Release Notes

Published
2025-12-22

Junos OS Evolved Release 25.4R1

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 25.4R1.

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

Table of Contents

Junos OS Evolved Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 25.4R1 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

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 25.4R1, click the Group by Release link. You can collapse and expand the list as needed.

- ACX7020
- ACX7024
- ACX7024X
- ACX7100-32C
- ACX7100-48L
- ACX7332
- ACX7348
- ACX7509

The following sections highlight the key features in this release.

Class of Service

- **Classification and rewrite support at SRv6 encapsulation, transit, and de-encapsulation nodes (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—You can classify and rewrite SRv6 traffic at encapsulation, transit, and de-encapsulation nodes to enforce CoS marking consistency for Layer 3 VPN (L3VPN) and EVPN traffic. Configure policies to map VLAN priority code point (PCP) and Differentiated Services code point (DSCP)/DSCP-IPv6 code points and

to rewrite VLAN tags and IPv6 headers. Policies support reduced and non-reduced modes, segment identifier (SID) depths, micro- or classic SID, dynamic tunnels, penultimate segment pop (PSP) or ultimate segment pop (USP), and Topology Independent Loop-Free Alternate (TI-LFA). You can enable DSCP propagation to copy access-side DSCP into the SRv6 IP header. DSCP propagation disables DSCP, DSCP-IPv6, and inet-precedence rewrite rules. You cannot classify traffic based on inner payload code points at de-encapsulation.

[See Understanding SRv6 Classification and Rewrite.]

- **CoS support on pseudowire subscriber logical tunnel interfaces (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—You can enforce QoS on pseudowire subscriber interfaces to shape, prioritize, and mark traffic for L2/L3 VPN services by leveraging logical tunnel-anchored pseudowire client logical interfaces. Apply classifiers and rewrite rules per logical interface, and use hierarchical schedulers and traffic-control profiles on pseudowire logical interfaces or logical interface sets. You can use six priorities and weighted random early detection (WRED). Traffic rates are determined by the recycle bandwidth; use the logical tunnel port shaper to control upstream traffic.

  [See CoS Support for Pseudowire Subscriber Logical Tunnel Interfaces.]

Device Security

- **IMA coverage update (ACX Series, PTX Series, and QFX Series)**—Integrity Measurement Architecture (IMA) coverage now includes the following additional file systems:
  - ISO9660
  - PROC
  - SYSFS
  - DEBUGFS
  - RAMFS
  - SECURITYFS
  - EFIVARFS
  - DEVPTS
  - BINFMTFS
  - SELINUX
  - CGROUP
  - NSFS

- TRACEFS

IMA now enforces signature verification for the `kexec` kernel and `initramfs` images. It also generates a nonrepudiable log for new key addition events to IMA keyrings. These enhancements strengthen runtime integrity protections against unauthorized changes to Junos OS Evolved.

[See File Security with IMA.]

EVPN

- **Convergence improvements for EVPN-VPWS (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We've introduced EVPN-VPWS convergence enhancements on the ACX Series and PTX Series routers. With this feature, the router now supports route acknowledgment for EVPN-VPWS routes, extending the existing mechanism currently used for BGP routes. The router advertises an EVPN-VPWS route to its neighbor only after it confirms that:

  - The route has been successfully installed in the forwarding table of the Packet Forwarding Engine
  - The route version in the routing information base (RIB) matches the version in theforwarding information base (FIB) .

  If there are any acknowledgment errors, the router withdraws the EVPN-VPWS routes from the Packet Forwarding Engine. This ensures that EVPN-VPWS routes are reliably installed in hardware and that the versions are synchronized across the RIB and the FIB.

  By default, required acknowledgment from the PFE prior to advertising the EVPN-VPWS route is automatically enabled. To disable the acknowledgment requirement, include the `evpn-vpws-ack-disable` statement at the `[edit routing-options forwarding-table]` hierarchy level. Disabling the acknowledgment requirement only affects new EVPN-VPWS instances and does not impact the network.

  [See forwarding-table.]

- **EVPN VPWS and FXC with BGP-LU PIC (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)** —Support is added for EVPN VPWS and Flexible Cross-Connect (FXC) services in VLAN-aware and VLAN-unaware modes. BGP-LU PIC enables fast convergence for single-homing, multihoming all-active, and multihoming single-active deployments. This feature is supported across RSVP-TE, LDP, ISIS FlexAlgo, and SR-MPLS transport modes, with or without Per-Next-Hop (PNH) optimization.

  NOTE: For VPWS and FXC, BGP-LU PIC with PNH is not supported when the transport path uses SR-TE. In addition, BGP-PIC edge is not applicable or qualified from the

control plane, and BGP-LU PIC convergence cannot be achieved without a PNH configuration.

[See Overview of VPWS with EVPN Signaling Mechanisms.]

- **NSR and HA for EVPN VPWS Using BGP-LU PIC (ACX7332, ACX7348, and ACX7509)**—Support is added for Nonstop Routing (NSR) and High Availability (HA) with EVPN VPWS, maintaining forwarding state and service continuity during Routing Engine switchover events.

  [See Overview of VPWS with EVPN Signaling Mechanisms.]

- **BGP-LU PIC with Entropy Label support (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)** —Support is added for entropy labels in BGP-LU to improve load balancing across equal-cost MPLS paths and optimize traffic distribution.

  [See Configuring Entropy Labels and entropy-label (protocols bgp)]

- **EVPN-VPWS service over SRv6 underlay with BGP-LU PIC (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—Configure single-active or all-active multi homed EVPN VPWS networks using segment routing over IPv6 (SRv6).

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

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

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

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

Forwarding Options

- **Custom profiles for hardware resources (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—Although using the `hw-db-profile` configuration you can statically reserve hardware resources, you cannot use this configuration for all hardware resources because there are resources that need to be managed independently. You can use the custom profile infrastructure instead to customize hardware resource allocation as per application requirements. Using the custom profile infrastructure enables you to create application-specific dynamic hardware resource allocations.

  [See Custom Profiles for Hardware Resources, counter-profiles, hw-profiles, and hw-profile.]

- **Support for traffic statistics information for IRB interface with custom profile configuration (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509)**—You enable

the display of integrated routing and bridging (IRB) statistics when you create a custom counter engine profile for IRB ingress and egress logical interfaces (IFL). To view the IRB statistics, issue the `run show interfaces IRB extensive` command.

[See Custom Profiles for Hardware Resources, counter-profiles, hw-profiles, and hw-profile.]

- **Support to assign counter resources for the policer-ingress application (ACX7332)**—Support has been extended to the policer-ingress application to be assigned counter engines to provide the following numbers of counter resources - 4,000, 8,000, and 16,000. This will enable higher scale for firewall filters and policers. For example, two counter engines of 16,000 counters for a total of 32,000 counters can be assigned to the policer-ingress application using the following configuration statement:

```
set system packet-forwarding-options custom-profiles counter-profiles policer_32k app policer-ingress counter-16k 2
```

[See Custom Profiles for Hardware Resources, counter-profiles, hw-profiles, and hw-profile.]

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, and QFX5241-64QD)**—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:

  - Slax processor enhancements including a new mode, additional options, and simplified argument parsing
  - New libslax extension library functions
  - Improved SLAX syntax options
  - New SLAX functions and enhancements to existing functions and statements
  - Hexadecimal numbers support in SLAX scripts

  This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

  [See libslax Distribution Overview.]

Junos Telemetry

- **Support for LSP traffic sensors (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X)**—Junos telemetry supports sensors to monitor LSP-RSVP and LDP-LSP traffic statistics. The following sensor paths are added to retrieve traffic information:

  Native sensor paths:

  - `/junos/services/label-switched-path/usage/`
  - `/junos/services/ldp/label-switched-path/ingress/usage/`
  - `/junos/services/ldp/label-switched-path/transit/usage/`

  Openconfig sensor path: `/network-instances/network-instance/mpls/lsps/constrained-path/tunnels/tunnel/`

  > **NOTE**:
  >
  > - Zero-value counters are not exported.
  > - Telemetry is not supported for Point-to-Multipoint (P2MP) and dynamically configured LSPs.
  >
  > Use the `'show agent sensors'` command to view all active sensors.
  >
  > For more information, see Junos YANG Data Model Explorer.

- **Stream telemetry data in gNMI-based message format over UDP (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, and QFX5240-64QD)**—Junos OS Evolved uses a dial-out mechanism to send telemetry data to a collector over UDP . The message format is defined in the `jnx_gnmi_over_udp.proto` file. Only `STREAM` mode with `SAMPLE` as subscription mode is supported. The message contains full key name and value pair information so the collector does not require data models for processing or consuming the telemetry data.

  [See No Link Title, No Link Title, No Link Title, No Link Title, No Link Title, and Junos YANG Data Model Explorer.]

Interfaces

- **Smart SFP transceivers to transport TDM line traffic (ACX7348)**— We support the following smart SFP transceivers on the ACX7348 routers:

  - DS3 smartSFP (SFP-GE-TDM-DS3)

- E1 smartSFP (SFP-GE-TDM-E1)
- T1 smartSFP (SFP-GE-TDM-T1)
- STM1 smart SFP (SFP-GE-TDM-STM1)
- STM4 smart SFP (SFP-GE-TDM-STM4)
- STM16 smart SFP+ (SFP-XGE-TDM-STM16)

Every pair of Smart transceivers implements a Time Division Multiplexing (TDM) Circuit Emulation Service where TDM lines are transported transparently across a packet-switched network (PSN).

You can use these transceivers to transport TDM lines traffic such as E1 or T1 or DS3 or STM-1/OC-3, STM-4/OC-12, or STM-16/OC-64 encapsulated into data packets across PSNs such as Ethernet, VLAN-based, or MPLS networks. At the receiver end, another Smart transceiver, paired with the first one and appropriately configured, re-assembles the packets into the original bit stream and delivers it on its TDM line interface.

[See tdm-options (Interfaces).]

- **Chromatic dispersion configuration support with media-id mapping (ACX7024)—**You can configure and monitor chromatic dispersion settings using CLI commands for enhanced optical management. Specify `cd-val` and `media-id` through the `set interfaces et-<> optics-options` command. View mapped CD range details using `show interface diagnostic optics chromatic-dispersion et-<>`. The system retains previous CD values if incorrect configurations are provided, ensuring network stability. This feature optimizes performance for 100G-ZR SFPs and is adaptable for future high-speed variants.

- **Chromatic Dispersion Compensation (ACX7024)**— This feature supports 100G-ZR SFPs. A new CLI `set interfaces et-<> optics-options cd cd-val <> media-id` is introduces to support configuration of cd-range mapped with media-id.

- **Support for resilient hashing and consistent hashing (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—

  Configure resilient hashing on ACX routers to minimize flow remapping across link aggregation on groups (LAGs) or equal cost paths. Resilient hashing works with the default static hashing algorithm. When you configure resilient hashing on LAGs, the configuration is applicable to a specific aggregated Ethernet interface (aex).

  [See Resilient Hashing on LAGs and ECMP groups.]

- **Channelized Smart SFP (ACX7024)**— You can transport and aggregate legacy T1/E1 services over Ethernet by provisioning channelized Smart SFPs as STM1 (63 E1) or OC3 (84 DS1). These TDM channels will be transported over the ethernet network.

  The transport network could be one of the below and can be transported over any user Ethernet network provisioned.

- Ethernet network with destination MAC address programmed on the smart SFPs
- VLAN tagging network
- VLAN tagging with destination MAC set
- MPLS tagging mode

[See tdm-options (Interfaces).]

Layer 2 VPN

- **Layer 2 VPN and Layer 2 Circuits over GRE / UDP with FTI Tunneling (ACX Series) —** Use Layer 2 virtual private network (L2VPN) and Layer 2 circuit (L2CKT) services over generic routing encapsulation (GRE) and user datagram protocol (UDP) tunnels that use FTI tunneling. This configuration extends your Layer 2 domain across flexible IP-based networks. You can use existing CLI commands to configure these services.

  [See Configuring Layer 2 Ethernet Services over GRE Tunnel Interfaces, Configuring Tunnel Interfaces on ACX Routers, and Configuring Flexible Tunnel Interfaces.]

- **Support for L2VPN and L2 circuit services over pseudowire interfaces (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7100-32C, ACX7332, ACX7348, and ACX7509)**—Use the PS interface to provision L2VPN and L2CKT services for unified Layer 2 connectivity across your infrastructure. This capability helps you extend customer segments, optimize resource utilization, and simplify service operations. Configure service instances and bind them to the PS interface, and specify encapsulation, VLAN membership, and CoS policies as required.

  [See Understanding Layer 2 VPNs, Layer 2 Circuit Overview, and MPLS Pseudowires Configuration.]

MACsec

- **Automatic adjustment of MTU for MACsec overhead (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Use this feature to automatically adjust the maximum transmission unit (MTU) for the Media Access Control Security (MACsec) overhead. Without this feature, you must adjust the interface MTU and the protocol MTU manually.

  Use this feature to ensure the interface or protocol MTU is adjusted properly to account for the MACsec overhead. This feature is disabled by default. To enable this feature, first enable MACsec. Then configure the `enable-auto-mtu-update` statement at the `[edit security macsec]` hierarchy level. This feature applies to physical interfaces, logical interfaces, and physical interfaces that are members of aggregated Ethernet interfaces.

  [See Media MTU and Protocol MTU.]

- **Support for a custom EAPoL EtherType to improve network tunneling of MACsec packets (ACX7100-32C, ACX7332, ACX7348, and ACX7509)**—MACsec uses Extensible Authentication Protocol over LAN (EAPoL) as a transport protocol to establish sessions. Some networks filter packets based on their EtherType value. By default, the EtherType for all EAPoL packets is 0x888e. To ensure the network tunnels the MACsec packets properly, you can set a custom EtherType for EAPoL packets.

  To configure an EAPoL profile with a custom EtherType, use the `ether-type` *`ether-type-value`* statement at the `[edit forwarding-options custom-eapol-ether-type-profiles` *`eapol-profile-name`*`]` hierarchy level. To apply the EtherType to MACsec packets, configure the `eapol-ethertype-profile` *`eapol-profile-name`* statement at the `[edit security macsec connectivity-association` *`ca-name`* `mka]` hierarchy level.

  [See Media Access Control Security (MACsec) over WAN, custom-eapol-ethertype-profiles, and mka.]

Network Management and Monitoring

- **AES-256 Encryption Algorithm Support for SNMPv3 (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can configure Advanced Encryption Standard (AES) 256 algorithm for an SNMPv3 user. To configure AES-256 algorithms for an SNMPv3 user, include the `privacy-aes256` statement at the `edit snmp v3 usm local-engine user username` hierarchy level. AES-256 is a symmetric encryption algorithm that uses a 256-bit key to encrypt or decrypt messages and provides high-level security for protecting sensitive information.

  [See Configure SNMPv3 Authentication Type and Encryption Type, show snmp v3, and usm.]

Network Time Protocol (NTP)

- **NTP time offset threshold support (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-36CD, PTX10002-60MR, 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, QFX5241-64OD, QFX5241-64QD, QFX5250-64OE, QFX5250-64OE-DOT2L, QFX5250-64OE-DAO, QFX5250-64OE-DO-T3, and QFX5250-64OE-L)**—Use the `ntp threshold` configuration to detect and handle NTP time discrepancies, mitigate spoofed NTP attacks, and preserve log accuracy. Use the CLI command `set system ntp threshold action accept|reject <value>` to configure a threshold and action to accept (and log), or to reject synchronization that exceeds the threshold value. You can also manually synchronize time with the NTP service if required.

  [See ntp threshold.]

Multicast

- **Configure multicast tunnel interfaces with bandwidth reservation (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)**—You can configure multicast tunnel interfaces on ACX7000 series routers and reserve bandwidth for tunnel services. Specify the bandwidth in Gbps using the tunnel-services multicast bandwidth statement at the `[edit chassis fpc| feb slot slot-number pfe number core number channel number]` hierarchy level for FPC or FEB-based systems. This setup allows you to efficiently manage multicast traffic through the multicast tunnel (MT) interface and support ROSEN and BGP-MVPN features.

  [See multicast (Tunnel Services), and Configure Multicast Tunnel Interfaces on ACX7000 Series Routers.]

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, PTX Series, and QFX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):

  - ML-DSA-87 PQC algorithm for digital signatures
  - SHA-512 for hashing

  For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

  [See PQC Signatures for Software Images.]

- **Support for Quantum Buffer in SSH (ACX Series, PTX Series, and QFX Series)**—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptoagility. The feature uses finite field cryptography (FFC) to extend the security life span of the current systems against quantum attacks. Quantum Buffer provides a phased approach to adopting post-quantum cryptography (PQC), thereby mitigating operational risks associated with the transition.

  To enable the feature, configure the following command:

  - `set system services ssh moduli type` *`name`* `refresh` *`frequency`* `count` *`count`*

    The configuration dynamically generates prime moduli for existing Diffie-Hellman (DH) group exchange algorithms, `group-exchange-sha1` and `group-exchange-sha2`. The `qbufd` process is responsible for generating the moduli.

  [See Quantum Buffer and moduli.]

- **Support for Shor-resistant and other default key exchange algorithms in SSH (ACX Series, PTX Series, and QFX Series)**—SSH supports the hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm, which is Shor-resistant and improves protection against quantum attacks.

  Configure `sntrup761x25519-sha512` at the [`edit system services ssh key-exchange`] hierarchy level.

  Additionally, SSH includes default support for the following Diffie-Hellman (DH) group key exchange algorithms that are available at the [`edit system services ssh key-exchange`] hierarchy level.

  - `dh-group16-sha512`
  - `dh-group18-sha512`

  [See key-exchange.]

Precision Time Protocol (PTP)

- **Support for frequency and phase offset relaxation (ACX7024 and ACX7024X)**—You can relax the frequency and phase offsets required for a Precision Time Protocol (PTP) lock. Configure the `frequency-lock-threshold` and `phase-adjust-threshold` options of the `ptp` statement, to relax these parameters. You can also relax the maximum phase offset to adjust in a phase-aligned state by using the `phase-adjust-threshold` statement in the PTP configuration.

  [See ptp.]

- **PTP passive and timeReceiver port monitoring with SNMP MIB support (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—You can enhance router performance monitoring through Precision Time Protocol (PTP) passive port monitoring, timeReceiver port monitoring, and SNMP MIB integration. This feature allows you to collect detailed metrics for active PTP timeReceiver ports using IEEE 1588-2019 standards, including timeTransmitter-timeReceiver delay and sync message transmission. Use `show ptp global-information` and `show ptp passive-port-monitor-status` commands to display performance statistics, ensuring efficient network management. Passive monitoring alerts you to potential fiber asymmetries or clock failures, improving network reliability by identifying issues when phase thresholds are exceeded.

  [See PTP Passive Port Performance Monitoring.]

- **Assisted Partial Timing Support G.8275.2 enhancements: configurable clock class threshold and support for multiple timeReceiver clocks (ACX7024 and ACX7024X)**—Deploy Assisted Partial Timing Support (APTS) to use Global Navigation Satellite System (GNSS) as the primary timing source with Precision Time Protocol (PTP) (G.8275.2) as backup. During GNSS outages, maintain phase and frequency using measured PTP delay asymmetry.

  With the newly introduced enhancements on APTS in this release, you can now configure the clock-class threshold. The system will consider the upstream PTP timeTransmitter as a backup source only if it's clock class is less than or equal to the set threshold value. You can also increase APTS backup capacity by enabling up to four PTP timeReceiver ports.

[See Assisted Partial Timing Support on ACX7024 and ACX7024X Routers.]

- **Assisted Full Timing Support with GNSS as primary and PTP and Synchronous Ethernet as backup sources (ACX7024, ACX7332, and ACX7348)**—Use Assisted Full Timing Support (AFTS) to enhance timing resilience by prioritizing Global Navigation Satellite System (GNSS) and use G.8275.1 (PTP and Synchronous Ethernet) as a backup source. This profile protects against spoofing, jamming, or atmospheric issues, helping you avoid costly holdover. During normal operation, source time from the local primary reference time clock (PRTC) or grandmaster sourced by GNSS. On GNSS failure, AFTS automatically falls back to PTP and Synchronous Ethernet. This AFTS configuration maintains phase and frequency without extended holdover and uses the physical layer clock if PTP is unavailable.

  [See Assisted Full Timing Support.]

Routing Protocols

- **IS-IS multi-instance support over a single interface (ACX7020, ACX7024, ACX7100, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, QFX5130, QFX5220, QFX5230-64CD, QFX5240, QFX5241-32OD, and QFX5700)**—We have enhaced the IS-IS multi-instance feature to support multiple IS-IS instances on the same logical interface with instance identifier TLV 7.

  Include the `instance-id` statement at the `[edit protocols isis-instance` *`name`* hierarchy level.

  [See How to Configure Multiple Independent IGP Instances of IS-IS.]

Services Applications

- **ITU-T Y.1564 Ethernet SLA validation** (ACX Series)—We support the ITU-T Y.1564 standard for turning up, installing, and troubleshooting Ethernet-based services. It is the only standard test methodology that allows for complete Ethernet service-level agreement (SLA) validation in a single test. You can configure tests and services and get a report about performance. Configure this feature using the `y1564` statement at the `[edit services monitoring]` hierarchy level. Use the following commands to:

  - Start a test: `test services monitoring y1564` *`test-name`* `start`
  - Check test results: `show services monitoring y1564`
  - Clear test result history: `clear services monitoring y1564`

- **Hardware-assisted inline active flow monitoring for IPv4 traffic with IPFIX and version 9 flow export (ACX7020, ACX7024, ACX7024X, ACX7100, ACX7332, ACX7348, and ACX7509)**—Monitor traffic behavior, utilization, and anomalies by sampling flows and exporting records to an IP Flow Information Export (IPFIX) or a version 9 collector using inline active flow monitoring. We now support hardware-assisted inline active flow monitoring for IPv4 traffic. Configure the `hw-assisted`

statement at the `[edit forwarding-options sampling instance instance-name family inet output inline-jflow]` hierarchy level to enable ASIC-based sampling. A separate data template contains the Information Elements (IEs) for hardware-assisted inline active flow monitoring.

[See Understand Inline Active Flow Monitoring.]

Software Installation and Upgrade

- **Load set-formatted and XML-based configuration files for ZTP (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5140-24CD8O, , QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, QFX5700, and QFX5700E)**—You can load set-formatted or XML-based configuration files when your device provisions for zero-touch provisioning (ZTP). Reuse existing set-style or XML-based configuration files for automated onboarding to avoid converting them to hierarchical syntax. Provide the configuration file in set format or in XML and specify the configuration file name under the DHCP vendor configuration options.

  [See Zero Touch Provisioning.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- [See profile-sharing.]

- **Support for UHP in IS-IS SR-MPLS (ACX7020, ACX7100, ACX7332, ACX7348, ACX7509,ACX7024, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, and PTX12008)** —Use Ultimate Hop Popping (UHP) with IS-IS or OSPF so the egress provider edge (PE) can process its own node SID. ISIS advertises a node SID with the P flag set and E flag unset. In controller-driven segment routing traffic engineering (SR-TE) the controller inserts the egress PE node SID beneath the SR-TE binding SID. If the Binding SID route fails on the penultimate hop, the egress PE might see its own node SID as the top label instead of penultimate hop popping (PHP). With the P flag set, the PE expects UHP and processes its MPLS label. Include the `ultimate-hop-popping` statement at the `[edit protocols isis source-packet-routing]` hierarchy level.

  [See ultimate-hop-popping.]

- **Delay normalization for OSPF Flexible Algorithm metrics and advertisements across IGP instances (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—Use delay normalization to compute and advertise a normalized delay metric for Flexible Algorithm, to improve path-selection consistency across all IGP instances. The device normalizes each received delay, compares each value with the

previously saved normalized value, and triggers link-state advertisement (LSA) generation when the values differ.

Delay normalization is disabled by default. To enable and configure delay normalization, use the `normalize` *`interval offset`* statement at the `[edit protocols ospf` *`area interface`* `delay-measurement]` hierarchy level.

[See delay-measurement (Protocols OSPF) and How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering.]

- **Selectively control per-prefix backup paths with OSPF import policy (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can selectively enable backup paths for specific prefixes to optimize redundancy and resource utilization. By default, a configured backup path applies to all prefixes. To exclude specific prefixes or ranges, create an OSPF import policy and configure the `no-backup` option in the `then` clause of the policy to suppress backup path installation for matching routes. You can reserve backup protection for critical prefixes while preventing unnecessary backups for others.

  [See Understanding Backup Selection Policy for OSPF Protocol.]

- **Preference-based Path Selection of L-OSPF Flexible Algorithm routes (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can control path selection by configuring the preference for L-OSPF Flexible Algorithm routes in `inetcolor.0` and `mpls.0`.

  Configure `flex-algorithm-preference` statement at the `[edit protocols ospf]` hierarchy level to prioritize desired routes and improve traffic engineering across IP and MPLS domains.

- **Policy-based redistribution of OSPF prefix SIDs across IGP instances (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can redistribute Segment Routing (SR) prefix-SIDs across OSPF IGP instances using route policy without explicitly specifying a prefix-segment index. This feature standardizes SR labels across instances and improves operational efficiency. Configure a policy with the `from prefix-segment` statement to match routes carrying prefix-segment information. In the `then` clause, use `prefix-segment redistribute` to inherit segment information from the matched route. We also support stitching `mpls.0` routes to enable interoperability between different IGP instances.

- **Non-router-ID endpoints as SR-TE destinations (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)**—Use non-router-ID endpoints as destinations in Segment Routing—Traffic Engineering (SR-TE) policies for Segment Routing for MPLS (SR-MPLS). Traditionally, these policies use router IDs, but you can specify anycast addresses to

enhance redundancy and load balancing in SR-MPLS networks. Use IPv4 and IPv6 anycast addresses as IGP-learned destinations with or without Segment Identifier (SID) stack compression. These anycast addresses are not redistributed (R-bit set). Use them as the to address for SR-TE policies with associated compute profiles..

[See Non-Router-ID Endpoints in Segment Routing Traffic Engineering.]

VPNs

- **L3VPN Support over pseudowire interface (ACX7020-AC, ACX7020-DC, ACX7024, ACX7024X, ACX7100-32C, ACX7332, ACX7348, and ACX7509)**—You can provision Layer 3 VPN (L3VPN) services over the PS interface and achieve parity with base L3VPN capabilities. Configure this feature using existing L3VPN workflows; you do not need additional CLI options.

