

Release Notes

Published
2025-01-28

Junos OS Evolved Release 24.4R1

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 24.4R1.

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

What's New | 1

Class of Service	2
High Availability	3
Interfaces	3
Junos Telemetry Interface	3
MPLS	4
Multicast	6
Network Management and Monitoring	7
Precision Time Protocol (PTP)	8
Routing Policy and layer2-policer Firewall Filters	8
Software Installation and Upgrade	8
Source Packet Routing in Networking (SPRING) or Segment Routing	9
Subscriber Management and Services	10
Additional Features	10

What's Changed | 14

Known Limitations | 18

Open Issues | 19

Junos OS Evolved Release Notes for PTX Series

What's New | 21

Hardware	22
Authentication and Access Control	75
Chassis	75

Class of Service	76
Ethernet Switching and Bridging	77
High Availability	78
Interfaces	78
Junos Telemetry Interface	78
Layer 2 VPN	82
MACsec	82
MPLS	83
Multicast	83
Multichassis Link Aggregation (MC-LAG)	84
Network Management and Monitoring	85
Routing Policy and layer2-policer Firewall Filters	86
Routing Protocols	87
Services Applications	87
Software Installation and Upgrade	89
Source Packet Routing in Networking (SPRING) or Segment Routing	89
Additional Features	90

What's Changed | 102

Known Limitations | 107

Open Issues | 107

Junos OS Evolved Release Notes for QFX Series

What's New | 109

Hardware	110
Authentication and Access Control	144
Additional Features Optimized for AI-ML Fabrics	145
Chassis	147

Class of Service	148
Forwarding Options	148
High Availability	148
Interfaces	149
Junos Telemetry Interface	149
Multicast	152
Network Management and Monitoring	152
Platform and Infrastructure	153
Precision Time Protocol (PTP)	153
Routing Policy and layer2-policer Firewall Filters	153
Routing Protocols	154
Software Installation and Upgrade	154
Additional Features	155

What's Changed | 157

Known Limitations | 161

Open Issues | 163

Upgrade Your Junos OS Evolved Software | 164

Documentation Updates | 165

Licensing | 165

Finding More Information | 165

Requesting Technical Support | 166

Revision History | 167

Junos OS Evolved Release Notes for ACX Series

IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 14](#)
- [Known Limitations | 18](#)
- [Open Issues | 19](#)

These release notes accompany Junos OS Evolved Release 24.4R1 for ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348 and ACX7509 devices. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- [Class of Service | 2](#)
- [High Availability | 3](#)
- [Interfaces | 3](#)
- [Junos Telemetry Interface | 3](#)
- [MPLS | 4](#)
- [Multicast | 6](#)
- [Network Management and Monitoring | 7](#)
- [Precision Time Protocol \(PTP\) | 8](#)
- [Routing Policy and layer2-policer Firewall Filters | 8](#)
- [Software Installation and Upgrade | 8](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 9](#)
- [Subscriber Management and Services | 10](#)

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 24.4R1, click the Group by Release link. You can collapse and expand the list as needed.

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

The following sections highlight the key features in this release.

Class of Service

- **Support for eight scheduler priority levels (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)** —The listed ACX7000 routers support eight scheduler priority levels for port scheduling: low-latency, strict-high, high, medium-high, medium-low, low-high, low-medium, and low. Low-high and low-medium are the two newly introduced priority levels. Configure scheduler priority at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.

[See [Schedulers Overview for ACX Series Router](#) and [priority \(Schedulers\)](#).]

- **Hierarchical CoS support at four levels (ACX7024, ACX7100-32C, ACX7100-48L, ACX7348, and ACX7509)** —The listed ACX7000 routers support four levels of hierarchical class of service (CoS). You can configure scheduling at four levels:
 - Physical interface level
 - Logical interface set level
 - Logical interface level
 - Queues

Define scheduler properties at each level by applying traffic control profiles at the physical, logical interface set, and logical interface levels and schedulers to the queues.

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

High Availability

- **Support for high availability (ACX7348)**—Use the high availability feature to achieve routing engine redundancy and reliability in packet-based communications. We support both graceful Routing Engine switchover (GRES), graceful restart (GR), and nonstop active routing (NSR). ACX7348 does not support PFE redundancy.



NOTE: If you alter the present flow or introduce a new flow during the Routing Engine switchover, the convergence does not take place until switchover completes. Topology changes during the switchover are applied only after switchover. Traffic loss and minor statistics loss is expected during switchover.

On ACX7348 device, if a configuration belonging to features like Broadband network gateway (BNG), VXLAN, sFlow, J-Flow, and port mirroring is detected during Routing Engine switchover, then datapath is reset and traffic reconvergence is seen.

[See [Understanding High Availability Features](#).]

Interfaces

- **Support for MC-LAG active/standby mode on L2CKT connections (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—You can configure multichassis link aggregation (MC-LAG) interfaces in active/standby mode for Layer 2 circuit (L2CKT) connections. This configuration enables you to set up a standby L2CKT connection on your interface, providing you with a backup link to maintain your service if the primary link fails. Use existing commands to configure this feature in the interface hierarchy.

[See [Getting Started with MC-LAG](#) and [Configuring Redundant Pseudowires for Layer 2 Circuits and VPLS](#).]

- **Support for MTU value of 9996 bytes (ACX7509)**—You can configure a maximum transmission unit (MTU) size of 9996 bytes to optimize network performance for large data packets. To configure the MTU value, use the set `interfaces interface-name mtu` command.

[See [show interfaces extensive](#).]

Junos Telemetry Interface

- **Sensor support for system, CPU, and memory statistics (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—We support periodic streaming of system

statistics, CPU parameters and memory-related parameters. Statistics support the health-monitoring application. The data model `openconfig-system.yang` version 0.10.0 supports this feature. Subscribe to statistics using either Juniper's proprietary remote procedure call (gRPC) service or gRPC Network Management Interface (gNMI).

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

- **Support for policers and ACLs in firewall filters (ACX7024 and PTX10003)**—The ACX7024 and PTX10003 support subscribable YANG data models for operational states. The genstate YANG models expose a subset of `show` command data through the gNMI subscribe RPC. A gNMI telemetry collector can subscribe to the resource paths defined in the published models to query for specific state data. This feature provides genstate YANG data model support for policers and ACLs in firewall filters.

[See [Junos Genstate YANG Data Models](#) and [gNMI Genstate Subscription](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **OpenConfig and state sensor support for Layer 2 (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—The listed ACX Series routers support OpenConfig configuration and state sensors for VLAN, multicast, and Spanning Tree Protocol (STP).

[See [Mapping OpenConfig VLAN Commands to Junos Configuration](#), [Mapping OpenConfig STP Commands to Junos Configuration](#), and [Mapping OpenConfig Multicast Commands to Junos Configuration](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **Support for Health Monitoring telemetry data for standby nodes and FPCs (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—We've expanded telemetry data to support Health Monitoring telemetry beyond the primary node to include standby nodes and Flexible PIC Concentrators (FPCs). You can stream statistics that include load average, process parameters, and component CPU utilization using either Juniper's proprietary remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) transport from the device to the collector.

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

MPLS

- **OAM support for labeled IS-IS and labeled OSPF flexible algorithm SR paths (ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024)**—Junos OS Evolved supports the following Operation, Administration, and Maintenance (OAM) capabilities for labeled IS-IS Flexible Algorithm (flex algo) segment routing paths:
 - IPv4 and IPv6 MPLS ping
 - IPv4 and IPv6 MPLS traceroute
 - Equal-cost multipath (ECMP) traceroute

Junos OS Evolved also supports IPv4 MPLS ping and IPv4 MPLS traceroute for labeled OSPF flex algo segment routing paths. The OAM functionality is used to detect data plane failures in segment routing paths for the purposes of fault detection and isolation.

To enable these OAM capabilities, we've introduced the `algorithm` option in the following commands:

- `ping mpls segment routing isis fec algorithm algorithm-id`
- `ping mpls segment routing ospf fec algorithm algorithm-id`
- `traceroute mpls segment routing isis fec algorithm algorithm-id`
- `traceroute mpls segment routing ospf fec algorithm algorithm-id`

[See [ping mpls segment routing isis](#), [ping mpls segment routing ospf](#), [traceroute mpls segment-routing ospf](#), and [traceroute mpls segment-routing isis](#).]

- **Support for MPLS Layer 2 Connections on subscriber stacking models using dynamic PWHT Interfaces (ACX7100-48L, ACX7332, ACX7348)**—Subscriber management and stacking models using PPPoE, DHCPv4, and DHCPv6 over dynamic PWHT interfaces are fully supported for L2CKT, L2VPN, and EVPN types. Use the existing configuration commands under the appropriate interface to configure these options.

[See [MPLS Pseudowire Configuration](#), [Pseudowire Headend Termination \(PWHT\) Configuration](#), [Pseudowire Termination at an EVPN](#), and [Layer 2 Circuit Overview](#).]

- **Support for L2VPN and L2CKT over SR-TE (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X)**—We now support L2VPN and L2CKT configurations and features over segment routing-traffic engineered (SR-TE) networks. Supported features include metro services, entropy labels, FAT labels, and pseudowire redundancy. You can see the L2VPN and L2CKT status for SR-TE configurations using the `show l2vpn connections` and `show l2circuit connections` commands.

[See [MPLS Traffic Engineering Configuration](#), [show l2vpn connections](#), and [show l2circuit connections](#).]

- **SRv6-TE tunnels with micro-SIDs in PCEP (ACX7100-32C, ACX7100-48L, ACX7024, ACX7332, ACX7348, ACX7509, and PTX10002-36QDD)**—This feature enhances traffic engineering and network optimization by enabling the reporting, delegation, and creation of these tunnels. You can report and delegate static SRv6-TE tunnels with micro-SID configurations to a PCE and initiate these tunnels through PCE, improving control and management. Key functionalities include reporting static SRv6-TE tunnels with micro-SIDs to the PCE, delegating their management, and creating them with proper SID structure and endpoint behavior checks. Existing CLI commands are extended to support these features, facilitating effective configuration and monitoring.

[See [SRv6-TE Tunnels with micro-SIDs in PCEP](#).]

- **PCEP multipath support for SR-TE (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)**—You can configure the multipath feature (primary or secondary paths) for Path Computation Element Protocol (PCEP) segment routing-traffic engineering (SR-TE) as defined in draft-ietf-pcemultipath-06. We support the following multipath capabilities:
 - When the PCEP multipath feature is enabled, you can configure multiple primary or secondary paths in a candidate path that you configure and control using Path Computation Client (PCC). Note that the PCEP multipath feature is enabled by default.
 - When the PCEP multipath feature is disabled, you can configure only one primary path in a candidate path. Note that a secondary path configuration is not allowed.

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



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

[See [PCEP Configuration](#).]

- **SRv6 and MPLS service interworking (ACX7100-32C, ACX7100-48L, ACX7024, ACX7348, and ACX7509)**—SRv6 and MPLS service interworking feature enables seamless integration and interoperability between SRv6 and MPLS within a single network, enabling the translation of service instructions and ensuring proper traffic encapsulation and forwarding. This functionality is pivotal for networks transitioning to SRv6 while maintaining their existing MPLS infrastructure, providing a pathway for incremental deployment. This interworking feature supports both BGP SRv6 and BGP MPLS-based Layer 3 services, ensuring redundancy and scalability.

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

[See [Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP](#).]

Multicast

- **Enhanced MVPN provider tunnel selection criteria (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We support the following enhanced MVPN provider tunnel selection criteria to finetune multicast path-selection across the core network.
 - Regular Expression for selecting RSVP tunnels for ingress replication.
 - Colored inet.3 table for ingress replication.
 - Root Address for MLDP P2MP tunnels.

[See [Provider Tunnel Selection In Ingress Replication](#).]

- **Enhanced L3 multicast operational commands (ACX7100-32C, PTX10004, and QFX5130-32CD)**—The show instance command is now extended to all routing instances for the following commands. Earlier, only specific PIM-enabled routing instances were displayed.

- show pim join instance all
- show pim rps instance all
- show pim statistics instance all
- show multicast route instance all
- show multicast statistics instance all

Additionally, the show pim statistics output will display V2 Sparse Join and V2 Sparse Prune counters.

The show igmp statistics output will also display the V1/V2/V3 Membership Query field.

[See [show pim statistics](#), [show multicast statistics](#), and [show igmp statistics](#).]

Network Management and Monitoring

- **Enhanced RSI timing support (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—Use the new CLI option timing-debug that we've introduced to the existing request support information with-components command to access component-specific commands and timing-specific information.

This option archives log files specific to components such as PTP, Synchronous Ethernet, and Global Navigation Satellite System (GNSS). This feature streamlines the data collection process to minimize network downtime and efficiently diagnoses timing-related issues by focusing on essential CLI and Packet Forwarding Engine outputs without requiring a full request support information (RSI) dump.

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

- **Dying Gasp support through SNMP (ACX7024 and ACX7024X)**—We've introduced dying gasp feature support through SNMP in case of power failure. When a device detects a power failure, it sends SNMP traps to the network management system (NMS) during power failures. The device can send dying gasp packets to up to five servers for each notification type.

This feature supports SNMPv1 and SNMPv2 but not SNMPv3, and is only available on AC power supplies. The hold time for dying gasp is 4 ms for a single AC power supply and doubles with two active AC power supplies.

[See [Dying Gasp Functionality](#) and [SNMP Traps and Informs](#).]

Precision Time Protocol (PTP)

- **Precision Time Protocol with MACsec encryption (ACX7332, ACX7348, and ACX7100-32C)**—You can use Precision Time Protocol (PTP) with Media Access Control Security (MACsec) encryption enabled on the same port simultaneously. This configuration ensures secure and precise time synchronization, enhancing network performance and security.

[See [Guidelines to Configure PTP over Ethernet.](#)]

- **Telecom Grandmaster functionality (ACX7332 and ACX7348)**—Use the inbuilt Global Navigation Satellite System (GNSS) receivers on these routers to provide the Telecom Grandmaster (T-GM) functionality. This feature ensures precise time synchronization, which is crucial for telecom networks to maintain accurate timing and coordination across various network elements.

[See [Integrated Global Navigation Satellite System \(GNSS\) on Routing Platforms.](#)]

Routing Policy and layer2-policer Firewall Filters

- **Support for profile categories (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, and ACX7024)**—Profile categories are a way to distinguish firewall filters based on the direction and interface type. The profile categories are namely, ingress-inet6-user-acl, ingress-inet6-lo0-acl, and egress-inet6-user-acl.

[See [Overview of Firewall Filter Profiles on ACX Series Routers \(Junos OS Evolved\).](#)]

- **Support for firewall filters (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509)**— Support for firewall filters on PFE for services, such as bridge, IPv4, IPv6, CCC, MPLS and so on, based on packet match conditions and actions on ACX series devices. You can enable firewall filters on ACX series routers to monitor and control the traffic transiting the router or destined for the routing engine. The types of filters supported are as follows:
 - **Interface-specific Filter**—Unique firewall filter instance per logical interface for both ingress and egress traffic stream and for all protocol families.
 - **Physical-interface Filter**—Unique firewall filter instance per physical interface applicable only for ethernet-switching, circuit cross-connect (CCC), inet, and inet6 family for ingress traffic.
 - **Global Filter**—Unique firewall filter instance applied globally at PFE level applicable only for ethernet-switching, circuit cross-connect (CCC), inet, and inet6 family for ingress traffic.

[See [Firewall Filters Overview.](#)]

Software Installation and Upgrade

- **ZTP on WAN Interfaces (ACX7100-32C, ACX7100-48L, and ACX7509)** — Zero-touch provisioning (ZTP) dynamically detects the port speed of WAN interfaces and uses this information to create ZTP client ports with the same speed. ZTP automatically cycles through the WAN ports until it receives

DHCP options from the DHCP server. The device uses the DHCP options to perform the bootstrap process.

[See [Zero Touch Provisioning](#).]

- **Static configuration of MAC-IP bindings (ACX7100-32C, ACX7100-48L, PTX10001-36MR, and PTX10008)**—You can configure MAC-IP bindings on interfaces to improve network management and host communication. This setup is similar to configuring static MAC addresses on an interface. Use this feature to streamline operations in static environments, such as Internet Exchange Points (IXPs), where Customer Edge (CE) routers remain fixed.

[See [Static Configuration of MAC-IP Bindings](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Application-specific link attribute support in OSPFv2 for segment routing–traffic engineering (SR-TE) (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, and ACX7348)**—Advertise traffic engineering (TE) attributes such as TE-metric, delay-metric, or admin-groups for RSVP and flexible algorithms on the same link using the application-specific link attribute as defined in RFC 8920, *OSPF Application-Specific Link Attributes*. Optimize network paths based on specific application requirements to enhance traffic management and efficiency.

Configure the flexible algorithm application-specific traffic engineering attribute by including the application-specific statement at the [edit protocols ospf area interface] hierarchy level and the strict-asla-based-flex-algorithm statement at the [edit protocols ospf source-packet-routing] hierarchy level.

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

- **Multi-instance OSPF with SR (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Configure and run multiple independent interior gateway protocol (IGP) instances of OSPFv2 with segment routing (SR) on a router. Use this feature to create two or more OSPF instances and apply SR-MPLS on each instance. Multiple instances of OSPF can advertise different prefix-segment IDs (SIDs), and other instances can use these SIDs for routing decisions.

Multi-instance OSPF combined with SR enhances network flexibility, scalability, and control over traffic engineering, especially in large and complex networks.



NOTE: Junos OS does not support the configuration of the same logical interface in multiple IGP instances of OSPFv2.

[See [Multiple Independent IGP Instances of OSPFv2 Overview](#) and [Example: Configure Multiple Independent Instances of OSPFv2 with Segment Routing](#).]

- **NSR Support for SRv6 IS-IS and SRv6 BGP (ACX7332, ACX7348, and ACX7509)**—Use IS-IS Nonstop Active Routing (NSR) for dynamic micro adjacency segment identifiers (SIDs) and dynamic classic adjacency End-x SIDs. Junos OS allocates the same dynamic SID on both the active and backup RE after a switch-over, ensuring that dynamically-allocated SIDs on the primary RE are not repurposed. You can also use BGP NSR for dynamic DT SIDs. Note that Junos OS currently does not support NSR for classic dynamic End SIDs.

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

Subscriber Management and Services

- **Subscriber stacking over pseudowire interfaces (ACX7100-48L, ACX7332, and ACX7348)**—We support subscriber access models using dual-stack configurations over DHCP Server/Relay and Point-to-Point Protocol over Ethernet (PPPoE) on Layer 2 pseudowire interfaces. This setup supports features such as firewall filters, hierarchical class of service (CoS), lawful intercept, and multicast .

Configure dual-stack options using the existing commands.

[See [Dual-Stack Access Models in a DHCP Network](#), [Dual-Stack Access Models in a PPPoE Network](#), and [MPLS Pseudowire Configuration](#).]

-

Additional Features

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

- **BGP autodiscovery underlay in EVPN-VXLAN (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)**

[See [BGP Auto-Discovered Neighbors](#).]

- **Enhanced OISM for IPv4 multicast traffic in EVPN-VXLAN fabrics (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509).** Support includes:

- Enhanced optimized intersubnet multicast (OISM) mode—the asymmetric bridge domains model, also called the bridge domains not everywhere (BDNE) model.

