Route Leak Prevention and Detection](#), [otc-local-role](#), and [Supported Standards for BGP](#).]

- **Enhanced service route resolution for BGP multipath with list next hop (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD)**—This feature improves how service routes resolve over BGP multipath routes that use list-next hop structures. The resolver now tracks all contributing paths, not just the active one. When a service route resolves to BGP multipath helper route, RPD builds internal dependencies across each indirect next hop in the list. This ensures that your service route automatically re-resolves whenever any contributing path changes.

Previously, inactive paths could change without triggering a re-resolution. To fix this issue, RPD now links multipath routes to their full set of contributing routes using a patricia tree structure. The resolver can now detect changes across the entire path set and update service routes as needed.

You do not need to configure new CLI settings, but you must ensure that the BGP multipath list-nexthop command statement is enabled. You can use the updated `show route resolution list-nh` and `show krt indirect-next-hop` commands to inspect dependencies, contributing next hops, and the resulting forwarding decisions.

This enhancement improves resolution accuracy, supports re-evaluation during inactive path changes, and strengthens overall routing consistency in hybrid internal BGP (IBGP) and external BGP (EBGP) environments.

[See [Understanding BGP Path Selection](#).]

Serviceability

- **Collecting and archiving system-state counters to aid in system debugging (QFX5240-64QD and QFX5240-64QD)**—We support a shell script that collects and archives system-state counters in log files. These log files provide a close record of the system state over time that can be useful in system debugging. The script starts running at switch bootup. You can use a CLI operational mode command to restart or stop the execution of the script. You can also edit some set-up values of the script execution, such as number of log files to be archived.

[See [Collect System-State Counters to Capture Information for System Debugging](#).]

Additional Features

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

- **Enhanced OISM in EVPN-VXLAN ERB overlay networks with an IPv6 underlay** (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)

[See [Enhanced OISM with an EVPN-VXLAN IPv6 Underlay Configuration](#), [EVPN-VXLAN with an IPv6 Underlay](#), and [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **NIST purge method for media sanitization** (ACX7024, ACX7024X, ACX7100-32C, ACX7332, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10008, QFX5700, and QFX5700E). We've extended support for NIST media sanitization for SATA hard disk drives to include:
 - Cryptographic scramble and block erase priorities for the purge method.
 - Enhanced secure erase priority for the clear method.

[See [NIST Special Publication 800-88, Guidelines for Media Sanitization](#) and [request system zeroize](#).]

- **Optimized EVPN-VXLAN DCI multicast support with enhanced OISM** (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)

[See [EVPN-VXLAN DCI Multicast with Enhanced OISM](#).]

- **Support for EVPN-VXLAN DCI stitching with an IPv6 underlay** (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)

[See [EVPN-VXLAN with an IPv6 Underlay](#).]

- **Support for Request Support Information (RSI) commands (QFX5130-32CD)**—We've introduced the following options at the [request support information] hierarchy level for QFX5130-32CD:

- cos
- l3-debug
- evpn-vxlan

[See [request support information](#).]

- **Support for internet-options commands** (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR|PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5220). We have added additional traffic management command options to the [set system internet-options] hierarchy:

- no-tcp-reset drop-all-tcp
- no-tcp-reset drop-tcp-with-syn-only
- ipv6-path-mtu-discovery-timeout *minutes*

[See [internet-options](#).]

- **Specify the install package name as a URL in the request system software add command** (ACX Series, PTX Series, and QFX Series)

[See [request system software add \(Junos OS Evolved\)](#).]

What's Changed

IN THIS SECTION

- General Routing | 92
- Class of Service (CoS) | 93
- EVPN | 93

- [Interfaces and Chassis | 94](#)
- [Junos XML API and Scripting | 94](#)
- [Routing Protocols | 94](#)
- [System Management | 94](#)
- [User Interface and Configuration | 95](#)

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

General Routing

- **Changes to default behavior under forwarding-table**— `ecmp-fast-reroute` and `indirect-next-hop-change-acknowledgements` are enabled by default under the `edit routing-options forwarding-table` hierarchy. You can verify these defaults by running the `show configuration routing-options forwarding-table` in operational mode.

[See [ecmp-fast-reroute](#) and [indirect-next-hop-change-acknowledgements](#).]