  [See Layer%203%20VPNs%20User%20Guide%20for%20Routing%20Devices and MPLS Pseudowires Configuration.]

Additional Features

We've extended support for the following features to the platforms shown in parentheses:

- **GRES support with PFE warm boot, state retention, and IP tunnels (ACX7332)**

  [See .]

- **Support forGRES, graceful restart, and nonstop active routing (NSR)** (ACX7332)

  [See Understanding High Availability Features on Juniper Networks Routers.]

- **Two-Way Active Measurement Protocol (TWAMP) monitoring service (RFC 5357) hardware timestamp support to enable Flex Algo and SR-MPLS support** (ACX7020, ACX7024, ACX7100, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See the `offload-type inline-timestamping` option of the test-session statement.]

- **Support for 10GBASE-T Copper SFP+ transceiver** (ACX7024 and ACX7024X). The ACX7024 and ACX7024X routers support the 10GBASE-T SFP+ transceiver, a copper based SFP+ module with an RJ-45 connector. It offers flexible speed support, operating at 100 Mbps, 1 Gbps, or 10 Gbps based on the network requirements.

  [See Hardware Compatibility Tool.]

- **Supported transceivers, optical interfaces, and DAC cables** (ACX7020, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5220-32CD, QFX5230-64CD, QFX5240-64QD, QFX5240-64OD, QFX5700, and QFX5700E). 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 firewall filters on pseudowire and logical tunnel interfaces** (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)

  [See Firewall Filters.]

- **Support for IRB global filters in ingress direction with family `inet` only** (ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509). Global IRB firewall filters are applicable only for traffic coming from SP-style L2 logical interfaces (IFLs) present in BD and are not applicable for EP-style configuration.

  [See Firewall Filters.]

- **Support for SRv6 Network Programming, SRv6-TE, and services over SRv6-TE** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X). We support the following SRv6 features:

  - Topology Independent Loop-Free Alternate (TI-LFA) backup path for SRv6
  - Layer 3 services over SRv6 in BGP (END.DT4 and END.DT6)
  - Operations, Administration and Management (OAM) ping for SRv6
  - SRv6 traceroute
  - Static SR-TE for SRv6
  - Dynamic tunnels for SRv6
  - SRv6-TE
  - Flex Algo for SRv6
  - SRv6 microloop avoidance
  - SRv6 header compression (uSID/micro-SID)
  - Service interworking between SRv6 and SR-MPLS
  - EVPN-VPWS
  - Service route resolution using SRv6-SIDs (BPG PIC Edge/Core for L3VPN)

  NOTE: By default, USP or USD flavors of END SID are enabled on ACX nodes for all the SIDs. The ACX Series device does not support the End SID functionality on a per SID

basis for USP and USD and instead provides support based on incoming packets, SL value, and number of SRHs in the packet.

[See Understanding SRv6 Network Programming in IS-IS Networks.]

- **Renaming OpenSSH implementation to JSSH** (All platforms). The OpenSSH implementation in Junos OS Evolved is renamed "JSSH" to highlight its customized nature.

  [See ssh.]

What's Changed

IN THIS SECTION

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

General Routing

- SSH key options for user account credentials. You can configure key-options *key-options* option at the `set system login user` *user* `authentication [ssh-rsa|ssh-ecdsa|ssh-ed25519] ssh key` hierarchy level.

  [See login.]

- Displays the event log of learned MAC addresses. By default mac-learning-logs are stored in UTC timestamps. To view the logs in system timezone, use the show ethernet-switching mac-learning-log use-system-timezone command. The show ethernet-switching mac-learning-log use-system-timezone command also prints the time zone abbreviations [IST, UTC, etc] in the timestamp. To view the logs in system timezone by default by using the show ethernet-switching mac-learning-log command, you need to configure the system-timezone statement at the [edit protocols l2-learning mac-learning-log] hierarchy level.

- When you run the request vmhost zeroize command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines.
- **Deprecated license trace (Junos OS Evolved)**—We've deprecated the CLI option `show system license liblicense-trace`.
- On the MPC7E-10G line card, when you configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet ports, use the speed statement at both the edit interfaces <interface name> gigether-options and edit interfaces interface <name hierarchy> levels.
- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked.
- **New option for debug collector data storage path**—We've included the option `outdir` to specify an output directory for storing debug collector data in a customised path. This allows you to organise and access diagnostic information more efficiently, adapting storage to your specific requirements.

  [See request system debug-info.]

Platform and Infrastructure

- Tacacs authorisation support for local authentication without password-Starting in Junos OS Evolved Release 25.4R1, you need not configure password under edit system authentication-order to enable password-options.
- **Commit validation for unique user IDs**—We have added support to validate the user configuration to ensure that each user is assigned a unique UID. A commit fails if duplicate UIDs are detected, ensuring stronger validation and preventing identity conflicts. Previously, a commit was successful even when multiple users shared the same UID, triggering only a warning and logging a syslog message.

User Interface and Configuration

- **Generate genstate YANG modules on Junos devices**—You can use `show system schema` operational command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

  [See show system schema.]

Known Limitations

IN THIS SECTION

Learn about limitations in this release for ACX Series routers.

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

General Routing

- When the original flow is egressing out through an 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 AE interface. For non-AE egress interface, this works fine and the sflow frame reflects the correct egress port.PR1647870
- With T-GM(Timing - Grand Master) enabled system like ACX7024, ACX7348 and ACX7332, when master port link comes up late, then PTP system servo state can move from INITIALIZING to PHASE ALIGNED. This will not cause any issue since when in INITIALIZING state, no PTP packets are sent to down streams devices and in PHASE ALIGNED state, PTP packets are sent with all correct clock parameters. When system moves to directly PHASE_ALIGNED, it is assured that the system is internally locked to phase and frequency of GNSS.PR1887028
- ACX7348 and ACX7332: Transient traffic loss is observed after PTP is enabled.PR1904131
- This PR is covering two behaviours of AFTS node. 1. During the system's transition from GNSS to hybrid mode, multiple phase spikes may occasionally be observed. However, these spikes remain within the acceptable threshold of 260 nanoseconds, and therefore, the performance is considered to be within the permissible operational range. 2. Rare Phase Jump Events In very rare scenarios, a phase jump exceeding 260 nanoseconds may occur. These events are momentary in nature, and observations indicate that the system recovers immediately.PR1908744
- As per GNSS receiver vendor, this issue has happened due to the sudden temperature changes affecting the GF-8801 unit which uses TCXO oscillator. Please refer to the attached plot (https:// gnats.juniper.net/web/default/1908802/attachments/GF_8801_Temp_vs_ Freq_Mode_Change.png) for detailed analysis. GF-8801 includes one TCXO to generate 10MHz and this TCXO can also be

drifted suddenly when sudden ambient temperature changes or strong air flow directly impact GF-8801.PR1908802

- In an infrequent case scenario, the system will be stuck in the acquiring state when asymmetry is configured on the primary LAG(400G). And the way to recover from the issue is to restart the fpc slot where the 400G port is present.PR1909225

Open Issues

IN THIS SECTION

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

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

General Routing

- On ACX7024 and ACX7100 platforms, the `show chassis fpc` command should not display temperature, CPU, and memory utilization values. These metrics are not valid for these platforms since FPC(Flexible PIC Concentrator) is a logical FRU (Field Replaceable Unit).PR1735274
- When ingress policer is configured, to drop ingress traffic, on an interface with upMep the CFM packets generated from the upMep will also be dropped due to the policer. This will lead to CFM session going down.PR1754938
- On Junos OS Evolved ACX7000 platforms configured with MPLS tunnel and Storm Control, the Layer 2 or MPLS tunnel-terminated known unicast traffic arriving at the ingress interface assigned with a high drop precedence by the interface-level classifier will get dropped on the interface where storm control profile is active. Due to this, MPLS tunnel traffic may get affected.PR1802525
- On all Junos OS Evolved platforms, MPLS (Multiprotocol Label Switching) payload type ether-pseudowire configured along with any other MPLS payload type does not work as expected and traffic will be impacted.PR1824219

- On all Junos OS Evolved platforms, host-originated L3 (Layer 3) traffic is marked with DSCP(Differentiated Services Code Point) value 48 [INET 110b] in the IP header, even if class-of-service host-outbound configuration is not present. Host-originated L2 (Layer 2) traffic is marked with IEEE 802.1p value 0 in the header, even if class-of-service host-outbound configuration is not present.PR1837443
- An LSI IFL remains in RPD even after being deleted by the interface manager daemon. It is visible in show interface routing but not in show interfaces, indicating that RPD still holds the IFL despite its removal elsewhere. rpd-agent does not send a delete message to RPD due to a reference count issue. Another daemon?likely l2ald?still holds a reference to the IFL. rpd-agent only sends the delete once all references are cleared, which doesn't happen in this case. The fix is to send a "delete pending" message from rpd-agent to RPD. RPD will treat this as a delete and remove the IFL, ensuring consistency across the system.PR1866522
- ACX7100-32C:While Verifying Accounting Flow inline-jflow-error-information of inline-as-lookup-failure is not getting as expected.PR1882224
- On Junos OS Evolved ACX7024 and ACX7024X, the **Active Disk Usage Exceeded** alarm is raised and not cleared even if the usage is less than 50%. This issue has no impact on traffic.PR1884419
- With traffic running and deletion and addtion of fti tunnel multiple times, IPv6 traffic drop is observed in the DUT. PR1896144
- During the transition from Slave 1 to Slave 2 in AFTS, there is a scenario where Slave 1 gets disabled, but the SPLL information from Slave 2 is received with a delay. As a result, the system initially remains in thePhase Alignedstate (Active, Ready) with Slave 2. Due to the delayed SPLL update, it then shifts to theAcquiringstate (Active, Not Ready) as it begins synchronisation with Slave 2, and then the system successfully re-aligns and returns to thePhase Alignedstate (Active, Ready) with Slave 2.PR1909821

Resolved Issues

IN THIS SECTION

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

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

General Routing

- [Clocking Solution]:ACX7348: PTP performance is not working as expected for asymmetry configuration. PR1793926
- ACX7509 might report SSD DMA error on secondary disk. PR1848468
- After GRES ptp-spll-lock-state is stuck in ACQUIRING instead of phase aligned. The timingd re core is also observed which is not seen in rerun. PR1855379
- Interfaces on ACX7509 go down after reboot with any FPC in an offline state. PR1857014
- Harmless error logs seen when EVPN Type-5 or IRB configured. PR1860489
- [ACX7000 Junos OS Evolved] syslog message "RPD_DYN_CFG_SMID_REG_FAILED: Failed to open session database: -1" pops every 5 seconds. PR1865702
- System stops responding to `show system applications app ndp detail` command or prevents applications from start/stop/restart. PR1866988
- The FEB might remain in fault state on ACX7509 after consecutive RE switchovers within 60 seconds. PR1869477
- Interfaces are not created after system reboot on Junos OS Evolved ACX platforms. PR1880571
- The debug-collection fails due to insufficient space. PR1883317
- Fan tray X incompatible alarm was raised when the fan unit was removed. PR1885345
- The application rpd-agent might restart with a coredump after interface related event changes. PR1885455
- ACX7000 DWDM optics wavelength not applied after reboot. PR1885817

- VSTP convergence failure on Junos OS Evolved ACX7000 when IRB shares bridge domain with VSTP enabled IFL. PR1886421
- Commit can fail if apply-path contains pattern-matching tokens other than the wildcard. PR1888201
- SSD firmware for both primary and secondary SSD is not displayed. PR1889258
- Traffic drop is observed in an EVPN multihoming as the MAC route points to the ESI interface when the CE (ESI) IFD flaps. PR1889335
- Framing error observed on AE interface after performing RE switchover on Junos OS Evolved ACX7509 platform. PR1891275
- Inaccurate FRU temperature readings reported via SNMP after switchover or reboot on ACX7509. PR1891630
- Firewall filter IP precedence matches incorrectly when a range of values is configured. PR1891684
- The fetched firmware version results an incorrect value on ACX7509. PR1891967
- Incorrect shared memory unit interpretation causes memory leak log errors. PR1892348
- On Junos OS Evolved ACX7509 system's clock break synchronization. PR1892751
- Classification of traffic is not working due to Forwarding Class to Queue mapping configuration failure. PR1893395
- Auto-negotiation mismatches leading to connectivity issues on Junos OS Evolved ACX7509, ACX7332, and ACX7348 platforms. PR1893886
- ACX7348 - Continuous log messages seen on backup routing engine after RE switchover. PR1894824
- System crashes due to a CPU glitch. PR1895749
- Adding a new key to authentication-key-chain causes kernel crash. PR1895827
- The other port in a port group goes down after an RE switchover if one of the port in the port group is already down on an ACX7509. PR1898908
- TPM key regeneration failure after zeroize affects the sZTP and gRPC functionality. PR1899669
- PCIe fatal error resulting in an unexpected device reboot and evo-pfemand crash on Junos OS Evolved ACX7322 platform. PR1900115
- The egress traffic control profile configured under hierarchical QoS does not take effect when applied to an L2IFL interface that is associated with an IRB. PR1900224

- Deactivating or deleting `auto-configure` configuration statement in an interface is not working properly. PR1902855
- In an EVPN Flex Algorithm environment, if multiple routing instances share the same color and the same VLAN, traffic might not be handled correctly. PR1905116
- System stops responding to `show system applications app ndp detail` command or prevents applications from start/stop/restart. PR1905807
- After RE switchover on the device, the capacity of DC PSM is observed to be zero after toggling the input power. PR1907310
- Next hop is missing when ECMP uses an IRB interface in an EVPN-VXLAN scenario. PR1908282
- The evo-pfemand crash is seen on ACX7000 platforms. PR1910126

Interfaces and Chassis

- ACX7024 see a "\" character in show chassis hardware output. PR1890751
- Performing Power-cycle from Junos OS Evolved after Bios Recovery (LKG/USB). PR1893274
- Unexpected commit failure will be observed when deactivating of unit <unit> or changing any configuration and performing a commit operation. PR1898055
- AE interfaces flap during GRES following an RE reboot on all Junos OS Evolved platforms. PR1898531
- IFL installation failure with the same outer VLAN tag and distinct inner VLAN lists. PR1911684

Layer 2 Features

- In VPLS scenario due to ungraceful switchover rpd process crash is observed. PR1882938
- Stale ARP details after RE switchover leading to traffic loss observed on ACX7000. PR1896785

Network Management and Monitoring

- Configuration under `set system trace application <app-name>` gets lost during rollback with certain sequence of steps. PR1869479

Routing Protocols

- BGP Prefix-SID Label collision causing RPD crash. PR1889749

Services Applications

- PAA installation failure on reboot due to **Not found default vrf** error. PR1886928

User Interface and Configuration

- Slow configuration commit observed on devices with a single Routing Engine (RE). PR1884781
- Traffic loss occurs due to missing apply-path handling during reboot or configd restart. PR1900883

Junos OS Evolved Release Notes for PTX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 25.4R1 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

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 25.4R1, 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.
- **PTX10K-LC1301-36DD line card (PTX10008)**—The PTX10K-LC1301-36DD line card features 36 ports, delivering a line rate throughput of 28.8 Tbps. The 36 high-density 800-Gigabit Ethernet (800 GbE) QSFP-DD ports support speeds of up to 800 Gbps. The line card houses two Juniper Networks' custom Express 5 ASICs, and each ASIC comprises two Packet Forwarding Engines.

  [See PTX10008 Line Card Components and Descriptions.]
- **New JNP10008-SF5 SIB (PTX10008)**—The JNP10008-SF5 Switch Interface Board (SIB) supports up to 28.8 Tbps of bandwidth per slot for the PTX10KLC1301-36DD line card installed in a PTX10008 router running Junos OS Evolved.

  [See PTX10008 Switch Fabric.]
- **Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers**

Feature	Description
Chassis	Packet Forwarding Engine resiliency . We provide resiliency feature support for the Packet Forwarding Engine, which enables the system to detect, report, and take action on Packet Forwarding Engine faults. Actions are taken based on the default configuration or a user configuration available for the errors.
	Fabric hardening and resiliency support on PTX10K-LC1301-36DD line cards. [See Fabric Hardening and Recovery on PTX10K Devices.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• Interoperability support and CLI enhancements. The PTX10008 router with the JNP10008-SF5 SIB supports default interoperability between the PTX10K-LC1301-36DD, PTX10K-LC1201-36CD, and PTX10K-LC1202-36MR line cards. Use the `set chassis interoperability express5-enhanced` command to bring up the system in the `express5` mode-specific functionalities. This disables the line-card interoperability feature. You can verify the interoperability status using the `show chassis interoperability` command. The existing commands for PTX10008 with PTX10K-LC1201-36CD line card will support for PTX10008 with PTX10K-LC1301-36DD line card as well. Following are the new CLI command updates: • The `show chassis fpc slot detail` command displays the Packet Forwarding Engine ASIC type. • In the `set chassis fpc` command, you must use `pfe-instance` instead of `pfe`. • The `show chassis fpc 5 pfe-instance all` command displays `pfe-instance` in the output. [See interoperability, show chassis interoperability, chassis, and show chassis fpc.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• Fabric resiliency support for JNP10008-SF5 SIB. The JNP10008-SF5 SIB supports fabric resiliency, enhancing fault management for fabric links. You can benefit from features including error detection, logging, alarm generation, SNMP trap sending, LED error indications, and self-healing. Use the CLI command `show system errors active detail` to view logged errors, ensuring comprehensive fault monitoring and increased system reliability. [See Fabric Resiliency and show system errors active.]
	• FPC fabric management for JNP10008-SF5 SIB. You can use the CLI command `set chassis fpc` to manage FPC online and offline states gracefully. Use the `set chassis fabric event reachability-fault` command to configure options for detecting fabric reachability faults and trigger automatic connectivity restoration. Additionally, use the `extended` keyword in `show chassis fabric fpcs` and `show chassis fabric sibs` commands to view detailed link information within planes, and identify partially enabled planes with the `Degraded` keyword in the `show chassis fabric fpcs` command. [See reachability-fault, show chassis fabric fpcs, and show chassis fabric sibs.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description

	• Support for JNP10008-SF5 SIB (PTX10008) — The PTX10008 supports the JNP10008-SF5 Switch Interface Board (SIB), which includes 18 fabric planes. You can use the `extended` keyword with the `show chassis sibs` command to view detailed plane information. Use the `set chassis sib` command to gracefully bring SIBs online or offline. Note that mixing JNP10008-SF3 SIB with JNP10008-SF5 SIB will result in compatibility errors indicated by specific CLI commands:
	  ◦ The `show chassis sibs` detail command displays `Incompatible with other SIBs` in the output.
	  ◦ The `show chassis alarms` command displays `SIB Incompatible` in the output.
	  ◦ The `request chassis sib online` command displays `Request failed since Fru is incompatible with other slots!` in the output.
	[See show chassis sibs, show chassis alarms, request chassis sib, and Fabric Management on PTX10K Devices.]

	• Optics EM policy support. The Environment Monitoring (EM) policy includes optics temperature sensors for PTX10008 routers with the PTX10K-LC1301-36DD line card. It ensures efficient thermal management of high-power optical modules. Key functionalities include temperature monitoring integration, automatic shutdown procedures, and CLI commands for managing and configuring the EM policy. [See Optics EM Policy Support.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
Class of service (CoS)	• Support for CoS features, including classifiers (behavior aggregate (BA), fixed, and multifield (MF)), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, drop profiles, HCoS, and policy map . • [See CoS Features and Limitations on PTX Series Routers and Class of Service.]
	• Support for on-chip queue buffer for PFC-enabled queues. A priority-based flow control (PFC)-enabled queue with a buffer-size less than 450 microseconds is viewed and installed as a PFC-enabled on-chip queue. When a queue is in PFC on-chip mode, the entire virtual output queue (VOQ) buffer is always on-chip and is not scaled based on bandwidth usage. [See **buffer-size (Schedulers)**.]

Dynamic Host Configuration Protocol (DHCP)	• Support for DHCPv4 relay agent and DHCPv6 relay agent, including:
	  • DHCP relay: Layer 3 (L3) interfaces
	  • DHCP relay: Option 82 for Layer 2 VLANs
	  • DHCP relay: Option 82 for L3 interfaces
	  • Extended DHCP relay agent
	  • Virtual router-aware DHCP (VR-aware DHCP)
	  [See Extended DHCP Relay Agent Overview.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
EVPN	• Support for EVPN-VXLAN Layer 2 (L2) gateways and Layer 3 (L3) gateways with EVPN Type 5 routes. [See EVPN User Guide.] • Support for ping and traceroute for EVPN-VXLAN [See Understanding Overlay ping and traceroute Packet Support.] • Support for Static VXLAN (L2 gateway). [See See Static VXLAN.]
	Support for EVPN-MPLS L2 and L3 features. [See EVPN Overview.]
	Support for EVPN-VPWS. [See Overview of VPWS with EVPN Signalig Mechanisms.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

Infrastructure	• We support the following IP and Infrastructure features:
	  • Junos telemetry interface (JTI) support for Packet Forwarding Engine sensors for usage, network processing unit (NPU) memory, NPU utilization, and pipeline NPU and ASIC. Using the JTI, you can export statistics using remote procedure call (gRPC) services, gRPC Network Management Interface (gNMI) services, and UDP transport.
	  Use these sensors:
	  • **/junos/system/linecard/packet/usage/**
	  • **/junos/system/linecard/npu/memory/**
	  • **/junos/system/linecard/npu/utilization/**
	  • **/components/component/integrated-circuit/state/**
	  • **/components/component/integrated-circuit/pipelinecounters/**
	  For pipeline sensors, the four packet and drop counter categories are interface, lookup, queuing, and host interface.
	  [See Junos YANG Data Model Explorer.]
	• Traffic drops classification based on trap classification.
	• Support for distributed denial-of-service (DDoS) IS-IS classification and higher DDoS bandwidth for Layer 2 and Layer 3 protocols.

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

	[See show ddos-protection protocols isis and protocols (DDoS).]
	• Support for load balancing under the `[edit forwarding-options enhanced-hash-key]` hierarchy. Load balancing includes:
	  • GRE key inclusion for transit IPv4 and IPv6 traffic
	  • IP Layer 3 fields
	  • IP Layer 4 fields
	  • IPv6 flow label inclusion
	  • MPLS labels
	  • MPLS port data
	  • MPLS pseudowire traffic
	  • Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets
	  • RSVP-TE load balancing in proportion to LSP bandwidth
	[See enhanced-hash-key.]
	Support for 128-way equal-cost multipath (ECMP) routing for MPLS transit cases.
	The following features do not support 128-way ECMP:
	• Multicast
	• P2MP

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• MC-LAG • Weighted unilist • Consistent hashing • Link protection (MPLS) • Adaptive load balancing • Class-based forwarding • Support for classification override configured under a forwarding policy. [See CoS Features and Limitations on PTX Series Routers and Overriding the Input Classification.] • You can configure passive monitoring on any interface on the PTX10008 routers to monitor MPLS-encapsulated packets. [See Passive Monitoring and passive-monitor-mode.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

Interfaces and chassis	• Support for VRRP. The following features are not supported for VRRP on Junos OS Evolved:
	  • ISSU
	  • Proxy ARP
	  • MC-LAG
	  • Distribution support on aggregated Ethernet interfaces
	  • IRB
	  • Inline delegation
	  [See Understanding VRRP.]
	• Support for the following protocols:
	  • LAG (aggregated Ethernet)
	  • LACP
	  • LLDP
	• Support for link fault management (LFM). We support IEEE 802.3ah OAM LFM to monitor point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The following LFM features are supported:
	  • Link discovery with active and passive modes
	  • Detect-LOC
	  • Remote loopback
	  • Loopback tracking
	  • Action profile

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• GRES and non-graceful Routing Engine switchover [See Introduction to OAM Link Fault Management (LFM).]
	We support the following optics: • 800GbE • 400GbE • 100GbE/2x100GbE • 10GbE/25GbE/40GbE • We support MAC address accounting for 10GbE, 40GbE, 100GbE, 200GbE, 400GbE, and 800GbE interfaces. • Support for MAC accounting for source and destination MAC addresses for Layer 3 interfaces and aggregated Ethernet interfaces. To enable MAC accounting, use the existing `mac-learn-enable` command at the `[edit interfaces interface-name gigether-options ethernet-switch-profile]` or `[edit interfaces aex aggregated-ether-options ethernet-switch-profile]` hierarchy level.

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

IP tunneling	• Support for the following Packet Forwarding Engine tunnel features:
	  • Filter-based GRE encapsulation and de-encapsulation and filter-based MPLS-in-UDP de-encapsulation. We've enabled the following encapsulation and de-encapsulation workflow:
	  1. An incoming packet matches a filter term with an encapsulate action. The packet is encapsulated in an IP+GRE header and is forwarded to the endpoint's destination.
	
	`set firewall tunnel-end-point` *`tunnel-name`* `ipv4|ipv6 source-address` *`address`*
	`set firewall tunnel-end-point` *`tunnel-name`* `ipv4|ipv6 destination-address` *`address`*
	`set firewall tunnel-end-point` *`tunnel-name`* `gre`
	`set firewall family inet|inet6 filter` *`name`* `term` *`name`* `from source-address` *`address`*
	`set firewall family inet|inet6 filter` *`name`* `term` *`name`* `then encapsulate` *`tunnel-name`*
	`set firewall family inet|inet6 filter` *`name`* `term last then accept`
	`set interfaces` *`interface-name`* `unit` *`number`* `family inet|inet6 filter input`
	`set interfaces` *`interface-name`* `unit` *`number`* `family inet|inet6 address` *`address`*
	`# This source address differs from the one for the tunnel endpoint.`

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

	2. At the destination, the packet matches a filter term with a de-encapsulate action. The GRE header or MPLS-in-UDP header is stripped from the packet. The inner packet is routed to its destination.
	
	`set firewall family inet|inet6 filter` *`name`* `term` *`name`* `from source-address` *`address`*
	`set firewall family inet|inet6 filter` *`name`* `term` *`name`* `from protocol gre`
	`set firewall family inet|inet6 filter` *`name`* `term` *`name`* `then decapsulate gre #`
	`Optionally de-encapsulate mpls-in-udp.`
	`set firewall family inet|inet6 filter` *`name`* `term last then accept`
	`set interfaces` *`interface-name`* `unit` *`number`* `family inet|inet6 filter input` *`filter-name`*
	`set interfaces` *`interface-name`* `unit` *`number`* `family inet|inet6 address` *`address`*
	`# This is the destination address.`
	
	[See Components of Filter-Based Tunneling Across IPv4 Networks and tunnel-end-point.]
	