With enhanced OISM on these devices, you must configure the vxlan-extended host profile at the [edit system packet-forwarding-options system-profile] hierarchy level.



NOTE: The Packet Forwarding Engine reboots when you change the system profile.

- MAC-VRF EVPN instances with vlan-based or vlan-aware service types only.

On these devices, in the EVPN instance (EVI) you must configure the `conserve-mcast-routes-in-pfe` option at the `[edit routing-instances name multicast-snooping-options oism]` hierarchy level.

- IPv4 multicast traffic with IGMPv2, IGMPv3, and IGMP snooping.
- Server leaf or border leaf OISM device roles.



NOTE: ACX Series OISM leaf devices can only have multihoming peers that are also ACX Series devices.

- External multicast source and receiver communication using classic Layer 3 (L3) interfaces only.

[See [Overview of Enhanced OISM](#) and [How Enhanced OISM Works](#).]

- **Fast reroute and egress link protection support for EVPN E-Tree** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)

[See [Fast Reroute for Egress Link Protection with EVPN-VXLAN Multihoming](#) and [Convergence in a Multihomed EVPN-MPLS Network](#).]

- **FlowTapLite support** (ACX7100, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)—In Junos OS Evolved FlowTapLite support, there are several differences from the Junos OS support.

[See [Understanding Flow-Tap Architecture](#).]

- **Optimized intersubnet multicast (OISM) for IPv4 multicast traffic in EVPN-VXLAN fabrics** (ACX7332, ACX7348, and ACX7509). Support on these devices includes:

- Regular OISM mode—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



NOTE: On these devices, in the EVPN instance (EVI) you must configure the `conserve-mcast-routes-in-pfe` option at the `[edit routing-instances name multicast-snooping-options oism]` hierarchy level.

- IPv4 multicast traffic with IGMPv2, IGMPv3, and IGMP snooping
- Server leaf or border leaf OISM device roles



NOTE: ACX Series OISM leaf devices can only have multihoming peers that are also ACX Series devices.

- External multicast source and receiver communication using classic Layer 3 (L3) interfaces only

[See [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **SMET Support in EVPN-MPLS** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)

[See [Overview of Selective Multicast Forwarding](#), [Configuring the number of SMET Nexthops](#) and [multicast-replication](#).]

- **QSFP-100G coherent ZR optics performance monitoring** (ACX7024, ACX7348, and PTX10001-36MR; and the PTX10004, PTX10008, and PTX10016 with the PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards installed). Monitor the performance of QSFP-100G coherent ZR optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Accumulate performance metrics into 15-minute and 1-day interval bins. Use the `show interfaces transport pm` command to view current and historical performance data.

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

- **QSFP-100G-LR coherent ZR optics support** (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, and ACX7348). Manage optical transport links efficiently with QSFP-100G-LR coherent ZR optics.

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

- **QSFP-100G-ER4L coherent ZR optics support** (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, and ACX7348). Manage optical transport links efficiently with QSFP-100G-LR coherent ZR optics.

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

- **Static route tracking using the results of RPM and TWAMP tests** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016). We've extended support for static route tracking to Junos OS Evolved and included Two-Way Active Measurement Protocol (TWAMP) test support as well. You use RPM or TWAMP probes to detect link status and to change the preferred-route state on the basis of the probe results. Tracked static routes can be IPv4 or IPv6, and each IPv4 and IPv6 tracked static route supports up to 16 next hops. You can also configure the metric, route preference, and tag values for each IPv4 or IPv6 destination prefix. However, you configure this feature differently on Junos OS Evolved devices; you configure the `sla-tracking` statement at the `[edit routing-options]` hierarchy level. For Junos OS, you would configure the `rpm-tracking` statement at the same hierarchy level.

[See [Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#), [Understand Two-Way Active Measurement Protocol](#), [sla-tracking](#), and [show route sla-tracking](#).]

- **Support for dynamic list next hop** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)

[See [Configuring Dynamic List Next Hop](#).]

- **Support for EVPN-MPLS egress link protection** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, and ACX7024X)

[See [Convergence in a Multihomed EVPN-MPLS Network](#).]

- **Support for hard interface shutdown when a device detects EVPN core isolation conditions** (ACX7024, ACX7100-32C, and ACX7100-48L)

[See [Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions](#), [network-isolation](#), and [network-isolation-profile](#).]

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

- **Support for SRv6 LSPs in PCEP** (ACX7348 and PTX10002-36QDD). The Path Computation Element Protocol (PCEP) supports all types of SRv6 LSPs, such as PCE-initiated, locally created, and delegated SRv6 LSPs.

[See [SRv6 LSP in PCEP](#).]

- **Support for VPLS multihoming and the mapping of VPLS traffic to specific LSPs** (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)

[See [Configuring VPLS Multihoming](#) and [Mapping VPLS Traffic to Specific LSPs](#).]

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

- **Support for policer and count actions** (ACX7332)

[See [Firewall Filter Nonterminating Actions](#).]

- **QSFP-100G coherent ZR optics performance monitoring** (ACX7100-32C and ACX7100-48L). Monitor the performance of QSFP-100G coherent ZR optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Accumulate performance metrics into 15-minute and 1-day interval bins. Use the `show interfaces transport pm` command to view current and historical performance data.

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

- **Support for L2 VPN and L2 Circuit** (ACX7100-32C, ACX7100-48L, ACX7024, ACX7024X, ACX7332, ACX7348, and ACX7509,)

[See [Understanding Layer 2 VPNs](#) and [Layer 2 Circuit Overview](#).]

- **Exclude hops in the RSVP LSP path (ACX7332, ACX7509, PTX10002-36QDD, PTX10008)**—You can configure a list of hops to be excluded in the label-switched path (LSP) so that RSVP LSPs avoid those hops and links in the traffic engineering (TE) domain. When an RSVP LSP is signaled in the network, the path message carries the excluded list of hops. When the downstream routers perform loose hop expansion, such as inter-domain LSP or abstract node expansion, the transit routers use the same excluded list of hops that the ingress router uses for path computation. This mechanism enables intermediate routers to avoid the routers included in the excluded hop list. The routers try alternative paths to help with the convergence of LSPs when a complete end-to-end path computation is not possible.

Additionally, ingress routers receive PathErr messages and when computing another path, the routers use a PathErr message sender's address to avoid the link or node that generates an error. Transit routers also need this error avoidance information during retry attempts. RFC4814 defines the exclude hop information and is accepted in RSVP signaling.

To configure LSPs to exclude a list of hops, include the exclude statement at the [edit protocols mpls path path-name next-hop] hierarchy level. The ingress routers exclude the hops in CSPF computation and are also included in RSVP LSP signaling.

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 15](#)
- [EVPN | 15](#)
- [General Routing | 15](#)
- [Junos Telemetry Interface | 16](#)
- [Junos XML API and Scripting | 16](#)
- [Network Management and Monitoring | 16](#)
- [PTP \(Precision Time Protocol\) | 17](#)
- [Routing Policies and Firewall Filters | 17](#)
- [Routing Protocols | 17](#)
- [User Interface and Configuration | 17](#)

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

Authentication and Access Control

- **Disabled CDN auto download (Junos OS Evolved)**— The PKI process periodically, by default every 24 hours, polls the CDN server for the latest default trusted CA bundle and updates the list for any changes to the trusted CAs in the bundle. If there are any changes, PKI process loads them in the background. The auto download of CA certificates might generate core files. We've disabled the service of PKI query to CDN server periodically to download the latest trusted CA bundle.
- On Junos OS Evolved, password authentication for SCP based configuration archival is supported.

EVPN

- **EVPN system log messages for CCC interface up and down events**—Devices will now log EVPN and EVPN-VPWS interface up and down event messages for interfaces configured with circuit cross-connect (CCC) encapsulation types. You can look for error messages with message types `EVPN_INTF_CCC_DOWN` and `EVPN_INTF_CCC_UP` in the device system log file (`/var/log/syslog`).

General Routing

- **Change to the commit process**—In prior Junos OS Evolved releases, if you use the `commit prepare` command and modify the configuration before activating the configuration using the `commit activate` command, the prepared commit cache becomes invalid due to the interim configuration change. As a result, you cannot perform a regular commit operation using the `commit` command. The CLI shows an error message: 'error: Commit activation is pending, either activate or clear commit prepare'. If you now try running the `commit activate` command, the CLI shows an error message: 'error: Prepared commit cache invalid, failed to activate'. You then must clear the prepared configuration using the `clear system commit prepared` command before performing a regular commit operation. From this Junos and Junos OS Evolved release, when you modify a device configuration after 'commit prepare' and then issue a 'commit', the OS detects that the prepared cache is invalid and automatically clears the prepared cache before proceeding with regular 'commit' operation.

[See [Commit Preparation and Activation Overview](#).]

Junos Telemetry Interface

- The `show agent sensors` command output for gRPC sensors is truncated on the Junos OS Evolved platform to align with the output format of the Junos OS platform.

Junos XML API and Scripting

- **Commit script input to identify software upgrades during boot time (ACX Series, PTX Series, and QFX Series)**—The `junos-context` node-set includes the `sw-upgrade-in-progress` tag. Commit scripts can test the `sw-upgrade-in-progress` tag value to determine if the commit is taking place during boot time and a software upgrade is in progress. The tag value is `yes` if the commit takes place during the first reboot after a software upgrade, software downgrade, or rollback. The tag value is `no` if the device is booting normally.

[See [Global Parameters and Variables in Junos OS Automation Scripts](#).]

Network Management and Monitoring

- In a firewall filter configured with a `port-mirror-instance` or `port-mirror` action, if `l2-mirror` action is also configured, then `port-mirroring` instance family should be any. In the absence of the `l2-mirror` action, `port-mirroring` instance family should be the firewall filter family.
- **Python 2 interpreter option deprecated for Juniper Extension Toolkit (JET) applications (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX10K-LC1202-36MR (line cards for PTX10016, PTX10008 and PTX10004), QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-64OD, QFX5240-QD, QFX5700, and QFX5700E)**—Python 2.7 is already not supported on Junos OS Evolved devices as of an earlier release. The `python` statement at the `edit system extensions extension-service application file <filename>` hierarchy level was used to interpret JET applications written in Python 2. This statement is now deprecated. To run daemonized on-device JET applications written in Python 3, use the `python3` statement.

[See [file \(JET\)](#).]

PTP (Precision Time Protocol)

- **Maximum limit of PTP local masters (PTX10008)**— You can configure up to 512 PTP masters at the `edit protocols ptp master interface interface-name multicast-mode hierarchy level` on PTX10008 series routers. Earlier the system was rejecting the commit while trying to configure more than 128 PTP masters.

Routing Policies and Firewall Filters

- Support added for source and destination port optimization for port ranges for IPv6 input firewall filters.

Routing Protocols

- **Update to IGMP snooping membership command options**— The `instance` option is now visible when issuing the `show igmp snooping membership ?` command. Earlier, the `instance` option was available but not visible when `?` was issued to view all possible completions for the `show igmp snooping membership` command.

[See [show igmp snooping membership](#).]

- **MLD snooping proxy and l2-querier source-address (ACX7024, ACX7100-32C, PTX10001-36MR, QFX5120-32C, and QFX5130-32CD)**— The `source-address` configured for proxy and l2-querier under the `mld-snooping` hierarchy should be an IPv6 link-local address in the range of `fe80::/64`. The CLI help text has been updated to "Source IPv6 link local address to use for proxy/L2 querier". In earlier releases, the CLI help text read, "Source IP address to use for proxy/L2 querier."

[See [source-address](#).]

User Interface and Configuration

- **Compact format deprecated for JSON-formatted state data (ACX Series, PTX Series, and QFX Series)**— We've removed the `compact` option at the `[edit system export-format state-data json]` hierarchy level because Junos devices no longer support emitting JSON-formatted state data in compact format.
- **Access privileges for request support information command (ACX Series, PTX Series, and QFX Series)**— The `request support information` command is designed to generate system information for

troubleshooting and debugging purposes. Users with the specific access privileges `maintenance`, `view`, and `view-configuration` can execute `request support information` command.

Known Limitations

IN THIS SECTION

- [General Routing | 18](#)

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

- G.8275.1- G.8273.2 1PPS cTE performance test may be marginally outside class-C for PTP BC on ACX7100-48L, especially for mixed speed port testing with combinations of 10G / 25G channelized ports and 100G ports. On each reboot, the 1PPS cTE measurement may be within the class-C measurement threshold, or may randomly be out of threshold by a few nanoseconds. [PR1607381](#)
- PTP to PTP Noise transfer will be failing for frequencies 1. 0.03125 HZ 2. 0.123125 HZ [PR1608786](#)
- During FPC offline CLI command execution, in case PIC is in onlining transition state, then user needs to wait for 5 minutes to get the FPC state. Alternatively, User should check PIC status along with FPC status for FPC online/offline CLI execution. For online CLI action, make sure FPC and PIC status are in offline state before triggering the online CLI. For offline CLI action, make sure FPC and PIC status are in online state before triggering the offline CLI. [PR1738954](#)
- ACX7024 ports support 10G/1G/25G multi-rate. When peering with other platform or other vendor devices, For example using SFP-LX10 for 1G connection, the link may remain physically down The reason is Auto-negotiation is not supported in ACX7024 PFE due to Broadcom limitation. In order to make it work, user has to explicitly configure speed/duplex on both sides, and disable auto-negotiation on the peer side [PR1759804](#)

- On ACX, among the cpu queues, few protocols share a single cpu queue. During queue sharing, any DDOS violation caused by a particular protocol traffic, triggers violations for all the protocols mapped to that CPU queue. This is a design limitation. [PR1810673](#)
- MTU config is only supported on IFL family level in ACX platform to generate ICMP packet too big error when MTU exception is HIT . Customer has to adapt the config on ACX box . For Ex: instead of configuring MTU on IFD like below set interfaces et-0/0/0 mtu 1400 we need to configure mtu like below on IFL family level delete interfaces et-0/0/0 mtu 1400 set interfaces et-0/0/0 unit 0 family inet mtu 1400 [PR1828681](#)

Open Issues

IN THIS SECTION

- [Dynamic Host Configuration Protocol | 19](#)
- [General Routing | 20](#)
- [Multicast | 20](#)
- [Precision Time Protocol \(PTP\) | 20](#)
- [Storm Control | 21](#)

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

Dynamic Host Configuration Protocol

- When DHCP trace options are enabled, there is a possibility that jdhcpd could generate a core file. In general, traceoptions should be enabled only for debugging. They should be disabled once debugging is done. [PR1771121](#)

General Routing

- When ingress policer is configured, to drop ingress traffic, on an interface with upMep the CFM packets generated from the upMep is also be dropped due to the policer. This leads to CFM session going down. [PR1754938](#)
- This is day-1 issue of Juniper server. It exists on all ACX Series platforms where Juniper Networks server runs. Once the asymmetry configured at the backup port, The error propagated to the down stream eventually causes this performance issue of spike. [PR1793926](#)
- SRv6 dynamic sid missing in Packet Forwarding Engine (PFE) and losing 100% VPN traffic after locator address is modified. The issue can be rectified by: Delete the locator/block: commit; re-configure the new locator/block prefix or restart rpd on the primary and backup Routing Engine. In case of ACX Series, please use the restart evo-pfem and on both Routing Engines after the changes to restore consistency. [PR1828806](#)
- Ingress interface-specific IPv6 filters might not work for logical interfaces like aggregated Ethernet/IRB/FTI after events like Packet Forwarding Engine (PFE)/router restart. [PR1855552](#)

Multicast

- On all Junos Evolved platforms, when MVPN (Multicast Virtual Private Network) configured device receives PIM join or IGMP report from the remote peer installs route in next-hop and not in multicast composite next-hop, causing the rpd process to crash. [PR1777774](#)
- When multicast packets transit ACX7100 devices through VXLAN VTEP or core interface without multicast configurations, errors are seen. Suggested to use DDOS configurations provided with the workaround. [PR1796501](#)

Precision Time Protocol (PTP)

- syncE to PTP and syncE to 1pps transient response marginally fails. This happens when the servo gets the initial 100ns jump in one measurement window and the next 100ns in the next measurement window adjusting less initially. [PR1611848](#)

Storm Control

- Storm Control is not working when it is configured on aggregated Ethernet. [PR1611848](#)

Junos OS Evolved Release Notes for PTX Series

IN THIS SECTION

- [What's New | 21](#)
- [What's Changed | 102](#)
- [Known Limitations | 107](#)
- [Open Issues | 107](#)

These release notes accompany Junos OS Evolved Release 24.4R1 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

- [Hardware | 22](#)
- [Authentication and Access Control | 75](#)
- [Chassis | 75](#)
- [Class of Service | 76](#)
- [Ethernet Switching and Bridging | 77](#)
- [High Availability | 78](#)
- [Interfaces | 78](#)
- [Junos Telemetry Interface | 78](#)

- [Layer 2 VPN | 82](#)
- [MACsec | 82](#)
- [MPLS | 83](#)
- [Multicast | 83](#)
- [Multichassis Link Aggregation \(MC-LAG\) | 84](#)
- [Network Management and Monitoring | 85](#)
- [Routing Policy and layer2-policer Firewall Filters | 86](#)
- [Routing Protocols | 87](#)
- [Services Applications | 87](#)
- [Software Installation and Upgrade | 89](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 89](#)
- [Additional Features | 90](#)

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 24.4R1, click the Group by Release link. You can collapse and expand the list as needed.

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

The following sections highlight the key features in this release.

Hardware

- **PTX10002-36QDD router (PTX Series)**—The PTX10002-36QDD is a fixed-configuration router that features 36 high-density and cost-efficient 800-Gigabit Ethernet (800GbE) ports network ports in a 2-U form factor. With 28.8 terabits per second (Tbps) of throughput, the PTX10002-36QDD is optimally designed for peering, core routing, and infrastructure edge routing roles in cloud provider, service provider, and content provider networks.

The router supports 2200-W or 3000-W high-voltage HVAC/HVDC and DC power supply units (PSUs) and front-to-back airflow.

You can channelize the ports on the PTX10002-36QDD and increase the number of interfaces.

To install the PTX10002-36QDD router and perform initial configuration, routine maintenance, and troubleshooting, see the [PTX10002-36QDD Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform.