- **Changes to show system alarms command output (QFX5130 and QFX5220)**—When the current version of the firmware is less than the minimum supported version, you can now see alarms for this mismatch in the output of the command. These alarms were not shown previously. For example, when you have a firmware version mismatch, you should now see output similar to the following:

```
user@host> show system alarms
18 alarms currently active
Alarm time Class Description
2024-09-09 04:55:00 PDT Minor CHASSIS 0 BIOS ROM
minimum supported firmware version mismatch
2024-09-09 04:55:20 PDT Minor CHASSIS 0 Fan CPLD
minimum supported firmware version mismatch
2024-09-09 04:55:19 PDT Minor CHASSIS 0 Optics
CPLD minimum supported firmware version mismatch
```

- On QFX5230 and QFX5240 devices, ECN counters at the port level did not properly account for ECN marked packets due to local congestion. This led to port-level ECN counters showing different numbers than the sum of queue-level ECN counters. This has been fixed starting in Junos OS Evolved 25.1R1.

- Support added for source and destination port optimization for port ranges for IPv6 input firewall filters.
- The `show subscribers extensive client-type dhcp | display xml validate` command has now been updated to display correct output instead of **Duplicate data element** error.
- **SFP Optics LOS alarms**— SFP Optics don't support Tx laser disabled alarm, Tx loss of signal functionality alarm, and Rx loss of signal alarm as diagnostics output.

See [show interfaces diagnostics optics](#).

- **Change in CLI output**—The CLI output for `show system license bandwidth`, `show system license bandwidth fpc`, and `show system license fpc` commands is updated.

See [Monitor Junos Licenses](#).]

- **Deprecated license trace (Junos OS Evolved)**—We've deprecated the CLI option `show system license liblicense-trace`.

Class of Service (CoS)

- In Junos OS Evolved, do not associate a default forwarding class name with a different queue number. Use a custom forwarding class name instead. Defining customer forwarding classes with factory default forwarding class names causes errors.

EVPN

- **Easy EVPN LAG (EZ-LAG) feature lightweight loop detection configuration change**— We have updated the configuration generated by the easy EZ-LAG features commit script to use logical interface names for the lightweight leaf to server loop detection feature. The EZ-LAG commit script previously generated the loop detection configuration with physical interface names, but the loop detection feature works only for logical interfaces. The lightweight loop detection configuration uses the `loop-detect` statement at the `edit protocols` hierarchy level.

[See [Easy EVPN LAG \(EZ-LAG\) Configuration](#).]

- **Duplicate MAC detection timeout (QFX5000 Series switches)**—The default setting for `auto-recovery-time` is 5 minutes on these platforms only.

[See [duplicate-mac-detection](#).]

Interfaces and Chassis

- **FEC statistics display (QFX5700)**—The `show interfaces interface-name extensive` command displays the FEC statistics on the host side because of the PHY introduction. This CLI display change is applicable to all PHY platforms.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, PTX Series, and QFX Series)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [cert-file](#).]

Routing Protocols

- **Extension of traceoptions support for VLANs in IGMP/MLD snooping**— The `traceoptions` option is supported under the `edit routing-instance protocols igmp-snooping vlan` and `edit routing-instance protocols mld-snooping vlan` hierarchy. `traceoptions` can be enabled for both specific and all vlans.

See [vlan \(IGMP Snooping\)](#) and [vlan \(MLD Snooping\)](#).]

System Management

- **Core files generated for related processes are bundled into a compressed TAR file (ACX Series, PTX Series, and QFX Series)**—When a process generates a core file and the system generates a live core for its related process, the system bundles the core files into a compressed TAR file. In earlier releases, the system does not bundle the files.

The process pairs that support this feature are: `bfdd` and `bfddagent`; `cfmd` and `cfmd-agent`; `dot1xd` and `dot1xd-agent`; `l2ald` and `l2ald-agent`; `l2cpd` and `l2cpd-agent`; `mcsnoopd` and `mcsnoopd-agent`; `ppmd` and `ppmdagent`; `routing` and `rpddagent`.