	• Support for FTI-based encapsulation and de-encapsulation of IPv4 and IPv6 packets. You can configure IP-IP encapsulation and de-encapsulation on flexible tunnel interfaces (FTIs). The default mode is loopback encap mode. Use the `bypass-loopback` statement at the `[edit interfaces fti` *`number`* `unit` *`logical-unit-number`* `tunnel encapsulation ipip]` hierarchy level to change into flattened encap mode to achieve line-rate performance.

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	[See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.] • Support for configuring MPLS protocols over FTI tunnels, thereby transporting MPLS packets over IP networks that do not support MPLS. GRE and UDP tunnels support the MPLS protocol for both IPv4 and IPv6 traffic. You can configure encapsulation and de-encapsulation for the GRE and UDP tunnels. To allow the MPLS traffic on the UDP tunnels, include the `mpls port-number` statement at the [`edit forwarding-options tunnels udp port-profile` *`profile-name`*] hierarchy level. To allow the MPLS traffic on the GRE tunnels, include the `mpls` statement at the [`edit interfaces fti0 unit` *`unit`* `family`] hierarchy level. [See Flexible Tunnel Interfaces Overview.] • Support for gress filter-based encapsulation. For an outgoing packet matching the filter term, the packet is encapsulated inside an IP + GRE header as specified by the tunnel configuration. IP lookup is performed on the outer header and the packet is forwarded accordingly. The IP lookup for GRE-encapsulation-capable route is limited to the implicit default routing instance. [See Understanding Filter-Based Tunneling Across IPv4 Networks.] • Support for configuring the output filter action with a nondefault routing instance or a specified routing instance. [See Firewall Filter Terminating Actions.] • Ingress filter-based de-encapsulation by using firewall filters for GRE and UDP tunnels

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	[See Configuring a Filter to De-Encapsulate GRE Traffic and decapsulate (Firewall Filter.]
Junos telemetry interface (JTI)	Junos telemetry interface (JTI) supports new platform sensors for the PTX10008. You can export platform-specific software and chassis component statistics using remote procedure call (gRPC) services, gRPC Network Management Interface (gNMI) services, and UDP transport. NewXpaths are added in the YANG data model. [For a complete list of Xpaths supported by the device, see Junos YANG Data Model Explorer.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	Packet Forwarding Engine'sINT for PTX Series routers. The Junos OS Evolved Packet Forwarding Engine introduces a framework in the data plane, called inband network telemetry (INT), which collects and reports network state information without the intervention of the control plane. The header in the INT model has telemetry instructions that instruct an INT-capable device the state it must collect. The network state information is exported by the data plane either to the telemetry monitoring system or is written into the packet. INT has source, transit, and sink support. The INT source embeds the INT metadata in the packet and the sink collects the metadata from the data packet for processing. We do not support INT source, sink, and all INT application modes on PTX10008 routers. The JNP10K-LC1301-36DD line card on PTX10008 supports only INT transit node in Junos OS Evolved Release 24.4R1. Among the three INT application modes INT-XD, INT-MX, and INT-MD, the JNP10K-LC1301 line card on PTX10008 supports only INT-MD mode and INT as a transit node. The `set forwarding-options` configuration command is updated with a new `inband-telemetry` option, to enable or disable this feature. [See Junos YANG Data Model Explorer.]
Layer 2 features	Support for Q-in-Q tunneling. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

	• The PTX10008 router supports the following Layer 2 basic learning, bridging and flooding features:
	  • Enterprise-style bridging (support both trunk and access mode)
	  • Service provider-style bridging (also known as sub-interface mode)
	  • BPDU block/filter
	  • xSTP
	  • Handle broadcast, unknown unicast and multicast (BUM) traffic, including split horizon
	  • MAC learning and aging
	  • Static MAC addresses
	  • Trunk port and VLAN membership
	  • 802.1Q EtherType—8100
	  • 802.1Q VLAN tagging—Single tagging with normalized to bridge domain tag at ingress
	  • Clearing all MAC address information
	  • Global MAC limit
	  • Global source MAC aging time
	  • MAC moves
	  • LACP and LLDP
	  • Disabling MAC learning at global and interface level

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• Native VLAN ID for Layer 2 logical interfaces • Single VLAN-tagged Layer 2 logical interfaces • Interface statistics **NOTE**: The router does not support the `show ethernet-switching statistics` command and child logical interface statistics for aggregated Ethernet. • Flexible Ethernet services **NOTE**: Enterprise-style Layer 2 logical interfaces aren't allowed under the `flexible-ethernet-services` encapsulation. • Virtual switch • Persistent MAC learning (sticky MAC) • Service provider bridging: • Multiple logical interfaces on the same physical interface that are part of the same bridge domain • Ethernet bridge encapsulation [See Layer 2 Bridging, Address Learning, and Forwarding User Guide.] • Support for IRB: • All Layer 2 protocols already supported on the router Layer 3 protocols: BGP, IGMP, IS-IS, OSPF, PIM, and RIP Per-IRB logical interface MAC and statistics IRB Layer 3 multicast support with flooding only

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	Address family support for IPv4 and IPv6, and support for IPv4 MTUs and IPv6 MTUs with different MTU values IRB interface in VRF routing instances Directed subnet broadcast support with IRB. [See Integrated Routing and Bridging.] • Support for interface MAC limit action. You can specify the action (`drop`, `drop and log`, `log`, or `shut down`) that Junos OS Evolved takes when packets with new source MAC addresses are received after the MAC address limit is reached. [See Configuring MAC Limiting and packet-action.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

Layer 3 features	• Support for 256-way ECMP. You can configure a maximum of 256 ECMP next hops for external BGP (EBGP) peers. This feature increases the number of direct BGP peer connections, which improves latency and optimizes data flow. However, we support 128 ECMP next hops for MPLS routes. Note that we do not support consistent load balancing (consistent hashing) for IPv4 or IPv6 with this feature.
	
	[See Understanding BGP Multipath.]
	
	• Support for the following Layer 3 forwarding features for IPv4, IPv6, MPLS, LAG, ECMP, MTU checks, ICMP, OSPF, IS-IS, ARP, NDP, BGP, BFD, LACP, LDP, RSVP, LLDP, VRF-lite, TTL expiry, IP options, IP fragmentation, DDoS
	
	• BFD support, including:
	
	  • Distributed BFD and BFD-triggered local repair (BFD authentication is not supported.)
	
	  • Independent micro-BFD sessions enabled on a per-member link basis for a LAG bundle
	
	  • Inline BFD
	
	  [See Understanding BFD.]
	
	• BGP flowspec signaling support. BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX10008 devices with LC1201, LC1202 and LC1301 line cards.Propagating firewall filter information as part of BGP enables you to propagate firewall filters against denial-of-service (DOS) attacks dynamically across autonomous systems. The following match conditions are not supported:

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• ICMP codes alone [inet/inet6] • Source/destination prefix with offset for inet6 • Flow label for inet6 fragment (for inet6) Junos OS Evolved doesn't support the traffic marking action on this router. To configure flow routes statically, configure the match conditions and actions at the [edit routing-options] hierarchy level.
MACsec	Media Access Control Security (MACsec) is supported on physical interfaces. [See Understanding Media Access control Security (MACsec).]
	Support for Media Access Control Security (MACsec) bounded delay protection. [See Configuring Bounded Delay Protection.]
Managing devices	Support for additional RPCs for the gNOI certificate management (cert) service. Junos OS Evolved supports the following gRPC Network Operations Interface (gNOI) cert service RPCs: • CanGenerateCSR()—Query if the target device can generate a certificate signing request (CSR) with the specified key type, key size, and certificate type. • RevokeCertificates()—Revoke certificates on the target device. [See gNOI Certificate Management (Cert) Service .]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

MPLS	• We support the following MPLs features:
	  • Support for MPLS FRR—MPLS fast reroute (FRR) provides faster convergence time (less than 50 milliseconds) for RSVP tunnels. The Routing Engine creates backup paths and the Packet Forwarding Engine installs the backup-path labels and next hops.
	  [See Fast Reroute Overview.]
	  • Support for 256-way ECMP. You can configure a maximum of 256 equal-cost multipath (ECMP) next hops for external BGP (EBGP) peers. This feature increases the number of direct BGP peer connections, which improves latency and optimizes data flow. However, we support 128 ECMP next hops for MPLS routes. Note that we do not support consistent load balancing (consistent hashing) for IPv4 or IPv6 with this feature.
	  [See Understanding BGP Multipath.]
	  • Support for MPLS features, including:
	  • CLI support for monitoring MPLS label usage
	  • Inline MPLS and IPv6 lookup for explicit null
	  • 32,000 transit LSPs
	  • Explicit null support for MPLS LSPs
	  • MPLS Label Block configuration
	  • MPLS over untagged Layer 3 interfaces

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description

	• MPLS OAM - LSP ping
	• JTI: OCST: MPLS operational state streaming (v2.2.0)
	• 2000 ingress LSP support
	• 2000 egress LSP support
	• Entropy label support
	• MPLS: JTI: Junos telemetry interface MPLS self-ping, TE++, and misc augmentation
	• Support for LDP, including:
	  • Configurable label withdraw delay
	  • Egress policy
	  • Explicit null
	  • Graceful restart signaling
	  • IGP synchronization
	  • Ingress policy
	  • IPv6 for LDP transport session
	  • Strict targeted hellos
	  • Track IGP metric
	  • Tunneling (LDP over RSVP)
	• RSVP++

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

	• Support for RSVP-TE, including:
	  • Bypass LSP static configuration
	  • Ingress LSP statistics in a file
	  • RSVP-TE hitless-MBB with no artificial delays
	  • 32,000 transit LSPs
	  • Automatic bandwidth
	  • Class-based forwarding (CBF) with 16 classes
	  • CBF with next-hop resolution
	  • Convergence and scalability
	  • Graceful restart signaling
	  • JTI interface statistics and LSP event export
	  • LSP next-hop policy
	  • LSP self-ping
	  • MPLS fast reroute (FRR)
	  • MTU signaling
	  • Optimize adaptive teardown
	  • Node/link protection
	  • Refresh reduction

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• Soft preemption • Shared Risk Link Group (SRLG) • Static LSPs with IPv4 next hop, IPv6 next hop, and IPv6 next hop with next-table support for bypass • Traffic engineering, including: • TE++: Dynamic ingress LSP splitting • Traffic engineering extensions (OSPF-TE and ISIS-TE) • Traffic engineering options `bgp`, `bgp-igp`, `bgp-igp-both-ribs`, and `mpls-forwarding` [See MPLS Applications User Guide .] • Segment routing support. You can configure the following Source Packet Routing in Networking (SPRING) or segment routing features on the router: • MPLS (segment routing using IS-IS): • Ping and traceroute for single IS-IS node or prefix segment • BGP Link State (BGP-LS): • Segment routing extensions for IS-IS • Segment routing extensions for OSPF • BGP:

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• Binding segment identifier (SID) for segment routing–traffic engineering (SR-TE) • Binding SID for SR-TE [draft-previdi-idr-segment-routing-te-policy] • Programmable routing protocol process APIs for SR-TE policy provisioning • Static SR-TE policy with mandatory color specification • Static SR-TE policy without color specification • IS-IS: • Adjacency SID • Advertising maximum link bandwidth and administrative color without RSVP-TE configuration • Anycast and prefix SIDs • Configurable segment routing global block (SRGB) • Node and link SIDs • Segment routing mapping server (SRMS) and client • Topology Independent Loop-Free Alternate (TI-LFA):

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• Link and node protection for IPv4 addressing (not required for IPv6 prefixes) • Link and node protection for IPv4 addressing (required for IPv6 prefixes) • Protection for SRMS prefixes • OSPF: • Advertising maximum-link bandwidth and administrative color without RSVP-TE configuration • Anycast SID • Configurable SRGB • Inter-area support • Node and link SID • Prefix SID • Segment routing mapping server (SRMS) and client • Static adjacency SID • TI-LFA: • Link and node protection • Protection for SRMS prefixes

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• MPLS ping and traceroute for single OSPF node or prefix segment • IGP adjacency SID hold time • Path Computation Element Protocol (PCEP) for segment routing LSPs • BGP IPv4 labeled-unicast resolution over: • BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF • Noncolored IPv4 SR-TE with segment routing using IS-IS and OSPF • Static colored IPv4 SR-TE with segment routing using IS-IS and OSPF • • BGP Layer 3 VPN over: • Colored SR-TE tunnels and IPv4 protocol next hops • Non-colored SR-TE tunnels and IPv4 protocol next hops • BGP-triggered dynamic SR-TE colored tunnels • Class-based forwarding and forwarding table policy LSP next-hop selection among noncolored SR-TE LSPs • First-hop label support for SID instead of an IP address

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

	• Path specification using router IP addresses (segment routing segment list path ERO support using IP address as next hop and loose mode)
	• SR-TE color mode:
	  • 00—Route resolution fallback to IGP path
	  • 01—Route resolution fallback to color only null routes
	• Static LSPs with member-link next hops for aggregated Ethernet bundles (also known as adjacent SID per LAG bundle or aggregated Ethernet member link)
	[See Understanding Source Packet Routing in Networking (SPRING).]
	• Support for Layer 2 VPN features, includeing:
	  • Transport of Layer 2 frames over MPLS (LDP signaling)
	  • Layer 2 VPNs over tunnels (BGP signaling)
	  • Simple Ethernet and VLAN-based cross-connect (also known as connections)
	  • Local and remote switching
	  • Ethernet and VLAN CCC
	  • Single-tagged CCC logical interfaces
	  • Control word
	  • Regular and aggregated Ethernet interfaces

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• Layer 2 protocol pass-through • Layer 2 circuit backup interface and backup neighbor • Layer 2 circuit statistics and CoS • VCCV with type 2 and type 3 [See **Layer 2 VPNs and VPLS User Guide for Routing Devices** and **TCC Overview**.] • VLAN ID lists for Layer 2 Circuits. VLAN ID lists allow you to link multiple VLAN IDs to a single logical interface for Layer 2 traffic. [See vlan-id-list (Ethernet VLAN Circuit), vlan-id-list, and Configuring VLAN Identifiers for VLANs and VPLS Routing Instances.] • MPLS-based Layer 3 VPNs support includes: • MPLS over Layer 3 VLAN-tagged subinterfaces • Per-next-hop label allocation • Mapping of the label-switched interface (LSI) logical interface label to the VPN routing and forwarding (VRF) routing table using the `vrf-table-label` statement • ICMP tunneling and MPLS traceroute • Disabling time-to-live (TTL) decrementing using `no-propagate-ttl` [See Layer 3 VPNs Feature Guide for Routing Devices.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• Support for IP-over-IP encapsulation to facilitate IP overlay construction over an IP transport network. An IP network contains edge devices and core devices. To achieve higher scale and reliability among these devices, use an overlay encapsulation to logically isolate the core network from the external network that the edge devices interact with. Static configuration or a BGP protocol configuration is used to distribute routes and signal dynamic tunnels. The dynamic-tunnels configuration creates IP-over-IP encapsulation-only tunnels in the Packet Forwarding Engine. The following are not supported: • Dynamic tunnel de-encapsulation operation • Next-hop-based statistics for dynamic tunnels • IP fragmentation at tunnel start point and path MTU discovery for IPv4/IPv6 [See Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation .] • Redistribution of IPv4 routes with IPv6 next hop into BGP. Devices can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.] • Link delay advertisement. You can get the measurement of various performance metrics in IP networks, which helps to distribute network-performance information in a scalable fashion. [See How to Enable Link Delay Measurement and Advertising in IS-IS.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
Multicast	• Support for multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows. MoFRR is supported for PIM sparse mode (SM) and source-specific multicast (SSM) modes only. Support does not extend to Multipoint LDP-based MoFRR. [See Understanding Multicast-Only Fast Reroute.] • Bidirectional Protocol Independent Multicast for multicast traffic. [See pim-snooping.] • Support for RSVP-based and LDP-based point-to-multipoint (P2MP) LSPs with graceful restart. In addition, the router supports IP unicast traffic in a label-edge router (LER) role and both IP unicast and multicast traffic in a label-switching router (LSR) role. [SeePoint-to-Multipoint LSP Configuration.] • Support for MPLS features P2MP ping and P2MP LSPs traceroute. MPLS ping and traceroute provide the mechanism to detect data-plane failure and isolate faults in the MPLS network. The traceroute or ping is initiated to validate LSP paths on P2MP. [See MPLS Applications User Guide.] • Optimized fast branch updates. The method of making fast branch updates to a multicast replication tree has been refined. Now, any membership changes in the tree trigger fast make-before-break (FMBB) re-optimization of the tree and ensure that there is no traffic loss. [See Multicast Shortest-Path Tree.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• Multicast support for next-generation MVPN including IR, RSVP-P2MP, and LDP-P2MP provider tunnel, inclusive and selective PMSI tunnel, rendezvous-point tree (RPT)-shortest-path tree (SPT) mode, turnaround provider edge (PE) device, rendezvous point (RP) mechanisms such as auto-RP, bootstrap router (BSR), and embedded RP. [See Multiprotocol BGP MVPNs Overview, Understanding Next-Generation MVPN Concepts, and Understanding Next-Generation MVPN Control Plane.] • Multicast support for next-generation MVPN including IR, RSVP-P2MP, and LDP-P2MP provider tunnel, inclusive and Selective PMSI tunnel, Rendezvous-point tree (RPT)-shortest-path tree (SPT) mode, turnaround provider edge (PE) device, RP mechanisms such as auto rendezvous point (RP), bootstrap router (BSR), and embedded RP. [See Multiprotocol BGP MVPNs Overview, Understanding Next-Generation MVPN Concepts, and Understanding Next-Generation MVPN Control Plane.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• MVPN BIER with MPLS encapsulation. Junos OS Evolved supports the Bit Index Explicit Replication (BIER) architecture to simplify control and forwarding planes by eliminating the need for multicast trees and per-flow states. With BGP-MVPN as an overlay, you can configure BIER-enabled provider tunnels for multicast VPNs. [See BIER Overview and bier.] • IS-IS as routing underlay for BIER. Junos OS Evolved supports the advertisement of BIER information of one or more BIER subdomains using IS-IS as the IGP underlay. Key BIER information such as BFR IDs and BFR prefixes in each subdomain are flooded through the IS-IS domain to generate the BIER forwarding table. [See IS-IS Extension for BIER and bier-sub-domain (Protocols IS-IS).]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

Network management and monitoring	• Local and remote port mirroring:
	  • Local port mirroring is used to copy the packet entering or leaving the system or port and send sampled packet through a predesignated port provided by configuration to remote devices or servers. Applications running on servers can analyze these packets and use the results based on the requirement.
	  • Remote port mirroring is used to send a sampled packet to a remote destination provided by configuration. The packet is encapsulated in a GRE header. Remote port mirroring makes use of the flexible tunnel interface (FTI) to encapsulate and send the packets out of the box. This feature also provides an option for configuring a policer for the given instance, so that the rate of sampling can be policed.
	• Port mirroring support for EVPN-VXLAN
	• Filter and mirror ingress and egress traffic on any network port to CPU. Junos devices support filtering and mirroring of incoming and outgoing packets, sending those packets to the CPU, and saving them to a file. This on-device packet capture feature can help you with protocol and application analysis, debugging, troubleshooting, network forensics, audit trails, and network attack detection. On-device packet capture (or “self-mirroring”) sends the sampled copy to a CPU and writes the copy into a packet capture (**.pcap**) file. The process does not require you to use any device connected to your network device.
	[See On-Device Packet Capture.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• Support for the sFlow technology, which is a monitoring technology for high-speed switched or routed networks. The sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. [See Understanding How to Use sFlow Technology for Network Monitoring.]
	Support for additional RPCs for the gNOI certificate management (cert) service. Junos OS Evolved supports the following gRPC Network Operations Interface (gNOI) cert service RPCs: • CanGenerateCSR() —Query if the target device can generate a certificate signing request (CSR) with the specified key type, key size, and certificate type. • RevokeCertificates()—Revoke certificates on the target device. [See gNOI Certificate Management (Cert) Service .]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description

	• Up maintenance association end points (MEPs) in distributed periodic packet management (PPM)
	• Distributed Y.1731 on synthetic loss measurement (SLM), delay measurement (DM), and loss measurement (LM)
	• Down MEPs on bridges, circuit cross-connect (CCC) , and EVPN
	• Distributed session support for CFM on aggregated Ethernet
	• Enhanced CFM mode
	• IPv4 (inet) support for Data Model (DM) and synthetic loss message (SLM)
	• Action profile for marking a link down, except for EVPN and bridge up MEP
	• LM colorless mode
	• DM and LM on aggregated Ethernet if all active child links are on the same Packet Forwarding Engine
	• Supported CFM protocol data units (PDUs), as follows:
	  • Continuity check messages (CCM)
	  • LBM
	  • LBR
	  • Link Trace Message (LTM)
	  • Link Trace Reply (LTR)

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• Delay measurement message (DMM) • Delay measurement reply (DMR) • LMM • LMR • Synthetic loss message (SLM) • Synthetic loss reply (SLR) • Enterprise and service provider configurations • VLAN normalization • VLAN transparency for CFM PDUs • CoS forwarding class (FC) and CoS packet loss priority (PLP) for CFM • CFM session on child physical interface in distributed mode • SNMP • Chassis ID or Send ID type, length, and value • Trunk mode • Maintenance association intermediate point (MIP)
Platform and infrastructure	Support for Synchronous Ethernet timing, Synchronous Ethernet over LAG, and Timing SNMP and MIB (SYNCE).

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	Platform resiliency support. PTX10008 routers with specific line cards support platform resiliency. Resiliency enables the router to handle failures and faults related to the hardware components such as line cards, switch fabric, Control Boards, fan trays, fan tray controllers, and power supply units. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See show system errors active.]
	Support for G.8273.2 and G.8275.1 profiles, hybrid mode with PTPoE (PTPoE and Synchronous Ethernet), one-step timestamping mode, and PTPoE support over LAG interoperability with child links spread across PTX10K-LC1301-36DD, PTX10K-LC1201-36CD, and PTX10000-LC1202-36MR line cards. [See Precision Time Protocol (PTP) Overview and PTP over Ethernet Overview.]
	Platform resiliency for PTX10K-LC1301-36DD . The PTX10K-LC1301-36DD line card supports platform resiliency. Resiliency includes handling faults pertaining to the line card hardware and transceivers. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See show system errors active.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
Segment routing	• • Support for SRv6 network programming in IS-IS. Use this feature to configure segment routing in a core IPv6 network without an MPLS data plane. • To enable SRv6 network programming in an IPv6 domain, include the `srv6` statement at the `[edit protocols isis sourcepacket- routing]` hierarchy level. • To advertise the Segment Routing Header (SRH) locator with a mapped flexible algorithm, include the `algorithm` statement at the `[edit protocols isis source-packet-routing srv6 locator]` hierarchy level. • To configure a Topology Independent Loop-Free Alternate (TI-LFA) backup path for SRv6 in an IS-IS network, include the `transitsrh- insert` statement at the `[edit protocols isis sourcepacket- routing srv6]` hierarchy level. See How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for SRv6 network programming and Layer 3 services over SRv6 in BGP. You can configure BGP-based Layer 3 services over an SRv6 core. You can enable Layer 3 overlay services with BGP as the control plane and SRv6 as the data plane. SRv6 network programming provides flexibility to leverage segment routing without deploying MPLS. [See **Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP**.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers ***(Continued)***

Feature	Description
	• Operations, Administration and Management (OAM) ping support for segment routing with IPv6 (SRv6) network programming. You can perform an OAM ping operation for any SRv6 segment identifier (SID) whose behavior allows upper layer header processing for an applicable OAM payload. As segment routing with IPv6 data plane (SRv6) adds only the new Type 4 routing extension header, you can use the existing ICMPv6-based ping mechanisms for an SRv6 network to provide OAM support for SRv6. Ping with O-Flag (segment header) is not supported. [See ITU-T Y.1731 Ethernet Service OAM Overview and How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for SRv6 traceroute. We support the traceroute mechanism for segment routing for IPv6 (SRv6) segment identifiers. You can use traceroute for both UDP and ICMP probes. By default, traceroute uses UDP probes. For ICMP probes, use the `traceroute` command with the `probe-icmp` option. [See How to Enable SRv6 Network Programming in IS-IS Networks.] • SRv6 support for static SR-TE policy. You can configure static segment routing–traffic engineering (SR-TE) tunnels over an SRv6 data plane. Use the following configuration commands to enable SRv6 support: • For an SR-TE policy: `set protocols source-packet-routing srv6`

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	• For an SR-TE tunnel: `set protocols source-packet-routing source-routing-path lsp name srv6` • For an SR-TE segment list: `set protocols source-packet-routing source-routing-path segment-list srv6` [See Understanding SR-TE Policy for SRv6 Tunnel.]
	Support for SRv6 micro-SIDs (uSIDs). You can compress multiple SRv6 addresses into a single IPv6 address (uSID). [See Micro SID support in SRv6, micro-sid, and block.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
Services applications	• Inline monitoring services support for packet mirroring with metadata. [See Inline Monitoring Services Configuration.] • Hardware-based IPFIX export for inline monitoring services. [See Understand Inline Active Flow Monitoring.] • Juniper Resiliency Interface (JRI) support. [See Juniper Resiliency Interface.] • HTTP and TCP probe types for RPM. You can configure the http-get, http-metadata-get, and tcp-ping probe types for real-time performance monitoring (RPM) probes. You must configure the offload-type none statement to be able to commit the configuration. [See probe-server, probe-type, and rpm.] • Inline active flow monitoring support, including support for egress sampling, for multiple BGP next-hop support, and for MPLS, MPLS-IPv4, and MPLS-IPv6 templates. [See Understand Inline Active Flow Monitoring.]
Software installation and upgrade	• Support for secure zero-touch provisioning (SZTP). [See Secure Zero Touch Provisioning.] • Support for ZTP using WAN interfaces. [See See Zero Touch Provisioning.]

Table 1: Features Supported on PTX10K-LC1301-36DD line card for PTX10008 Routers *(Continued)*

Feature	Description
	Firmware upgrade support. [See request system firmware upgrade.]
Additional feature support	Firewall filter support. [See Firewall filter support.]
	Policer and policer overhead interoperability. [See Routing Policies, Firewall Filters, and Traffic Policers User Guide.]