Table 1: PTX10002-36QDD Feature Support

Feature	Description
Chassis	<ul style="list-style-type: none"> Support for the following chassis management functionalities: <ul style="list-style-type: none"> The presence of two ASIC packages enables you to take a Flexible PIC Concentrator (FPC) offline or bring it online to restart the FPC without impacting the power to the FPC. When you connect 3000-watt (W) power supply units (PSUs), the system operates in normal power mode. You can change the operating power mode from normal to power-optimized by using the <code>set chassis mode power-optimized</code> command. The <code>show chassis fpc</code> command displays both PFE and PFE-Instance details. On the router, when you run the <code>request chassis fpc</code> command, you must use <code>pfe</code> instead of <code>pfe-instance</code> to control the FPC operations. Also, when you run the <code>request chassis fpc</code> command, you must commit the command for both the Packet Forwarding Engines that are present. <p>[See Power Mode Management on PTX10002-36QDD, chassis, request chassis fpc, and show chassis fpc.]</p> <ul style="list-style-type: none"> Support for resiliency features to manage fabric faults, including but not limited to: <ul style="list-style-type: none"> Auto-heal functionality to recover the faulty link by fixing the errors automatically. All Packet Forwarding Engines disabled when the number of fabric link errors exceeds four in the system. <p>You can use the existing CLI commands for the fabric management. The following commands display new or different fields in their outputs:</p> <ul style="list-style-type: none"> <code>show chassis fabric fpcs</code> displays peer FPC and PFE details because the Packet Forwarding Engines are directly connected to each other.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • show chassis fabric topology displays only physical link connectivity. <p>[See Chassis-Level User Guide, show chassis fabric fpcs, and show chassis fabric topology.]</p> <ul style="list-style-type: none"> • Optics EM policy support. We've extended the Junos Environment Monitoring (EM) policy to include optics temperature sensors for PTX10002-36QDD routers. It includes the following features: <ul style="list-style-type: none"> • The Optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed. • Junos OS Evolved automatically triggers optics shutdown for 100GbE, 400GbE, and 800GbE optics when the Fire Shutdown threshold is breached. Auto-recovery is not supported for optics disabled by the EM policy. To re-enable the optics, use the request interface optics-reset command or perform optics online insertion and removal (OIR). • EM policy is enabled by default on all 100GbE, 400GbE, and 800GbE optics that are Multi-source Agreements (MSA)-compliant and support diag EEPROM with temperature monitoring. This policy is not applicable for loopback optics and direct attach copper (DAC) cables. <p>To disable EM policy or view temperature threshold values, use the following CLI commands:</p> <ul style="list-style-type: none"> • set chassis fpc <i>fpc_slot</i> pic <i>pic_slot</i> port <i>port_no</i> no-temperature-monitoring explicitly disables the EM policy on specific WAN ports. • show chassis temperature-thresholds displays the optics temperature threshold values. • show chassis environment displays the optics temperature. <p>[See chassis-adc-temperature-sensor.]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> Routing Engine support. The fixed-configuration PTX10002-36QDD router supports an inbuilt Routing Engine represented by the model number RE-JNP10002-36QDD in the CLI. <p>The router does not support:</p> <ul style="list-style-type: none"> A pluggable Routing Engine GRES, as the router does not have a redundant Routing Engine The following operational commands: <ul style="list-style-type: none"> request chassis routing-engine master acquire request chassis routing-engine master release <p>[See show chassis hardware.]</p> <ul style="list-style-type: none"> Routing Engine resiliency. We've enabled Routing Engine resiliency for the faults related to CPU memory and DIMM. The Routing Engine supports fault-handling actions such as logging errors, raising alarms, sending SNMP traps, and providing indication about an error through the LEDs. <p>[See show system errors active.]</p> <ul style="list-style-type: none"> Support for fabric platform resiliency includes resiliency functionality to manage hardware components such as the FPCs, PSUs, and fans. <p>[See show chassis power detail, show chassis fpc, and show chassis fan.]</p> <ul style="list-style-type: none"> Packet Forwarding Engine resiliency. The software detects, reports, and takes action on Packet Forwarding Engine faults. Actions are taken based on the default configuration or user configuration available for the errors. <p>[See show system errors active.]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Class of service	<ul style="list-style-type: none"> Support for class-of-service (CoS) features, including classifiers (behavior aggregate (BA), fixed, and multifield (MF)), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, and drop profiles. <p>[See CoS Features and Limitations on PTX Series Routers.]</p> <ul style="list-style-type: none"> Support for priority-based flow control (PFC) watchdog, which detects and mitigates PFC pause storms received for PFC-enabled queues. <p>We've added the <code>jnxCosWatchdogTxQueueTable</code> table to the SNMP class-of-service (CoS) MIB to show statistics for transmitting PFC queues related to the PFC watchdog. Table entries are indicated by <code>jnxCosWatchdogTxQueueEntry</code> and contain the following objects:</p> <ul style="list-style-type: none"> <code>jnxCosWatchdogIfIndex</code>—The index of an interface on which PFC and PFC watchdog are enabled. <code>jnxCosWatchdogTxQueueId</code>—The ID of the queue of the PFC-enabled interface. <code>jnxCosWatchdogTxQueueRecoveredCount</code>—The number of times a queue recovered after a PFC pause storm. <code>jnxCosWatchdogTotalPktDrop</code>—The total number of packets dropped due to PFC pause storm mitigation since the device was started. <code>jnxCosWatchdogLastPktDrop</code>—The number of packets dropped due to the last PFC pause storm. <p>[See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved and PFC Watchdog.]</p> <ul style="list-style-type: none"> Support for importing existing classifier and rewrite rules to form new rules. Support for priority-based flow control (PFC) at Layer 3 for untagged traffic and explicit congestion notification (ECN).

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p data-bbox="755 352 1398 422">[See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic and CoS Explicit Congestion Notification.]</p> <ul data-bbox="719 457 1421 663" style="list-style-type: none"> • Queue-depth monitoring support for virtual output queues. Virtual output queue (VOQ) queue-depth monitoring, or latency monitoring, measures peak queue occupancy of a VOQ. Junos OS Evolved supports VOQ queue-depth monitoring to report peak queue length for a given physical interface for each Packet Forwarding Engine. <p data-bbox="755 695 1143 722">[See VOQ Queue-depth Monitoring.]</p> <ul data-bbox="719 758 1404 1199" style="list-style-type: none"> • Support for export of physical interface queue statistics to an outside collector. Use UDP (native) streaming, remote procedure call (gRPC) services, or gRPC network management interface (gNMI) services by using the sensor/junos/system/linecard/interface/queue/. Each physical interface has eight queues. The following counters are exported as part of this sensor for all configured physical interfaces: <ul data-bbox="755 1037 1235 1199" style="list-style-type: none"> • Transmitted packets and transmitted bytes • Red drop packets and bytes • Tail drop packets and bytes <p data-bbox="755 1234 1365 1335">This feature includes zero suppression support. It does not include support for summed-up counters on aggregated Ethernet (ae) interfaces.</p> <p data-bbox="755 1367 1360 1430">[See sensor (Junos Telemetry Interface) and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p> <ul data-bbox="719 1465 1421 1709" style="list-style-type: none"> • Hierarchical CoS support. The router supports up to four levels of scheduling on an interface (physical interfaces, logical interface sets, logical interfaces, and queues). The router does not support hierarchical CoS on integrated routing and bridging (IRB) or aggregated Ethernet interfaces. Also, hierarchical CoS schedulers should not include buffer or drop profile configurations.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p>To enable hierarchical scheduling, set <code>hierarchical-scheduler</code> at the <code>[edit interfaces interface-name]</code> hierarchy level.</p> <p>[See Hierarchical Class of Service in ACX Series Routers.]</p> <ul style="list-style-type: none"> • Support for classification override configured under a forwarding policy. <p>[See CoS Features and Limitations on PTX Series Routers and Overriding the Input Classification.]</p>
Dynamic Host Configuration Protocol	<ul style="list-style-type: none"> • DHCPv4 relay agent and DHCPv6 relay agent are supported. The router supports the following DHCP features: <ul style="list-style-type: none"> • 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) <p>[See Extended DHCP Relay Agent Overview.]</p>
Hardware	<ul style="list-style-type: none"> • Supported transceivers, optical interfaces, and DAC cables. Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available. <p>[See Hardware Compatibility Tool .]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
High availability and resiliency	<ul style="list-style-type: none"> • BFD support, including: <ul style="list-style-type: none"> • 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 <p>[See Understanding BFD .]</p> • 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. <p>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.</p> <p>The router does not support the following features:</p> <ul style="list-style-type: none"> • 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 <p>[See Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation .]</p> <ul style="list-style-type: none"> • Support for VRRP. <p>The following features are not supported for VRRP on Junos OS Evolved:</p> <ul style="list-style-type: none"> • ISSU

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<ul style="list-style-type: none">• Proxy ARP• MC-LAG• Distribution support on aggregated Ethernet interfaces• IRB• Inline delegation <p>[See Understanding VRRP .]</p>

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
Interfaces	<ul style="list-style-type: none"> Interface support. The PTX10002-36QDD supports multiple port speeds and various channels under each port. The router supports a maximum port speed of 400 Gbps in low power mode. In standard power mode, it supports a port speed of 800 Gbps. If a port speed is configured (or a port speed is determined by default) without configuring number-of-sub-ports (at the [edit interfaces <i>interface-name</i>] hierarchy level), the port operates in nonchannelized mode. [See Port Speed on PTX Routers.] 400G-ZR and 400G-ZR+ support enhancements. We support 400G-ZR and 400G-ZR+ optics enhancements on the PTX10002-36QDD. The enhancements include application selection and configuration of target output power. You can view the advertised applications and switch between the applications. [See Features of 400ZR and 400G OpenZR+.] Support for performance monitoring and TCA. We support performance monitoring for the PTX10002-36QDD optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics using the show interfaces transport pm command and manage optical transport links efficiently. [See show interfaces transport pm.] Support for timing and synchronization. The PTX10002-36QDD supports Synchronous Ethernet compliant with the following ITU recommendations: <ul style="list-style-type: none"> G.8262/G.8262.1—Specifies timing characteristics of Synchronous Ethernet equipment clock (EEC). G.8264—Describes the Ethernet Synchronization Message Channel (ESMC). [See Synchronous Ethernet Overview.]

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<ul style="list-style-type: none"> Support for load balancing under the [edit forwarding-options enhanced-hash-key] hierarchy. <p>Load balancing includes:</p> <ul style="list-style-type: none"> 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 <p>[See enhanced-hash-key.]</p> <ul style="list-style-type: none"> Support for 128-way equal-cost multipath (ECMP) routing for MPLS transit cases. <p>The following features do not support 128-way ECMP:</p> <ul style="list-style-type: none"> Multicast P2MP MC-LAG Weighted unilist Consistent hashing Link protection (MPLS)

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<ul style="list-style-type: none"> • Adaptive load balancing • Class-based forwarding • 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 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 <i>number</i> unit <i>logical-unit-number</i> tunnel encapsulation ipip] hierarchy level to change the mode to flattened encap mode to achieve line-rate performance. [See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.] • Support for configuring UDP tunnel encapsulation on FTIs. You can configure encapsulation by using the tunnel encapsulation udp source <i>address</i> destination <i>address</i> statement at the [edit interfaces fti unit <i>unit</i>] hierarchy level. Keep in mind the following when configuring this feature: <ul style="list-style-type: none"> • Adding tunnel-termination makes the tunnel a de-encapsulation-only tunnel and encapsulation is disabled. • Specifying both the source and destination address is mandatory when you do not configure tunnel-termination. • Configuring a variable prefix mask on the source address is not allowed.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p data-bbox="755 352 1117 384">[See encapsulation (interfaces-fti).]</p> <ul style="list-style-type: none"> <li data-bbox="719 420 1409 636">• GRE tunnel encapsulation using loopback-based interface. You can configure GRE tunnel encapsulation on flexible tunnel interfaces (FTIs) using the loopback interface. Configure encapsulation by using the <code>tunnel encapsulation gre source <i>address</i> destination <i>address</i></code> statement at the [edit interfaces <i>fti0</i> unit <i>unit</i>] hierarchy level. <p data-bbox="755 667 1117 699">[See encapsulation (interfaces-fti).]</p> <ul style="list-style-type: none"> <li data-bbox="719 735 1417 951">• Support for GRE tunnel de-encapsulation using FTIs. Flexible tunnel interfaces (FTIs) support GRE tunnel de-encapsulation. When you enable the <code>tunnel-termination</code> statement at the [edit interfaces <i>fti0</i> unit <i>unit-number</i>] hierarchy level, tunnels are terminated on the WAN interface before any other actions—such as sampling, port mirroring, or filtering—are applied. <p data-bbox="755 982 1409 1045">[See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.]</p> <ul style="list-style-type: none"> <li data-bbox="719 1081 1393 1287">• Support for configuring MPLS protocols over FTI tunnels, thereby transporting MPLS packets over IP networks that do not support MPLS. Generic routing encapsulation (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. <p data-bbox="755 1323 1409 1507">To allow the MPLS traffic on the UDP tunnels, include the <code>mpls port-number</code> statement at the [edit forwarding-options tunnels udp port-profile <i>profile-name</i>] hierarchy level. To allow the MPLS traffic on the GRE tunnels, include the <code>mpls</code> statement at the [edit interfaces <i>fti0</i> unit <i>unit</i> family] hierarchy level.</p> <p data-bbox="755 1539 1193 1570">[See Flexible Tunnel Interfaces Overview.]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Junos telemetry interface	<ul style="list-style-type: none"> • JTI support for Packet Forwarding Engine sensors for usage, network processing unit (NPU) memory, NPU utilization, and pipeline NPU and ASIC. Using the Junos telemetry interface (JTI), you can export statistics using remote procedure call (gRPC) services, gRPC Network Management Interface (gNMI) services, and UDP transport. <p>Use these sensors:</p> <ul style="list-style-type: none"> • <code>/junos/system/linecard/packet/usage/</code> • <code>/junos/system/linecard/npu/memory/</code> • <code>/junos/system/linecard/npu/utilization/</code> • <code>/components/component/integrated-circuit/state/</code> • <code>/components/component/integrated-circuit/pipeline-counters/</code> <p>For pipeline sensors, the four packet and drop counter categories are interface, lookup, queuing, and host interface.</p> <p>[See Junos YANG Data Model Explorer.]</p> <ul style="list-style-type: none"> • JTI support for platform sensors. Using the Junos telemetry interface (JTI), 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. <p>Use these sensors:</p> <ul style="list-style-type: none"> • <code>/junos/system/cmerror/</code> • <code>/junos/system/linecard/</code> • <code>/components/components/</code> • <code>/system/alarms/</code> • <code>/state/interfaces/</code>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none">• /state/chassis/ <p>[See Junos YANG Data Model Explorer.]</p>

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> Support for flow-aware transport (FAT) for pseudowires labels on ingress routers, with parsing that includes all the payload fields in the hash calculation. These flow labels are supported: <ul style="list-style-type: none"> L2circuit, LDP-signaled pseudowires L2VPN, BGP-signaled pseudowires L2VPN with FEC129 (BGP autodiscovery) <p>[See flow-label-receive and flow-label-transmit.]</p> Support for VLAN tag manipulation: pop, push, and swap. <p>[See Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough.]</p> Support for virtual circuit connection verification (VCCV) protocol, which transfers control packets from one provider edge (PE) router to another PE router by creating a separate channel in the pseudowires. The pseudowires set up signaling peers and use a control word to maintain proper sequencing of pseudowire packets over the packet-switched network. <p>[See BFD Support for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS, Configuring BFD for VCCV for Layer 2 Circuits, MPLS Pseudowires Configurations, show ldp database, and show route instance.]</p> Support for inner VLAN transparency. We support the pop, push, swap, pop-pop, pop-swap, swap-push, push-push, and swap-swap operations on port-based and VLAN-based Metro Ethernet Forum (MEF) Layer 2 services. VLAN transparency refers to preserving inner VLANs in the packet that are not subject to manipulation and are not used for forwarding. Based on the scenarios, VLAN transparency works on up to four VLAN tags. <p>[See Understanding VLAN Manipulation (Normalization and VLAN Mapping) on Ethernet Services.]</p> Support for the following protocols:

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ol style="list-style-type: none"><li data-bbox="756 352 1065 384">1. LAG (aggregated Ethernet)<li data-bbox="756 415 846 447">2. LACP<li data-bbox="756 478 846 510">3. LLDP

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • Support for the following Layer 3 forwarding features: <ul style="list-style-type: none"> • 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

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none">• IP fragmentation• DDoS

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
MACsec	<ul style="list-style-type: none"> • MACsec support in static CAK mode on physical interfaces with dynamic power management. Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for traffic on Ethernet links. This device supports MACsec in static connectivity association key (CAK) mode. This device supports MACsec on physical interfaces to enable you to secure your network using any of the following encryption types: <ul style="list-style-type: none"> • GCM-AES-128 • GCM-AES-256 • GCM-AES-XPB-128 • GCM-AES-XPB-256 <p>This device supports the following MACsec features:</p> <ul style="list-style-type: none"> • Configurable security association key (SAK) rekey period • MACsec Key Agreement (MKA) protocol fail-open mode • Preshared key (PSK) chains and hitless rollover • PSK password encryption using single password • Fallback PSK • Extended packet numbering (XPB) • Jumbo frames <p>[See Understanding Media Access Control Security (MACsec).]</p> <ul style="list-style-type: none"> • MACsec dynamic power management support. Use MACsec to secure your network with the knowledge that your device is working to optimize power usage. To save power, the device dynamically powers MACsec blocks on and off based on the MACsec configuration. You might experience minimal traffic loss during the power block transition. <p>[See Understanding Media Access Control Security (MACsec).]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
MPLS	<ul style="list-style-type: none"> 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. <p>[See Fast Reroute Overview.]</p> <ul style="list-style-type: none"> Support for MPLS features, including: <ul style="list-style-type: none"> 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 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 and TE++ LDP, including: <ul style="list-style-type: none"> Configurable label withdraw delay Egress policy

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • 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++ • RSVP-TE, including: <ul style="list-style-type: none"> • Bypass LSP static configuration • Ingress LSP statistics in a file • RSVP-TE hitless-MBB with no artificial delays • 32,000 transit LSPs • Auto 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

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • LSP self-ping • MPLS fast reroute (FRR) • MTU signaling • Optimize adaptive teardown • Node/link protection • Refresh reduction • 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: <ul style="list-style-type: none"> • 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 <p>[See MPLS Applications User Guide .]</p> <ul style="list-style-type: none"> • Support for an increased scale of transit RSVP-TE–signaled MPLS label-switched paths (LSPs) that are enabled with link protection. • Enhanced scaling for the following MPLS features: <ul style="list-style-type: none"> • RSVP transit LSPs with link and node protection • RSVP ingress and egress LSPs with ultimate-hop popping (UHP) and penultimate-hop popping (PHP) • LDP-over-RSVP LSPs

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • Packet Forwarding Engine statistics • Fast reroute (FRR) and make before break (MBB) • Weighted ECMP • Ping and traceroute • Clone route • Transit statistics <p>[See MPLS Applications User Guide .]</p> <ul style="list-style-type: none"> • 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. <p>[See Point-to-Multipoint LSPs Overview .]</p> <ul style="list-style-type: none"> • 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. <p>[See MPLS Applications User Guide .]</p> <ul style="list-style-type: none"> • Optimized fast branch updates. We've refined the method of making fast-branch updates to a multicast replication tree. 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. <p>[See Multicast Shortest-Path Tree .]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Multicast	<ul style="list-style-type: none"> • 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).] • IPv4 and IPv6 multicast support including MSDP, support for PIM-SM as the first-hop router (FHR) or last-hop router (LHR), and support for anycast, static, or local rendezvous point (RP). • Support for multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows. MoFRR minimizes multicast packet loss in PIM domains when there are link failures. 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.] • Support for bidirectional Protocol Independent Multicast (PIM) for multicast traffic. [See pim-snooping.]