User Interface and Configuration

- **Access privileges for request support information command (ACX Series, PTX Series, and QFX Series)**
— The request support information command is designed to generate system information for troubleshooting and debugging purposes. Users with the specific access privileges maintenance, view, and view-configuration can execute request support information command.
- **Changes to the show system storage command output (ACX Series, PTX Series, and QFX Series)**—We've updated the show system storage command output to include only true (physical) storage and exclude any host/hypervisor level storage. In earlier releases, the output also includes a container/jail storage which does not have a separate storage of its own.

[See [show system storage](#).]

- **Option to view combined disk space usage statistics for all configuration databases (ACX Series, PTX Series, and QFX Series)**—The show system configuration database usage command provides the merge option. When you include the merge option, the command output displays combined disk space usage statistics for all configuration databases, including the static configuration database and all ephemeral configuration database instances.

[See [show system configuration database usage](#).]

- **Changes to show system statistics command output (ACX Series, PTX Series, and QFX Series)**—We've updated the output fields for the following commands:
 - show system statistics mpls—We've added a new field: Packets with IPv6 explicit NULL tag.
 - show system statistics jtd—We've renamed two fields to better reflect the underlying data representation. The MPLS IPv4 packet received field is renamed to MPLS IPv4 explicit NULL packets received, and the MPLS IPv6 packets received field is renamed to MPLS IPv6 explicit NULL packets received.

Known Limitations

IN THIS SECTION

- [General Routing | 96](#)
- [Interfaces and Chassis | 96](#)

Learn about limitations in this release for the 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

- On Junos Evolved TD4 platform (QFX5130-48CM), intermittent link flaps causing the hardware link scan to malfunction, resulting in incorrect high values in the Correction Field (CF) of Precision Time Protocol (PTP) packets, causing improper timestamping. This impacts system time inaccuracies across the network and the overall timing synchronization of the network until remediated.[PR1829734](#)
- On the QFX5230/QFX5240 platforms, host injected control protocol frames are accounted under the egress port's best-effort queue (Q0 by default) even though they are sent on the network control queue. this is only a display issue of the show command output and does not affect the actual prioritization of traffic.[PR1835046](#)

Interfaces and Chassis

- This is unsupported configuration for channelized interfaces.[PR1858941](#)

Open Issues

IN THIS SECTION

- [General Routing | 97](#)
- [Interfaces and Chassis | 97](#)

Learn about open issues 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

- QFX5700 MacSec: Minor packet drops (0.0000001%) observed when MacSec is enabled.[PR1816407](#)
- USB disks with Junos OS Evolved images from Junos OS Evolved 23.4R2 onwards might not be detectable by Windows. They still have valid images, and can be used for Junos OS Evolved installs. The only issue is that new images cannot be installed on these USB disks because Windows no longer recognizes these USB drives.[PR1819846](#)
- On Junos Evolved platform QFX5230 and QFX5240, any type of interface flap will cause the Priority-based Flow Control (PFC) generation stops from the switch under storm condition after neighbor link flap.[PR1827068](#)
- Egress port speed field in the IFA 2.0 Metadata stack is not supported on QFX5230 and it is always '0' irrespective of the egress port. speed.[PR1860574](#)
- On QFX5000 Junos OS Evolved platforms interface counter sensor names are modified as per OpenConfig YANG model details. https://gnats.juniper.net/web/default/1860876/attachments/interface_counter_s_cvbc_details.docx[PR1860876](#)
- On Junos Evolved OS platforms like QFX5000 (QFX5230/QFX5240), ungraceful swapping a 400G DAC with a 400G DR4 optic (or vice-versa) might lead to traffic loss.[PR1862711](#)
- Graceful JOJI of FPC is the recommended way instead of ungraceful.[PR1874712](#)
- There are only the logs seen related to getDieTemperature. There is no functionality effected and CPU is not effected. System is also stable. [PR1878950](#)
- On QFX5130 or QFX5700 platforms, in Ethernet VPN-Virtual Extensible LAN (EVPN-VxLAN) scenario, the Explicit Congestion Notification (ECN) Congestion Experienced (CE) bit is not be set for ECN-capable traffic when there is congestion. Since the CE bit is not set, packets are dropped. [PR1880166](#)