Chassis

- **PIC support for optical module as FRU (PTX Series)**—You can manage optical modules independently of the Packet Forwarding Engine forwarding state at the Physical Interface Module (PIC) level. This change enhances flexibility and reliability as optical modules remain accessible in the hardware inventory even if the Packet Forwarding Engine is powered off. Ports and optic modules associated with a disabled Packet Forwarding Engine remain powered off and inactive, with their LED status reflecting this state. Changes in port speed through CLI doesn't affect powered-off ports. When a Packet Forwarding Engine is taken offline, the associated ports are merely disabled, not removed. This functionality improves monitoring of optical modules even during PFE power transitions.

  Use the `show interfaces diagnostics optics` command to view detailed diagnostics, including temperature, voltage, and power status, of optical modules. The `set chassis fpc pfe power off` command manages state visibility and diagnostics for optical modules. This command also manages the operational state of the Packet Forwarding Engines. The `show chassis hardware` command now displays optical modules associated with offlined PFEs.

  [See No Link Title, No Link Title, No Link Title, No Link Title .]

- [See No Link Title.]
- [See No Link Title.]

Class of Service

- **Support for on-chip queue buffer for PFC-enabled queues (PTX10008)**—The PTX10008 views and installs a priority-based flow control (PFC)-enabled queue with a buffer-size less than 450 microseconds as a PFC-enabled on-chip queue. When a queue is in PFC on-chip mode, the entire virtual output queue (VOQ) buffer is always on-chip and is not scaled based on bandwidth usage.

  [See buffer-size (Schedulers).]

- **HCoS interoperability support (PTX10008 with the Express 5-based PTX10K-LC1301-36DD line card)**—By default, PTX10008 routers boot with Express 5-based PTX10K-LC1301-36DD line cards in interop mode. This mode does not support hierarchical class-of-service (HCoS) configuration on the PTX10K-LC1301-36DD line card. To enable HCoS on this line card, run the `set chassis interoperability express5-enhanced` command. Commit the change and reboot the router. After the reboot, the PTX10008 with the PTX10K-LC1301-36DD line card supports HCoS configuration.

  [See Hierarchical Class of Service Overview.]

- **Per-queue accounting of ECN packets (PTX10002-36QDD and PTX10008)**—You can track explicit congestion notification (ECN)-marked traffic with ECN-CE packet and byte counters in queue statistics, helping you monitor congestion and validate ECN behavior. Enable support by setting `class-of-service options hierarchical-scheduler-disable`, and then reboot to operate in non-HCoS mode.

  [See ECN Packets per Queue and show interfaces queue.]

Device Security

- **IMA coverage update (ACX Series, PTX Series, and QFX Series)**—Integrity Measurement Architecture (IMA) coverage now includes the following additional file systems:

  - ISO9660
  - PROC
  - SYSFS
  - DEBUGFS
  - RAMFS
  - SECURITYFS
  - EFIVARFS
  - DEVPTS
  - BINFMTFS
  - SELINUX

- CGROUP
- NSFS
- TRACEFS

IMA now enforces signature verification for the `kexec` kernel and `initramfs` images. It also generates a nonrepudiable log for new key addition events to IMA keyrings. These enhancements strengthen runtime integrity protections against unauthorized changes to Junos OS Evolved.

[See File Security with IMA.]

EVPN

- **Convergence improvements for EVPN-VPWS (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We've introduced EVPN-VPWS convergence enhancements on the ACX Series and PTX Series routers. With this feature, the router now supports route acknowledgment for EVPN-VPWS routes, extending the existing mechanism currently used for BGP routes. The router advertises an EVPN-VPWS route to its neighbor only after it confirms that:

  - The route has been successfully installed in the forwarding table of the Packet Forwarding Engine
  - The route version in the routing information base (RIB) matches the version in theforwarding information base (FIB) .

  If there are any acknowledgment errors, the router withdraws the EVPN-VPWS routes from the Packet Forwarding Engine. This ensures that EVPN-VPWS routes are reliably installed in hardware and that the versions are synchronized across the RIB and the FIB.

  By default, required acknowledgment from the PFE prior to advertising the EVPN-VPWS route is automatically enabled. To disable the acknowledgment requirement, include the `evpn-vpws-ack-disable` statement at the `[edit routing-options forwarding-table]` hierarchy level. Disabling the acknowledgment requirement only affects new EVPN-VPWS instances and does not impact the network.

  [See forwarding-table.]

- **Lightweight PE-CE loop detection in EVPN-VXLAN fabrics (PTX10002-36QDD)**—You can enhance network reliability on EVPN-VXLAN fabrics by implementing lightweight PE-CE loop detection. This feature assists in preventing routing loops between provider edge (PE) and customer edge (CE) devices, ensuring efficient data flow and maintaining network integrity. By detecting loops early, you minimize disruptions and optimize network performance, supporting stable and consistent connectivity.

  [See EVPN-VXLAN Lightweight Leaf to Server Loop Detection.]

- **OISM with a default IRB routing instance (PTX10001-36MR)**—You can now configure optimized intersubnet multicast (OISM) using a default Layer 3 (L3) integrated routing and bridging (IRB) routing instance. We support this feature with:

  - IPv4 underlay peering in the EVPN-VXLAN network.
  - IPv4 multicast traffic with IGMPv1/v2 and IGMP snooping.

  With a default L3 routing instance, you configure L3 OISM parameters at the global level instead of at a VRF instance level. For example, configure the L3 EVPN protocol statements at the `[edit protocols evpn ...]` hierarchy level instead of at the `[edit routing-instances` *`name`* `protocols evpn ...]` hierarchy level. Also, to prevent routing loops when you use a default routing instance, you must set a forwarding table export policy that includes the `install-nexthop except overlay-vxlan-interfaces` policy option.

  [See OISM Support with Tenant L3 VRF Instances or the Default L3 Routing Instance.]

- **Weighted ECMP for EVPN-MPLS Type 5 routes (PTX10002-36QDD, PTX10004, and PTX10008)**—Use weighted ECMP (WECMP) for EVPN-MPLS Type 5 routes by advertising generalized weights from gateway counts with the EVPN link bandwidth extended community. This advertisement distributes traffic across next hops and prevents congestion when leaf capacities differ. To configure an export policy, use `set policy-options policy-statement` *`policy-name`* `then aggregate-weight [multiplier one through 32]` and `set routing-instances` *`instance-name`* `protocols evpn ip-prefix-routes export` *`policy-name`*. Use `show evpn ip-prefix-database extensive` to verify weights, balance factors, and feature state. If weights are missing or inconsistent from the Type 5 route of any leaf device, the border leaf device generates a debug log and follows regular ECMP forwarding to the leaf devices.

  [See `policy-statement`.]

High Availability

- **Support for node-by-node upgrade by disabling GRES (PTX10004, PTX10008, and PTX10016)**—By default, Junos OS Evolved upgrades the entire cluster during a maintenance window, including both Routing Engines and all Flexible PIC Concentrators (FPCs). You can now upgrade the device one Routing Engine at a time by temporarily disabling GRES, upgrading the hardware backup Routing Engine first, switching over to it, then upgrading the new backup Routing Engine and re-enabling GRES.

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

Interfaces

- **Application Selection on 400ZR and 400 G OpenZR optics modules (PTX10003)**—Use application selection to configure different operational modes and optimize performance for your network. You can define how the optics behaves regarding reach, capacity, and line system compatibility. Configure application selection with the `set interfaces <interface> optics-options application hostid <hostid> mediaid`

`<mediaid> [domainid <domainid>]` command under the `interface optics-options` hierarchy. See Features of 400ZR and 400G OpenZR+.

- **Support for 800 G Open ZR+ pluggable modules (PTX10008)**—Enhance data center and infrastructure connectivity with high-capacity 800G-ZR+ pluggable modules. These modules support multiple optical modes and high-performance modulation formats, enabling transmission to reach up to 450 km at 800 Gbps. PTX 10008 LC1301 does not support the Nx200G speeds (where N is one to four). The optics modules also support features such as application selection, wavelength configuration, optical loopback, and configuration of target output power. See 800ZR and 800G OpenZR+ Optical Transceivers.

- **WAN interface support on PTX10K-LC1301 (PTX10008)**—You can configure the WAN interface MTU up to 16000 bytes for transit traffic. The device supports host traffic up to 9500 bytes, reducing fragmentation and enhancing throughput. The default MTU is 1514 bytes. The device supports a minimum packet size of 597 bytes at line rate.

  Configure the interface using the `set interfaces <IFD name> mtu <MTU>` command. Verify the value using the `show interfaces et-0/0/11` command

- **Extension of ae interface ID range (PTX10002-36QDD and PTX10008)**—You can configure aggregate Ethernet interface names in the range 0–999999, expanding the previous upper bound of ae4091. You gain naming flexibility for large deployments without changing ae bundle behavior or protocols. The supported ae bundle count per platform remains unchanged.

  Configure the interface using the `set interfaces ae40000` command to adopt broader numbering schemes consistent with your operational conventions.

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, and QFX5241-64QD)**—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:

  - Slax processor enhancements including a new mode, additional options, and simplified argument parsing
  - New libslax extension library functions
  - Improved SLAX syntax options
  - New SLAX functions and enhancements to existing functions and statements
  - Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See libslax Distribution Overview.]

Junos Telemetry

- **Export FIB data to external collectors using Junos telemetry (PTX10008, PTX10002-36QDD)**—Use Junos telemetry to stream or export Forwarding Information Base (FIB) statistics to external collectors. Telemetry data is streamed in "ON_CHANGE" mode. This feature supports the OpenConfig YANG model OC-AFT. To enable and manage FIB streaming, use the following CLI commands on the client device:

  - To enable tracing, use the `set system trace application` *`application-name`* `node` *`node-name`* `level` *`level`* command at the `[edit system]` hierarchy. For example, `set system trace application` *`fibd`* `node` *`re0`* `level` *`info`*.

  - Use `set system fib-streaming` and `delete system fib-streaming` statements at the `[edit]` hierarchy level to launch or terminate the process.

  - Use the `set system fib-streaming traceoptions file` *`file-name`* statement at the `[edit]` hierarchy level to configure a logging file.

  - Use the `set system fib-streaming traceoptions flag` *`flag-name`* statement at the `[edit]` hierarchy level to configure various trace parameters.

  - Use the `set system fib-streaming traceoptionslevel` *`level-name`* statement at the `[edit]` hierarchy level to configure log levels.

  - Use the `restart fib-streaming` command to restart the process.

  To view FIB streaming information, use the following operational mode commands:

  - `show fib-streaming`

  - `show fib-streaming next-hop-groups`

  - `show fib-streaming next-hops`

  - `show fib-streaming routes ipv4-unicast`

  - `show fib-streaming routes ipv6-unicast`

  - `show fib-streaming routes mpls`

  To view the trace data for the specified application use `show trace application` *`app-name`* command. For example, `show trace application` *`fibtd`*.

Navigate to the sensor path */network-instances/network-instance/afts/* on the Junos YANG Data Model Explorerto view all the supported sensors. For information on telemetry modes, see Telemetry Modes. For sensor and sensor path information, see Junos YANG Data Model Explorer.

[See Junos YANG Data Model Explorer, No Link Title, No Link Title, and No Link Title.]

- **Sensor path support for the table connection `disable-metric-propagation` leaf (PTX10001-36MR, PTX10008, PTX10016)**—By default, the OpenConfig protocol sets the destination protocol metric based on the source protocol metric. Set the `disable-metric-propagation` leaf to true to stop this behavior. The device then sets the metric to zero or a policy-defined value. Subscribe to the following resource paths to obtain sensor-specific information:

  - `/network-instances/network-instance/table-connections/table-connection/config/disable-metric-propagation`
  - `/network-instances/network-instance/table-connections/table-connection/state/disable-metric-propagation`

  For more information, see Junos YANG Data Model Explorer.

- **OpenConfig IS-IS enhancements including graceful restart and maximum ECMP paths (PTX10004, PTX10008, PTX10016)**—To achieve compliance with OpenConfig data models for IS-IS protocols, configure the missing leaves for graceful restart and maximum ECMP paths. This step ensures remote management across different vendors without requiring device login, by using specific CLI commands to set configurations that include OpenConfig network instances and graceful restart options.

  For more information, see Junos YANG Data Model Explorer.

- **Export 400G ZR and ZR+ transceiver, terminal device, and optics data to external collectors using Junos telemetry (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Use Junos Telemetry to stream terminal device, transceiver, and optical channel (state and performance) statistics from PTX10001-36MR, PTX10004, PTX10008, and PTX10016 devices to external collectors. This information facilitates enhanced performance management and monitoring. To configure a telemetry subscription to deliver data to your collector, see Telemetry Data Subscriptions over gNMI and Establish a Dial-in Telemetry Connection.

  Subscribe to the following resource paths for terminal data, transceiver, and optical channel statistics:

  - `/components/component/optical-channel/config/operational-mode`
  - `/components/component/optical-channel/state/operational-mode`
  - `/components/component/transceiver/physical-channel/channel/state/output-power/max`
  - `/components/component/transceiver/physical-channel/channel/state/output-power/min`
  - `/components/component/transceiver/physical-channel/channel/state/output-power/avg`
  - `/components/component/transceiver/physical-channel/channel/state/output-power/interval`

- /components/component/transceiver/physical-channel/channel/state/output-power/max-time
- /components/component/transceiver/physical-channel/channel/state/output-power/min-time
- /terminal-device/logical-channels/channel/config/loopback-mode
- /terminal-device/logical-channels/channel/state/loopback-mode

View the complete list of supported sensors in the Junos YANG Data Model Explorer.

[See Establish a Dial-in Telemetry Connection and Junos YANG Data Model Explorer.]

- **Stream telemetry data in gNMI-based message format over UDP (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, and QFX5240-64QD)**—Junos OS Evolved uses a dial-out mechanism to send telemetry data to a collector over UDP . The message format is defined in the `jnx_gnmi_over_udp.proto` file. Only `STREAM` mode with `SAMPLE` as subscription mode is supported. The message contains full key name and value pair information so the collector does not require data models for processing or consuming the telemetry data.

  [See No Link Title, No Link Title, No Link Title, No Link Title, No Link Title, and Junos YANG Data Model Explorer.]

MACsec

- **Automatic adjustment of MTU for MACsec overhead (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Use this feature to automatically adjust the maximum transmission unit (MTU) for the Media Access Control Security (MACsec) overhead. Without this feature, you must adjust the interface MTU and the protocol MTU manually.

  Use this feature to ensure the interface or protocol MTU is adjusted properly to account for the MACsec overhead. This feature is disabled by default. To enable this feature, first enable MACsec. Then configure the `enable-auto-mtu-update` statement at the `[edit security macsec]` hierarchy level. This feature applies to physical interfaces, logical interfaces, and physical interfaces that are members of aggregated Ethernet interfaces.

  [See Media MTU and Protocol MTU.]

- **MACsec support during GRES and NSR (PTX10004, PTX10008, and PTX10016)**—The GRES feature enables a switch or router with redundant routing engines to continue forwarding packets, even if one Routing Engine fails. Nonstop active routing (NSR) is an enhancement on GRES that does not rely on helper routers (or switches) to assist the routing platform in restoring routing protocol information. You can configure MACsec to provide uninterrupted MACsec service and secure your traffic during a Routing Engine switchover.

[See Configuring Advanced MACsec Features.]

- **Support for a custom EAPoL EtherType to improve network tunneling of MACsec packets for Layer 2 traffic (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**—MACsec uses Extensible Authentication Protocol over LAN (EAPoL) as a transport protocol to establish sessions. Some networks filter packets based on their EtherType value. By default, the EtherType for all EAPoL packets is 0x888e. To ensure the network tunnels the MACsec packets properly, you can set a custom EtherType for EAPoL packets.

  To configure an EAPoL profile with a custom EtherType, use the `ether-type` *`ether-type-value`* statement at the `[edit forwarding-options custom-eapol-ether-type-profiles (EAPOL_ETHERTYPE1 | EAPOL_ETHERTYPE2)]` hierarchy level. By default, the EtherType value for the EAPOL_ETHERTYPE1 profile is 0x876f and the EtherType value for the EAPOL_ETHERTYPE2 profile is 0xb860. If you configure a different value, you must use an EtherType that isn't already reserved for another use. To apply the EtherType to MACsec packets, configure the `eapol-ethertype-profile` *`eapol-profile-name`* statement at the `[edit security macsec connectivity-association` *`ca-name`* `mka]` hierarchy level.

  [See Media Access Control Security (MACsec) over WAN, custom-eapol-ethertype-profiles, and mka.]

Network Management and Monitoring

- **Ephemeral database default commit synchronize model changed to synchronous (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We've changed the default commit synchronize model for the ephemeral database from the asynchronous model to the synchronous model. The synchronous model ensures better synchronization of ephemeral configurations across Routing Engines or Virtual Chassis members by processing commits synchronously. With this change, only the synchronous model supports synchronizing ephemeral data on devices that have graceful Routing Engine switchover (GRES) or nonstop active routing (NSR) enabled.

  [See Understanding Ephemeral Database Commit Synchronize Models.]

- **Proactive resource monitoring and alarms for data plane stability (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—The Resource Monitoring Infrastructure on Junos OS devices enables proactive and continuous monitoring of critical hardware resources, including Chash, Jencap, Counters, Firewall, and Datapath components. Threshold limits in platform settings ensure real-time detection of abnormal usage without relying on external telemetry collectors. If resource utilization exceeds configured thresholds, the system generates syslog messages, raise alarms through the CMerror module, and forwards events to platform daemons. This infrastructure stabilizes the data plane and reduces traffic loss.

  To view CMerror information, use the `show system errors active` command.

  [See show system errors active.]

- **Ingress and egress sFlow support for L2 traffic (PTX10002-36QDD)**—Use sFlow to monitor ingress or egress data and control traffic forwarded through L2 logical interfaces associated with a bridge domain (BD). This feature supports both enterprise and service provider style L2 configurations.

  [See sFlow Technology Overview.]

- **AES-256 Encryption Algorithm Support for SNMPv3 (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can configure Advanced Encryption Standard (AES) 256 algorithm for an SNMPv3 user. To configure AES-256 algorithms for an SNMPv3 user, include the `privacy-aes256` statement at the `edit snmp v3 usm local-engine user username` hierarchy level. AES-256 is a symmetric encryption algorithm that uses a 256-bit key to encrypt or decrypt messages and provides high-level security for protecting sensitive information.

  [See Configure SNMPv3 Authentication Type and Encryption Type, show snmp v3, and usm.]

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, PTX Series, and QFX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):

  - ML-DSA-87 PQC algorithm for digital signatures
  - SHA-512 for hashing

  For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

  [See PQC Signatures for Software Images.]

- **Support for Quantum Buffer in SSH (ACX Series, PTX Series, and QFX Series)**—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptoagility. The feature uses finite field cryptography (FFC) to extend the security life span of the current systems against quantum attacks. Quantum Buffer provides a phased approach to adopting post-quantum cryptography (PQC), thereby mitigating operational risks associated with the transition.

  To enable the feature, configure the following command:

  - `set system services ssh moduli type` *`name`* `refresh` *`frequency`* `count` *`count`*

    The configuration dynamically generates prime moduli for existing Diffie-Hellman (DH) group exchange algorithms, `group-exchange-sha1` and `group-exchange-sha2`. The `qbufd` process is responsible for generating the moduli.

[See Quantum Buffer and moduli.]

- **Support for Shor-resistant and other default key exchange algorithms in SSH (ACX Series, PTX Series, and QFX Series)**—SSH supports the hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm, which is Shor-resistant and improves protection against quantum attacks.

  Configure `sntrup761x25519-sha512` at the [`edit system services ssh key-exchange`] hierarchy level.

  Additionally, SSH includes default support for the following Diffie-Hellman (DH) group key exchange algorithms that are available at the [`edit system services ssh key-exchange`] hierarchy level.

  - `dh-group16-sha512`
  - `dh-group18-sha512`

  [See key-exchange.]

Precision Time Protocol (PTP)

- **Transparent clock support (PTX10008 with JNP10K-LC1301 line cards)**—Use the transparent clock feature to improve synchronization between the timeTransmitter and timeReceiver clocks. Transparent clocks ensure that the timeTransmitter and timeReceiver clocks are not affected by packet delay variation, enhancing overall network performance and reliability.

  [See PTP Transparent Clocks.]

Routing Policy and Firewall Filters

- **Support for new options for the show route command (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)**—We've introduced enhancements to the show route command to enable network operators understand and troubleshoot routing tables more easily. Use the `prefix-length-distribution` option to display counts of prefix lengths across routing tables for each instance. You can run the `show route` *`destination`*`covering-subnets` command to walk up the radix tree and list all routes covering a destination. These options provide deeper insight for troubleshooting and optimize routing behavior.

  [See show route.]

- **Increased firewall filter scale (PTX10002-36QDD)**—You can apply the `prefix-scale-mode` configuration on a firewall filter to increase its scale to up to 225,000 prefixes per term, and in total up to 1 million prefixes for IPv4, 512,000 for 64-bit IPv6, and 256,000 for 128-bit IPv6 as match conditions.

  [See prefix-scale-mode.]

- **Apply selective class-based filtering without enabling SCU accounting and/or DCU accounting (PTX10002-36QDD, PTX10004, and PTX10008)**—Use the `class-based-filtering`configuration to filter traffic without enabling source class usage (SCU) or destination class usage (DCU) accounting.

Applying this configuration improves line-rate performance because SCU and/or DCU lookup is not performed for SCU and/or DCU accounting and only selective class-based filtering is performed.

[See class-based-filtering.]

Routing Protocols

- **AS loop check in BGP Networks (PTX10016)**—We have enabled Autonomous systems (AS) path loop check for external BGP (EBGP) and internal BGP (IBGP) sessions by default. The loop check is made in the BGP peer's AS path domain. Use this feature to configure and manage AS path loop detection on Junos devices.

  You can disable AS path loop check for IBGP sessions including all routing instances using the statement `no-loop-check` statement at the `[edit protocols bgp defaults ibgp]` hierarchy level.

  [See no-loop-check.]

- **Configure BGP keepalive value independent of holdtimer value (PTX10001-36MR)**— BGP uses hold time to terminate unresposive sessions, which is reset each time a BGP message is received on a BGP connection. In the absence of any BGP message a keepalive timer runs that triggers a keepalive message type when it expires. By default, this keepalive timer is one third of the negotiated BGP hold time, if greater than zero. You can now configure the keepalive timer independently of the hold time between 1 second through 21845 seconds. Include the `keepalive-time` statement at the `[edit protocols bgp]` hierarchy level.

  [See keepalive-time .]

- **Generate static RT-Constrain route based on community/wildcard (PTX10016)**—When the RT-Constrain feature is partially deployed in a network, the resource saving benefit is lost. We have extended the static RT-Constrain feature to generate host static RT-Constrain entries from fully qualified route targets configured in the routing policy. You can assign BGP communities or a wildcard route target on the static RT-Constrain route. You can also configure the static RT-Constrain route's origin AS in the NLRI while retaining the global AS number.

- **IS-IS multi-instance support over a single interface (ACX7020, ACX7024, ACX7100, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, QFX5130, QFX5220, QFX5230-64CD, QFX5240, QFX5241-32OD, and QFX5700)**—We have enhaced the IS-IS multi-instance feature to support multiple IS-IS instances on the same logical interface with instance identifier TLV 7.

  Include the `instance-id` statement at the `[edit protocols isis-instance` *`name`* hierarchy level.

  [See How to Configure Multiple Independent IGP Instances of IS-IS.]

- **Selectively disable NH validation based on Community/RT/RD (PTX10016)**—Define a BGP import policy to selectively disable next hop resolution. The policy action sets the next hop to fictitious

instead of indirect next hop and avoids resolving the next hop for routes that match the community specified in the policy.

- **BGP prefix limit based on route target to limit VPN prefixes (PTX10016)**—Typically L3VPN deployments limit routes at the customer edge peer level with the `prefix-limit` configuration for a BGP peer family. We have shifted this control to a central location such as the route reflector or a ASBR so that routes originating at all sites in a VPN are taken into account. BGP maintains and enforces the prefix limit as specified by the route target communities originating at various VPN sites to limit the number of prefixes a BGP peer can advertise or receive to conserve resources.

  [See prefix-limit.]

- **Replace BGP AS path to maintain network interoperability (PTX10008)**—Define a routing policy to match and replace a list of autonomous systems (AS) numbers or private AS numbers with the local AS number of the BGP peering session to maintain network interoperability. This configuration works only on AS sequences and not on AS sets. In addition to using external BGP (EBGP), enable internal BGP (IBGP) to leverage this capability in route-reflector scenarios. Include the policy action `as-path-replace as-list | private` statement at the `[edit policy-options policy-statement statement-name then]` hierarchy level to activate the feature.

  [See Autonomous Systems for BGP Sessions.]

Services Applications

- **Inline active flow monitoring using IP Flow Information Export (IPFIX) and version 9 templates for IPv6 BGP next-hop addresses within IPv4 data templates** (PTX10002-36QDD). We support including Information Element 63, IPv6 BGP NextHop Address, in the IPv4 and MPLS-IPv4 templates. This support is to enable the correct reporting of the BGP next-hop address for the ingress direction of sampling enabled on an SRv6 tunnel's head end node.

  [See Understand Inline Active Flow Monitoring.]

Software Installation and Upgrade

- **Separate stage and activate options for software installations (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—You have the option to install software in separate staging and activation phases, which provides finer control over software management. The `request system software add stage` command enables you to install a new software image without setting it as the next boot version, supporting multiple installations. The `request system software stage` command enables you to validate and save the active configuration for an installed software image. Additionally, you can activate an installed software image at any time by using the `request system software activate` command, which sets the image as the next boot version. By staging and validating software images in advance, you can shorten the maintenance window for software upgrades, which is reduced to activating the image and rebooting the system.

[See Stage and Activate Junos OS Evolved Software, request system software add, request system software stage, and request system software activate.]

- **Secure boot and common BIOS support (JNP10008-SF5 switch interface board (PTX10008))**—The Secure boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms.

  [See Junos OS Evolved Overview and request system firmware upgrade (Junos OS Evolved).]

- **Include certificates in bundled ISOs (PTX10002-36QDD and PTX10008)**—You can include certificates in bundled ISOs so that the certificates required for custom packages inside the ISO are installed automatically. To install a bundled ISO that includes certificates, use the `request system software add restart` **`allow-bundled-keys`** *`bundled-iso-path`* command:

  [See Creating a Bundled ISO.]