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Routing policy and firewall filters	<ul style="list-style-type: none"> Support added to hierarchical policers for applying user-selectable bandwidth for premium and non-premium traffic. Use the firewall filter action policer-charge to subtract available bandwidth credits and make bandwidth available to the aggregate policer. Firewall output filtering support using Fast Lookup Filter (FFT) block for line-rate performance of up to 2 billion PPS. The fast-lookup-filter statement from the CLI filter configuration prioritizes output filtering (but not input filtering) on the FFT block. FFT enables support for 128 unique output filters across IPv4, IPv6, or MPLS families. [See fast-lookup-filter (PTX).] SNMP MIB support for the jnxFirewallCounterTable object. Junos OS Evolved SNMP extends support for the jnxFirewallCounterTable and its objects: <ul style="list-style-type: none"> jnxFirewallCounterEntry jnxFWCounterPacketCount jnxFWCounterByteCount jnxFWCounterDisplayFilterName jnxFWCounterDisplayName jnxFWCounterDisplayType [See SNMP MIB Explorer.] Firewall filter support. IPv4 and IPv6 firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on the router from a source address to a destination address. [See Firewall Filter Match Conditions and Actions (PTX Series Routers).] Support for filter-based forwarding.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p data-bbox="755 352 1333 422">[See Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address.]</p> <ul data-bbox="719 457 1421 674" style="list-style-type: none"> • Support for SDN-based networks to configure certain router interfaces to pass traffic toward an SDN controller. Use firewall filters to match and redirect packets defined at the [edit services inline-monitoring instance] hierarchy level. Supported match criteria includes IPv4, IPv6, and family any (destination), VLAN ID, and certain traceroute redirect packets. <p data-bbox="755 701 922 730">[See controller.]</p> <ul data-bbox="719 766 1421 940" style="list-style-type: none"> • Support for firewall filters on discard interfaces. You can apply firewall filters on a discard interface. The action specified by the filter (log or count) is executed before the traffic is discarded. Firewall filters are supported only for IPv4 and IPv6 traffic in the egress direction of the interface. <p data-bbox="755 968 1101 997">[See Configuring Firewall Filters.]</p> <ul data-bbox="719 1033 1162 1459" style="list-style-type: none"> • Support for firewall features, including: <ul data-bbox="755 1098 1063 1459" style="list-style-type: none"> • Forwarding IPv4 and IPv6 • Firewall filter • Load balancing • MPLS fast reroute • Host path • Egress peer engineering <p data-bbox="755 1493 1398 1562">[See Firewall Filter Match Conditions and Actions (PTX Series Routers).]</p> <ul data-bbox="719 1591 1388 1766" style="list-style-type: none"> • Support for input-chain and output-chain CLI filters. Use multiple levels of CLI filters. The filter chain helps in logically grouping filters with a specific pattern of rules, instead of evaluating all the filter terms in one filter and deciding at the filter's last term. The feature provides you flexibility in

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<p>modeling the filter as and when it is applicable in the solution. You can configure up to eight filters in both input chains and output chains.</p> <p>[See Example: Using Firewall Filter Chains, output-chain, and input-chain.]</p> <ul style="list-style-type: none"> • Support for nested filters, which enable you to reference a common firewall filter by attaching it to multiple firewall policies (a filter being one or more match conditions and corresponding actions). You can bind nested filters to the following interface types: <ul style="list-style-type: none"> • <code>inet</code>—Both input and output directions • <code>inet6</code>—Both input and output directions • <code>mpls</code>—Input direction only <p>You can also bind the filters to routing instances, and in the input direction, in the output direction, or in both directions.</p> <p>[See Guidelines for Nesting References to Multiple Firewall Filters and Example: Nesting References to Multiple Firewall Filters.]</p> • Support for matching <code>ip-options</code> in IPv4 packet headers. Use the <code>ip-options</code> any match condition to match fields in the IPv4 header and create firewall filter rules to handle the matched packets. Specifying <code>ip-options</code> provides a finer level of control, so for example, you can create a rule to drop any IPv4 packets that do not include at least one IP option in the header. Configure the match condition at the <code>[edit firewall family inet filter <i>name</i> term <i>name</i> from ip-options any]</code> hierarchy level. <p>[See Firewall Filter Match Conditions for IPv4 Traffic .]</p> <ul style="list-style-type: none"> • Support for labeling interfaces with specified group IDs from 1 through 255 and matching the interface-group ID on the firewall filter. The filter recognizes which interface the packet comes from and performs actions only specified for a certain interface group.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p data-bbox="755 352 1369 384">[See Understanding BGP Flow Routes for Traffic Filtering .]</p> <ul style="list-style-type: none"> <li data-bbox="719 415 1421 483">• Firewall filter support for bitwise logical operations for TCP flag match. <p data-bbox="755 514 1341 579">[See Firewall Filter Match Conditions Based on Bit-Field Values .]</p> <ul style="list-style-type: none"> <li data-bbox="719 611 1421 1245"> • MPLS filter payload match. IPv4 and IPv6 payload fields match conditions are available for MPLS traffic. Additionally, the following match conditions are available: <ul style="list-style-type: none"> <li data-bbox="755 747 1421 856">• MPLS header EXP match conditions for MPLS traffic—exp0, exp1, exp0-except, exp1-except. Existing match conditions exp and exp-except will be deprecated. <li data-bbox="755 888 1421 997">• MPLS header Label match conditions for MPLS traffic—label0, label1, label0-except, label1-except. Existing match conditions label and label-except will be deprecated. <li data-bbox="755 1029 1421 1138">• MPLS header TTL match conditions for MPLS traffic—ttl0, ttl1, ttl0-except, ttl1-except. Existing match conditions ttl and ttl-except will be deprecated. <li data-bbox="755 1169 1421 1245">• MPLS header Bottom of Stack match conditions for MPLS traffic—bottom-of-stack0 and bottom-of-stack1 <p data-bbox="755 1276 1336 1308">[See Firewall Filter Match Conditions for MPLS Traffic .]</p> <ul style="list-style-type: none"> <li data-bbox="719 1339 1349 1371">• Unicast RPF support for both IPv4 and IPv6 traffic flows. <p data-bbox="755 1402 1211 1434">[See Configuring Unicast RPF Loose Mode .]</p> <ul style="list-style-type: none"> <li data-bbox="719 1465 1377 1533">• Enhanced scaling for DoS and protection offers loose mode unicast RPF on IPv4 and IPv6. <p data-bbox="755 1564 1211 1596">[See Configuring Unicast RPF Loose Mode .]</p> <ul style="list-style-type: none"> <li data-bbox="719 1627 1377 1766">• Support for DCU and SCU accounting. Source class usage (SCU) accounting provides a breakdown of output interface traffic statistics that originates from specific prefixes. Destination class usage (DCU) accounting provides a

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p>breakdown of input interface traffic statistics that is destined for specific prefixes.</p> <p>[See Understanding Source Class Usage and Destination Class Usage Options.]</p> <ul style="list-style-type: none"> • Class-based firewall filters. You can apply firewall filters actions such as drop, reject, sample, and police on packets classified by destination class usage (DCU) and source class usage (SCU) accounting. You can use this feature, for example, as part of a design to provide distributed denial-of-service (DDoS) protection to specific customers. <p>[See Configure the Filter Profile.]</p> <ul style="list-style-type: none"> • Support for forwarding class and packet loss priority (PLP) as policer actions. You can use forwarding class (FC), and both FC and PLP together, as policer actions in policer policy configurations. This includes both ingress and egress directions. • Support for two-color Layer 3 interface policers (ingress and egress). <p>[See Basic Two-Rate Three-Color Policers.]</p> <ul style="list-style-type: none"> • Support for packet-rate policers. You can use a count of packets as the threshold for traffic policers. Per-packet policers can better mitigate low-and-slow types of denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks. <p>You can apply packet-level policers in the ingress or egress interface direction. These policers support both two-color and three-color policers. The following families are supported: inet, inet6, mpls, and ethernet-switching .</p> <p>Configure per-packet policer rates using the pps-limit (packets per second) and packet-burst-size-limit (packets) configuration statements at the [edit firewall policer <i>policer-name</i>] hierarchy level.</p> <p>[See Packets-Per-Second (pps)-Based Policer Overview and pps-limit (Policer).]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> Shared-bandwidth and percentage policer. Use the shared-bandwidth policer for instances where policers are attached to aggregated Ethernet interface bundles with child legs spanning different Packet Forwarding Engine or Flexible Port Concentrator (FPC) instances. The bandwidth policers program the policer token bucket with weighted bandwidth or burst (depending on the number of child legs per Packet Forwarding Engine). <p>The percentage policer feature enables you to configure the bandwidth policer relative to the physical-interface speed where you configure the class-of-service (CoS) shaping rate. After the configuration, the egress policer can then use this base CoS shaping rate instead of the physical-interface speed.</p> <p>[See Configure the Filter Profile.]</p> <ul style="list-style-type: none"> Two-color and three-color traffic policers for input and output traffic. The supported actions are discard, forwarding-class, and loss-priority (high and low). You can attach policers to logical interfaces and the protocol families mpls, inet, and inet6. <p>[See Basic Two-Rate Three-Color Policers.]</p> <ul style="list-style-type: none"> 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: <ol style="list-style-type: none"> 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. <pre> set firewall tunnel-end-point <i>tunnel-name</i> ipv4 ipv6 source-address <i>address</i> set firewall tunnel-end-point <i>tunnel-name</i> ipv4 ipv6 destination-address <i>address</i> set firewall tunnel-end-point <i>tunnel-name</i> gre set firewall family inet inet6 filter <i>name</i> term <i>name</i> from source-address <i>address</i> </pre>

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<pre> set firewall family inet inet6 filter <i>name</i> term <i>name</i> then encapsulate <i>tunnel-name</i> set firewall family inet inet6 filter <i>name</i> term last then accept set interfaces <i>interface-name</i> unit <i>number</i> family inet inet6 filter input set interfaces <i>interface-name</i> unit <i>number</i> family inet inet6 address <i>address</i> # This source address differs from the one for the tunnel endpoint. 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 <i>name</i> term <i>name</i> from source-address <i>address</i> set firewall family inet inet6 filter <i>name</i> term <i>name</i> from protocol gre set firewall family inet inet6 filter <i>name</i> term <i>name</i> then decapsulate gre # Optionally de-encapsulate mpls-in-udp. set firewall family inet inet6 filter <i>name</i> term last then accept set interfaces <i>interface-name</i> unit <i>number</i> family inet inet6 filter input <i>filter-name</i> set interfaces <i>interface-name</i> unit <i>number</i> family inet inet6 address <i>address</i> # This is the destination address. [See Components of Filter-Based Tunneling Across IPv4 Networks and tunnel-end-point.] • Support for tunnel de-encapsulation using firewall filters for GRE and UDP tunnels. [See Configuring a Filter to De-Encapsulate GRE Traffic and decapsulate (Firewall Filter).] </pre>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • BGP flow specification. BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX10002-36QDD devices with 14.4 Tbps 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. <p>The following match conditions are not supported:</p> <ul style="list-style-type: none"> • ICMP codes alone [inet/inet6] • Source/destination prefix with offset for inet6 • Flow label for inet6 fragment [for inet6] <p>Junos OS Evolved running on this router doesn't support the traffic marking action.</p> <p>To configure flow routes statically, configure the match conditions and actions at the [edit routing-options] hierarchy level.</p> <ul style="list-style-type: none"> • Forwarding IPv6 transit statistics. <p>[See BGP User Guide.]</p>

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> Local port mirroring support. You can use port mirroring to copy packets entering or exiting a port or entering a VLAN and to send the copies to a local interface for local monitoring. <p>The following features are included:</p> <ul style="list-style-type: none"> Interface filter on ingress and egress Forwarding table filter (FTF) on ingress Families inet and inet6 Aggregated Ethernet interfaces at both ingress and egress <p>Use the following CLI hierarchies to configure port mirroring:</p> <ul style="list-style-type: none"> [edit interfaces] [edit forwarding-options port-mirroring] [edit firewall filter] <p>You can configure family inet and family inet6 in the [edit interfaces] and the [edit forwarding-options port-mirroring] hierarchies for this feature. This feature applies to global port mirroring only.</p> <p>[See Understanding Port Mirroring and Analyzers.]</p> <ul style="list-style-type: none"> Remote port mirroring with ToS or DSCP settings. You can send sampled copies of incoming packets to remotely connected network management software. You send the packets using GRE, which is supported by flexible tunnel interfaces (FTIs). You can set ToS and DSCP values to provide necessary priorities in the network for these packets. You can also apply policing to sampled packets that are leaving the FTI. Configure the settings you need in the [edit forwarding-options port-mirroring instance <i>instance-name</i> output] hierarchy. <p>[See instance (Port Mirroring).]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> Support for additional family any in port mirroring You can configure family any (as well as the earlier family options, inet, and inet6) for local port mirroring and remote port mirroring. You can use the family any configuration option to process the families any, ccc, ethernet-switching, or mpls. <p>NOTE: You use the family any configuration option to process all four families.</p> <p>Use [edit forwarding-options port-mirroring] for local port mirroring or [edit forwarding-options port-mirroring instance <i>instance-name</i>] for remote port mirroring, with both configurations also requiring a firewall filter.</p> <p>The following configuration statements are no longer part of the port mirroring configuration on PTX Series devices:</p> <ul style="list-style-type: none"> next-hop for family any family vpls no-filter-check hosted-service server-profile <p>[See port-mirroring.]</p> <ul style="list-style-type: none"> Support for EVPN-VXLAN filtering and port mirroring based on VNI match conditions. You can construct a firewall filter to filter EVPN-VXLAN traffic by using the VXLAN network identifier (VNI) values in the match condition on ingress and egress interfaces. This feature supports redirecting traffic to a global port-mirroring instance. <p>To filter traffic based on the VNI, use the following commands:</p> <pre>set firewall filter <i>filter-name</i> term <i>term-name</i> from vxlan vni <i>vni-value</i></pre>

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<p data-bbox="755 359 1370 420">set firewall filter <i>filter-name</i> term <i>term-name</i> from vxlan vni-except <i>vni-value</i></p> <p data-bbox="755 464 1386 489"><i>vni-value</i> can be a numeric value or range of numeric values.</p> <p data-bbox="755 525 1398 585">[See Firewall Filter Match Conditions and Actions (PTX Series Routers).]</p> <ul style="list-style-type: none"> <li data-bbox="719 625 1390 758">• 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. <p data-bbox="755 791 1127 816">[See sFlow Monitoring Technology.]</p> <ul style="list-style-type: none"> <li data-bbox="719 856 1390 917">• sFlow technology support for MPLS interfaces to sample and report MPLS traffic on the routers. <p data-bbox="755 951 1114 976">[See sFlow Technology Overview.]</p> <ul style="list-style-type: none"> <li data-bbox="719 1016 1401 1329">• Ingress and egress sFlow functionalities for transit nodes are supported for IPv4-in-IPv4, IPv6-in-IPv4, and regular IPv4/IPv6 traffic. In transit-only devices, the IP-in-IP encapsulated packet can transit through the device without any change or might get de-encapsulated and forwarded or de-encapsulated and encapsulated and forwarded based on the next-hop configuration. Additionally, the packet might traverse through multiple VRF instances while getting forwarded. The router supports ingress and egress sFlow for all those variations. <p data-bbox="755 1362 1127 1388">[See sFlow Monitoring Technology.]</p> <ul style="list-style-type: none"> <li data-bbox="719 1428 1417 1740">• sFlow technology support for exporting extended IPv4 and IPv6 tunnel egress structure. sFlow technology supports the export of the Extended Tunnel Egress Structure fields for traffic entering IPv4 or IPv6 GRE tunnels. These additional attributes provide information about the GRE tunnel into which a packet entering the device will get encapsulated. The GRE tunnel could be IPv4 or IPv6. The feature is supported only when sFlow is enabled in the ingress direction wherein firewall-based GRE happens on IPv4 or IPv6 packets.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p>The device supports the feature for the following traffic scenarios when ingress sFlow sampling is enabled:</p> <ul style="list-style-type: none"> • Incoming IPv4 traffic that undergoes IPv4 GRE • Incoming IPv6 traffic that undergoes IPv4 GRE • Incoming IPv4 traffic that undergoes IPv6 GRE • Incoming IPv6 traffic that undergoes IPv6 GRE <p>[See sFlow Monitoring Technology.]</p> <ul style="list-style-type: none"> • Sample size support in sFlow. You can configure the sFlow sample size of the raw packet header to be exported as part of the sFlow record to the collector. The configurable range of sample size is from 128 bytes through 512 bytes. <p>[See sFlow Monitoring Technology.]</p> <ul style="list-style-type: none"> • Support for passive monitoring, including support for passive monitoring on MPLS-encapsulated packets. You can configure passive monitoring on any interface on the PTX Series routers, and you can use this feature to monitor MPLS-encapsulated packets. After you enable passive monitoring, the router accepts and monitors traffic on the interface and forwards those packets to monitoring tools such as IDS servers and packet analyzers, or to other devices such as other routers or end-node hosts. <p>[See Passive Monitoring and passive-monitor-mode.]</p> <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Link discovery with active and passive modes • Detect-LOC • Remote loopback

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none">• Loopback tracking• Action profile• GRES and non-graceful Routing Engine switchover <p>[See Introduction to OAM Link Fault Management (LFM).]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Segment routing	<ul style="list-style-type: none"> • Segment routing support. You can configure the following Source Packet Routing in Networking (SPRING) or segment routing features on the router: <ul style="list-style-type: none"> • MPLS (segment routing using IS-IS): <ul style="list-style-type: none"> • Ping and traceroute for single IS-IS node or prefix segment • BGP Link State (BGP-LS): <ul style="list-style-type: none"> • Segment routing extensions for IS-IS • Segment routing extensions for OSPF • BGP: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Adjacency SID • Advertising maximum link bandwidth and administrative color without RSVP-TE configuration • Anycast and prefix SIDs • Configurable segment routing global block (SRGB)

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • Node and link SIDs • Segment Routing Mapping Server (SRMS) and client • Topology Independent Loop-Free Alternate (TI-LFA): <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Link and node protection • Protection for SRMS prefixes

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF • Non-colored 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: <ul style="list-style-type: none"> • 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 non-colored SR-TE LSPs • First-hop label support for SID instead of an IP address • 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: <ul style="list-style-type: none"> • 00—Route resolution fallback to IGP path • 01—Route resolution fallback to color only null routes

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> Static LSPs with member-link next hops for aggregated Ethernet bundles (also known as adjacent SID per LAG bundle or aggregated Ethernet member link) <p>[See Understanding Source Packet Routing in Networking (SPRING).]</p> <ul style="list-style-type: none"> Support for scaled-up static and BGP segment routing policies, where each policy contains eight segment routing paths with five labels per path without make-before-break (MBB). <p>[See egress-chaining and fib-next-hop-split.]</p> <ul style="list-style-type: none"> SPRING statistics sensor support for JTI supports export of SPRING statistics to an outside collector by using remote procedure call (gRPC) services. The feature provides the segment-identifier (SID)-level and interface-level traffic counts for SPRING traffic. These statistics reflect the SPRING LSP utilization in the traffic engineering database, which aids in correctly rerouting the RSVP LSPs. <p>To enable SPRING statistics, include the following statements on the client device:</p> <ul style="list-style-type: none"> For egress (per-interface egress), use <code>set protocols isis source-packet-routing sensor-based-stats per-interface per-member-link egress</code> For egress (per-SID egress), use <code>set protocols isis source-packet-routing sensor-based-stats per-sid egress</code> For ingress (per-SID ingress), use <code>set protocols isis source-packet-routing sensor-based-stats per-sid ingress</code>. <p>Use the following sensors to export statistics by means of gRPC services to an outside collector:</p> <ul style="list-style-type: none"> <code>/junos/services/segment-routing/interface/egress/usage/</code> for egress (per-interface egress) aggregate SPRING traffic.