Interfaces and Chassis

- On QFX5230-64CD platform, 400G DAC cable of 2.5m length and 4X100G DAC BO might not link up with some peer devices. This issue is not seen with all peer devices. If this happens, please replace the 2.5m cable with a 1m DAC cable or use supported 400G Optics instead.[PR1747315](#)
- Link with the DAC (SFP56-50G-DAC-3M) comes up with 25G default speed configuration when switch is rebooted. Hence, When 50G speed config is applied , peer side sees the link flap.[PR1836697](#)

- Issue is not related to QFX5220 and not related to optics (QSFP56-DD-400GBASE-DR4) as mentioned in the PR synopsis and description. We encounter this issue is happening on the Marvel Line-side TX on the QFX5700. Input error and framing are small in number it should not have significant impact on traffic.[PR1848109](#)
- Junk values are shown for Carrier transition counter sometimes on initiating `clear interfaces statistics all` command and `picd` app restart. This issue is observed on QFX5220-32CD, QFX5230-64CD, QFX5240-64OD.[PR1864375](#)
- To ensure proper functionality of hitless restart/ISSU, do not mark any port as "unused".[PR1868377](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 98](#)
- [EVPN | 99](#)

Learn about the issues fixed in this release for QFX 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

- When a VXLAN encapsulated IP packet, or an IP packet with UDP port matching the VXLAN UDP port, is received on a trunk interface that does not have VXLAN configuration it is dropped. The Issue is not seen with untagged L3 interfaces [PR1805922](#)
- DHCPv4/v6 packets may be dropped because DHCP packets are not routed to kernel after initial `jdhcpd` starts [PR1816246](#)
- In high-scale, high-load environments, the `l2ald` process may experience hangs during Apstra polling. [PR1828741](#)
- Packets are incorrectly directed to an inappropriate filter and are lost due to TCAM exhaustion. [PR1840632](#)

- Stale MAC-IP entries are not cleared in an EVPN-VXLAN scenario when encapsulate-inner-vlan or decapsulate-accept-inner-vlan or both knobs are present. [PR1844623](#)
- Packet Forwarding Engine restart and high CPU utilization for evo-pfemamd process in large-Scale VXLAN Fabrics. [PR1845230](#)
- QFX5240-64 OD/QD: PSM upgrade failure. [PR1848912](#)
- VXLAN packets with inner VLAN tag are dropped in transit device. [PR1849807](#)
- OSPFv3 authentication with IPSEC doesn't work on Junos OS Evolved QFX platforms. [PR1862778](#)
- Traffic forwarding fails to recover after manually removing and reinserting the Direct Attached Copper cable. [PR1862893](#)
- ECN copy support from VXLAN header to inner Header on Type-5 decap tunnels. [PR1865671](#)
- The command output for `show system snapshot` randomly returns **Error: Snapshot is in progress ... Please retry later!** [PR1868309](#)
- QFX5241-64OD: Media type not displayed correctly for OSFP DAC and DACBO cables. [PR1873970](#)
- Rare assert with an Junos OS Evolved process. [PR1874151](#)
- EVACL does not work on TH4 and TH5 when bound on access ports. [PR1876492](#)
- Optics description is displayed as UNKNOWN for 3PO optics on Junos OS Evolved platforms. [PR1882916](#)

EVPN

- Duplicate traffic will occur in an EVPN-VxLAN DCI configuration with translation VNI. [PR1853110](#)

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

- [Self-Help Online Tools and Resources | 102](#)
- [Creating a Service Request with JTAC | 103](#)

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net/>
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

Revision History

8 December 2025—Revision 9, Junos OS Evolved Release 25.2R1

23 October 2025—Revision 8, Junos OS Evolved Release 25.2R1

25 September 2025—Revision 7, Junos OS Evolved Release 25.2R1

14 August 2025—Revision 6, Junos OS Evolved Release 25.2R1

6 August 2025—Revision 5, Junos OS Evolved Release 25.2R1

5 August 2025—Revision 4, Junos OS Evolved Release 25.2R1

15 July 2025—Revision 3, Junos OS Evolved Release 25.2R1

10 July 2025—Revision 2, Junos OS Evolved Release 25.2R1

30 June 2025—Revision 1, Junos OS Evolved Release 25.2R1

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2025 Juniper Networks, Inc. All rights reserved.