- **Load set-formatted and XML-based configuration files for ZTP (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5140-24CD8O, , QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, QFX5700, and QFX5700E)**—You can load set-formatted or XML-based configuration files when your device provisions for zero-touch provisioning (ZTP). Reuse existing set-style or XML-based configuration files for automated onboarding to avoid converting them to hierarchical syntax. Provide the configuration file in set format or in XML and specify the configuration file name under the DHCP vendor configuration options.

  [See Zero Touch Provisioning.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for S-BFD over SRv6-TE paths with Classic and Micro SIDs (PTX10002-36QDD, PTX10004, PTX10008, and PTX12008)**—You can validate and control Segment Routing for IPv6—Traffic Engineering (SRv6-TE) path liveness with Seamless-BFD (S-BFD) for Classic and Micro Segment Identifiers (SIDs). S-BFD helps improve path convergence and reliability by ensuring that a SRv6-TE path is usable only when its S-BFD session is up. S-BFD sessions run in distributed mode and support non-stop routing (NSR) and GRES.

  To interoperate with responders that expect the IPv6 local-host destination, use the `set protocols source-packet-routing source-routing-path` *`lsp-path-name`* `primary` *`segment-list-name`* `bfd-liveness-detection sbfd destination-ipv6-local-host` configuration statement.

  You can display SRv6 S-BFD session details with the `show spring-traffic-engineering sbfd` command.

[See S-BFD for SRv6 TE Paths.]

- **Color-based forwarding (CBF) for SRv6-TE (PTX10002-36QDD, PTX10004, PTX10008, PTX10016, and PTX12008)**—We extend the existing color-based forwarding (CBF) functionality to SRv6 traffic engineering enabling you to adapt to complex network demands. Use CBF for SRv6-TE to steer traffic across multiple transport tunnels based on class of service (CoS). This approach improves route selection, next-hop resolution, and service quality. Resolver enhancements support SRv6 Segment Identifiers (SID) across multipath routes, and the `preserve-next-hop-hierarchy` configuration prevents misrouting.

  [See preserve-nexthop-hierarchy (SRv6-TE).]

- [See profile-sharing.]

- **Support for UHP in IS-IS SR-MPLS (ACX7020, ACX7100, ACX7332, ACX7348, ACX7509,ACX7024, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, and PTX12008)** —Use Ultimate Hop Popping (UHP) with IS-IS or OSPF so the egress provider edge (PE) can process its own node SID. ISIS advertises a node SID with the P flag set and E flag unset. In controller-driven segment routing traffic engineering (SR-TE) the controller inserts the egress PE node SID beneath the SR-TE binding SID. If the Binding SID route fails on the penultimate hop, the egress PE might see its own node SID as the top label instead of penultimate hop popping (PHP). With the P flag set, the PE expects UHP and processes its MPLS label. Include the `ultimate-hop-popping` statement at the `[edit protocols isis source-packet-routing]` hierarchy level.

  [See ultimate-hop-popping.]

- **SRv6-TE route resolution over BGP without IGP (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**—An SRv6 tunnel consists of paths with segment identifiers (SIDs) over IGP that steer traffic to a traffic-engineering (TE) path. If IGP is not available, configure these SIDs statically and advertise them through External BGP (EBGP). This feature is supported on both classic and micro SRv6 SIDs

  Networks that deploy network orchestrator to steer transit traffic onto a TE path and advertise these transit prefixes using BGP color community don't have a service SID. In this case, the last SID must not be removed, and the ingress SRv6 TE tunnel acts as a transit tunnel to forward transit traffic with SRv6 encapsulation.

  Include the `no-remove-srv6-last-sid` statement at the `[edit protocols source-packet-routing]` hierarchy level and the `use-ingress-routes-as-transit` statement at the `[edit protocols source-packet-routing srv6]` hierarchy level.

  [See no-remove-srv6-last-sid.]

- **Delay normalization for OSPF Flexible Algorithm metrics and advertisements across IGP instances (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD,**

PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)—Use delay normalization to compute and advertise a normalized delay metric for Flexible Algorithm, to improve path-selection consistency across all IGP instances. The device normalizes each received delay, compares each value with the previously saved normalized value, and triggers link-state advertisement (LSA) generation when the values differ.

Delay normalization is disabled by default. To enable and configure delay normalization, use the `normalize` *`interval offset`* statement at the `[edit protocols ospf` *`area interface`* `delay-measurement]` hierarchy level.

[See delay-measurement (Protocols OSPF) and How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering.]

- **Selectively control per-prefix backup paths with OSPF import policy (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can selectively enable backup paths for specific prefixes to optimize redundancy and resource utilization. By default, a configured backup path applies to all prefixes. To exclude specific prefixes or ranges, create an OSPF import policy and configure the `no-backup` option in the `then` clause of the policy to suppress backup path installation for matching routes. You can reserve backup protection for critical prefixes while preventing unnecessary backups for others.

  [See Understanding Backup Selection Policy for OSPF Protocol.]

- **Preference-based Path Selection of L-OSPF Flexible Algorithm routes (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can control path selection by configuring the preference for L-OSPF Flexible Algorithm routes in `inetcolor.0` and `mpls.0`.

  Configure `flex-algorithm-preference` statement at the `[edit protocols ospf]` hierarchy level to prioritize desired routes and improve traffic engineering across IP and MPLS domains.

- **Policy-based redistribution of OSPF prefix SIDs across IGP instances (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can redistribute Segment Routing (SR) prefix-SIDs across OSPF IGP instances using route policy without explicitly specifying a prefix-segment index. This feature standardizes SR labels across instances and improves operational efficiency. Configure a policy with the `from prefix-segment` statement to match routes carrying prefix-segment information. In the `then` clause, use `prefix-segment redistribute` to inherit segment information from the matched route. We also support stitching `mpls.0` routes to enable interoperability between different IGP instances.

- **Non-router-ID endpoints as SR-TE destinations (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR,**

PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—Use non-router-ID endpoints as destinations in Segment Routing—Traffic Engineering (SR-TE) policies for Segment Routing for MPLS (SR-MPLS). Traditionally, these policies use router IDs, but you can specify anycast addresses to enhance redundancy and load balancing in SR-MPLS networks. Use IPv4 and IPv6 anycast addresses as IGP-learned destinations with or without Segment Identifier (SID) stack compression. These anycast addresses are not redistributed (R-bit set). Use them as the to address for SR-TE policies with associated compute profiles..

[See Non-Router-ID Endpoints in Segment Routing Traffic Engineering.]

- **Static SID configuration in SRv6 Manager (PTX10002-36QDD, PTX10004, PTX10008, PTX10016, and PTX12008)**—Configure SRv6 classic, micro node, adjacency SIDs, along with classic END and END-X SIDs, and install them in the routing table without using interior gateway protocols (IGPs) such as IS-IS. Advertise these static routes through a BGP export policy for path computation. This configuration enables controllers to receive static node and adjacency SIDs over BGP and compute SR-TE paths across a domain that does not use IS-IS.

  [See Understanding SRv6 Static Segment Identifier.]

Additional Features

We've extended support for the following features to the platforms shown in parentheses:

- **CoS interface telemetry support** (PTX10008). Support for gRPC Network Management Interface (gNMI) streaming of CoS interface queue statistics. To stream statistics, use the resource path **/qos/interfaces/interface/output/queues/queue/state/**.

  [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

- **On-box aggregation support** (PTX10008)

  [See Junos YANG Data Model Explorer.]

- **OpenConfig QoS operational state sensors** (PTX10008)

  [See Telemetry Sensor Explorer.]

- **IGMP snooping and MLD snooping** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards). The device supports the features while operating in either of the following chassis modes:

  - (Default mode) Both Express 5 and Express 4 line cards interoperating on the device, called interop mode, where device operation is limited to the functions and scale of the Express 4 line card.

- Only Express 5 line cards operating on the device. In this case, to enable the device to operate using the enhanced functions and scale available on the Express 5 line cards, enable the enhanced Express 5 mode on the device as follows:

  ```
  set chassis interoperability express5-enhanced
  ```

  You must reboot the device for this setting to take effect.

Note that Express 5 line cards don't support multicast snooping route counters, which record the number of packets that use a snooping route toward the destination. However, Express 4 line cards do update the snooping route counters. As a result, on devices running in interop mode, the packet statistics reported in the `show multicast snooping route extensive` CLI command might not match the actual packet count for snooping routes.

[See IGMP Snooping Overview and Understanding MLD Snooping.]

- **Interoperability support for EVPN-VPWS service** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD, PTX10K-LC1202-36MR line cards). The PTX10008 router supports the deployment of EVPN-VPWS on the PTX10K-LC1201-36C, PTX10K-LC1202-36MR, and PTX10K-LC1301-36DD line cards within the same chassis. The control word is automatically enabled and must remain enabled on all three line cards. We also do not recommend configuring a custom tag protocol identifier (TPID) on the PTX10K-LC1301-36DD line card because the PTX10K-LC1201-36C and PTX10K-LC1202-36MR line cards do not support custom TPID and the router might drop these packets.

  [See Overview of VPWS with EVPN Signaling Mechanisms.]

- **Interconnecting EVPN-VXLAN data centers with EVPN-MPLS in WAN using gateway nodes** (PTX10002-36QDD).

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

- **OISM for IPv4 multicast traffic in EVPN-VXLAN fabrics** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards). The device supports the feature while operating in either of the following chassis modes:

  - (Default mode) Both Express 5 and Express 4 line cards interoperating on the device, called interop mode, where device operation is limited to the functions and scale of the Express 4 line card.

- Only Express 5 line cards operating on the device. In this case, to enable the device to operate using the enhanced functions and scale available on the Express 5 line cards, enable the enhanced Express 5 mode on the device as follows:

  ```
  set chassis interoperability express5-enhanced
  ```

  You must reboot the device for this setting to take effect.

Note that Express 5 line cards don't support multicast snooping route counters, which record the number of packets that use a snooping route toward the destination. However, Express 4 line cards do update the snooping route counters. As a result, on devices running in interop mode, the packet statistics reported in the `show multicast snooping route extensive` CLI command might not match the actual packet count for snooping routes.

In either mode, OISM support on this device includes:

- Regular OISM mode only—The original symmetric bridge domains model, also called the bridge domains everywhere (BDE) model
- MAC-VRF EVPN instances with `vlan-based` or `vlan-aware` service types only
- IPv4 multicast traffic with IGMPv2, IGMPv3, and IGMP snooping
- Server leaf, border leaf, or lean spine OISM device roles
- External multicast source and receiver communication using any of the following methods:
  - Classic L3 interfaces
  - EVPN multicast VLAN (M-VLAN) integrated routing and bridging (IRB) interfaces
  - Non-EVPN IRB interfaces

[See Optimized Intersubnet Multicast in EVPN Networks.]

- **Support for multiple bridge domains and local switching in VPLS** (PTX10002-36QDD). We support the following VPLS features:
  - Multiple bridge domains within a single VPLS routing instance

    [See Introduction to Configuring VPLS.]
  - Local switching with hierarchical VPLS

    [See Local Switching (VPLS) and Example: Configuring LDP-Based H-VPLS Using a Single Mesh Group to Terminate the Layer 2 Circuits.]

- **Transport class tunnel support for EVPN** (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016). We support EVPN-VPWS, EVPN-ELAN, and EVPN E-Tree services over transport class tunnels on PTX Series devices. These devices use the unified next-hop hierarchy, which takes advantage of improvements in the RPD infrastructure to provide improved scalability and faster convergence performance. The PTX Series platform enables unified next-hop hierarchy by default.

  [See Configuring EVPN over Transport Class Tunnels.]

- **Two-Way Active Measurement Protocol (TWAMP) monitoring service (RFC 5357) hardware timestamp support to enable Flex Algo and SR-MPLS support** (ACX7020, ACX7024, ACX7100, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)

  [See the `offload-type inline-timestamping` option of the test-session statement.]

- **MVPN bud node support with the looping back interface (LBI)** (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016).

  [See MVPN mLDP bud node support with Looping back interface.]

- **Support for performance monitoring and TCA** (PTX10008). Use performance monitoring for PTX10008 routers to measure current and historical metrics. The router accumulates these metrics into 15-minute and 1-day intervals. You can configure the 15-minute interval length. You can view these metrics by using the show interfaces transport pm command. This approach helps you manage optical transport links more efficiently.

- **Firewall filter support** (PTX10008)

  [See Understanding Firewall Filter Match Conditions.]

- **Firewall filter interop support**) (PTX10008 with the Express 5-based PTX10K-LC1301-36DD line card and the Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards)

  [See Firewall Filters.]

- **Policer and policer overhead interop support** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards)

  [See Firewall Filters and Traffic Policers.]

- **Support for ZTP using WAN interfaces** (PTX10008)

  See Zero Touch Provisioning.

- **Support for Secure ZTP (SZTP) using WAN interfaces** (PTX10008)

  See Zero Touch Provisioning.

- **MPLS feature support (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards)**. [See Fast Reroute Overview Understanding BGP MultipathMPLS Applications User Guide Understanding Source Packet Routing in Networking (SPRING)TCC OverviewLayer 2 VPNs and VPLS User Guide for Routing Devicesvlan-id-list (Ethernet VLAN Circuit)vlan-id-listConfiguring VLAN Identifiers for VLANs and VPLS Routing InstancesLayer 3 VPNs Feature Guide for Routing Devices Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGPHow to Enable Link Delay Measurement and Advertising in IS-IS.]

- **Support for Packet Forwarding Engine tunnel features**(PTX10008 routers with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

  [See Components of Filter-Based Tunneling Across IPv4 Networkstunnel-end-point, Tunnel and Encryption Services Interfaces User Guide for Routing Devices, Flexible Tunnel Interfaces Overview, Understanding Filter-Based Tunneling Across IPv4 Networks, Firewall Filter Terminating Actions, Configuring a Filter to De-Encapsulate GRE Traffic, and decapsulate (Firewall Filter.]

- **Support for multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

  [See Understanding Multicast-Only Fast Reroute, pim-snooping, Point-to-Multipoint LSP Configuration,MPLS Applications User Guide, Multicast Shortest-Path Tree, Multiprotocol BGP MVPNs Overview, Understanding Next-Generation MVPN Concepts, and Understanding Next-Generation MVPN Control Plane.]

- **Local and remote port mirroring** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

  [See On-Device Packet Capture.]

- **Support for the sFlow technology** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

  [See Understanding How to Use sFlow Technology for Network Monitoring.]

- **CFM Support** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards)

- **Support for basic Layer 2 learning, bridging, forwarding, and flooding, MAC Limit Actions, and IRB Support** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards)

  See Layer 2 Bridging, Address Learning, and Forwarding User Guide, Integrated Routing and Bridging, Configuring MAC Limiting, and packet-action.

- **Support for Q-in-Q tunneling** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards)

  [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.]

- **Support for MAC address accounting for 10GbE, 40GbE, 100GbE, 200GbE, 400GbE, and 800GbE interfaces** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

- **Support for VRRP, LFM, LAG and LACP** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

  [See Understanding VRRP and Introduction to OAM Link Fault Management (LFM).].]

- **Support for Layer 3 unicast routing and forwarding, BFD, BGP flowspec signaling, and 256-way ECMP** (PTX10008 with the Express 5-based PTX10K-LC1301-36DD and Express 4-based PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards).

  [See Understanding BGP Multipath and Understanding BFD.]

- **Renaming OpenSSH implementation to JSSH** (All platforms). The OpenSSH implementation in Junos OS Evolved is renamed "JSSH" to highlight its customized nature.

  [See ssh.]

What's Changed

IN THIS SECTION

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

General Routing

- SSH key options for user account credentials. You can configure key-options *key-options* option at the `set system login user` *user* `authentication [ssh-rsa|ssh-ecdsa|ssh-ed25519] ssh key` hierarchy level.

  [See login.]

- Displays the event log of learned MAC addresses. By default mac-learning-logs are stored in UTC timestamps. To view the logs in system timezone, use the show ethernet-switching mac-learning-log use-system-timezone command. The show ethernet-switching mac-learning-log use-system-timezone command also prints the time zone abbreviations [IST, UTC, etc] in the timestamp. To view the logs in system timezone by default by using the show ethernet-switching mac-learning-log command, you need to configure the system-timezone statement at the [edit protocols l2-learning mac-learning-log] hierarchy level.

- [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.]

- When you run the request vmhost zeroize command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines.

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

- On the MPC7E-10G line card, when you configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet ports, use the speed statement at both the edit interfaces <interface name> gigether-options and edit interfaces interface <name hierarchy> levels.

- **SMAC accounting mismatch (PTX10002-36QDD, and PTX10008)**—Source MAC (SMAC) accounting over accounts the byte counter by including the L2 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]

- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked.

Forwarding and Sampling

- IPv6 packets with link-local source address dropped per RFC 4291 (PTX12008,PTX10002-36QDD)-IPv6 packets that use a link-local source address and a global destination address won't be forwarded on a PTX12008 Router or PTX10002-36QDD. This behavior complies with RFC 4291 requirements. These IPv6 packets are dropped, and the system sends an "ICMPv6 destination unreachable with code 2" error message to the packet's source.

Interfaces and Chassis

- Vlan Tagging 1. For PTX Junos OS Evolved 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 Junos OS Evolved 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.

Network Management and Monitoring

- **Ephemeral database default commit synchronize model changed to synchronous (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)**—We've changed the default commit synchronize model for the ephemeral database from the asynchronous model to the synchronous model. With this change, we've deprecated the `allow-commit-synchronize-with-gres` statement and only the synchronous model supports synchronizing ephemeral data on devices that have graceful Routing Engine switchover (GRES) or nonstop active routing (NSR) enabled.

  [See Understanding Ephemeral Database Commit Synchronize Models.]

Platform and Infrastructure

- **Tacacs authorisation support for local authentication without password**—Starting in Junos OS Evolved Release 25.4R1, you need not configure password under `edit system authentication-order` to enable password-options.

- **Commit validation for unique user IDs**—We have added support to validate the user configuration to ensure that each user is assigned a unique UID. A commit fails if duplicate UIDs are detected, ensuring stronger validation and preventing identity conflicts. Previously, a commit was successful even when multiple users shared the same UID, triggering only a warning and logging a syslog message.

User Interface and Configuration

- **Generate genstate YANG modules on Junos devices**—You can use `show system schema` operational command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

  [See show system schema.]

Known Limitations

IN THIS SECTION

Learn about limitations in this release for PTX Series routers.

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

General Routing

- When the low priority queue is oversubscribed, we may see a page timeout error under following conditions. The reason for this is detailed below . 1. If a queue is configured with shared buffer configuration we allocate interface equivalent of buffers to that queue. 2. When traffic in such queue is moving very slow or starved for a longer duration due to strict-priority-scheduling the PE ASIC marks pages staying long as Timed out. 3. Current Page time out value is ~3.5 to 7 Sec. 4. Such time out packet when reached Head of the queue will all get dropped leading to page time out count increase. 5. Since shared buffer config leads to large queue size, and if the queue is moving very

slowly, we never hit queue full if the input to such slow-moving queue is very less. 6. This is what being tested in this particular senario. 7. When that happen we will only see Page timeout and no Queue TAIL Drops. PR1581490

Open Issues

IN THIS SECTION

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

- For PTX10001-36mr, the software driver reads the voltage threshold erroneously causing the "Host 0 Voltage Threshold Crossed" alarm to present on the device. PR1592258
- Scapa : EVO CB1, SIB0 and FPC0 (JNP10K-LC1201) went to fault after system power outage in lab.PR1749793
- 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
- On all Junos OS Evolved platforms, host-originated L3 (Layer 3) traffic is marked with DSCP(Differentiated Services Code Point) value 48 [INET 110b] in the IP header, even if class-of-service host-outbound configuration is not present. Host-originated L2 (Layer 2) traffic is marked with IEEE 802.1p value 0 in the header, even if class-of-service host-outbound configuration is not present.PR1837443
- When multiple SIBs are offlined in succession, at times it can result in continuous traffic drops.PR1849563
- 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
- An LSI IFL remains in RPD even after being deleted by the interface manager daemon. It is visible in show interface routing but not in show interfaces, indicating that RPD still holds the IFL despite its removal elsewhere. rpd-agent does not send a delete message to RPD due to a reference count issue. Another daemon?likely l2ald?still holds a reference to the IFL. rpd-agent only sends the delete once all references are cleared, which doesn't happen in this case. The fix is to send a "delete pending" message from rpd-agent to RPD. RPD will treat this as a delete and remove the IFL, ensuring consistency across the system.PR1866522
- A PTX10003-160C with high protocols scaling number, its FPC might not restart properly. An example of a high scaling number is a L2VPN gateway with more than 9000 outbound labels.PR1881324
- In Junos OS Evolved, ARP resolution requests are throttled on FPC per logical interface level that is if resolution request on a logical interface expires then a throttle timer is started on that logical interface and no other resolution request could be generated on that logical interface when throttle timer is running. In this issue RE netstack is sending packets to FPC with hint to resolution request for already resolved IP address and triggering throttle timer. So resolution request for second IP address on same logical interface could not be generated for some time which is triggering phone-home application 10 seconds time out.PR1883158
- The aggregate DDOS statistics for ISIS and BFD might not reflect accurate values due to inherent design limitations.PR1887330
- For any given filter term, the maximum term size is limited to 512KB. When the cumulative size of an individual term?including match conditions, match values, and actions?exceeds this limit, the firewallD compiler cannot process the filter, resulting in an error and preventing the filter from being programmed in the PFE. The actual size of each term depends on various combinations of the configured match conditions and values. Therefore, it is not possible to provide a specific size limit that would prevent these issues, as the required size varies significantly based on the particular configuration.PR1906528

Resolved Issues

IN THIS SECTION

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

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

General Routing

- Adjust gNMI SetRequest Update handling for list keys + value payload. PR1691474
- IS-IS Level 2 disabled after upgrade due to OpenConfig YANG model change. PR1798178
- Non-Impactful CFMMAN Parse Error on First CFM CCM PDU During FPC/PFE Online Transition. PR1810549
- The l2cpd-agent crashes due to its state cleaned up during shutdown. PR1852221
- Ad hoc swapping of 400G DAC with 400G DR4 (or vice-versa) optic leads to traffic loss. PR1862711
- AE re-anchoring Impacting CFM on Untagged AE Interfaces. PR1864491
- System stops responding to `show system applications app ndp detail` command or prevents applications from start/stop/restart. PR1866988
- Software upgrade aborts or FPC does not come online on specific Junos OS Evolved PTX platforms while upgrading with an option validate-restart or add restart. PR1867902
- Junos OS Evolved: OS command injection vulnerabilities fixed (CVE-2025-60006). PR1870684
- 60A DIP switch setting on 3KW DC PSU limits output power capacity and system could shutdown based on system power usage. PR1870875

- Ethernet-switching flood filter stops working when policer action is added. PR1877486
- Restarting the evo-aftmand-bt or evo-cda-bt process disrupts PHY synchronization within the timingd service resulting in timing errors. PR1878029
- Link flapping occurs on stable subinterfaces when using QDD400GZR optics if any one subinterface in the channelized group is down. PR1878198
- Duplicate entries are observed on Junos OS Evolved platforms when subscribing leaf level sub-interface path. PR1879832
- bf_crcv_x_intr_cru_err_pg_psc transient interrupt may occur during FPC offline. PR1880275
- Firewall policy configured to match IP payloads fail matching on MPLS packets. PR1882315
- Packet drops are observed when MACsec with bounded-delay is configured due to key rollover. PR1883473
- IGMP or MLD packets associated with CCC services will be dropped instead of being forwarded. PR1885670
- Few telemetry paths are not exported after router reboot. PR1886043
- In scale EVPN scenario with IRB, host traffic via IRB will be impacted ICMP, OSPF, BGP and L3 protocols on Junos OS Evolved PTX series platform. PR1890125
- On PTX Junos OS Evolved platforms in EVPN-VXLAN scenario traffic will not egress out of the irb next-hop when the bridge-domain doesn't have VXLAN VNI configured. PR1890878
- The hardware-assisted-keepalives CFM transmit packets are queued into Q4 instead of Q3/ Network-control. PR1891363
- Micro BFD packets egress through incorrect queue on PTX platforms. PR1893165
- LACP packets are getting dropped because it is misclassified as IPv4 packets in the L2Circuit/L2VPN scenario. PR1894156
- IPv4 and IPv6 traffic is dropped over VXLAN tunnel on PTX router. PR1895937
- Interfaces take 30 seconds to come up for a short LOS event. PR1896329
- PFEs fail to come online after a rapid restart or power cycle. PR1898307
- Host loopback wedge error is seen when committing firewall filter leading to major FPC errors. PR1899071
- Junos OS Evolved PTX PFE instance stuck in fault state after restart. PR1900735
- Anomalies observed on master RE during switchover. PR1901341

- FPC unresponsive due to memory exhaustion on PTX platforms. PR1901728
- System stops responding to `show system applications app ndp detail` command or prevents applications from start/stop/restart. PR1905807
- Interfaces take 14 seconds to come up for a short LOS event. PR1913948
- MPLS explicit null label packets on PTX10003 skip ipv4/ipv6 filter processing. PR1914060

Class of Service (CoS)

- Application of per-unit-scheduler on a channelized interface results in commit failure. PR1890490
- CCL:DCG: subscription to /qos/interfaces/interface/output/queues/queue is streaming of data on pfh- interfaces in scaled DCGate setup. PR1911456

Flow-based and Packet-based Processing

- [PTX-10008-evo] ingress mpls IPFIX flows have TopLabelAddr V6: :: PR1874737
- J-Flow Export Record Bug in PTX10000 with Junos OS Evolved devices. PR1893450

Forwarding and Sampling

- Inconsistent QoS marking in MPLS domain due to mismatch between traffic classification and rewrite policy. PR1873317
- PFEs will enter FAULT state when multiple PFEs are on-lined / off-lined / restarted at the same time. PR1886060

Interfaces and Chassis

- AE interfaces flap during GRES following an RE reboot on all Junos OS Evolved platforms. PR1898531

MPLS

- More bandwidth admitted onto a TE link when Label Switched Paths (LSPs) undergoing make-before-break re-route over the same link carrying the bypass LSP during local repair. PR1896022

Network Management and Monitoring

- Syslog forwarding intermittently stops post DUT reboot on virtual devices. PR1853209

Routing Policy and Firewall Filters

- Traffic drop observed for two policer with same initial name and term names. PR1896496

User Interface and Configuration

- The xnm-ssl feature fails without loopback interface configuration on Junos OS Evolved platforms. PR1882996
- Slow configuration commit observed on devices with a single Routing Engine (RE). PR1884781
- Traffic loss occurs due to missing apply-path handling during reboot or configd restart. PR1900883

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 25.4R1 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

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 25.4R1, 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
- QFX5241-32OD
- QFX5241-64OD
- QFX5241-64QD
- QFX5700
- QFX5700E

Hardware

- **QFX5241-32OD Switch (QFX Series)**—The Juniper Networks® QFX5241-32OD Switches are fixed-configuration devices with 32 octal small form-factor pluggable (OSFP) ports that support speeds of up to 800 Gigabit Ethernet (GbE). Features such as 25.6 terabits per second (Tbps) throughput and 1-

U shallow buffer design make these switches optimal as end-of-row, leaf, or spine devices in IP fabric architectures. The switches support 2400-watt (W) AC power supply units (PSUs) and front-to-back airflow.