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • <code>/junos/services/segment-routing/sid/usage/</code> for egress (per-SID egress) and ingress (per-SID ingress) aggregate SPRING traffic. <p>[See source-packet-routing and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p> <ul style="list-style-type: none"> • BGP and statically configured SR-TE traffic statistics sensor support for JTI. <p>[See source-packet-routing, Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface), and Understanding OpenConfig and gRPC on Junos Telemetry Interface.]</p> <ul style="list-style-type: none"> • Support for segment routing over UDP. Configure the <code>udp-tunneling encapsulation</code> statements at the <code>[edit protocols isis source-packet-routing]</code> hierarchy level to enable SR-MPLS routers and IP-only routers to seamlessly coexist, by encapsulating SR-MPLS label stacks in IP/UDP encapsulation. This feature also supports: <ul style="list-style-type: none"> • Entropy in the UDP source port • Underlay and overlay ECMP at the start of the tunnel • Policy control to resolve dynamic tunnels <p>SR-over-UDP supports tunnels without a loopback stream in the Packet Forwarding Engine, thereby reducing additional bandwidth consumption.</p> <p>[See Next-Hop-Based Dynamic Tunnels and source-packet-routing (Protocols IS-IS).]</p> • SPRING : JTI : Ingress SR-TE statistics per binding SID and segment list (static, BGP, PCEP paths). Use this feature to provide route statistics for segment routing-traffic engineering (SR-TE) per label-switched path (LSP). Junos OS Evolved uses Junos telemetry interface (JTI) and gRPC services to provide the statistics. <p>Supported resource paths (sensors) include:</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • <code>/junos/services/segment-routing/traffic-engineering/tunnel/lsp/ingress/usage/</code> • <code>/junos/services/segment-routing/traffic-engineering/tunnel/lsp/transit/usage/</code> <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface, source-packet-routing, and show spring-traffic-engineering.)]</p> <ul style="list-style-type: none"> • SPRING statistics sensor support for JTI supports export of SPRING statistics to an outside collector by using remote procedure call (gRPC) services and gRPC Network Management Interface (gNMI) services. This feature provides interface-level and segment identifier (SID)-level ingress statistics. The feature also provides egress statistics for each child member at the physical interface level. <p>To enable SPRING statistics, include the following statements on the client device:</p> <ul style="list-style-type: none"> • For egress (per-child member at the physical interface level), use the set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress command. • For ingress (per-SID ingress and per-interface ingress), use the set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress command. <p>Use the following sensors to export statistics by means of gRPC or gNMI services to an outside collector:</p> <ul style="list-style-type: none"> • <code>/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/in-octets/</code> for ingress (per-SID ingress and per-interface ingress) SPRING traffic. • <code>/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/</code>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p>out-octets/ for egress (per-child member at the physical-interface level) SPRING traffic.</p> <ul style="list-style-type: none"> • /network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/out-pkts/ for egress (per-child member at the physical-interface level) SPRING traffic. <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface and source-packet-routing.]</p> <ul style="list-style-type: none"> • Support for segment-routing telemetry sensor enhancements. We support segment routing sensor enhancements for SID-level and interface-level traffic counts. These enhancements comply with the current supported sensors in the OpenConfig models openconfig-segment-routing.yang and openconfig-mpls.yang. • SR-TE colored policy RIB5 and SR-TE colored telemetry sensor support. We support JTI streaming and ON-CHANGE sensors that deliver operational state statistics for SR-TE colored policy RIB5 and SR-TE colored telemetry sensors. Statistics are delivered to an outside collector using gRPC or gNMI. The feature includes new OpenConfig resource paths for existing and new SR-TE policy (tunnel) and SR-TE per-LSP colored statistics. <p>[See Telemetry Sensor Explorer.]</p> <ul style="list-style-type: none"> • Support for SRv6 network programming in IS-IS. Use this feature to configure segment routing in a core IPv6 network without an MPLS dataplane. <p>To enable SRv6 network programming in an IPv6 domain, include the srv6 statement at the [edit protocols isis source-packet-routing] hierarchy level.</p> <p>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.</p>

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<p>To configure a TI-LFA backup path for SRv6 in an IS-IS network, include the <code>transit-srh-insert</code> statement at the [edit protocols isis source-packet-routing srv6] hierarchy level.</p> <p>[See How to Enable SRv6 Network Programming in IS-IS Networks.]</p> <ul style="list-style-type: none"> Support for SRv6 network programming and Layer 3 Services over SRv6 in BGP. You can configure BGP-based Layer 3 service 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. Such networks depend only on the IPv6 headers and header extensions for transmitting data. <p>To configure IPv4 and IPv6 transport over an SRv6 core, include the <code>end-dt4-sid</code> <i>sid</i> and the <code>end-dt6-sid</code> <i>sid</i> statements at the [edit protocols bgp source-packet-routing srv6 locator name] hierarchy level.</p> <p>To configure IPv4 VPN and IPv6 VPN service over an SRv6 core, include the <code>end-dt4-sid</code> <i>sid</i> and the <code>end-dt6-sid</code> <i>sid</i> statements at the [edit routing-instances <i>routing-instance-name</i> protocols bgp source-packet-routing srv6 locator <i>name</i>] hierarchy level.</p> <p>[See Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP.]</p> <ul style="list-style-type: none"> OAM ping support for segment routing with IPv6 (SRv6) network programming. You can perform an Operations, Administration and Management (OAM) ping operation for any SRv6 segment identifier (SID) whose behavior allows upper layer header processing for an applicable OAM payload. <p>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.</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<p data-bbox="755 352 1398 420">[See ITU-T Y.1731 Ethernet Service OAM Overview and How to Enable SRv6 Network Programming in IS-IS Networks.]</p> <ul data-bbox="719 457 1370 667" style="list-style-type: none"> • 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. <p data-bbox="755 697 1344 764">[See How to Enable SRv6 Network Programming in IS-IS Networks.]</p> <ul data-bbox="719 802 1393 898" style="list-style-type: none"> • SRv6 support for static SR-TE policy. You can configure static segment routing-traffic engineering (SR-TE) tunnels over an SRv6 data plane. <p data-bbox="755 928 1365 995">Use the following configuration commands to enable SRv6 support:</p> <ul data-bbox="755 1033 1390 1306" style="list-style-type: none"> • For an SR-TE policy: <code>set protocols source-packet-routing srv6</code> • For an SR-TE tunnel: <code>set protocols source-packet-routing source-routing-path lsp <i>name</i> srv6</code> • For an SR-TE segment list: <code>set protocols source-packet-routing source-routing-path segment-list srv6</code> <p data-bbox="755 1335 1284 1360">[See Understanding SR-TE Policy for SRv6 Tunnel.]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
Services applications	<ul style="list-style-type: none"> • Inline active flow monitoring support. [See Understand Inline Active Flow Monitoring.] • Juniper Resiliency Interface support. [See Juniper Resiliency Interface.] • Inline monitoring services support for packet mirroring with metadata. [See Inline Monitoring Services Configuration.] • 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: <ul style="list-style-type: none"> • 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.] • CFM support: <ul style="list-style-type: none"> • 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 Ethernet VPN (EVPN) • Distributed session support for connectivity fault management (CFM) on aggregated Ethernet • Enhanced CFM mode

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Continuity check messages (CCM) • LBM • LBR • Link Trace Message (LTM) • Link Trace Reply (LTR) • 1DM (one-way delay measurement) • 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

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • 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) <p>[See Connectivity Fault Management (CFM).]</p> <ul style="list-style-type: none"> • Support for enhanced CFM. The feature extends CFM support to inline mode. Support includes: <ul style="list-style-type: none"> • Up and down maintenance association end points (MEPs) on bridges, circuit cross-connect (CCC), and Ethernet VPN (EVPN) in inline mode • ITU-T Y.1731 on synthetic loss measurement (SLM) and delay measurement (DM) • Inline session support for connectivity fault management (CFM) on aggregated Ethernet • Enhanced CFM mode by default • Supported inline performance monitoring (PM) sessions, as follows: <ul style="list-style-type: none"> • PM Tx • PM Rx • PM responder • IPv4 (inet) and IPv6 (inet6) support for continuity check messages (CCM), delay measurement (DM), and synthetic loss message (SLM)

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • DM on aggregated Ethernet with at least one child link on the anchor Packet Forwarding Engine • Action profile for marking a link down, except for EVPN and bridge up MEP • Supported CFM protocol data units (PDUs) for inline handling, as follows: <ul style="list-style-type: none"> • CCM • Delay measurement message (DMM) • Delay measurement reply (DMR) • Synthetic loss message (SLM) • Synthetic loss reply (SLR) • Enterprise and service provider configurations • VLAN normalization • VLAN transparency for CFM PDUs • Combination of up MEP, down MEP, or maintenance association intermediate point (MIP) configuration over the same interface <p>[See Connectivity Fault Management (CFM).]</p>
Security services	<ul style="list-style-type: none"> • Support for DDoS IS-IS classification and higher DDoS bandwidth for Layer 2 and Layer 3 protocols. <p>[See show ddos-protection protocols isis and protocols (DDoS) (ACX Series, PTX Series, and QFX Series).]</p>
Software installation and upgrade	<ul style="list-style-type: none"> • Support for secure BIOS and secure boot implementation based on the UEFI 2.4 standard. <p>[See Secure Boot.]</p>

Table 1: PTX10002-36QDD Feature Support *(Continued)*

Feature	Description
VPNs	<ul style="list-style-type: none"> • MPLS-based Layer 3 VPNs support includes: <ul style="list-style-type: none"> • 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 <code>vrf-table-label</code> statement • ICMP tunneling and MPLS traceroute • Disabling time-to-live (TTL) decrementing using <code>no-propagate-ttl</code> <p>[See Layer 3 VPNs Feature Guide for Routing Devices.]</p> • Carriers-of-carriers and inter-AS VPN supported features include: <ul style="list-style-type: none"> • Carrier-of-carriers VPN service • Interprovider Layer 3 VPN Option A • Interprovider Layer 3 VPN Option B • Interprovider Layer 3 VPN Option C <p>However, traffic statistic collection for BGP labeled unicast is not supported for carrier-of-carrier VPNs and interprovider traffic.</p> <p>[See Carrier-of-Carrier VPNs.]</p> • Layer 2 VPN feature support includes: <ul style="list-style-type: none"> • 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)

Table 1: PTX10002-36QDD Feature Support (*Continued*)

Feature	Description
	<ul style="list-style-type: none"> • Local and remote switching • Ethernet and VLAN CCC • Single-tagged CCC logical interfaces • Control word • Regular and aggregated Ethernet interfaces • 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 <p>[See Layer 2 VPNs and VPLS User Guide for Routing Devices and TCC Overview.]</p>

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

Authentication and Access Control

- **Support for outbound SSH through HTTP proxy servers (PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5230-64CD)**—You can establish outbound SSH connections through an HTTP proxy server to enable secure remote management of devices, even when firewalls block outbound SSH connections.

Use the outbound-ssh client *name* proxy-server configuration statement to configure proxy server details.

[See [outbound-ssh](#).]

Chassis

- **Support for powering on, powering off, or restarting Packet Forwarding Engine (PTX10002-36QDD)**—You can power off, power on, or restart the Packet Forwarding Engines in the PTX10002-36QDD router by following these steps:

1. Configure the *pair* of Packet Forwarding Engines that you want to restart, power off, or power on—for example:

- `set chassis fpc 0 pfe 0 power on`
- `set chassis fpc 0 pfe 1 power on`



NOTE: The four Packet Forwarding Engines are numbered 0–3. You configure them in pairs—0 and 1; 2 and 3.

2. Issue the `power on`, `power off`, or `restart` command—for example:

- `request chassis fpc slot 0 pfe 0 power on`

3. Enter **yes** when the following question appears on the screen:

- Warning: pfe 1 will also be offlined. Do you wish to continue?
[yes,no]



NOTE: You can also set the `reset-pfe` action to reset a Packet Forwarding Engine when a chassis error occurs. Configure the action statement at `[edit chassis fpc slot-number error error-severity-level]` hierarchy level.

[See [request chassis fpc](#) and [action \(chassis error\)](#).]

Class of Service

- **Support for policy maps (PTX10002-36QDD)**—Use policy maps on PTX10002-36QDD routers to assign rewrite rules on a per-customer basis. You can use any packet field to identify a given flow and specify a rewrite value for that flow. PTX10002-36QDD routers support the following types of packet marking: INET-Precedence, DSCP, IEEE 802.1p, and IEEE 802.1ad..

You can define a policy map by including the `policy-map` statement at the `[edit class-of-service]` hierarchy level. You enable `policy-map-marking` at the egress interface at the `[edit class-of-service interfaces interface-name unit unit-number]` hierarchy level.

On PTX10002-36QDD routers, you can also use policy maps to pass meta data between firewall filters. That is, a policy map value set in an ingress filter can be matched on in an egress filter.

[See [Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps](#).]

- **Support for ECN copy on EVPN-VXLAN tunnels (PTX10002-36QDD)** —PTX10002-36QDD routers that originate or terminate EVPN-VXLAN tunnels and have explicit congestion notification (ECN) enabled automatically copy the ECN bits from the inner header to the outer header. The router

copies the ECN bits from the outer header to the inner header if the inner header has the ECT bit set. If the router experiences congestion, it sets the CE bits if the ECT bit is enabled.

[See [CoS Support on EVPN VXLANs](#).]

- **Support for low-threshold ECN (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—You can define a buffer rate, which is the base rate for buffer size calculation. The buffer rate is the target rate of a virtual output queue (VOQ), which is the intended egress queue rate during typical congestion.

Configure the buffer-rate statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy level.

You can also define more granular fill-level percentages for drop profiles. That is, you can now set fill-level percentages to tenths of a percent instead of just whole percentages (for example, 50.2 percent versus just 50 percent).

Define drop profiles at the [edit class-of-service drop-profiles *profile-name*] hierarchy level.

Setting the buffer rate and defining more granular drop profiles can improve the buffer resolution on lower-end buffers, which helps implement low-threshold explicit congestion notification (ECN). (Low-threshold ECN triggers ECN marking as soon as the buffer starts filling up).

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

- **Support for shared VOQ queue-depth monitor profiles across ae- interfaces (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—By default, a monitoring profile that you assign to an aggregated Ethernet (ae-) interface replicates across all members of the ae- interface. The monitoring profile also reports virtual output queue (VOQ) depth individually on each interface. On large systems, this process can quickly consume the maximum supported hardware monitoring profile IDs. To conserve monitoring profile IDs, include the shared option at the [set class-of-service interfaces *ae-interface* monitoring-profile *profile-name*] hierarchy level. The configured shared option creates only one monitoring profile ID to share across all member interfaces. The option also reports the largest peak on a member interface as the common peak for the ae- interface.

[See [VOQ Queue-depth Monitoring](#).]

Ethernet Switching and Bridging

- **Support for interface MAC limit action (PTX10002-36QDD)**—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](#).]

High Availability

- **Adaptive load balancing for ECMP (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can use the adaptive load balancing (ALB) feature for ECMP to address traffic distribution imbalances by monitoring packet and byte counts for each hash bucket. This feature employs a feedback mechanism to detect and correct overload conditions, ensuring fair traffic distribution across links.

Use the adaptive CLI option at the [edit policy-options policy-statement *policy* then load-balance] hierarchy level to enable ALB. To adjust the tolerance percentage, use the `ecmp-alb tolerance percentage` CLI option at the [edit chassis] hierarchy level.

[See [Configuring Adaptive Load Balancing](#).]

Interfaces

- **Support for MAC accounting for source and destination MAC addresses for L3 interfaces (PTX10002-36QDD)**—We support media access control (MAC) accounting for source and destination MAC addresses for Layer 3 (L3) interfaces and aggregated Ethernet interfaces. To enable MAC accounting, use the `mac-learn-enable` configuration statement at the [edit interfaces *interface-name* `gigether-options ethernet-switch-profile`] or the [edit interfaces `aex aggregated-ether-options ethernet-switch-profile`] hierarchy level.

[See [show interfaces mac-database](#).]

Junos Telemetry Interface

- **Support for gNMI telemetry subscriptions using the Opstated application and genstate YANG data models (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved supports the ODL-annotation-based YANG data model with a new application called Opstated. The Opstated application supports telemetry subscriptions that use the genstate YANG data model. Genstate YANG data models expose a subset of `show` command data through the gRPC Network Management Interface (gNMI) subscribe RPC. Opstated works with the command and remote procedure call (`cmd.dd`) data model within the Junos OS Evolved software to provide the available state found in a CLI operational mode command. Data is then efficiently streamed to a collector. [See [Junos Genstate YANG Data Models](#), [Juniper Github](#), and [gNMI Genstate Subscription](#). For sensors, see [Junos YANG Data Model Explorer](#).]
- **Support for policers and ACLs in firewall filters (ACX7024 and PTX10003)**—The ACX7024 and PTX10003 support subscribable YANG data models for operational states. The genstate YANG models expose a subset of `show` command data through the gNMI subscribe RPC. A gNMI telemetry collector can subscribe to the resource paths defined in the published models to query for specific state data. This feature provides genstate YANG data model support for policers and ACLs in firewall filters.

[See [Junos Genstate YANG Data Models](#) and [gNMI Genstate Subscription](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **Support for device hardware capacity statistics (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved supports improved statistics that inform you about the hardware components that you can add to a device. The resource path `/components/component/state/empty` now returns different values. In previous releases, the resource path streamed a value of `False` to indicate a field replaceable unit (FRU) slot, whether populated or unpopulated with hardware. This behavior has changed. There are now two values: `True` and `False`. The resource path streams the value `True` when the FRU slot is available, but not populated. The value `False` is streamed when the FRU slot is populated.

This feature supports the components Flexible PIC Concentrator (FPC), SFP transceiver, power supply module (PSM), Routing Engine, Switch Interface Board (SIB), and fan tray. Streaming and ON-CHANGE subscriptions support this resource path using the transport's gRPC Network Management Interface (gNMI) or Juniper's proprietary remote procedure call (gRPC). The feature also supports both INITIAL_SYNC and target-defined mode.

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

- **Support for ARP table, IPv6 neighbor discovery, and NTP sensors in genstate YANG data models (PTX10003)**—The PTX10003 supports subscribable YANG data models for operational states. The genstate YANG models expose a subset of `show` command data through the gNMI subscribe RPC. A gNMI telemetry collector can subscribe to the resource paths defined in the published models to query for specific state data. This feature supports state data for the Address Resolution Protocol (ARP) table, IPv6 neighbor discovery, and Network Timing Protocol (NTP). The genstate YANG data model supports the resource paths `genstate:/genstate/ipv6-nd-information/`, `genstate:/genstate/arp-table-information/`, and `genstate:/genstate/ntp-status/`.