To install the QFX5241-32OD switch and perform initial configuration, routine maintenance, and troubleshooting, see the QFX5241-32OD Switch Hardware Guide. See Feature Explorer for the complete list of features for any platform.

Table 2: QFX5241-32OD Feature Support

Feature	Description

Class of service (CoS)	• Support for CoS features on Layer 2 and Layer 3 interfaces, including:
	  • IPv4 and IPv6 unicast routing
	  • Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p
	  • Port scheduling
	  • Shared buffer
	  • Priority-based flow control (PFC) based on IEEE-802.1p for VLAN-tagged traffic. Protocols such as Remote Direct Memory Access (RDMA) over Converged Ethernet version 2 (RoCEv2) require DSCP-based PFC at Layer 3 for untagged traffic.
	  • Weighted random early detection (WRED) and explicit congestion notification (ECN)
	  • Telemetry support for CoS queue statistics exported using the sensor `/junos/system/linecard/qmon-sw/`.
	[See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description

EVPN	• Support for Layer 2 gateway and Address Resolution Protocol (ARP) suppression on EVPN-VXLAN.
	
	[See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy ARP and ARP Suppression, and overlay (Packet Forwarding Options).]
	
	• Support for CoS, and firewall filtering and policing on EVPN-VXLAN.
	
	[See CoS Support on EVPN VXLANs and Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]
	
	• Support for Wake-on-LAN (WOL) targeted broadcast on EVPN-VXLAN.
	
	[See Targeted Broadcast and targeted-broadcast.]
	
	• Support for EVPN-VXLAN Layer 2 gateway, including:
	
	  • Multihoming
	
	  • ARP suppression
	
	  • Layer 3 IPv4 underlay with integrated routing and bridging (IRB) and LAG
	
	  • Core isolation
	
	  • Broadcast, unknown unicast, and multicast (BUM) traffic forwarding by ingress replication only
	
	  • MAC move limits
	
	[See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy Arp and Arp Suppression, and Proxy NDP and NDP Suppression, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Edge-Routed Bridging Overlay Design and Implementation, Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions, and mac-move-limit.]
	
	• Support for EVPN-VXLAN Layer 3 gateway, including:

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
	• Layer 3 VXLAN gateway in edge-routed bridging fabric • Up to 256 VLANs with IRB enabled • Layer 3 underlay that supports IRB and LAG • ECMP in the underlay • IPv4 and IPv6 virtual gateway MAC address support for IRB interfaces • In-service software upgrade (ISSU) for Layer 3 gateway functionality [See Understanding EVPN with VXLAN Data Plane Encapsulation, Example: Configuring an EVPN-VXLAN Edge-Routed Bridging Fabric with a Virtual Gateway, Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network, and IP Fabric Underlay Network Design and Implementation.] • Support for EVPN-VXLAN Type 5 stitching, including: • Overlay and underlay ECMP • Type 5 stitching • Type 2 and Type 5 route coexistence • Symmetric IRB • In-service software upgrade (ISSU) [See Understanding EVPN with VXLAN Data Plane Encapsulation, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Understanding EVPN Type 5 Routes, EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN, NSR and Unified ISSU Support for EVPN, and irb-symmetric-routing.] • Support for sFlow technology on EVPN-VXLAN. [See Overview of sFlow Technology.]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
	• Support for port mirroring and analyzers on EVPN-VXLAN. [See Port Mirroring and Analyzers in an EVPN-VXLAN Environment.] • Support for forwarding EVPN data traffic on the spine device without any traffic loss while the leaf device performs a unified in-service software upgrade (unified ISSU) in spine-and-leaf topologies with external BGP (EBGP) connections. [See Understanding Unified ISSU.]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
Features optimized for AI-ML fabrics	• Support for priority-based flow control (PFC) watchdog. [See PFC Watchdog and congestion-notification-profile.] • Telemetry support for streaming IPv4 and IPv6 transit statistics using the native resource paths **/state/interfaces/interface[name=' ']/counters/ipv4/** and **/state/interfaces/interface[name=' ']/counters/ipv6/**. [See Junos YANG Data Model Explorer and route-accounting.] • Support for enabling or disabling dynamic load balancing (DLB). You can use the `dynamic-load-balance` statement to selectively enable or disable DLB based on `rdma-opcode` match or any match available in firewall filters. The optimal link is determined based on the modified port load and port queue metrics when DLB is enabled. [See rdma-opcode, dynamic-load-balance-selective, and egress-quantization.] • Support for PFC using Differentiated Services code points (DSCPs) at Layer 3 for untagged IPv6 traffic. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.] • Support for global load balancing (GLB). [See Global Load Balancing (GLB).] • Support for reactive path rebalancing. [See Reactive Path Rebalancing.] • SNMP and telemetry support for PFC, explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved, show snmp mib, and Guidelines for gRPC and

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
	gNMI Sensors (Junos Telemetry Interface). For sensors, see Junos YANG Data Model Explorer.] • Support for PFC X-ON threshold. [See xon (Input Congestion Notification).] • Extended sFlow monitoring functionality support to export sFlow sample packets through the mgmt_junos interface and nondefault virtual routing and forwarding (VRF) WAN ports. [See collector, show sflow collector, and System Logging and Routing Instances.] • Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool. [See buffer-dynamic-threshold.] • Support for increased global shared buffer pool of up to 147 MB. [See Configuring Ingress and Egress Dedicated Buffers.] • Support for Inband Flow Analyzer (IFA) 2.0 transit node. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description

Layer 2 features	• Support for Layer 2 unicast forwarding and VRRP.
	[See Understanding VRRP.]
	• Support for IGMP snooping, including:
	  • IGMPv1, IGMPv2, and IGMPv3
	  • IGMP proxy
	  • IGMP querier at Layer 2
	  • Any-source multicast (ASM) and source-specific multicast (SSM) modes
	  • Virtual router (VRF-lite) support
	  • Integrated routing and bridging (IRB) support
	[See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description

Layer 3 features	• Support for Layer 3 (L3) unicast forwarding and GRE tunneling. We support both IPv4 and IPv6 unicast routing.
	
	[See Generic Routing Encapsulation (GRE).]
	
	• Support for L3 multicast forwarding, including:
	
	  • PIM first-hop router rendezvous point (RP) functionality
	
	  • Multicast Source Discovery Protocol (MSDP)
	
	  • Make-before-break (MBB) support for multicast receivers on existing L3 aggregated Ethernet (ae*x*) or link aggregation group (LAG) interfaces. Support includes member addition, member deletion, link up, and link down events.
	
	  • PIM source-specific multicast (SSM)
	
	  • PIM sparse mode (PIM SM)
	
	  • PIM dense mode (PIM DM)
	
	  • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including:
	
	  • IPv4 and IPv6 multicast
	
	  • IGMPv1, IGMPv2, and IGMPv3
	
	  • Multicast Listener Discovery (MLD) versions 1 and 2
	
	  • Any-source multicast (ASM) and source-specific multicast (SSM) modes
	
	[See Multicast Routing Protocols and PIM Overview.]
	
	• Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6.
	
	[See DHCP Relay Agent.]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
Network management and monitoring	• Support for sFlow technology. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers. The switch can support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]
Platform and infrastructure	• Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.]
Precision Time Protocol (PTP)	• Support for Precision Time Protocol (PTP) transparent clock. [See Precision Time Protocol (PTP) Overview.]
Protection against DDoS attacks	• Support for configuration and installation of policers at the Packet Forwarding Engine level for defense from distributed denial-of-service (DDoS) attacks. By default, DDoS protection is enabled for many protocols on the QFX5241-32OD switch. [See Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers, show ddos-protection statistics, and show ddos-protection version.]
Routing policy and firewall filters	• Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]

Table 2: QFX5241-32OD Feature Support *(Continued)*

Feature	Description

Services applications	• Support for GRE features, including:
	  • GRE tunnels over Gigabit Ethernet, LAG, and VLAN
	  • Tagged subinterfaces
	  • Payload protocol for IPv4 and IPv6
	  • Delivery protocol for IPv4
	  • Multicast over GRE tunnels
	  • Tunnel statistics
	  • VRF with GRE
	  • Time-to-live (TTL)
	[See Generic Routing Encapsulation (GRE).]

Software installation and upgrade	• Support for zero-touch provisioning (ZTP) over IPv4 and IPv6 on the management and WAN interfaces. [See Zero Touch Provisioning.]

- **QFX5241-64OD and QFX5241-64QD switches (QFX Series)**—The Juniper Networks® QFX5241-64OD Switch and Juniper Networks® QFX5241-64QD Switch are fixed-configuration devices with 64 octal small form-factor pluggable (OSFP) or QSFP-DD ports that support speeds up to 800 Gigabit Ethernet (GbE). Features such as 51.2 terabits per second (Tbps) throughput and 2-U shallow buffer design make these switches optimal as end-of-row, leaf, or spine devices in IP fabric architectures. The switches support 3000-watt (W) AC and DC power supply units (PSUs) and front-to-back airflow. A key difference between the QFX5240 and QFX5241 is that QF5241 supports Secure BIOS and Secure Boot.

  To install the QFX5241-64OD and QFX5241-64QD switches and perform initial configuration, routine maintenance, and troubleshooting, see the QFX5241-64OD and QFX5241-64QD Switches Hardware Guide. See Feature Explorer for the complete list of features for any platform.

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support

Feature	Description

Class of service (CoS)	• Support for CoS features on Layer 2 and Layer 3 interfaces, including:
	  • IPv4 and IPv6 unicast routing
	  • Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p
	  • Port scheduling
	  • Shared buffer
	  • Priority-based flow control (PFC) based on IEEE-802.1p for VLAN-tagged traffic. Protocols such as Remote Direct Memory Access (RDMA) over Converged Ethernet version 2 (RoCEv2) require DSCP-based PFC at Layer 3 for untagged traffic.
	  • Weighted random early detection (WRED) and explicit congestion notification (ECN)
	  • Telemetry support for CoS queue statistics exported using the sensor `/junos/system/linecard/qmon-sw/`.
	[See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description
EVPN	• Support for Layer 2 gateway and Address Resolution Protocol (ARP) suppression on EVPN-VXLAN. [See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy ARP and ARP Suppression, and overlay (Packet Forwarding Options).] • Support for CoS, and firewall filtering and policing on EVPN-VXLAN. [See CoS Support on EVPN VXLANs and Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • Support for Wake-on-LAN (WOL) targeted broadcast on EVPN-VXLAN. [See Targeted Broadcast and targeted-broadcast.] • Support for EVPN-VXLAN Layer 2 gateway, including: • Multihoming • ARP suppression • Layer 3 IPv4 underlay with integrated routing and bridging (IRB) and LAG • Core isolation • Broadcast, unknown unicast, and multicast (BUM) traffic forwarding by ingress replication only • MAC move limits [See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy Arp and Arp Suppression, and Proxy NDP and NDP Suppression, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Edge-Routed Bridging Overlay Design and Implementation, Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions, and mac-move-limit.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description

	• Support for EVPN-VXLAN Layer 3 gateway, including:
	  • Layer 3 VXLAN gateway in edge-routed bridging fabric
	  • Up to 256 VLANs with IRB enabled
	  • Layer 3 underlay that supports IRB and LAG
	  • ECMP in the underlay
	  • IPv4 and IPv6 virtual gateway MAC address support for IRB interfaces
	  • In-service software upgrade (ISSU) for Layer 3 gateway functionality
	  [See Understanding EVPN with VXLAN Data Plane Encapsulation, Example: Configuring an EVPN-VXLAN Edge-Routed Bridging Fabric with a Virtual Gateway, Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network, and IP Fabric Underlay Network Design and Implementation.]
	• Support for EVPN-VXLAN Type 5 stitching, including:
	  • Overlay and underlay ECMP
	  • Type 5 stitching
	  • Type 2 and Type 5 route coexistence
	  • Symmetric IRB
	  • In-service software upgrade (ISSU)
	  [See Understanding EVPN with VXLAN Data Plane Encapsulation, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Understanding EVPN Type 5 Routes, EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN, NSR and Unified ISSU Support for EVPN, and irb-symmetric-routing.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description
	• Support for sFlow technology on EVPN-VXLAN. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers on EVPN-VXLAN. [See Port Mirroring and Analyzers in an EVPN-VXLAN Environment.] • Support for forwarding EVPN data traffic on the spine device without any traffic loss while the leaf device performs a unified in-service software upgrade (unified ISSU) in spine-and-leaf topologies with external BGP (EBGP) connections. [See Understanding Unified ISSU.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description
Features optimized for AI-ML fabrics	• Support for priority-based flow control (PFC) watchdog. [See PFC Watchdog and congestion-notification-profile.] • Telemetry support for streaming IPv4 and IPv6 transit statistics using the native resource paths **/state/interfaces/interface[name=' ']/counters/ipv4/** and **/state/interfaces/interface[name=' ']/counters/ipv6/**. [See Junos YANG Data Model Explorer and route-accounting.] • Support for enabling or disabling dynamic load balancing (DLB). You can use the `dynamic-load-balance` statement to selectively enable or disable DLB based on `rdma-opcode` match or any match available in firewall filters. The optimal link is determined based on the modified port load and port queue metrics when DLB is enabled. [See rdma-opcode, dynamic-load-balance-selective, and egress-quantization.] • Support for PFC using Differentiated Services code points (DSCPs) at Layer 3 for untagged IPv6 traffic. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over Converged Ethernet version 2 (RoCEv2). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.] • Support for global load balancing (GLB). [See Global Load Balancing (GLB).] • Support for reactive path rebalancing. [See Reactive Path Rebalancing.] • SNMP and telemetry support for PFC, explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved, show snmp mib, and Guidelines for gRPC and

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description
	gNMI Sensors (Junos Telemetry Interface). For sensors, see Junos YANG Data Model Explorer.] • Support for PFC X-ON threshold. [See xon (Input Congestion Notification).] • Extended sFlow monitoring functionality support to export sFlow sample packets through the mgmt_junos interface and nondefault virtual routing and forwarding (VRF) WAN ports. [See collector, show sflow collector, and System Logging and Routing Instances.] • Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool. [See buffer-dynamic-threshold.] • Support for increased global shared buffer pool of up to 147 MB. [See Configuring Ingress and Egress Dedicated Buffers.] • Support for Inband Flow Analyzer (IFA) 2.0 transit node. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description

Layer 2 features	• Support for Layer 2 unicast forwarding and VRRP.
	[See Understanding VRRP.]
	• Support for IGMP snooping, including:
	  • IGMPv1, IGMPv2, and IGMPv3
	  • IGMP proxy
	  • IGMP querier at Layer 2
	  • Any-source multicast (ASM) and source-specific multicast (SSM) modes
	  • Virtual router (VRF-lite) support
	  • Integrated routing and bridging (IRB) support
	[See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description

Layer 3 features	• Support for Layer 3 (L3) unicast forwarding and GRE tunneling. We support both IPv4 and IPv6 unicast routing.
	
	[See Generic Routing Encapsulation (GRE).]
	
	• Support for L3 multicast forwarding, including:
	
	  • PIM first-hop router rendezvous point (RP) functionality
	
	  • Multicast Source Discovery Protocol (MSDP)
	
	  • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (ae*x*) or link aggregation group (LAG) interfaces. Support includes member addition, member deletion, link up, and link down events.
	
	  • PIM source-specific multicast (SSM)
	
	  • PIM sparse mode (PIM SM)
	
	  • PIM dense mode (PIM DM)
	
	  • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including:
	
	  • IPv4 and IPv6 multicast
	
	  • IGMPv1, IGMPv2, and IGMPv3
	
	  • Multicast Listener Discovery (MLD) versions 1 and 2
	
	  • Any-source multicast (ASM) and source-specific multicast (SSM) modes
	
	[See Multicast Routing Protocols and PIM Overview.]
	
	• Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6.
	
	[See DHCP Relay Agent.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description
Network management and monitoring	• Support for sFlow technology. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers. The switches can support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]
Platform and infrastructure	• Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.]
Protection against DDoS attacks	• Support for configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from distributed denial-of-service (DDoS) attacks. By default, DDoS protection is enabled for many protocols on the QFX5241-64OD and QFX5241-64QD switches. [See Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers, show ddos-protection statistics, and show ddos-protection version.]
Routing policy and firewall filters	• Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]

Table 3: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description

Services applications	• Support for GRE features, including:
	  • GRE tunnels over Gigabit Ethernet, LAG, and VLAN
	  • Tagged subinterfaces
	  • Payload protocol for IPv4 and IPv6
	  • Delivery protocol for IPv4
	  • Multicast over GRE tunnels
	  • Tunnel statistics
	  • VRF with GRE
	  • Time-to-live (TTL)
	[See Generic Routing Encapsulation (GRE).]

Software installation and upgrade	• Support for zero-touch provisioning (ZTP) over IPv4 and IPv6 on the management and WAN interfaces. [See Zero Touch Provisioning.]

Authentication and Access Control

•

Class of Service

- **Support for drop congestion notification (QFX5240-64OD, QFX5240-64QD, QFX5241-64OD, QFX5241-64QD, and QFX5241E-64OD)**—QFX5240 and QFX5241 switches support drop congestion notification (DCN) to enhance congestion management. DCN is a congestion management technique based on packet trimming. Rather than drop a packet when congestion occurs, the device trims the packet's payload, which results in a smaller packet. The device then transmits this smaller packet through a high-priority queue toward its destination. Subsequent hops in the network must recognize DCN-drop marked packets and direct them to high-priority queues as well. End hosts must process the trimmed DCN packets, identify the packets dropped due to congestion, and request retransmission of those lost packets.

[See Drop Congestion Notification (DCN).]

- **Rewrite outer DSCP, preserve ECN, and copy inner DSCP by default in EVPN-VXLAN (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)**—You can control DSCP marking on network-facing underlay interfaces in EVPN-VXLAN by assigning one rewrite rule per interface for IPv4 and IPv6. By default, the device copies inner DSCP to the outer VXLAN header. A DSCP rewrite overrides the type-of service (ToS) copy and rewrites only the outer IP DSCP, preserving inner IP DSCP and ECN bits. The device does not support 802.1p rewrite. You can enable explicit congestion notification (ECN) per output queue through scheduler configuration. With weight random early detection (WRED), congested queues mark packets CE (congestion experienced) when all intermediate devices enable ECN. Use this capability to maintain QoS consistency across overlays and underlays.

  [See CoS Support on EVPN VXLANs.]

- **DCBX support for PFC (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—We support data center bridging and exchange capability (DCBX) support for priority-based flow control (PFC) on the listed switches.

  [See Understanding DCBX and priority-flow-control.]

- **Dynamic threshold profiles for shared buffer pools (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can tune the buffer allocated per priority group or port based on the dynamic threshold setting, also known as the alpha value. Existing design supports this feature through a global configuration that is applicable to all the ports on the device regardless of the port configuration or properties. Having a global alpha value is not effective when the device has ports operating at various speeds. With this feature, you can create dynamic thresholds per priority group and then associate the dynamic threshold profile on an ingress interface. This configuration creates interface-specific threshold values.

  You can configure per priority group dynamic thresholds for each interface to optimize shared buffer allocation and protect lossless traffic during congestion. Egress lossless queues inherit ingress priority-group alpha values.

  [See Dynamic Threshold Profiles for Shared Buffer Pools.]

Device Security

- **IMA coverage update (ACX Series, PTX Series, and QFX Series)**—Integrity Measurement Architecture (IMA) coverage now includes the following additional file systems:

  - ISO9660
  - PROC

- SYSFS
- DEBUGFS
- RAMFS
- SECURITYFS
- EFIVARFS
- DEVPTS
- BINFMTFS
- SELINUX
- CGROUP
- NSFS
- TRACEFS

IMA now enforces signature verification for the `kexec` kernel and `initramfs` images. It also generates a nonrepudiable log for new key addition events to IMA keyrings. These enhancements strengthen runtime integrity protections against unauthorized changes to Junos OS Evolved.

[See File Security with IMA.]

EVPN

- **Copy and preserve 802.1p CoS priority bits between VLAN S-tags and C-tags with EVPN-VXLAN (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)**—Use this feature to copy and preserve the 802.1p priority between customer VLAN headers (C-tags) and service provider VLAN headers (S-tags) in EVPN-VXLAN traffic. We support this feature with service provider-style interface configurations.

  - When a leaf device adds an S-tag to a packet, this feature copies the class-of-service (CoS) priority bits from the C-tag to the S-tag.
  - When a leaf device removes an S-tag from a packet, this feature retains the priority bits from the C-tag.
  - The traffic priority remains unchanged in the VXLAN tunnel.

  Use the `vxlan-enable-dot1p-copy` option at the `[edit forwarding-options]` hierarchy level to enable this behavior. This setting overrides VLAN rewrite settings.

  [See forwarding-options.]

- **EVPN multihoming and multitenancy support over colored IP fabric with BGP DPF (QFX5130-32CD, QFX5240-64OD, and QFX5240-64QD)**—You can leverage EVPN-VXLAN over colored IP fabric using BGP deterministic path forwarding (DPF) to support multihoming and multitenancy configurations for AI/ML applications. This functionality facilitates EVPN for Layer 3 networks with EVPN Type 5, enhancing network segmentation and resource allocation. By using a colored logical fabric, you can achieve flexible routing as uncolored routes integrate seamlessly with all color-coded sessions, optimizing network efficiency and adaptability.

  [See BGP Deterministic Path Forwarding in a CLOS Network.]

- **IGMP snooping and MLD snooping in EVPN-VXLAN bridged overlay networks with IPv4 or IPv6 underlays (QFX5130-32CD, QFX5130-48, QFX5130-48CM, and QFX5700)**—With IGMP snooping and MLD snooping, the devices in an EVPN-VXLAN network forward multicast traffic flows only to the hosts that subscribe to those flows, rather than flooding the traffic to all hosts. Now you can enable Layer 2 (L2) multicast with IGMP snooping for intra-VLAN IPv4 multicast traffic and MLD snooping for IPv6 multicast traffic in an EVPN-VXLAN bridged overlay network with:

  - IPv4 or IPv6 underlay peering in the EVPN-VXLAN network with both IPv4 and IPv6 data traffic.
  - MAC-VRF EVPN instances with `vlan-based` or `vlan-aware` service types.
  - Enterprise-style interfaces.

  You must configure the bridge domains (VLANs) symmetrically on EVPN multihoming peer devices.

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

- **Hardware-assisted inline BFD for EVPN-VXLAN types 2 and 5 with 3x 100-ms timers (QFX5130-32CD, QFX5130-48C, QFX5130E-32CD, QFX5700, and QFX5700E)**—You can use hardware-assisted inline BFD over VXLAN tunnels for rapid, deterministic failure detection with 100 x 3 millisecond timers on supported platforms. Use IPv4 and IPv6 multihop BFD for Type 2 (L2/L3) with ECMP or multihomed VTEPs. Use Type 5 with ECMP, including pure Type 5 routing instances. To enable hardware-assisted inline BFD, configure bfd on overlay bgp sessions, peer overlays between loopbacks, and apply the specified timers.

  [See Understanding How BFD Detects Network Failures.]

- **Egress rate limiting per VNI for VXLAN unicast and BUM traffic (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)**—You can enforce per-VNI egress rate limits on VXLAN tunnel-initiated traffic to prevent congestion, mitigate denial-of-service (DoS) risk, and prioritize critical services. This configuration targets VXLAN traffic and preserves locally switched or routed flows. We added a new egress VLAN ACL filter profile to support the *egress rate limit per VNI* feature. You enable this profile with `set system packet-forwarding-options firewall profiles ethernet-switching egress profile1`. Changing the filter profile triggers a Packet Forwarding Engine restart. Create the filter using `set firewall family ethernet-switching filter` *`filter-name`* `term` *`term-name`* `from`

`vxlan tunnel-initiated` and `traffic-type known-unicast` for unicast traffic or `traffic-type-except known-unicast` for BUM traffic. Set two-color or three-color policers with the `discard` action and attach the filters per VLAN with `set routing-instances` *`instance-name`* `vlans` *`vlan-name`* `forwarding-options filter output` *`filter-name`*. You use `show firewall` to view policer statistics.

[See ethernet-switching and tunnel-initiated.]

- **sFlow support for EVPN-VXLAN multicast (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)**—You can use sFlow technology to sample EVPN-VXLAN multicast traffic configured at the interface level on customer-facing ports. The software replication of samples is also enabled. The software replicates each sample to all egress-sampling-enabled interfaces that are part of the multicast group.

  The sFlow collector can be reached through a standard Layer 3 gateway (underlay), a management IP address (reachable through the default or a nondefault routing instance), or a VXLAN tunnel (overlay).

  To enable known multicast sampling (disabled by default), use the `set protocols sflow egress-multicast` command.

  To set the maximum packet replication rate in software (the default is 2000), use `set protocols sflow max-replication-rate 1000`.

  Limitations:

  - The `out-priority` field for a VLAN is always set to 0 in sFlow samples.
  - IPv6 underlay transport for the EVPN-VXLAN sFlow use case is not supported.
  - EVPN-VXLAN unknown multicast traffic is not sampled; only known multicast traffic is supported for sampling.

  [See sFlow Technology Overview.]

- **Symmetric IRB for IPv6 underlay with EVPN Type-2 MAC-IP routes for intersubnet routing (QFX5130-32CD and QFX5130-48C)**—You can now use symmetrical integrated routing and bridging (IRB) to forward traffic consistently across all provider edge (PE) devices. The device carries EVPN Type 2 MAC-IP routes across an IPv6 underlay that provides an ECMP transport for scalable, multitenant L2 and L3 connectivity. For each routing instance, use the `irb-symmetric-routing vni` statement under the `[edit protocols evpn]` hierarchy to maintain symmetrical bridging and routing.

  [See EVPN-VXLAN with an IPv6 Underlay and Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics.]

- **Automatically synchronize multicast router interface status among EVPN multihoming peers in EVPN-VXLAN bridged overlays (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)**—

Multihoming peer EVPN-VXLAN border leaf (BL) devices that connect to external multicast receivers will now automatically synchronize the multicast router (mrouter) interface status for the EVPN segment identifiers (ESIs) the peer devices share. This feature ensures that multicast traffic reaches external receivers when a non-designated forwarder (NDF) BL peer device receives IGMP or MLD queries from the external PIM domain.

The peer BL devices:

- Use EVPN Type 7 Join Sync routes to communicate the mrouter interface status among their ESI peers.
- Advertise EVPN Type 3 Inclusive Multicast Ethernet Tag (IMET) routes that enable the other EVPN devices to learn how to reach the mrouter interfaces.

This behavior happens by default, so you no longer need to manually configure the `multicast-router-interface` setting on the multihoming peer BL devices. If needed, you can disable this behavior using the `skip-multicast-router-port-sync-nlri` statement at the **`[edit protocols evpn]`** hierarchy level.

We support this feature in EVPN-VXLAN bridged overlay (BO) networks with IPv4 or IPv6 underlay peering.

- **Enable scaling for stretched VXLAN campus networks (QFX5130-32CD and QFX5700)**—To support large-scaled stretched VXLAN campus networks, we provide new routing policy options, sample routing policies, and new statements to optimize how host routes are managed across the access, distribution, and core layers. With this feature, you can configure the network to install host routes in the core layer but not advertise the host routes to the distribution and access layers. The core devices advertise only subnet routes (using EVPN Type 5 routes) to the distribution devices. The distribution devices then advertise the subnet routes to the access layer. The configuration includes policies to ensure the EVPN Type 5 subnet routes are the preferred routes on the distribution and access layer devices. This design reduces the route table burden on access and distribution devices, enabling greater scalability.

Forwarding Options

- **Use the `no-queue-pair` configuration to exclude `queue-pair` from the hash calculation (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Previously, the 5 tuple approach was used for hash calculation to identify a flow. But for remote direct memory access (RDMA) traffic, the 5 tuple hashing mechanism wasn't providing the desired entropy. So considering `queue-pair` into the hash calculation provided better entropy. Currently, `queue-pair` is by default included in the hash calculation You can use the `no-queue-pair` configuration to exclude `queue-pair` from the hash calculation.

  [See enhanced-hash-key.]

- **Symmetric hashing support (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD,**

QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, and QFX5700)—Symmetric hashing ensures that bidirectional traffic flows use the same path across ECMP routes or LAGs.

When enabled, symmetric hashing normalizes Layer 3 and Layer 4 fields before computing the hash. This normalization guarantees that forward and reverse traffic produce the same hash value, mapping to the same link or ECMP path. This behavior is important for services that require flow symmetry, such as Network Address Translation (NAT), firewalls, distributed denial-of-service (DDoS) protection, and data collection.

To enable symmetric hashing, use the `set forwarding-options enhanced-hash-key symmetric-hash` *`inet/inet6`* command.

[See enhanced-hash-key.]

- **Unknown unicast drop configuration for VLAN interfaces (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)**—You can enhance network performance and prevent traffic storms by configuring your switch to drop unknown unicast packets. This action prevents the flooding of unicast packets with unknown destination MAC addresses across VLAN interfaces.When you enable this feature, the switch learns and adds the source MAC address to the MAC address table. The switch drops packets with unlearned destination MAC addresses. This approach ensures efficient network resource usage and optimal network performance.

  [See Understanding and Preventing Unknown Unicast Forwarding.]

High Availability

- **Inline micro-BFD support (QFX5130-32CD, QFX5130E-32CD, QFX5700, and QFX5700E)**—You can use inline micro-BFD sessions to monitor the status of individual member links in a LAG bundle. To configure micro-BFD, use the `set interfaces` *`interface`* `aggregated-ether-options bfd-liveness-detection` configuration command.

  The BFD local address is the loopback address of the source of the micro-BFD sessions. Micro-BFD sessions do not support ISSU or authentication.

  [See Understanding Independent Micro BFD Sessions for LAG.]

- **In-service hitless graceful reboot (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can perform a hitless graceful reboot to minimize traffic loss and prevent interface flapping. Initiate the reboot with the `request system reboot hitless` command, or use the `request system reboot hitless no-confirm` option to bypass prompts. To review reboot details, use the `show system software hitless-reboot` command. The reboot operation will be canceled if another reboot is pending, upgrade tasks are active, graceful restart is not configured, or application configuration validation fails. Hitless reboot works only for protocols and configurations that are supported by unified in-service software upgrade (unified ISSU).

  [See Unified ISSU for Junos OS Evolved.]

- **Fast Fast Boot-enabled unified ISSU (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5130E-32CD)**—You can perform unified ISSU on with minimal disruption using Fast Fast Boot (FFB). FFB maintains forwarding with minimal traffic drop and control-plane outage. You can enable unified in-service software upgrade (unified ISSU) with the `request system software add` *`package-name`* `restart` command.

  [See Understanding Unified ISSU.]

Interfaces

- **Application selection enhancements on ZR and ZR-M optics (QFX5130-32CD and QFX5220)**—Use these application selection enhancements to view the advertised applications for R and ZR-M optics and switch among those applications as needed. [See 400ZR and 400G OpenZR+ Optical Transceivers.]

- **Port speed and channelization (QFX5241-32OD)—** You can tailor port speeds and channelize interfaces on QFX5241-32OD switches. On the QFX5241-32OD, you can configure native 800-Gbps speed on the OSFP ports and 10-Gbps speed on the SFP28 ports.

  You can also channelize the OSFP ports into the following speeds:

  - 1x800 Gbps
  - 2x400 Gbps
  - 4x200 Gbps

  Even-numbered ports support 8x100G or 8x50G; the next odd-numbered port supports up to 2x400G and 2x100G. [See Port Speed on QFX Switches.]

- **Monitor pre-FEC BER on optical interfaces (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can assess raw optical signal quality by measuring the pre–forward error correction (FEC) bit error rate (BER), which reflects errors before correction. Use pre-FEC to evaluate the effectiveness of cable performance and error correction. These metrics help you identify impairments early and guide link tuning or maintenance decisions. You collect both metrics in monitoring workflows to support objective path selection, capacity planning, and service assurance. Use the `show interfaces` *`interface-name`* command to display the pre-FEC BER. [See show interfaces.]

- **FEC histogram and statistics on optical interfaces (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Use the FEC histogram to evaluate link quality. The histogram shows symbol error corrections per FEC codeword for precise assessment of transmission integrity. Run `show interfaces` *`interface-name`* to display FEC counters, corrected and uncorrected codeword counts, pre-FEC BER, and the FEC histogram. This visibility enables you to detect degradation and guide troubleshooting and optimization. [See show interfaces.]

- **Interface flap damping with up/down hold timers to suppress transient routing updates (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Use interface damping to prevent

transient interface flaps during transport link switching from causing unwanted routing updates. Configure hold-down and hold-up timers with the `interface hold-time` statement to delay advertising state changes. The damping mechanism ignores transitions within these timers and advertises the interface only when the state persists at timer expiration. Apply damping with routing protocols to stabilize route convergence during millisecond-scale flaps. [See Physical Interface Damping Overview and Hold timer for LAG interfaces].

- **Support for native 800 Gbps and SFP28 10/25 G with expanded channelization options for OSFP and QSFP-DD (QFX5241-64OD and QFX5241-64QD)**—Optimize port density and migration by selecting native speeds or channelization on QFX5241-64OD and QFX5241-64QD switches.

  On the QFX5241-64OD, you can configure native 800-Gbps speed on the OSFP ports and 10-Gbps speed on the SFP28 ports. You can also channelize OSFP ports into 2x400 Gbps or 8x100 Gbps.

  On the QFX5241-64QD:

  - Configure the following native speeds on the QSFP-DD ports:
    - 800 Gbps
    - 400 Gbps
    - 100 Gbps
  - Use 10-Gbps speed on the SFP28 ports.

  You can also channelize the QSFP-DD ports into the following speeds:

  - 2x400 Gbps
  - 4x100 Gbps
  - 8x100 Gbps
  - 8x50 Gbps

  Do not mix 8x100 Gbps or 8x50 Gbps on odd and even ports. See Port Speed on QFX Series Switches. Also, 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.

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, and**

QFX5241-64QD)—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:

- Slax processor enhancements including a new mode, additional options, and simplified argument parsing
- New libslax extension library functions
- Improved SLAX syntax options
- New SLAX functions and enhancements to existing functions and statements
- Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See libslax Distribution Overview.]

Junos Telemetry

- **Export packet capture statistics to external collectors using Junos telemetry (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5700, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—The packet capture feature records the first configured number of host-bound packets on each physical interface and exports them to an external collector over the Junos telemetry infrastructure. You can use this data to debug and fix network or performance issues. Subscribe to the packet-capture sensor at `/junos/system/linecard/packet-capture`. The device captures 50 ingress packets when an interface transitions from the DOWN state to the UP state. The data is encoded in Google Protocol Buffer format and streamed over gRPC with SSL encryption. To enable packet capture, configure:

  ```
  edit system packet-forwarding-options packet-capture packet-capture-enable
  ```

  Use the *show agent sensors* command to view the packet capture sensor information.

  For more information, see Junos YANG Data Model Explorer.

- **Stream telemetry data in gNMI-based message format over UDP (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, and QFX5240-64QD)**—Junos OS Evolved uses a dial-out mechanism to send telemetry data to a collector over UDP . The message format is defined in the `jnx_gnmi_over_udp.proto` file. Only `STREAM` mode with `SAMPLE` as subscription mode is supported. The message contains full key name and value pair information so the collector does not require data models for processing or consuming the telemetry data.

[See No Link Title, No Link Title, No Link Title, No Link Title, No Link Title, and Junos YANG Data Model Explorer.]

Layer 2 VPN

- **Support for rewriting the L2PT destination MAC address (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)**—In Layer 2 Protocol Tunneling (L2PT), the device rewrites the original multicast destination MAC address of the packet. The packet then travels across the provider network transparently to the other end of the tunnel, where the destination device restores the original multicast destination MAC address. By default, the device rewrites the multicast destination MAC address with the predefined multicast tunneling MAC address 01:00:0C:CD:CD:D0 in the MAC table. You can optionally specify a different multicast MAC address.

  To specify the MAC address, use the `tunnel-destination-mac` *`mac-address`* statement at the `[edit protocols layer2-control layer2-control mac-rewrite]` hierarchy level. You can set any non-reserved multicast MAC address.

  [See mac-rewrite and Layer 2 Protocol Tunneling (L2PT).]

Multicast

- **Hash-based PIM RPF selection (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)**—Use hash-based Protocol Independent Multicast (PIM) reverse path forwarding (RPF) selection to route multicast traffic for a specific source and group (S, G) consistently through the same upstream node. The device uses multi-level hashing of PIM neighbor attributes such as router ID, cluster ID, and interface ID for consistent routing. For configuration, include the `hash-based-rpf-selection` statement under the `edit protocols pim` hierarchy.

  [See Hash-based PIM RPF selection, and hash-based-rpf-selection.]

Network Management and Monitoring

- **Ingress inline sFlow and sFlow collector statistics (QFX5130-32CD, QFX5130-48CM, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—Ingress inline sFlow samples packets on ingress local ports in real time, providing timely visibility into live traffic without relying on control plane processing. The device handles sampling decisions, statistics collection, and sample-destination selection in the Packet Forwarding Engine after ingress processing and before queuing, minimizing forwarding impact. The device copies selected packet and mirrors to a configured analyzer (mirror-to-port) for detailed analysis while traffic moves normally. Inline operation ensures accurate flow insight at high speeds, with overhead controlled by a configurable sampling rate and consistent behavior regardless of the egress port.

- **AES-256 Encryption Algorithm Support for SNMPv3 (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5700E,**

QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—You can configure Advanced Encryption Standard (AES) 256 algorithm for an SNMPv3 user. To configure AES-256 algorithms for an SNMPv3 user, include the `privacy-aes256` statement at the `edit snmp v3 usm local-engine user username` hierarchy level. AES-256 is a symmetric encryption algorithm that uses a 256-bit key to encrypt or decrypt messages and provides high-level security for protecting sensitive information.

[See Configure SNMPv3 Authentication Type and Encryption Type, show snmp v3, and usm.]

- **Timestamp option for tap-aggregation packets (QFX5220-32CD and QFX5220-128C)**—You can configure the tap-aggregation feature to insert a timestamp in packets at data capture—before the packets are sent to the tool ports for analysis. The timestamp shows exactly when the data packet was captured.

  Configure the Precision Time Protocol (PTP) reference clock on the tap-aggregation switch and ensure that PTP is running when the timestamp is inserted. Your tap-aggregation switch must also sync the PTP FPGA's recovered time-of-day with the system chip's time-of-day. The command you use to enable the timestamp option is:

  `[edit] user@switch# set interfaces` *`interface-name`* `timestamp ingress`

- **Ingress ACL UDF filtering function on tap ports on TAP-aggregation switches (QFX5130-32CD, QFX5130-48C, QFX5700, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can apply ingress access control list (ACL) with user-defined fields (UDF) on tap interfaces. This helps you choose the traffic to be sent to the tool interfaces on a tap-aggregation switch. If the ACL match conflicts with the tap-aggregation rule, the ACL match takes precedence.

  Also on these switches, you can now configure the TAP-aggregation interfaces under the `[edit interfaces]` hierarchy level as follows:

  - Add an interface to a tap group with

    `[edit]`

    ```
    user@switch# set interfaces interface-name unit 0 mode tap group tap-group-name
    ```

  - Add an interface to a tool group with

    `[edit]`

    ```
    user@switch# set interfaces interface-name unit 0 mode tool group tool-group-name
    ```

- **Dropped-packet notification (QFX5240-64OD and QFX5240-64QD)**—Packet drops are common occurrences on network switches and routers. Debugging packet drops can be complex and time-consuming. The packet-processing pipeline supports a limited set of drop counters, but these

counters are insufficient for debugging complex packet-drop issues. Debugging difficulties can result in high mean times to recovery (MTTRs).

A feature called *dropped-packet notification*, also referred to as mirror on drop (MoD), can help you debug packet drops in real time. The areas of packet drop monitored include:

- Packets dropped due to processing in the ingress pipeline
- Packets dropped due to congestion in the MMU

Dropped-packet notification on the platforms named in this description is stateless and flow unaware.

You configure much of the dropped-packet notification feature at the `[edit forwarding-options mirror-profile]` hierarchy level.

- **1:N port mirroring for sending a source packet to multiple Layer 2 destinations (QFX5130-32CD|QFX5130E-32CD|QFX5130-48C|QFX5130-48CM|QFX5700|QFX5700E|QFX5220|QFX5230-64CD|QFX5240-64OD|QFX5240-64QD)**—You can use the 1:N port mirroring feature to mirror traffic to multiple Layer 2 destinations. This feature requires either one or both of the following configurations:

  - A port-mirroring instance that is based on a firewall filter. Use the configuration statements in the `[edit forwarding-options port-mirroring instance]` hierarchy.
  - A native analyzer. Use the configuration statements in the `[edit forwarding-options analyzer]` hierarchy.

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

  [See 1:N Port Mirroring to Multiple Destinations on Switches.]

Platform and Infrastructure

- **NIST purge method for media sanitization (QFX5240-64OD and QFX5240-64QD)**—We've extended support for NIST media sanitization for Non-Volatile Memory Express (NVMe) solid-state drives to include:

  - Cryptographic scramble and block erase priorities for the purge method
  - NVMe format with user data erase for the purge method if the previous step fails
  - NVMe format for the clear method

  For example, you can use this high level of data destruction when you pull a device from production. To maintain data security, sanitize any disk drives in the device before they leave your premises. The NIST Special Publication 800-88 specifies the priority levels for sanitizing disk drives. In Junos OS

Evolved, sanitize a disk drive using the `request system zeroize (disk1 | disk2)` command. The sanitization process starts at the highest NIST sanitization priority that the NVMe drive supports. If the attempt fails, the process uses the method associated with the next lowest NIST priority level, and so on, until the disk is sanitized either using one of the NIST methods or using the Linux dd command.

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

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, PTX Series, and QFX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):

  - ML-DSA-87 PQC algorithm for digital signatures
  - SHA-512 for hashing

  For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

  [See PQC Signatures for Software Images.]

- **Support for Quantum Buffer in SSH (ACX Series, PTX Series, and QFX Series)**—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptoagility. The feature uses finite field cryptography (FFC) to extend the security life span of the current systems against quantum attacks. Quantum Buffer provides a phased approach to adopting post-quantum cryptography (PQC), thereby mitigating operational risks associated with the transition.

  To enable the feature, configure the following command:

  - `set system services ssh moduli type` *`name`* `refresh` *`frequency`* `count` *`count`*

    The configuration dynamically generates prime moduli for existing Diffie-Hellman (DH) group exchange algorithms, `group-exchange-sha1` and `group-exchange-sha2`. The `qbufd` process is responsible for generating the moduli.

  [See Quantum Buffer and moduli.]

- **Support for Shor-resistant and other default key exchange algorithms in SSH (ACX Series, PTX Series, and QFX Series)**—SSH supports the hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm, which is Shor-resistant and improves protection against quantum attacks.

  Configure `sntrup761x25519-sha512` at the [`edit system services ssh key-exchange`] hierarchy level.

  Additionally, SSH includes default support for the following Diffie-Hellman (DH) group key exchange algorithms that are available at the [`edit system services ssh key-exchange`] hierarchy level.

- dh-group16-sha512
- dh-group18-sha512

[See key-exchange.]

Routing Policy and Firewall Filters

- **Configure multiple flex match ranges within the same term (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can configure multiple flex match ranges within the same term using the `flexible-match-ranges` configuration statement, allowing traffic to be filtered based on several criteria simultaneously. This feature supports family inet, inet6, and Ethernet switching.

  [See flexible-match-ranges.]

Routing Protocols

- **IS-IS multi-instance support over a single interface (ACX7020, ACX7024, ACX7100, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, PTX10016, QFX5130, QFX5220, QFX5230-64CD, QFX5240, QFX5241-32OD, and QFX5700)**—We have enhaced the IS-IS multi-instance feature to support multiple IS-IS instances on the same logical interface with instance identifier TLV 7.

  Include the `instance-id` statement at the `[edit protocols isis-instance` *`name`* hierarchy level.

  [See How to Configure Multiple Independent IGP Instances of IS-IS.]

- **GLB multi-link support on IP Fabric (QFX5240-64OD and QFX5240-64QD)**—We are extending Global Load Balancing (GLB) to support multiple paths between spine and top-of-rack switches on a 3-stage Clos IP fabric.

  To enable GLB for multi-link on a 3-Clos IP fabric, include the `glb-multilink-mode` *`max-val|avg-val`* statement at the `[edit forwarding-options enhanced-hash-key]` hierarchy level. By default, the spine advertises the average quality of all links. Enable GLB globally at `[edit protocols bgp]` hierarchy level.

  [See Configure GLB on 3-CLOS IP Fabric with Multilinks.]

- **Support for BGP bandwidth unequal load balancing for EVPN-VXLAN routes (QFX5130 and QFX5700)**—We have extended the asymmetric traffic distribution support to EVPN-VXLAN routes. You can enable this feature based on path cost. This approach optimizes BGP traffic across paths with different bandwidths. The control plane calculates balance values for each path, and the Packet Forwarding Engine uses these values to distribute traffic proportionally ensuring efficient utilization of network resources.

  [See **BGP Link-Bandwidth Community**.]

- **Support for 256-way ECMP (QFX5130-32CD, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Use this feature to increase the number of direct BGP peer connections, improve latency, and optimize data flow by configuring up to 256 ECMP next hops for external BGP peers.

  [See Example: Load Balancing BGP Traffic.]

- **Proxy ARP and NDP Proxy Support with VRRP (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008 and PTX10016)**—Proxy ARP and NDP Proxy support with VRRP allows a router to use the VRRP virtual MAC address when responding to proxied ARP, NDP, or DAD requests. When an interface is configured with both proxy and VRRP, only the VRRP master responds to proxy requests, and the response uses the VRRP virtual MAC instead of the physical interface MAC. This behavior applies to restricted and unrestricted modes for Proxy ARP, NDP proxy, and DAD proxy. By using the virtual MAC address, the feature ensures consistent address resolution and seamless connectivity during VRRP failover events, preventing disruption when mastership changes occur. This feature is supported on Junos OS Evolved platforms.

Services Applications

- **IFA 2.0 initiator and terminator support with new live and probe modes (QFX5240-64OD and QFX5240-64QD)**—Collect per-flow metrics in the data plane using Inband Flow Analyzer (IFA) 2.0 to monitor latency and congestion, to trace paths and flows, and to analyze per-hop and end-to-end telemetry without control-plane involvement or forwarding impact. The new modes are as follows:

  - In live mode, IFA 2.0 uses the original packets as IFA probe packets.
  - In probe mode, IFA 2.0 uses copies of the original packets as the IFA probe packets.

  Configure an IFA initiator node to initiate IFA probe packets:

  - In live mode, the initiator inserts IFA headers and metadata into the live packets.
  - In probe mode, the initiator samples the flow and copies that sample to create the probe.

  Configure an IFA terminator node to parse the IFA header in the packet to determine if it is a live or probe packet:

  - If it is a live packet, then the terminator strips the IFA layer and forwards the modified live packet to its destination.
  - If the packet is a probe packet, then the terminator drops the packet and sends a mirror copy to the collector.

  To specify the mode, configure the `mode (live | probe)` statement at the `[edit services inband-flow-telemetry]` hierarchy level. The default mode is `probe`. Use the `show services inband-flow-telemetry global` command to see probe results that include the mode information.

[See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring, inband-flow-telemetry, and show services inband-flow-telemetry.]

Software Installation and Upgrade

- **Firmware upgrade support (QFX5241-32OD, QFX5241-64OD, and QFX5241-64QD)**—We support the `request system firmware upgrade` command to upgrade firmware.

  [See request system firmware upgrade (Junos OS Evolved).]

- **Secure boot and common BIOS support (QFX5241-64OD and QFX5241-64QD)**—The Secure boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms.

  [See Junos OS Evolved Overview and request system firmware upgrade (Junos OS Evolved).]

- **Load set-formatted and XML-based configuration files for ZTP (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5140-24CD8O, , QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, QFX5700, and QFX5700E)**—You can load set-formatted or XML-based configuration files when your device provisions for zero-touch provisioning (ZTP). Reuse existing set-style or XML-based configuration files for automated onboarding to avoid converting them to hierarchical syntax. Provide the configuration file in set format or in XML and specify the configuration file name under the DHCP vendor configuration options.

  [See Zero Touch Provisioning.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Path quality profile sharing for GLB multi-link support on IP Fabric (QFX5240-64OD, and QFX5240-64QD)**—While we support larger Clos networks and more GPUs, the TH5 chipset can only support 64 profiles. In Clos networks with five or more stages, some nodes, like super spines exceed 64 next-next-hop nodes. We can reuse the profiles under specific conditions to support more than 64 next-next-hop nodes. We support GLB for Clos networks with profile sharing in hyperscaler artificial intelligence/machine learning (AI/ML) fabrics containing tens of thousands of leaves or GPUs.

  [See profile-sharing.]

- **Delay normalization for OSPF Flexible Algorithm metrics and advertisements across IGP instances (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD,**

PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)—Use delay normalization to compute and advertise a normalized delay metric for Flexible Algorithm, to improve path-selection consistency across all IGP instances. The device normalizes each received delay, compares each value with the previously saved normalized value, and triggers link-state advertisement (LSA) generation when the values differ.

Delay normalization is disabled by default. To enable and configure delay normalization, use the `normalize` *`interval offset`* statement at the `[edit protocols ospf` *`area interface`* `delay-measurement]` hierarchy level.

[See delay-measurement (Protocols OSPF) and How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering.]

- **Selectively control per-prefix backup paths with OSPF import policy (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can selectively enable backup paths for specific prefixes to optimize redundancy and resource utilization. By default, a configured backup path applies to all prefixes. To exclude specific prefixes or ranges, create an OSPF import policy and configure the `no-backup` option in the `then` clause of the policy to suppress backup path installation for matching routes. You can reserve backup protection for critical prefixes while preventing unnecessary backups for others.

  [See Understanding Backup Selection Policy for OSPF Protocol.]

- **Preference-based Path Selection of L-OSPF Flexible Algorithm routes (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can control path selection by configuring the preference for L-OSPF Flexible Algorithm routes in `inetcolor.0` and `mpls.0`.

  Configure `flex-algorithm-preference` statement at the `[edit protocols ospf]` hierarchy level to prioritize desired routes and improve traffic engineering across IP and MPLS domains.

- **Policy-based redistribution of OSPF prefix SIDs across IGP instances (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008)**—You can redistribute Segment Routing (SR) prefix-SIDs across OSPF IGP instances using route policy without explicitly specifying a prefix-segment index. This feature standardizes SR labels across instances and improves operational efficiency. Configure a policy with the `from prefix-segment` statement to match routes carrying prefix-segment information. In the `then` clause, use `prefix-segment redistribute` to inherit segment information from the matched route. We also support stitching `mpls.0` routes to enable interoperability between different IGP instances.

Storm Control

- **Storm control (QFX5220, QFX5230, and QFX5240)**—Storm control enables a switch to monitor traffic levels and drop broadcast, multicast, and unknown unicast packets when traffic exceeds a specified level, called the storm control level. This feature prevents packet proliferation and network degradation.

  [See Understanding Storm Control.]

System Management

- **Display password expiration information for users (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Configure the `maximum-lifetime` password option so that user passwords expire after the specified number of days. Root users can view other users' password expiration details on the device, and non-root users can view their own. Use the `show system login password-expiry-information` command to view password expiration information.

  [See password (Login).]

Additional Features

We've extended support for the following features to the platforms shown in parentheses:

- **Support for Precision Time Protocol (PTP) transparent clock** (QFX5240-64OD and QFX5240-64QD)

  [See Precision Time Protocol (PTP) Overview.]