[See [Junos Genstate YANG Data Models](#) and [gNMI Genstate Subscription](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **Support for Health Monitoring telemetry data for standby nodes and FPCs (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—We've expanded telemetry data to support Health Monitoring telemetry beyond the primary node to include standby nodes and Flexible PIC Concentrators (FPCs). You can stream statistics that include load average, process parameters, and component CPU utilization using either Juniper's proprietary remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) transport from the device to the collector.

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

- **OpenConfig sensor support for ZR and ZR+ optical transceivers (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved supports data streaming for ZR and ZR+

optics. You can create a subscription in INITIAL_SYNC or TARGET_DEFINED mode using Juniper's proprietary remote procedure call (gRPC) service or gRPC Network Management Interface (gNMI). Use these resource paths in a subscription to stream data:

- `/components/component/optical-channel/state/target-output-power`
- `/components/component/optical-channel/config/target-output-power`
- `/components/component/transceiver/state/supply-voltage/` new leaves instant, avg, min, max, interval, min-time, and max-time
- `/components/component/transceiver/physical-channels/channel/state/input-power/` new leaves avg, min, max, interval, min-time, and max-time
- `/components/component/state/temperature/` new leaves instant, avg, min, max, interval, min-time, and max-time, alarm-status, alarm-threshold, and alarm-severity

This feature is based on data models `openconfig-terminal-device.yang` (version 1.9.0), `openconfig-platform-transceiver.yang` (version 0.13.0), and `openconfig-platform.yang` (version 0.21.0).

[For sensors, see [Junos YANG Data Model Explorer](#). For CLI operational mode commands, see [show interfaces diagnostics optics](#) (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) and [show interfaces extensive](#).]

- **Periodic streaming of selected prefixes using IP oc-aft (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved supports periodic streaming of specific prefixes under the IP OpenConfig Abstract Forwarding Table (oc-aft) sensor child path `/network-instances/network-instance/afts/`. To enable prefix filtering on the target (source) device, include the prefix statement at the `[edit fib-streaming prefix-list table table-name family family-name]` hierarchy level. After you enable this feature, only interface data with the required prefixes and their corresponding next hops and next-hop group containers are exported to the oc-aft collector. Reducing the set of interfaces to only the ones of interest exported to the collector decreases the overall CPU and resource usage on Routing Engines, Flexible PIC Concentrators (FPCs), and Modular Port Concentrators (MPCs). The recommended periodic interval for streaming resource paths under `/network-instances/network-instance/afts/` is 5 minutes.

[See [Configuring Prefix Filtering](#), [prefix-list](#), [show fib-streaming state](#), and [Junos YANG Data Model Explorer](#).]

- **Support for backup next-hop group sensor (PTX10008 and PTX10016)**—With this feature, you can send telemetry data for the backup next-hop group from your device to the collector using both streaming and ON_CHANGE subscriptions. These subscriptions utilize Juniper's proprietary gRPC or gRPC Network Management Interface (gNMI). Use the resource path `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/state/backup-next-hop-group` in a subscription.

To enable this feature, add the backup-next-hop-group statement at the `[edit system fib-streaming model ocaft]` hierarchy level in the configuration mode.

Deleting the configuration disables the feature:

```
delete system fib-streaming model ocraft backup-next-hop-group
```

[See [Configuring Prefix Filtering](#), [prefix-list](#), [show fib-streaming state](#), and [Junos YANG Data Model Explorer](#).]

- **Support for genstate YANG data models (PTX10003)**—The PTX10003 supports genstate YANG data models for operational state. You can subscribe to these models, which expose a subset of `show` and `show extension-service` command data through the gNMI subscribe RPC. A gNMI telemetry collector can subscribe to the resource paths in the data models to query specific state data. Periodic streaming telemetry is also supported. We've added the following resource paths:

- `genstate:/genstate/interface-information/physical-interface`
- `genstate:/genstate/interface-information/physical-interface/logical-interface`
- `genstate:genstate/request-response-client-information`
- `genstate:genstate/request-response-server-information`

We've removed the `genstate:/genstate/interface-information/logical-interface` resource path

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

- **Support for genstate YANG data models (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can subscribe to genstate YANG models on Junos OS Evolved devices to access a subset of `show` command data. This feature allows a gNMI telemetry collector to subscribe to resource paths in the models, enabling you to query specific state data. This feature supports the `show agent sensors` command. The supported root resource path is `genstate:genstate/sensor-information`.

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

- **IPFIX telemetry support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—You can stream Inline J-Flow for IP Flow Information Export (IPFIX) operational states on the `/state/sampling/flow-monitoring/` native model path.

Test gRPC, gNMI, and UDP transport protocols on PTX Series platforms with fixed-form-factor and multi-line-card chassis. Data export varies based on J-Flow configuration on single or multiple lines. If you configure J-Flow on IPv4, IPv6, and MPLS families, all supported leaves are exported with the correct values for corresponding families. After configuring J-Flow on all families, if you delete one family and add it back, the values are exported for all families initially.

Data is not exported for the deleted families. If you configure the family again, data export resumes for the newly added family. Junos OS Evolved does not allow telemetry subscriptions for a few unsupported leaves.

[For a complete list of all other sensors available, see [Junos YANG Data Model Explorer](#).]

- **UDP streaming support (PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved introduces User Datagram Protocol (UDP) streaming support on the PTX10003, PTX10004, PTX10008, and PTX10016 routers with the following sensors:
 - `/system/state/hostname`
 - `/system/state/...`
 - `/system/clock/state/timezone-name`
 - `/system/mount-points/mount-point/state/...`
 - `system/processes/process[pid]`
 - `/pidsystem/processes/process[pid]/state/...`
 - `/components/component[RE0:CPU0]/name`
 - `/components/component[RE0:CPU0]/cpu/utilization/state/...`
 - `/components/component[RE0:CPU0]/state/description`
 - `/components/component[RE0:CPU0]/state/...`

[For a complete list of sensors available, see [Junos YANG Data Model Explorer](#).]

Layer 2 VPN

- **VLAN ID lists for Layer 2 circuits (PTX10001-36MR, PTX10004, and PTX10008)**—We support VLAN ID lists for Layer 2 circuits on the listed PTX Series routers. With VLAN ID lists, you can 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](#).]

MACsec

- **MACsec bounded delay protection (PTX10002-36QDD)**—You can enable Media Access Control Security (MACsec) bounded delay protection to protect your network against man-in-the-middle attacks. When you enable MACsec bounded delay protection, the device ensures that a frame is not sent after a delay of two seconds or more. MACsec periodically compares the number of frames transmitted to the number received. If a frame is sent but not received within two seconds, such as during a man-in-the-middle-attack, MACsec drops the packet.

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

MPLS

- **SRv6-TE tunnels with micro-SIDs in PCEP (ACX7100-32C, ACX7100-48L, ACX7024, ACX7332, ACX7348, ACX7509, and PTX10002-36QDD)**—This feature enhances traffic engineering and network optimization by enabling the reporting, delegation, and creation of these tunnels. You can report and delegate static SRv6-TE tunnels with micro-SID configurations to a PCE and initiate these tunnels through PCE, improving control and management. Key functionalities include reporting static SRv6-TE tunnels with micro-SIDs to the PCE, delegating their management, and creating them with proper SID structure and endpoint behavior checks. Existing CLI commands are extended to support these features, facilitating effective configuration and monitoring.

[See [SRv6-TE Tunnels with micro-SIDs in PCEP](#).]

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

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



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

[See [PCEP Configuration](#).]

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

[See [Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP](#).]

Multicast

- **Enhanced MVPN provider tunnel selection criteria (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We support the following enhanced MVPN provider tunnel selection criteria to finetune multicast path-selection across the core network.
 - Regular Expression for selecting RSVP tunnels for ingress replication.
 - Colored inet.3 table for ingress replication.
 - Root Address for MLDP P2MP tunnels.

[See [Provider Tunnel Selection In Ingress Replication](#).]

- **Enhanced L3 multicast operational commands (ACX7100-32C, PTX10004, and QFX5130-32CD)**—The show instance command is now extended to all routing instances for the following commands. Earlier, only specific PIM-enabled routing instances were displayed.

- show pim join instance all
- show pim rps instance all
- show pim statistics instance all
- show multicast route instance all
- show multicast statistics instance all

Additionally, the show pim statistics output will display V2 Sparse Join and V2 Sparse Prune counters.

The show igmp statistics output will also display the V1/V2/V3 Membership Query field.

[See [show pim statistics](#), [show multicast statistics](#), and [show igmp statistics](#).]

- **Multicast support for Next-Generation MVPN (NGMVPN) (PTX10002-36QDD)**—Support includes:
 - IR, RSVP-P2MP, and LDP-P2MP provider tunnel
 - Inclusive and selective PMSI tunnel
 - Rendezvous-point tree (RPT)-shortest-path tree (SPT) mode
 - Restart individual PFE instances
 - Turnaround provider edge (PE) device
 - RP mechanisms, including 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](#).]

Multichassis Link Aggregation (MC-LAG)

- **Multichassis Link Aggregation Groups (MC-LAGs) support (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**—We support MC-LAGs, which enable a client device to form a logical LAG interface between two MC-LAG peers. An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running STP.

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

Network Management and Monitoring

- **sFlow OpenConfig translation and telemetry enhancement (PTX10008)**—We've extended the sFlow OpenConfig translation and telemetry functionality. When you configure the first path in an OpenConfig container, all the associated OpenConfig default settings for that container are automatically enabled. The default values are available in the `openconfig-sampling-sflow.yang` DM.

This functionality simplifies the configuration process by ensuring that all necessary defaults are in place without requiring additional manual configuration.

The default sample size is 128 bytes but you can also set an explicit sample size. You can verify the sample size using the `run show sflow` command.

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

- **sFlow support for filter-based UDP and GRE tunnel de-encapsulation (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Use sFlow to monitor ingress and egress traffic profiles for filter-based GRE and UDP tunnel de-encapsulation. sFlow support for UDP tunnel de-encapsulation enhances network monitoring with statistical sampling of network packets through firewall filters. This feature supports IPv4, IPv6, and MPLS payloads, where de-encapsulation and firewall actions are enabled on the ingress side. Key functionalities include de-encapsulation action mapping for preserving payload headers, handling VRF configuration for accurate packet lookup, and determining sample content based on packet structure. The device supports a sampling rate of up to 35,000 samples per second.

The behavior of sFlow for UDP de-encapsulation is consistent with that of GRE de-encapsulation, providing a unified approach to monitoring across different tunneling methods.

- **Chunked framing support in NETCONF sessions (PTX10008)**—Junos devices support the chunked framing mechanism for messages in a NETCONF session. Chunked framing is a standardized framing mechanism that ensures that character sequences within XML elements are not misinterpreted as message boundaries. If you enable RFC 6242 compliance, and both peers advertise the `:base:1.1` capability, the NETCONF session uses chunked framing for the remainder of the session. Otherwise, the NETCONF session uses the character sequence `]]>]]>` as the message separator.

[See [Configure RFC-Compliant NETCONF Sessions](#).]

- **Ingress and egress sFlow support for L2 traffic (PTX10001-36MR, PTX10004, PTX10008, PTX10016)**— 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](#).]

Routing Policy and layer2-policer Firewall Filters

- **Support for including layer 2 header in policer overhead calculation (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016)**—Support is added to include layer 2 header in the policer overhead calculation using the `layer2-policer` configuration statement. By default, policer overhead calculation is - layer 3 header length + payload length. After setting this configuration statement, policer overhead accounting calculation is - layer 2 header length + layer 3 header length + payload length. Supported for any, INET, INET6, and MPLS firewall filter families.

[See [layer2-policer](#).]

- **Support for counting the number of BGP large communities (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016)**—You can use `large-community-count` to count the number of BGP large communities.

[See [large-community-count](#).]

- **Support added to hierarchical policers for applying user-selectable bandwidth for premium and non-premium traffic (PTX10002-36QDD)**—You can use the new firewall filter action `policer-charge` to subtract available bandwidth credits and make it available to the aggregate policer.

[See [policer-charge](#).]

- **Scalable source address based forwarding (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Source based forwarding is used for providing value added services where paths through the network are selected based on the packet's source address. In some deployments, the destination node of the packet, identified by the protocol next hop, stays the same, but the path through the network may vary. However, in some deployments, the destination of the packet changes too. In this scenario, separate forwarding tables are used for traffic forwarding, and traffic is steered into these tables based on the packet's source address. `source-based-forwarding` and `source-lookup` configuration statements have been introduced to enable this functionality.

[See [source-based-forwarding](#) and [source-lookup](#).]

- **Support for matching first byte of payload value (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**

[See [Firewall Filter Match Conditions for IPv4 Traffic](#).]

- **Support for matching IPv6 flow label field (PTX10008)**—Support is added for matching the 20-bit `flow-label` field in the header of an IPv6 packet. We've added two new match conditions for this feature—`flow-label flow-label value` and `flow-label flow-label value mask mask value`.

[See [Firewall Filter Match Conditions for IPv6 Traffic](#).]

- **Support for increasing firewall filter scale (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—We support two new configuration statements—`scale-mode` and `no-incremental-update`. Use `scale-mode` to

accommodate more firewall filter terms, when you're focused more on scale than on performance. Use `no-incremental-update` to prevent the firewall filter from undergoing incremental update; the filter undergoes make-before-break (MBB).

[See [scale-mode](#) and [no-incremental-update](#).]

Routing Protocols

- **Improved handling of Packet Forwarding Engine install errors (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016)—**

We have extended the mechanism whereby the forwarding plane can inform the control plane of any errors in next hops. If a route error is received from the Packet Forwarding Engine, the route error is directly added to the route-error list. For next-hop errors, the error is propagated to the parent route and the parent route is added to the route-error list. Once the route pointing to the next hop with an error is determined, the protocols can then take corrective action. To enable this error checking, configure the `pfe-install-error` statement at the `[edit routing-options forwarding-table]` hierarchy level, then restart the `rpd` process. Two commands are available to check on these route errors after you have configured this feature:

- If the forwarding plane is not forwarding traffic, use the `show route pfe-install-error` command to check if the Packet Forwarding Engine has sent an error for the route that is dropping traffic.
- To check for next-hop errors, use the `show routing nhdb pfe-install-error` command. We support errors for routes pointing to a limited set of next hops: `RT_NH_INDIRECT`, `RT_NH_ROUTER`, `RT_NH_CHAIN`, `RT_NH_INDXD`, `RT_NH_LIST`, `RT_NH_TUNNEL_COMP`, and `RT_NH_FRR_INDIRECT`.

[See [forwarding-table](#), [show route pfe-install-error](#), [show routing nhdb pfe-install-error](#), and [Packet Forwarding Engine Install Error Checking](#).]

- **Enhanced Routing Policies and Multi-Instance IS-IS Support (ACX-Series, QFX-series, and PTX-Series)—**We've introduced enhancements to simplify routing policies and improve IS-IS multi-instance support. You can now tag local and direct routes with tag and tag2 values, match multiple tag2s in a single policy term, and set IS-IS Down bits during inter-instance route redistribution for precise control. Policy configurations support regex for dynamic matching of multiple IS-IS instances, while wildcard patterns streamline operational commands. Additionally, administrators can reuse the same Micro SID Locator and Node-SID across IS-IS instances, enhancing SRv6 scalability. These updates reduce complexity, improve flexibility, and provide greater control for efficient network management.

Services Applications

- **Inline active flow monitoring IP Flow Information Export (IPFIX) and version 9 template support for CoS policy-map name reporting in the ingress direction (PTX10002-36QDD)—**We support a new Juniper-specific enterprise Information Element ID, 32765, in the data record templates `ip4-template`

and ipv6-template. This IE ID is four bytes long and contains the first four characters of the policy-map name. Ensure the first four letters of your policy-map are unique.

Configure this IE ID with the `include-policy-map-name` statement at the `[edit services flow-monitoring (version-ipfix|version9) template-name data-record-fields]` hierarchy level. Configure policy maps at the `[edit class-of-service policy-map]` hierarchy level.

[See [data-record-fields](#), [Understand Inline Active Flow Monitoring](#), and [Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps](#).]

- **Inline active flow monitoring multiple BGP next-hop support (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—We have added support for reporting an accurate BGP next-hop address for the load-balanced traffic over multiple BGP peers in the ingress direction. Previously, we reported only the first address in a list of BGP next hops. To contain the accurate BGP next-hop address, we use the IPv4 BGP Nexthop Address (IE 18) field in the IPv4 and MPLS-IPv4 templates and the IPv6 BGP Nexthop Address (IE 63) field in the IPv6 and MPLS-IPv6 templates, for both the IPFIX and the version 9 formats.

To configure this feature, include the `nexthop-learning enable` statement at the `[edit services flow-monitoring (version-ipfix | version9) template]` hierarchy level. When you enable this feature, the `fragmentIdentification` (IE 54) field reports a value of 0.

[See [Understand Inline Active Flow Monitoring](#) and [nexthop-learning](#).]

- **RFC 8402 SR support for TWAMP probes (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**—We have added support in Two-Way Active Measurement Protocol (TWAMP) for segment routing (SR) as defined in RFC 8402, which broadly specifies the SR architecture. We support two types of SR for TWAMP probes:

- **SR-MPLS:** Uses a list of labels, each representing a segment end node.
- **SRv6:** Uses a type 4 IPv6 routing header introduced in RFC 8754, with each segment end node represented as an IPv6 address or IPv6 segment identifier (SID).

You can specify the list of SR-MPLS or SRv6 segments for a TWAMP probe to reach the reflector. In addition, you can specify the same information for the return path from the reflector to the client. This return path information is embedded in the probe itself by using the extensions proposed in *Simple TWAMP (STAMP) Extensions for Segment Routing Networks*, draft-ietf-ippm-stamp-srpm, namely the return path TLV and its return segment list sub-TLVs, as appropriate depending on the segment routing type. The TWAMP probes are timestamped in either the Routing Engine or the Packet Forwarding Engine.

To configure this feature, include the `source-routing` statement at the `[edit services monitoring twamp client control-connection name test-session session-name]` hierarchy level.

[See [Understand Two-Way Active Measurement Protocol](#) and [source-routing](#).]

Software Installation and Upgrade

- **Support for SZTP (PTX10002-36QDD)**—Use RFC-8572-based secure zero-touch provisioning (SZTP) to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before initiating ZTP.

To enable mutual authentication, the system generates a unique digital voucher based on the Digital Device ID or Cryptographic Digital Identity (DevID) of the network device. The DevID is embedded inside Trusted Platform Module (TPM) 2.0 chip on the network device. We issue a digital voucher to customers for each eligible network device.

[See [Secure Zero Touch Provisioning](#) and [Generate Secure ZTP Vouchers](#).]

- **Switching between SZTP and ZTP on secure platforms (PTX10002-36QDD)**—You can switch between using secure zero-touch provisioning (SZTP) and zero-touch provisioning (ZTP) on secure platforms. To override the default behavior of your secure device, you can issue the `request system zeroize ztp-option secure disable` command. When you issue this command, the CLI checks to see if the default platform behavior is secure. If the default platform is secure, the device will run ZTP after you reboot. If the default platform is not secure, the process ends. When you issue the `request system zeroize ztp-option secure enable` command, the CLI checks to see if the platform behavior is secure. If the default platform is secure, the process ends. If the platform isn't secure, you will receive an error that says the platform is not secure and cannot switch to SZTP. The process ends.

[See [Switching between Secure Zero Touch Provisioning and Zero Touch Provisioning](#).]

- **ZTP on WAN Interfaces (PTX10002-36QDD)**—Zero-touch provisioning (ZTP) dynamically detects the port speed of WAN interfaces and uses this information to create ZTP client ports with the same speed. ZTP automatically cycles through the WAN ports until it receives DHCP options from the DHCP server. The device uses the DHCP options to perform the bootstrap process.

[See [Zero Touch Provisioning](#).]

- **Static configuration of MAC-IP bindings (ACX7100-32C, ACX7100-48L, PTX10001-36MR, and PTX10008)**—You can configure MAC-IP bindings on interfaces to improve network management and host communication. This setup is similar to configuring static MAC addresses on an interface. Use this feature to streamline operations in static environments, such as Internet Exchange Points (IXPs), where Customer Edge (CE) routers remain fixed.

[See [Static Configuration of MAC-IP Bindings](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Multi-instance OSPF with SR (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)**—Configure and run multiple independent interior gateway protocol (IGP) instances of OSPFv2 with segment routing (SR) on a router. Use this feature to create two or more OSPF

instances and apply SR-MPLS on each instance. Multiple instances of OSPF can advertise different prefix-segment IDs (SIDs), and other instances can use these SIDs for routing decisions.

Multi-instance OSPF combined with SR enhances network flexibility, scalability, and control over traffic engineering, especially in large and complex networks.



NOTE: Junos OS does not support the configuration of the same logical interface in multiple IGP instances of OSPFv2.

[See [Multiple Independent IGP Instances of OSPFv2 Overview](#) and [Example: Configure Multiple Independent Instances of OSPFv2 with Segment Routing](#).]

Additional Features

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

- **Additional family in port mirroring** (PTX10002-36QDD). You can configure family any (as well as the earlier family options, inet and inet6) for local port mirroring and remote port mirroring. You use family any for family any, ccc, ethernet-switching, or mpls.



NOTE: You use the family any configuration option to process all four families.

You no longer configure port mirroring by using the [edit forwarding-options port-mirroring analyzer] hierarchy on the PTX devices. You now use [edit forwarding-options port-mirroring] for local port mirroring or [edit forwarding-options port-mirroring instance *instance-name*] for remote port mirroring, with both of those configurations also requiring a firewall filter.

The following configuration statements are no longer part of the port mirroring configuration on PTX:

- next-hop for family any
- family vpls
- no-filter-check
- hosted-service
- server-profile

[See [Example: Configure Port Mirroring with Family any and a Firewall Filter](#) and [port-mirroring](#).]

- **Avoid microloops in IS-IS SRv6 networks** (PTX10002-36QDD). You can enable post-convergence path calculation on a device to avoid microloops if a link or metric changes in an SRv6 network.

[See [How to Configure Microloop Avoidance for IS-IS in SRv6 Networks.](#)]

- **BGP autodiscovery underlay in EVPN-VXLAN** (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)

[See [BGP Auto-Discovered Neighbors.](#)]

- **Connectivity fault management and enhanced CFM** (PTX10002-36QDD). Support includes:
 - Up and down maintenance association end points (MEPs) on bridges, circuit cross-connect (CCC), and Ethernet VPN (EVPN) in inline mode
 - ITU-T Y.1731 on synthetic loss measurement (SLM) and delay measurement (DM)
 - Inline session support for connectivity fault management (CFM) on aggregated Ethernet
 - Enhanced CFM mode by default
 - Supported inline performance monitoring (PM) sessions include:
 - PM Tx
 - PM Rx
 - PM responder
 - IPv4 (inet) and IPv6 (inet6) support for continuity check messages (CCM), delay measurement (DM), and synthetic loss message (SLM)
 - DM on aggregated Ethernet with at least one child link on the anchor Packet Forwarding Engine
 - Action profile for marking a link down, except for EVPN and bridge up MEP
 - Supported CFM protocol data units (PDUs) for inline handling include:
 - CCM
 - Delay measurement message (DMM)
 - Delay measurement reply (DMR)
 - Synthetic loss message (SLM)
 - Synthetic loss reply (SLR)
 - Enterprise and service provider configurations
 - VLAN normalization

- VLAN transparency for CFM PDUs
- CoS forwarding class and CoS packet loss priority (PLP) for CFM
- Combination of up MEP, down MEP, or maintenance association intermediate point (MIP) configuration over the same interface

[See [Connectivity Fault Management \(CFM\)](#).]

- **CoS interface telemetry support** (PTX10002-36QDD). 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 \(Junos Telemetry Interface\)](#).]

- **EVPN-VXLAN fabric with an IPv6 underlay** (PTX10002-36QDD). Support includes:
 - Quality of service (QoS) and class of service (CoS) classification with explicit congestion notification (ECN) copy upon VXLAN tunnel encapsulation and de-encapsulation.

Priority-based flow control (PFC), differentiated services code point (DSCP) copy, and IEEE 802.1p rewrite are not supported.

 - Dynamic Host Configuration Protocol (DHCP) relay with DHCPv4 and DHCPv6.

[See [EVPN-VXLAN with an IPv6 Underlay](#) and [Example: Configure an IPv6 Underlay for Layer 2 VXLAN Gateway Leaf Devices](#).]

- **EVPN-VXLAN L2 gateways and L3 gateways with EVPN Type 5 routes** (PTX10002-36QDD).

Support includes:

- Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) Layer 2 (L2) and Layer 3 (L3) gateway operations in edge-routed bridging (ERB) overlay and centrally routed bridging (CRB) overlay fabrics
- EVPN instances using the MAC-VRF instance type with VLAN-based, VLAN-bundle, or VLAN-aware bundle service types
- Pure EVPN Type 5 (IP prefix) route virtual routing and forwarding (VRF) model
- Integrated routing and bridging (IRB) for IPv4 and IPv6 data traffic
- Q-in-Q dual tagging for VXLAN network identifier (VNI) mapping with service provider-style logical interface configurations only
- Overlapping VLAN IDs across MAC-VRF instances
- Underlay reachability over ECMP

- Active/active multihoming with Ethernet segment identifiers (ESIs) per physical interface
- Proxy Address Resolution Protocol (ARP) and proxy Network Discovery Protocol (NDP), and ARP or NDP suppression
- Multicast IRB support without IGMP snooping or MLD snooping
- IEEE 802.1p and Differentiated Services code point (DSCP) class of service (CoS) on EVPN-VXLAN tunnel interfaces, with both service provider-style or enterprise-style interface configurations (including classification and rewrite operations, but not DSCP copy support)

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

- **Firewall filtering using flood policer, IRB, and service provider egress filtering** (PTX10002-36QDD). You can use the flood policer feature to control flooding of the network with broadcast, unknown unicast, and multicast (BUM) traffic, and this control includes the EVPN flood policer.



NOTE: EVPN-MPLS configurations also support flood policers.

- **IGMP snooping and MLD snooping** (PTX10002-36QDD)

[See [IGMP Snooping Overview](#) and [Understanding MLD Snooping](#).]

- **Inline active flow monitoring for IPv4 and IPv6 traffic using IP Flow Information Export (IPFIX) and version 9 templates** (PTX10002-36QDD)

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

- **Inline active flow monitoring using IP Flow Information Export (IPFIX) and version 9 templates for IRB interfaces and BGP next-hop addresses** (PTX10002-36QDD). We now support:
 - IPv4 and IPv6 traffic on IRB interfaces.
 - A BGP next-hop address in the IPv6 and MPLS-IPv6 templates. Information Element 63, IPv6 BGP NextHop Address, is now available.

[See [Inline Active Flow Monitoring on IRB Interfaces](#) and [Understand Inline Active Flow Monitoring](#).]

- **IP-IP tunnel stitching** (PTX10002-36QDD).

[See [Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation](#) and [Example: Configuring Next-Hop-Based IP-Over-IP Dynamic Tunnels](#).]

- **Layer 2 and Layer 3 support for flood policers** (PTX10002-36QDD). You can configure firewall filters for flood policers on Layer 2 (family ccc) and Layer 3 (family any) traffic, in both the ingress and egress directions. Most match conditions (except packet-length) and most actions are supported.

- **Next-hop-based dynamic tunnels with IPv6 in the underlay network** (PTX10002-36QDD). You can encapsulate IPv4 and IPv6 packets inside the IPv6 packets between two IPv6 nodes. This encapsulation mechanism helps to create next-hop-based dynamic tunnels with IPv6 in the underlay network.

[See [Next-Hop-Based Dynamic Tunnels](#) and [show dynamic-tunnels database](#).]

- **OAM ping support for SRv6 network programming** (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016). You can perform an Operations, Administration and Management (OAM) ping operation for any Segment Routing with IPv6 (SRv6) segment identifier (SID) whose behavior allows upper layer header processing for an applicable OAM payload.

[See [ITU-T Y.1731 Ethernet Service OAM Overview How to Enable SRv6 Network Programming in IS-IS Networks](#).]

- **On-box aggregation support** (PTX10002-36QDD).

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

- **OpenConfig QoS operational state sensors** (PTX10002-36QDD).

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

- **OpenConfig QoS queue management profile and ECN configuration** (PTX10002-36QDD).

[See [Mapping OpenConfig QoS Commands to Junos Configuration](#).]

- **Optimized intersubnet multicast (OISM) for IPv4 multicast traffic in EVPN-VXLAN fabrics** (PTX10002-36QDD). 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 Layer 3 (L3) interfaces
 - EVPN multicast VLAN (M-VLAN) integrated routing and bridging (IRB) interfaces
 - Non-EVPN IRB interfaces

[See [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **QoS configuration and streaming with OpenConfig** (PTX10002-36QDD).

[See [Mapping OpenConfig QoS Commands to Junos Configuration](#) and [Junos YANG Data Model Explorer](#).]

- **Static VXLAN L2 gateway overlays** (PTX10002-36QDD)

[See [Static VXLAN](#), [remote-vtep-list](#), and [static-remote-vtep-list](#).]

- **QSFP-100G coherent ZR optics performance monitoring** (ACX7024, ACX7348, and PTX10001-36MR; and the PTX10004, PTX10008, and PTX10016 with the PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards installed). Monitor the performance of QSFP-100G coherent ZR optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Accumulate performance metrics into 15-minute and 1-day interval bins. Use the `show interfaces transport pm` command to view current and historical performance data.

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

- **Redistribution of IPv4 routes with IPv6 next hop into BGP** (PTX10002-36QDD). Devices can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, IPv4 traffic is tunneled from CPE devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers, which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop.

[See [Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP](#).]

- **SRv6 flexible algorithms in TED and BGP-LS** (PTX10002-36QDD). The router supports Segment Routing for IPv6 (SRv6) flexible algorithms in the traffic engineering database (TED) and in BGP Link State (BGP-LS).

[See [Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering](#) and [BGP Link-State Extensions for Source Packet Routing in Networking \(SPRING\)](#).]

- **Static route tracking using the results of RPM and TWAMP tests** (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, ACX7024, ACX7024X, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016). We've extended support for static route tracking to Junos OS Evolved and included Two-Way Active Measurement Protocol (TWAMP) test support as well. You use RPM or TWAMP probes to detect link status and to change the preferred-route state on the basis of the probe results. Tracked static routes can be IPv4 or IPv6, and each IPv4 and IPv6 tracked static route supports up to 16 next hops. You can also configure the metric, route preference, and tag values for each IPv4 or IPv6 destination prefix. However, you configure this feature differently on Junos OS Evolved devices; you configure the `sla-tracking` statement at the `[edit routing-options]` hierarchy level. For Junos OS, you would configure the `rpm-tracking` statement at the same hierarchy level.

[See [Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#), [Understand Two-Way Active Measurement Protocol, sla-tracking](#), and [show route sla-tracking](#).]

- **Support for automatic ingress LSP Policing in P2MP LSPs (PTX10002-36QDD).**

[See [Configuring Automatic Policers](#).]

- **Support for basic Layer 2 features (PTX10002-36QDD).** The PTX10002-36QDD 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)
 - Handle BUM (broadcast, unknown unicast and multicast) 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
 - Native VLAN ID for Layer 2 logical interfaces
 - Single VLAN-tagged Layer 2 logical interfaces
 - Interface statistics



NOTE: The `show ethernet-switching statistics` command and child logical interface statistics for aggregated Ethernet are not supported.

- 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 flexible algorithms for SRv6 in IS-IS** (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016). Configuration of segment routing in a core IPv6 network without an MPLS data plane supports these flexible algorithm (flex algo) features in IS-IS networks:
 - Advertisement of a Segment Routing Extension Header (SRH) locator with a mapped flexible algorithm
 - Use of application-specific link attributes (ASLA) for flexible algorithms
 - Configuration of a TI-LFA backup path for SRv6
 - Use of application-specific link attributes (ASLA) for flexible algorithms
 - Compression of SRv6 addresses into a single IPv6 address (micro-SID) in flex algo path computations

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

- **Support for G.8275.1 profile, PTP over Ethernet encapsulation, and hybrid mode over LAG with PTP over Ethernet** (PTX10002-36QDD).

[See [G.8275.1 Telecom Profile](#), [Guidelines for Configuring PTP over Ethernet](#), and [Hybrid Mode](#).]

- **Support for the gRPC Network Operations Interface (gNOI) certificate management cert service** (PTX10002-36QDD). You can execute supported cert service remote procedure calls (RPCs) to manage certificates on the network device. Using gNOI operations enables you to use the same suite of microservices to efficiently manage large-scale multivendor networks.

[See [gNOI Certificate Management \(Cert\) Service](#).]

- **Support for IRB** (PTX10002-36QDD). You can use integrated routing and bridging (IRB) to route Layer 3 traffic between a bridge domain and another routed interface. The PTX10002-36QDD router supports the following IRB features:
 - 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
 - Address family support for IPv4 and IPv6, and support for IPv4 maximum transmission units (MTUs) and IPv6 MTUs with different MTU values
 - IRB interface in virtual routing and forwarding (VRF) routing instances
 - Directed subnet broadcast support with IRB
 - Support for VRRP on IRB

[See [Integrated Routing and Bridging](#), [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port](#), and [Understanding VRRP](#).]

- **Support for LLDP, xSTP, and BPDU protection.** (PTX10002-36QDD) Support includes:
 - Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), Multiple Spanning Tree Protocol (MSTP), Spanning Tree Protocol (STP), root protection for STP, concurrent configuration of RSTP and VSTP, virtual switch, and BPDU protection for spanning-tree protocols
 - Bridge protocol data unit (BPDU) protection for EVPN-VXLAN
 - LLDP support, including:
 - LLDP on em0 interfaces
 - Disabling of LLDP time, length, and value (TLV) messages

[See [Configuring STP](#), [Understanding BPDU Protection for EVPN-VXLAN](#), and [Device Discovery Using LLDP](#).]

- **Support for micro-SIDs in TI-LFA, microloop avoidance, flex algo, and IS-IS MT** (PTX10002-36QDD). We extend the support of compressing SRv6 addresses into a single IPv6 address (micro-SID) in Topology Independent Loop-Free Alternate (TI-LFA), microloop avoidance, and Flexible Algorithm (flex algo) path computations. We also support IPv6 unicast topology (part of IS-IS MT) in TI-LFA, microloop avoidance, and flex algo computations.

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

- **Support for Q-in-Q tunneling** (PTX10002-36QDD).

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

- **Support for SRv6 LSPs in PCEP** (ACX7348 and PTX10002-36QDD). The Path Computation Element Protocol (PCEP) supports all types of SRv6 LSPs, such as PCE-initiated, locally created, and delegated SRv6 LSPs.

[See [SRv6 LSP in PCEP](#).]

- **Support for SRv6 traceroute** (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016). We support the traceroute mechanism for Segment Routing for IPv6 (SRv6) segment identifiers.

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

- **Support for EVPN-VPWS**(PTX10002-36QDD)

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

- **Support for transit traffic rates in bits per second (bps) and packets per second (pps) for both IPv4 and IPv6 at a logical interface level** (PTX10002-36QDD). Support includes:

- Support for classification override configured under a forwarding policy
- Support for VRRP
- Unicast RPF support for both IPv4 and IPv6 traffic flows
- Loose-mode unicast RPF on IPv4 and IPv6
- Support for destination class usage (DCU) and source class usage (SCU) accounting
- Class-based firewall filters
- Forwarding IPv6 transit statistics
- Operational state statistics for IPv6 logical interfaces

[See [Configuring IPv4 and IPv6 Accounting](#), [CoS Features and Limitations on PTX Series Routers](#), [Overriding the Input Classification](#), [Understanding VRRP](#), [Configuring Unicast RPF Loose Mode](#), [Understanding Source Class Usage and Destination Class Usage Options](#), [Configure the Filter Profile](#), [BGP User Guide](#), and [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

- **Support for VPLS** (PTX10001-36MR, PTX10002-36QDD, PTX1000, PTX10008, and PTX10016). — You can configure VPLS on the PTX10000 line of routers running Junos OS Evolved.

- To configure VPLS, configure the instance-type virtual-switch statement at the [edit routing-instances *routing-instance-name*] hierarchy level.

- In this release, we support single bridge domains. You must configure service-type single statement at the [edit routing-instances *routing-instance-name* vpls] hierarchy level.
- You must enable control-word at the [edit routing-instances *routing-instance-name* protocols vpls] hierarchy level.
- Encapsulation of ethernet-vpls and vlan-vpls is not supported on CE interfaces.
- To display VPLS MAC address information, use the show ethernet-switching table command.

[See [Introduction to Configuring VPLS](#)

- **Supported transceivers, optical interfaces, and DAC cables (PTX10004, PTX10008, and PTX10016)**—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.
- **Support for tunnel decapsulation using firewall filters for GRE and UDP tunnels** (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)

[See [Configuring a Filter to De-Encapsulate GRE Traffic](#) and [decapsulate \(Firewall Filter\)](#).]

- **Support for filter-based forwarding to a specific outgoing interface or destination IP address** (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016). Support for next-interface, next-ip and next-ip6.

[See [Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address](#).]

- **Support for per-packet random spray load balancing** (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)

[See [Configuring Per-Packet Load Balancing](#).]

- **Support for flexible firewall filter match conditions** (QFX5220, QFX5230, QFX5240, and PTX10002-36QDD)

[See [Firewall Filter Flexible Match Conditions](#).]

- **Fast lookup filter** (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)—Support extended to ethernet switching, any, mpls, and ccc firewall filter families for fast lookup filters on PTX Series routers.

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

- **Support for VPLS** (PTX10002-36QDD). We support the following VPLS features:

- Multihoming

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

- Hierarchical VPLS (H-VPLS)

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

- Flow Aware Transport (FAT) labels.