- **EVPN-VXLAN Layer 3 gateway** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD). Support includes:

  - Layer 3 (L3) VXLAN gateway in edge-routed bridging (ERB) fabric
  - Up to 256 integrated routing and bridging (IRB) enabled VLANs
  - Layer 3 underlay that supports IRB and link aggregation groups (LAGs)
  - ECMP in the underlay
  - IPv4 and IPv6 virtual gateway MAC address support for IRB interfaces
  - In-service software upgrade (ISSU) for L3 gateway functionality

  [See Understanding EVPN with VXLAN Data Plane Encapsulation, Example: Configuring an EVPN-VXLAN Edge-Routed Bridging Fabric with a Virtual Gateway, Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network, and IP Fabric Underlay Network Design and Implementation.]

- **EVPN-VXLAN Type 5 route** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD). Support includes:

- ECMP in underlay and overlay
- L3 underlay with integrated routing and bridging (IRB) and LAGs
- IPv6 user traffic support
- In-service software upgrade (ISSU) support
- Type 5 seamless stitching
- Symmetric IRB
- Coexistence of Type 2 and Type 5 routes

[See Understanding EVPN with VXLAN Data Plane Encapsulation, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Understanding EVPN Pure Type 5 Routes, EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN, NSR and Unified ISSU Support for EVPN, and irb-symmetric-routing.]

- **EVPN-VXLAN Layer 2 gateway** (QFX5240-64OD and QFX5240-64QD). Support includes:
  - Up to 508 VXLAN-enabled VLANs
  - L3 underlay with integrated routing and bridging (IRB) and LAGs
  - ECMP in the underlay
  - VXLAN tunnel endpoint (VTEP) ingress and egress counters
  - Broadcast, unknown unicast, and multicast (BUM) traffic forwarding by ingress replication only
  - MAC move functionality
  - Core isolation support as a control-plane feature
  - `no-mac-ip-learning` statement to disable `mac-ip` learning
  - MAC move limit enforcements with port disable and VLAN member disable actions
  - In-service software upgrade (ISSU) for L3 gateway functionality

  [See Understanding EVPN with VXLAN Data Plane Encapsulation, IP Fabric Underlay Network Design and Implementation, Overview of Hierarchical ECMP Groups on QFX5200 Switches, NSR and Unified ISSU Support for EVPN, and irb-symmetric-routing.]

- **Support for CoS and firewall filtering and policing on EVPN-VXLAN networks** (QFX5240-64OD and QFX5240-64QD)

  [See CoS Support on EVPN VXLANs and Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]

- **Support for sFlow technology on EVPN-VXLAN** (QFX5240-64OD and QFX5240-64QD)

  [See Overview of sFlow Technology.]

- **Support for port mirroring and analyzers on EVPN-VXLAN** (QFX5240-64OD and QFX5240-64QD)

  [See Port Mirroring and Analyzers in an EVPN-VXLAN Environment.]

- **Fast reroute for egress link protection (ELP) in EVPN-VXLAN multihoming environments** (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)

  [See Fast Reroute for Egress Link Protection with EVPN-VXLAN Multihoming and reroute-address.]

- **Simplified configuration for ESI LAGs with EVPN dual-homing (EZ-LAG)** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)

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

- **Inband Flow Analyzer (IFA) 2.0 transit node support** (QFX5230-64CD)

  [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.]

- **Support for performance monitoring and TCA** (QFX5130-32CD). Use performance monitoring for QFX5130-32CD switches to measure current and historical metrics. The switch accumulates these metrics into 15-minute and 1-day intervals. You can configure the 15-minute interval length. You can view these metrics by using the show interfaces transport pm command. This approach helps you manage optical transport links more efficiently.

- **Support for performance monitoring and TCA** (QFX5700). Use performance monitoring for QFX5700 switches to measure current and historical metrics. The switch accumulates these metrics into 15-minute and 1-day intervals. You can configure the 15-minute interval length. You can view these metrics by using the show interfaces transport pm command. This approach helps you manage optical transport links more efficiently. Also, 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. See show interfaces transport pm.

- **Supported transceivers, optical interfaces, and DAC cables** (ACX7020, ACX7100-32C, ACX7100-48L, ACX7348, ACX7509, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5220-32CD, QFX5230-64CD, QFX5240-64QD, QFX5240-64OD, QFX5700, and QFX5700E). 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.

- **Separate firmware installation packages** (QFX5130, QFX5220, QFX5230, QFX5240, and QFX5700). You can manage firmware upgrades using standalone firmware installation packages. The names for these packages begin with the prefix `jfirmware-junos-evo-install*`. You add a firmware package to the

system with the `request system software add` command. After you've added the firmware package to your system, you update the firmware for a hardware component using the `request system firmware upgrade` command.

[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).]

- **Support for wildcard mask match condition for source-address/destination-address match conditions for `inet6` address family** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)

  [See Understanding Firewall Filter Match Conditions and IPv6 Wildcard Mask Match Conditions.]

- **Support for SRv6 network programming** (QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, and QFX5241E-64OD). We support the following SRv6 features:

  - Topology Independent Loop-Free Alternate (TI-LFA) backup path for SRv6
  - Layer 3 Services over SRv6 in BGP (END.DT4 and END.DT6)
  - Operations, Administration and Management (OAM) ping for SRv6
  - SRv6 traceroute
  - Static SR-TE for SRv6
  - SRv6-TE
  - SRv6 header compression (uSID or micro-SID)

  [See **Understanding SRv6 Network Programming in IS-IS Networks**.]

- **Per port (IFL) `no-mac-learning` support** (QFX5130-32CD, QFX5130-48C, and QFX5700)

  [See MAC Learning and `no-mac-learning`.]

- **Resilient hashing support for overlay ECMPs with Type 5 routes** (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)

  [See Configure Resilient Hashing on ECMP Groups.]

- **Overlapping VLAN support for ERB EP style logical interfaces** (QFX5130-32CD, QFX5130-48C, and QFX5700)

  [See Overlapping VLAN Support Using Multiple Forwarding Instances or VLAN Normalization and `forwarding-instance`.]

- **Renaming OpenSSH implementation to JSSH** (All platforms). The OpenSSH implementation in Junos OS Evolved is renamed "JSSH" to highlight its customized nature.

  [See ssh.]

What's Changed

IN THIS SECTION

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

EVPN

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

General Routing

- SSH key options for user account credentials. You can configure key-options *key-options* option at the `set system login user` *user* `authentication [ssh-rsa|ssh-ecdsa|ssh-ed25519] ssh key` hierarchy level.

  [See login.]

- **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> <b>show system alarms</b> 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
```

- Displays the event log of learned MAC addresses. By default mac-learning-logs are stored in UTC timestamps. To view the logs in system timezone, use the show ethernet-switching mac-learning-log use-system-timezone command. The show ethernet-switching mac-learning-log use-system-timezone command also prints the time zone abbreviations [IST, UTC, etc] in the timestamp. To view the logs in system timezone by default by using the show ethernet-switching mac-learning-log command, you need to configure the system-timezone statement at the [edit protocols l2-learning mac-learning-log] hierarchy level.

- When you run the request vmhost zeroize command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines.

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

- On the MPC7E-10G line card, when you configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet ports, use the speed statement at both the edit interfaces <interface name> gigether-options and edit interfaces interface <name hierarchy> levels.

- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked.

- **New option for debug collector data storage path**—We've included the option `outdir` to specify an output directory for storing debug collector data in a customised path. This allows you to organise and access diagnostic information more efficiently, adapting storage to your specific requirements.

  [See request system debug-info.]</p>

- **High-power optics support with CLI configuration option (QFX5240 and QFX5241)**—You can enable high-power optics across all ports by configuring the high-power-mode option for each port. This feature supports up to 32 high-power modules, allowing you to benefit from enhanced connection capabilities. Ensure you configure the necessary settings to initialize and utilize high-power optics effectively, optimizing your network's performance.

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.
- **Default :0 sub interface for single subport (QFX5220-32CD)**—When you configure number-of-sub-ports 1 using the `set chassis fpc` *`fpc number`* `pic` *`pic number`* `port` *`port number`* `speed` *`speed`* `number-of-sub-ports 1` command, the `:0` sub-interface is created automatically. This configuration delivers deterministic sub-interface naming, simplifying provisioning, automation templates, and monitoring across platforms, maintaining parity with other implementations.

Platform and Infrastructure

- **Tacacs authorisation support for local authentication without password**—Starting in Junos OS Evolved Release 25.4R1, you need not configure password under `edit system authentication-order` to enable password-options.
- **Commit validation for unique user IDs**—We have added support to validate the user configuration to ensure that each user is assigned a unique UID. A commit fails if duplicate UIDs are detected, ensuring stronger validation and preventing identity conflicts. Previously, a commit was successful even when multiple users shared the same UID, triggering only a warning and logging a syslog message.

User Interface and Configuration

- **Stale ui-state.db data in persistent NETCONF sessions post-mgd restart**— Existing NETCONF sessions might fetch stale data from ui-state.db after mgd -N restart. New sessions correctly map the refreshed database. Scripts must establish new sessions post-restart to access updated values. Functional configuration remains unaffected. [Script failures monitoring "local-host" NETCONF sessions]-Scripts might fail when including "local-host" NETCONF sessions in monitoring operations. Internal sessions are now excluded from tracking. Scripts must filter out "local-host" sessions. No impact to internal application functionality.
- **Generate genstate YANG modules on Junos devices**—You can use `show system schema` operational command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

  [See show system schema.]

Known Limitations

IN THIS SECTION

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

- When 40G and 100G xcvrs are used together on the 16x100G PIC of QFX5700 such that they are connected to the same Broadcom Serdes Core on the PFE side, then one or more of the links from these interfaces might not come up. With this PIC, ports 0-3, 4-7, 8-11 and 12-15 share one Broadcom Serdes Core each.PR1867341

Routing Protocols

- With IGMP Snooping in EVPNoVXLAN with Mrouter sync feature enabled with "clear bgp neighbor all" in router where the igmp query is learnt locally when this AE goes from up to down state and all its mrouter states are removed and then the AE comes up [after core isolation] and the traffic coming into the system from another esi-lag will get dropped till all the query is learnt in AE and it becomes mrouter. This is the design limitation as it has core isolation in place and in this scenario we expect the traffic loss [12 sec] i.e till the time AE has learnt all the incoming queries from all vlan and marks its ifl as mrouter interface in corresponding VLAN.PR1876740

Open Issues

IN THIS SECTION

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

- 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
- 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
- Insufficient headroom cells are allocated per Priority-Group(PG) when cable-length is configured <100mts(default length) under 'congestion-notification-profile' configuration and entire traffic frame size is <128B. Because of this one can observe the Ingress drops on the Port/Interface. PR1846357
- Hitting with input-errors and framing-errors on channelizing QSFP56-DD-400GBASE-DR4. PR1848109
- During USB boot with a fully populated system, error messages were observed in the boot logs indicating IRQ was not handled. This has no functional impact, as the driver is not used on the QFX5241-32/64 platforms. PR1868237
- Graceful JOJI of FPC is the recommended way instead of ungraceful.PR1874712
- With a back-to-back connection using 800G ZR* optics, if one end is configured with native speed (default) and other end is configured wit 8x100G mode, then the expectation is that link is down on both ends because of the configuration mismatch between the ends; but one of the two ends shows as up. The same behaviour is presented by 400G ZR* optics also. PR1884559

- QFX5130 macsec: on macsec enabled interfaces, after evo-pfemand restart with CLI `restart evo-pfemand`, ping might not work sometimes.PR1891444
- A little higher convergence (~1-2 seconds) seen when following conditions are met. 1. A MH leaf is connected to more than 1 spine. 2. In 2 spines scenario, initially one of uplinks is down on the leaf. 3. BFD packet must hash to this path. 4. Wait for (4-6 seconds) and then make second uplink down from another spine (now leaf is isolated). 5. Convergence of ~1-2 seconds is seen. Higher convergence is not seen when 1. Both spine uplinks are down at the same time. 2. when remote leaf is rebooted.PR1894826
- Convergence of 5-13 seconds seen when uplinks from remote leafs to core are brought up. This is for type5 tunneled traffic only.PR1895977
- QFX5700: 16x100G PIC: Commit check added when 40G and 100G xcvrs are put to share a Broadcom serdes core. PR1900911
- This issue is seen intermittently on few devices. Probably due to some other kernel modules/ processes using the i2c lines for these SFP voltage controllers. Impact: During firmware upgrade of sfpg1/g2/g3. once in a few iterations there will be write failures, which is expected in situations where another driver is interacting with device on same i2c line. Retry will make it work. Recovery: There is no other impact on the device except upgrade failure, which is rare if customer wants to upgrade once. in case encountered, retry will solve the issue.PR1908672
- On a QFX5241/QFX5240 platform, changing the MTU value of an interface does not take effect when the UFT profile is also being changed in the same commit. The MTU configuration should be applied in a separate commit.PR1914797

Resolved Issues

IN THIS SECTION

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

- QFX5700 might report SSD DMA error on secondary disk. PR1848468
- Source subnet mask shows as 0 in sFlow samples on Junos OS Evolved platforms. PR1854029
- Ad hoc swapping of 400G DAC with 400G DR4 (or vice-versa) optic leads to traffic loss. PR1862711
- Interfaces won't come up after loading a configuration file that contains multiple forwarding profiles and MTU changes. PR1877286
- ECN CE bit is not set in an EVPN-VxLAN scenario. PR1880166
- One physical member link of ESI LAG interface flapping causes mac-moves and packet loss. PR1884214
- Static ARP/ND not supported on EVPN enabled IRB interfaces. PR1887492
- The egress queue got stuck, causing traffic drops even without congestion, after deactivate and activate class-of-service. PR1888213
- Enable high-power optics on all ports, limited to a maximum of 32 high-power modules. PR1890645
- Host-bound packet drops when using decapsulate filter action in non-default routing-instances. PR1896239
- [QFX5K-EVO]Spine which has VGA does not reply to traceroute in EVPN-VXLAN CRB environment. PR1896299
- QFX5241-32OD: Mismatch with optic description for "OSFP-2x400G-LR4-10 740-174937" and "OSFP-2x400G-FR4 740-174935". PR1904999

Interfaces and Chassis

- Allow `number-of-sub-ports 1` configuration to create the ':0' sub-interface. PR1900918
- Traffic loss and slow convergence might occur when rebooting a QFX5240 spine device with mixed 400G/800G links. PR1905737

Routing Protocols

- Scaled BGP sessions stuck in idle after interface rollback. PR1880630

Junos OS Evolved Release Notes for Third-Party Whitebox

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 25.4R1 for UfiSpace-S9500-22XST, UfiSpace-S9600-32X, UfiSpace-S9600-72XC, and UfiSpace-S9600-102XC devices.

NOTE: Junos OS Evolved Release 25.4R1 for UfiSpace products is an early limited support release for trial purposes only. Contact Juniper Networks to request a trial of this product.

What's New

IN THIS SECTION

Learn about new features introduced in this release for Third-Party Whitebox devices.

Hardware

- **UfiSpace S9600 series (Third-party white box)**—The S9600 series Open Aggregation Routers from UfiSpace are third-party white boxes that support Junos OS Evolved for a broadband network

gateway (BNG) application. These white boxes enable you to deploy a BNG solution with Junos OS Evolved without hardware vendor lock-in. The following UfiSpace products are supported in this release:

- **UfiSpace-S9600-32X**—A 2-U appliance that features 4 SFP28 and 32 QSFP28 interfaces to allow high bandwidth connectivity in the core. The UfiSpace-S9600-32X supports Junos OS Evolved when the white box functions as a spine router in a spine-leaf network architecture for a BNG application.
- **UfiSpace-S9600-72XC**—A 2-U appliance that features 64 SFP28 and 8 QSFP28 interfaces to allow high fanout connectivity to top-of-rack equipment. The UfiSpace-S9600-72XC supports Junos OS Evolved when the white box functions as a high fanout leaf router in a spine-leaf network architecture for a BNG application.
- **UfiSpace-S9600-102XC**—A 2-U appliance that features 96 SFP28 and 6 QSFP28 interfaces to allow very high fanout connectivity to top-of-rack equipment. The UfiSpace-S9600-102XC supports Junos OS Evolved when the white box functions as a very high fanout leaf router in a spine-leaf network architecture for a BNG application.

The UfiSpace white boxes come preloaded with the UfiSpace ONIE bootloader from the factory. You configure the bootloader to boot Junos OS Evolved from a repository that you prepopulate with the Junos OS Evolved image or ISO downloaded from the Juniper Networks software download site. After Junos OS Evolved boots, you can configure the Junos OS Evolved BNG features through CLI or REST API.

> **NOTE**: Junos OS Evolved is supported on these UfiSpace white boxes only for the BNG application and only as described above. You cannot run Junos OS Evolved on these white boxes for any other application or in any other manner including as a general-purpose router.

[See Feature Explorer for the complete list of features for this or any other platform.]

[See UfiSpace for information on the UfiSpace S9600 series Open Aggregation Routers.]

Junos Telemetry

- **Export sensor data to external collectors using Junos telemetry (UFISPACE-S9600-32X, UFISPACE-S9600-72XC, and UFISPACE-S9600-102XC)**—Use Junos telemetry to stream device and operational statistics from UFISPACE-S9600-32X, UFISPACE-S9600-72XC, and UFISPACE-S9600-102XC devices to external collectors. Monitor system health by subscribing to sensor paths for chassis temperature, CPU and memory utilization, fan speed, transceiver parameters, interface statistics, timex statistics, and subscriber details. Obtain platform, routing, process, system resource, port, optical, H-policer, PSM fan, and event metrics by subscribing to the respective resource paths.

  For more information, see Junos YANG Data Model Explorer.

Layer 2 VPN

- **Support for anycast gateways in Layer 2 VPNs (UFISPACE-S9600-32X, UFISPACE-S9600-72XC, and UFISPACE-S9600-102XC)**—We support Layer 2 VPNs services using an anycast gateway address. The anycast address is a single IP Address that is configured on all the multiple devices. This feature provides resiliency in the network and load balancing among the devices that share the anycast address.

  For more information, see Understanding Layer 2 VPNs.

- **Stitching a Layer 2 VPN to VPWS (UFISPACE-S9600-32X, UFISPACE-S9600-72XC, and UFISPACE-S9600-102XC)**—In a spine-leaf topology, the spine device is a central point where L2 VPN sevices from the backbone and vpws from the access layer converge. We support stitching Layer 2 VPN tunnels to the VPWS tunnels without requiring a physical hardware interface. To enable this feature, configure the spine devices to stitch Layer 2 VPN services to VPWS by configuring Layer 2 interwork.

  For more information, see Understanding Layer 2 VPNs.

Precision Time Protocol (PTP)

- **Synchronous Ethernet support (UfiSpace-S9600-72XC and UfiSpace-S9600-102XC)**—The UfiSpace-S9600-72XC and UfiSpace-S9600-102XC routers support synchronous Ethernet compliant with ITU-T G.8262/G.8262.1 with clock recovery on all 10/25/40/100 Gbps ports, including support for enhanced Ethernet Synchronization Messaging Channel (ESMC), LAG, and the Synchronous Ethernet MIB.

  [See Synchronous Ethernet Overview.]

Subscriber Management and Services

- **Subscriber management (UFISPACE-S9600-72XC and UFISPACE-S9600-102XC)**—Use the `set system services subscriber-management enable` command to enable the following subscriber management services for DHCP, PPPoE, and L2TP LAC subscribers:

  - Firewall filters
  - CoS
  - Lawful intercept
  - Multicast
  - HTTP-redirect
  - RADIUS
  - Resource monitoring (RSMON)

- Subscriber and service accounting.

  We introduce the following new CLI commands:

  - `show system subscriber-management health <clients | endpoints | services>` command to view client, endpoint, and service status
  - New options for `show subscribers` command: `interface`, `class-of-service`, and `firewall`.
  - `clear subscribers` command to clear subscribers firewall counters.

  [See enable (Enhanced Subscriber Management, Broadband Subscriber Services User Guide, Broadband Subscriber Management Getting Started Guide, and CLI Reference.]

- **Support for direction field in mirror header for subscriber secure policy (UFISPACE-S9600-72XC and UFISPACE-S9600-102XC)**—Use subscriber secure policy to send mirrored traffic to a mediation device on a per-subscriber basis. The mediation device uses a mirror header to differentiate multiple mirrored streams from different sources. You can use the `set system services radius-flow-tap extended-mirror-header` command to enable direction field in the mirror header. The direction field helps the mediation device to quickly identify the traffic direction and apply the appropriate rules defined by the law enforcement agency.

  [See Subscriber Secure Policy Overview, Mirror Header Format (Junos OS Evolved with Payload Direction), and radius-flow-tap.]

- **Support for initiating DHCP relay session over static IFL (UFISPACE-S9600-32X, UFISPACES9600-72XC, and UFISPACE-S9600-102XC)**—You can initiate a DHCP relay session over a static IFL (logical interface). To do this, use the `set interfaces interface_name auto-configure static-vlan-identity interface-tag tag_name` and `set interfaces interface-name unit logical-unit-number interface_tag tag_name` commands. The *tag_name* and the mapped dynamic profile properties allow the system to identify the subscriber interface and create a dynamic VLAN to initiate a DHCP relay session. You can use the `show subscriber (detail | extensive)` command to view the initiated sessions.

  [See interface-tag (Junos OS Evolved), auto-configure (Interfaces), unit (Interfaces), and show subscribers.]

- **Filters and filter services support (UFISPACE-S9600-32X, UFISPACES9600-72XC, and UFISPACE-S9600-102XC)**—You can configure filters and filter services to control the subscriber's data traffic.

  The supported filters include:

  - Ascend-Data-Filter
  - RPF
  - Classic filters.

  The supported filter services include:

- Prefix list match to filter traffic based on matching IP prefixes
- Sample filter action to apply sampling for selected traffic
- HTTP redirect to redirect HTTP traffic
- Allowlist or walled garden to permit only specified traffic.

[See Ascend-Data-Filter Policies for Subscriber Management Overview, Unicast RPF in Dynamic Profiles for Subscriber Interfaces, Classic Filters Overview, Parameterized Filter Match Conditions for IPv6 Traffic, Parameterized Filter Match Conditions for IPv4 Traffic, Parameterized Filter Nonterminating and Terminating Actions and Modifiers, HTTP Redirect Service Overview, and Configuring a Walled Garden as a Firewall Service Filter.]

- **Support for L2TP CSUN and counters (UFISPACES9600-72XC and UFISPACE-S9600-102XC)**—You can enable or disable Layer 2 Tunneling Protocol (L2TP) connection speed update notifications (CSUNs) for L2TP network servers (LNSs) that do not support receiving CSUNs from L2TP access concentrators (LACs), as defined in RFC5515. This enhancement allows overriding the `connection-speed-update` configuration by sending the 26-159 VSA attribute `L2tp-Csun-enable` in the access-accept message.

  Additionally, accounting messages include new VSA attributes to enable counters for policers and queues, providing improved monitoring capabilities. Use the following CLI commands to allow the system to send the counter statistics in RADIUS accounting messages:

  - `set dynamic-profiles` *`profile-name`* `firewall family` *`family`* `filter` *`filter-name`* `service-accounting` for policer counter statistics
  - `set dynamic-profiles` *`profile-name`* `class-of-service interfaces` *`interface-name`* `unit` *`unit-name`* `service-accounting` for queue counter statistics

  [See Configuring the Reporting and Processing of Subscriber Access Line Information, Juniper Networks VSAs Supported by the AAA Service Framework, AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS, Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile, class-of-service (Dynamic Profiles), and filter (Dynamic Profiles Filter Creation).]

- **Support for enhanced hierarchical policer under dynamic profile (UFISPACES9600-72XC and UFISPACE-S9600-102XC)**—You can configure an enhanced hierarchical policer under a dynamic profile to rate-limit subscriber traffic. Traffic policing is supported at four levels of hierarchies with respect to the traffic priority: High, Medium-High, Medium-Low, and Low. Use the `set dynamic-profiles` *`profile_name`* `firewall enhanced-hierarchical-policer` *`policer_name`* command to create the enhanced hierarchical policer. You can apply the policer to a subscriber interface as a filter action for aggregate traffic levels. Include the `logical-interface-policer` statement to rate-limit across multiple protocol families (inet and inet6), without requiring separate policer instances for each family.

Additionally, we support new VSA tags and variables to define values for the enhanced hierarchical policer default parameters.

[See Enhanced Hierarchical Policer Overview (Junos OS Evolved), Configure Enhanced Hierarchical Policer (Junos OS Evolved), enhanced-hierarchical-policer (Dynamic Profiles), Juniper Networks VSAs Supported by the AAA Service Framework, and Predefined Variables in Dynamic Profiles.]

- **PAO REST Interface support for DT-A4 Service Edge routers (UFISPACE-S9600-32X)**—PAO REST Interface support allows the POD Access Orchestrator (PAO) to manage and monitor DT-A4 routers through a REST-based interface. On Service Edge (leaf) routers, PAO uses this interface to dynamically create and delete empty-session VLANs that support subscriber onboarding workflows. Operators enable the feature by configuring the external-management hierarchy under system services subscriber-management. Spine routers support REST-based management and monitoring operations but do not support empty-session VLAN handling. This feature simplifies orchestration, supports automated recovery scenarios, and improves operational visibility in DT-A4 environments.

Upgrade Your Junos OS Evolved Software

For products impacted, see Feature Explorer.

Follow these steps to upgrade your Junos OS Evolved software:

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

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

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

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that Juniper Networks has developed over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see Juniper Flex Program.

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.

  https://apps.juniper.net/feature-explorer/

- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.

  https://prsearch.juniper.net/InfoCenter/index?page=prsearch

- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.

  https://apps.juniper.net/hct/home

> **NOTE:** To obtain information about the components that are supported on the devices and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6.

  https://pathfinder.juniper.net/compliance/

Requesting Technical Support

IN THIS SECTION

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at https://www.juniper.net/content/dam/www/assets/resource-guides/us/en/jtac-user-guide.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://support.juniper.net/support/

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://support.juniper.net/support/
- Call 1-888-314-5822 (toll free, US & Canada). If outside the US or Canada, use a country number listed from one of the regional tabs listed on the Contact Support page.
- Federal Government Support: 1-833-900-1454.

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

Revision History

22 December 2025—Revision 1, Junos OS Evolved Release 25.4R1