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

- Entropy labels

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

- Flood policer

[See [Configuring Firewall Filters and Policers for VPLS.](#)]

- Mac limits and actions

[See [mac-move-limit.](#)]

- CFM support for VPLS topology

[See [Configure a MEP to Generate and Respond to CFM Protocol Messages.](#)]

- **Exclude hops in the RSVP LSP path (ACX7332, ACX7509, PTX10002-36QDD, PTX10008)**—You can configure a list of hops to be excluded in the label-switched path (LSP) so that RSVP LSPs avoid those hops and links in the traffic engineering (TE) domain. When an RSVP LSP is signaled in the network, the path message carries the excluded list of hops. When the downstream routers perform loose hop expansion, such as inter-domain LSP or abstract node expansion, the transit routers use the same excluded list of hops that the ingress router uses for path computation. This mechanism enables intermediate routers to avoid the routers included in the excluded hop list. The routers try alternative paths to help with the convergence of LSPs when a complete end-to-end path computation is not possible.

Additionally, ingress routers receive PathErr messages and when computing another path, the routers use a PathErr message sender's address to avoid the link or node that generates an error. Transit routers also need this error avoidance information during retry attempts. RFC4814 defines the exclude hop information and is accepted in RSVP signaling.

To configure LSPs to exclude a list of hops, include the exclude statement at the [edit protocols mpls path path-name next-hop] hierarchy level. The ingress routers exclude the hops in CSPF computation and are also included in RSVP LSP signaling.

- **Support for VNI based match for EVPN-VXLAN (PTX10002-36QDD)**

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

- **Support for output filter-based GRE (PTX10002-36QDD)**—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 packet is forwarded accordingly. The IP lookup for GRE-encap capable route is limited to the implicit default routing-instance.

[See [Understanding Filter-Based Tunneling Across IPv4 Networks.](#)]

- **Support for configuring output filter action with non-default routing instance or a specified routing instance (PTX10002-36QDD)**

[See [Firewall Filter Terminating Actions.](#)]

- **Support for filter-based forwarding (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**

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

- **Firewall filter support for bitwise logical operations for TCP Flag match (PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)**

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

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 103](#)
- [Class of Service \(CoS\) | 103](#)
- [EVPN | 103](#)
- [General Routing | 103](#)
- [Interfaces and Chassis | 104](#)
- [Junos Telemetry Interface | 104](#)
- [Junos XML API and Scripting | 104](#)
- [Multicast | 105](#)
- [Network Management and Monitoring | 105](#)
- [PTP \(Precision Time Protocol\) | 106](#)
- [Routing Policies and Firewall Filters | 106](#)
- [Routing Protocols | 106](#)
- [User Interface and Configuration | 106](#)

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

Authentication and Access Control

- Disabled CDN auto download (Junos OS Evolved)— The PKI process periodically, by default every 24 hours, polls the CDN server for the latest default trusted CA bundle and updates the list for any changes to the trusted CAs in the bundle. If there are any changes, PKI process loads them in the background. The auto download of CA certificates might generate core files. We've disabled the service of PKI query to CDN server periodically to download the latest trusted CA bundle.
- On Junos OS Evolved, password authentication for SCP based configuration archival is supported.

Class of Service (CoS)

- Previously, the Junos OS Evolved system default scheduler was named "default" (no brackets), while the Junos OS system default scheduler is named "<default>" (with brackets). Now, the Junos OS Evolved system default scheduler is also named "<default>" (with brackets).

EVPN

- **EVPN system log messages for CCC interface up and down events**—Devices will now log EVPN and EVPN-VPWS interface up and down event messages for interfaces configured with circuit cross-connect (CCC) encapsulation types. You can look for error messages with message types EVPN_INTF_CCC_DOWN and EVPN_INTF_CCC_UP in the device system log file (/var/log/syslog).

General Routing

- The system now checks the port number value (z) in the 'set interfaces et-x/y/z:n' configuration for a valid port range on PTX10002-36QDD. Previously, configurations with invalid port numbers were committed successfully. With this update, the system displays a UI error message and prevents committing configurations with invalid port numbers, ensuring configuration accuracy and preventing potential issues.
- Change to the commit process—In prior Junos OS Evolved releases, if you use the commit prepare command and modify the configuration before activating the configuration using the commit activate

command, the prepared commit cache becomes invalid due to the interim configuration change. As a result, you cannot perform a regular commit operation using the commit command. The CLI shows an error message: 'error: Commit activation is pending, either activate or clear commit prepare'. If you now try running the commit activate command, the CLI shows an error message: 'error: Prepared commit cache invalid, failed to activate'. You then must clear the prepared configuration using the clear system commit prepared command before performing a regular commit operation. From this Junos and Junos OS Evolved release, when you modify a device configuration after 'commit prepare' and then issue a 'commit', the OS detects that the prepared cache is invalid and automatically clears the prepared cache before proceeding with regular 'commit' operation.

[See [Commit Preparation and Activation Overview](#).]

Interfaces and Chassis

- **Disable power redundancy alarms for JNP10K-PWR-DC2 PSM (PTX10008 and PTX10016)**- The JNP10K-PWR-DC2 PSM supports power redundancy across two DIP switches. When all input feeds are not connected to power supplies, it triggers a chassis alarm such as PSM 5 Input B0 and B1 Failed. Starting in Junos OS Evolved Release 24.2R1, you can disable this chassis alarm by using the set chassis alarm psm *psm number* input *input number* ignore command.

[See [JNP10K-PWR-DC2 Power Supply](#).]

- **DDoS protection protocols statistics update (PTX Series)**—Starting in Junos OS Evolved Release 23.2R2, the show ddos-protection protocols statistics displays the Max arrival rate and Arrival rate output values as expected. Earlier to this release, the Max arrival rate and Arrival rate output values were displayed larger than expected.

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

Junos Telemetry Interface

- The show agent sensors command output for gRPC sensors is truncated on the Junos OS Evolved platform to align with the output format of the Junos OS platform.

Junos XML API and Scripting

- **Commit script input to identify software upgrades during boot time (ACX Series, PTX Series, and QFX Series)**—The junos-context node-set includes the sw-upgrade-in-progress tag. Commit scripts can test

the `sw-upgrade-in-progress` tag value to determine if the commit is taking place during boot time and a software upgrade is in progress. The tag value is `yes` if the commit takes place during the first reboot after a software upgrade, software downgrade, or rollback. The tag value is `no` if the device is booting normally.

[See [Global Parameters and Variables in Junos OS Automation Scripts](#).]

Multicast

- **Non-revertive switchover for sender based MoFRR**— In earlier Junos releases, source-based MoFRR ensured that the traffic reverted to the primary path from the backup path, when the primary path or session was restored. This reversion could result in traffic loss. Starting in Junos OS 22.4R3-S1, source-based MoFRR will not revert to the primary path, i.e. traffic will continue to flow through the backup path as long as the traffic flow rate on the backup path does not go below the configured threshold set under protocols `mvpn hot-root-standby min-rate`.

[See [min-rate](#).]

Network Management and Monitoring

- In a firewall filter configured with a `port-mirror-instance` or `port-mirror` action, if `l2-mirror` action is also configured, then `port-mirroring` instance family should be any. In the absence of the `l2-mirror` action, `port-mirroring` instance family should be the firewall filter family.
- **Python 2 interpreter option deprecated for Juniper Extension Toolkit (JET) applications** (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX10K-LC1202-36MR (line cards for PTX10016, PTX10008 and PTX10004), QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-64OD, QFX5240-QD, QFX5700, and QFX5700E)—Python 2.7 is already not supported on Junos OS Evolved devices as of an earlier release. The `python` statement at the `edit system extensions extension-service application file <filename>` hierarchy level was used to interpret JET applications written in Python 2. This statement is now deprecated. To run daemonized on-device JET applications written in Python 3, use the `python3` statement.

[See [file \(JET\)](#).]

PTP (Precision Time Protocol)

- **Maximum limit of PTP local masters (PTX10008)**— You can configure up to 512 PTP masters at the `edit protocols ptp master interface interface-name multicast-mode hierarchy level` on PTX10008 series routers. Earlier the system was rejecting the commit while trying to configure more than 128 PTP masters.

Routing Policies and Firewall Filters

- Support added for source and destination port optimization for port ranges for ipv6 input firewall filters.

Routing Protocols

- **MLD snooping proxy and I2-querier source-address (ACX7024, ACX7100-32C, PTX10001-36MR, QFX5120-32C, and QFX5130-32CD)**— The source-address configured for proxy and I2-querier under the `mld-snooping` hierarchy should be an IPv6 link-local address in the range of `fe80::/64`. The CLI help text has been updated to "Source IPv6 link local address to use for proxy/L2 querier". In earlier releases, the CLI help text read, "Source IP address to use for proxy/L2 querier."

[See [source-address](#).]

User Interface and Configuration

- **Compact format deprecated for JSON-formatted state data (ACX Series, PTX Series, and QFX Series)**
—We've removed the `compact` option at the `[edit system export-format state-data json]` hierarchy level because Junos devices no longer support emitting JSON-formatted state data in compact format.
- **Access privileges for request support information command (ACX Series, PTX Series, and QFX Series)**
—The `request support information` command is designed to generate system information for troubleshooting and debugging purposes. Users with the specific access privileges `maintenance`, `view`, and `view-configuration` can execute `request support information` command.

Known Limitations

IN THIS SECTION

- [General Routing](#) | 107

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

- It occurs only in case of hardware issues where SCL is shorted to the GND. In 8 8-slot chassis, the SCL clock of I2C shorted on the PSM0, affects the I2C lines of PSM2,3,4. Similarly, PSM4 affects PSM5. This is an existing issue on the RCB I2C master. All slaves on the same master get affected in case the I2C clock is shorted to GND. Issue seen in the field should be RMA of the PSM0 if the short is in PSM0, else chassis which could be short in the short in I2C Clock in chassis/midplane. In case the issue is seen in the field where all the 4 PSM failures, Investigation should be done for the PSM0 failure which can cause the other PSM to fail. Then investigate on the other PSM or midplane/.chassis. [PR1807924](#)
- When a SIB/FPC/PFE is offlined gracefully or ungracefully while traffic is flowing, it can result in bf_sm_sch_dat_intr_oresource_drop CM major error on the SIBs which are online. Similarly the CM major error can be observed during SIB/FPC/PFE online operation while traffic is flowing. In the output of "show system errors inactive detail" this alarm can be observed. This alarm is a transient alarm observed only during offline/online operation. [PR1835365](#)

Open Issues

IN THIS SECTION

- [General Routing](#) | 108

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

- An improper handling of syntactically invalid structure vulnerability in Object Flooding Protocol (OFP) service of Juniper Networks Junos OS Evolved allows an unauthenticated, network-based attacker to cause a Denial of Service (DoS). Please refer to <https://supportportal.juniper.net/JSA75753> for more information. [PR1714333](#)
- Many issues are observed in Junos OS Evolved library regarding the DNS resolution and these were fixed in this PR. [PR1733616](#)
- When DHCP trace options are enabled, there is a possibility that jdhcpd could generate a core file. In general, traceoptions should be enabled only for debugging. They should be disabled once debuffing is done. [PR1771121](#)
- On PTX10001-36MR and PTX10004/PTX10008/PTX10016 Junos OS Evolved BT Packet Forwarding Engine based platforms, same priority and weights are given to strict-high queue and other non-strict high queues for same transmit-rate set in scheduler CLI configuration which leads to no preference to strict-high queue. This causes tail-drops packets count to increment on strict-high queue. There is no impact on system due to this issue. [PR1773709](#)
- On Junos OS Evolved PTX platforms in the EVPN-VxLAN (Ethernet VPN-Virtual Extensible LAN) DCI (Data Center Interconnect) multihoming scenario, due to source MAC update/modification destination MAC route might be deleted that results in traffic drop. The issue occurs when the MAC is learnt over ESI-LAG (Ethernet Segment Identifier - Link Aggregation Group). [PR1783935](#)
- The support to export key leaves was done through PR 1818150 recently and it is done only in the latest releases. The support to export key leaves is unavailable uniformly across pre-GNMI/ GNMI and UDP on the 23.4 release. [PR1831799](#)
- This is a transient log which sometimes is seen when LSI is recreated. This has no functional impact. It's generated because of additional dependency of LSI with BD. System takes care of cleaning the token for which this error is generated. This can be confirmed via VTY command show sandbox token. [PR1834443](#)
- On Junos Evolved PTX10003 platform the command indirect-next-hop-change-acknowledgements is required in the junos-default configuration. If the said command is missing, it will result in packet loss. [PR1836337](#)
- When AH is in progress and GRES happens, there is a possibility that AH might fail. In case of AH Failure, following recovery mechanisms can be tried: Try executing the AH through CLI to see if it

helps. If CLI AH command does not recover the link, PFE offline/online can be tried to recover the link. If PFE offline/online also does not recover the link, then FPC offline/online can be tried.

[PR1843391](#)

- The **RTI Await LSI Cookie** in longevity test during routing instance is added and deleted. [PR1843627](#)

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

- [What's New | 109](#)
- [What's Changed | 157](#)
- [Known Limitations | 161](#)
- [Open Issues | 163](#)

These release notes accompany Junos OS Evolved Release 24.4R1 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

- [Hardware | 110](#)
- [Authentication and Access Control | 144](#)
- [Additional Features Optimized for AI-ML Fabrics | 145](#)
- [Chassis | 147](#)
- [Class of Service | 148](#)
- [Forwarding Options | 148](#)

- [High Availability | 148](#)
- [Interfaces | 149](#)
- [Junos Telemetry Interface | 149](#)
- [Multicast | 152](#)
- [Network Management and Monitoring | 152](#)
- [Platform and Infrastructure | 153](#)
- [Precision Time Protocol \(PTP\) | 153](#)
- [Routing Policy and layer2-policer Firewall Filters | 153](#)
- [Routing Protocols | 154](#)
- [Software Installation and Upgrade | 154](#)
- [Additional Features | 155](#)

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

Hardware

- **QFX5130-48C and QFX5130-48CM switches (QFX Series)**—The Juniper Networks® QFX5130-48C is our first 1-U fixed form-factor switch that is completely optimized for 100GbE server connections. The QFX5130-48C switch offers high-density 100GbE access ports in an SFP-DD form factor optimized for servers. The switch also has high-density 400GbE ports in a QSFP-56 form factor optimized for easy uplinks to data centers. The QFX5130-48C provides a throughput of 8 Tbps by means of:
 - Forty-eight high-density 100GbE access ports that support SFP-DD transceivers optimized for servers.
 - Eight high-density 400GbE ports that support QSFP56 transceivers optimized for easy uplinks to the spine layer in data centers.

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

QFX5130-48CM is our first 1-U fixed form-factor switch that is completely optimized for 100GbE server connections. The Juniper Networks® QFX5130-48CM switch offers high-density 100GbE access ports in an SFP-DD form factor optimized for servers, along with high-density 400GbE ports in a QSFP-56 form factor optimized for easy uplinks to data centers. The QFX5130-48CM switch provides a throughput of 8 terabit per second (Tbps) by means of:

- Forty-eight high-density 100GbE access ports that support SFP-DD transceivers optimized for servers.
- Eight high-density 400GbE ports that support QSFP56 transceivers optimized for easy uplinks to the spine layer in data centers.
- Support for Media Access Control Security (MACsec) feature.

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

To install the QFX5130-48C and QFX5130-48CM switch and perform initial configuration, routine maintenance, and troubleshooting, see the [QFX5130 System Overview](#). See [Feature Explorer](#) for the complete list of features for any platform.

Table 2: QFX5130-48C and QFX5130-48CM Feature Support

Feature	Description
Chassis	<ul style="list-style-type: none"> • QFX5130-48C is a high-density 100GbE Ethernet/IP switch based on a single chip that offers 8-Tbps forwarding capacity. • Support for the next-generation, high-density, and cost-efficient 100GbE and 400GbE optimized fixed system <p>The QFX5130-48CM switch features:</p> <ul style="list-style-type: none"> • Forty-eight SFP56-DD 100GbE ports for server connectivity • Eight QSFP-DD 400GbE uplink ports • Up to 16-Tbps (bidirectional)/2.7-bpps throughput • Using breakout cables, you can increase the total number of supported 100/25/10GbE ports per switch to 72. • To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported , see the Hardware Compatibility Tool.

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Class of service	<ul style="list-style-type: none"> • CoS support on EVPN VXLAN networks. [See CoS Support on EVPN VXLANs.] • Support for priority-based flow control (PFC) of untagged traffic at Layer 3 using Differentiated Services Code Points (DSCP). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.]
Ethernet switching and bridging	<ul style="list-style-type: none"> • Support for Q-in-Q tunneling with a service-provider-style configuration. [See Configuring Q-in-Q Tunneling.] • LLDP support. [See Device Discovery Using LLDP.] • Support for MAC move limit with EVPN-VXLAN. [See Understanding MAC Move Limiting.]
Forwarding options	<ul style="list-style-type: none"> • Support for port mirroring in EVPN-VXLAN environments. [See How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics.]
High availability	<ul style="list-style-type: none"> • VRRP support on Packet Forwarding Engine. [See VRRP Overview.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Interfaces	<ul style="list-style-type: none"> Support for BGP flowspec. [See BGP.] Support for 48 SFP-DD and 8 QSFP-DD ports. Each switch also supports two 10GbE SFP+ ports. We support the following port configurations on each switch: <ul style="list-style-type: none"> 48x100/50/25/10GbE SFP-DD ports 8x400/200/100/40GbE QSFP-DD ports 2x10GbE SFP+ ports [See Port Settings (Interface Guide for Switches).] Support for MACsec on physical interfaces on QFX5120-48CM. This platform supports MACsec in dynamic connectivity association key (CAK) mode with GCM-AES-128, GCM-AES-256, GCM-AES-XPN-128, and GCM-AES-XPN-256 encryption. MACsec is supported on physical interfaces for switch-to-host and switch-to-switch links. [See Configuring MACsec.] Each switch has 48 SFP-DD and 8 QSFP-DD ports. QFX5130-48C/ 48CM also supports two SFP+ ports with 10 GbE. It supports the following port configurations: <ul style="list-style-type: none"> 48x100G / 50GbE / 25GbE / 10GbE on SFP-DD ports 8x400G / 200GbE / 100GbE / 40GbE on QSFP-DD ports 2x10GbE on SFP+ ports [See Port Speed on QFX5130-48C and QFX5130-48CM Switches.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Junos Telemetry Interface (JTI)	<ul style="list-style-type: none"> • JTI streaming support for hardware Routing Engine-based sensors. Subscribe to / components/sensor to stream hardware operational stages. Statistics include Routing Engine, power supply unit (PSU), Control Board (CB), FPC, and PIC states. <p>[See Junos YANG Data Model Explorer.]</p>
Multicast	<ul style="list-style-type: none"> • MLD snooping and IRB stitching support . <p>[See Understanding MLD Snooping.]</p> <ul style="list-style-type: none"> • Support for multicast forwarding. <p>[See Multicast Overview.]</p> <ul style="list-style-type: none"> • IGMP snooping support. <p>[See PIM Overview.]</p> <ul style="list-style-type: none"> • IGMP, MLD multicast snooping, and IRB elaboration with MBB. <p>[See IGMP Snooping Overview.]</p> <p>[See Understanding MLD Snooping.]</p>

Table 2: QFX5130-48C and QFX5130-48CM Feature Support (*Continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> • Support for sFlow. [See Overview of sFlow Technology.] • Support for analyzers and port mirroring. [See Understanding Port Mirroring and Analyzers.] • IPsec support for OSPFv2 and OSPFv3. [See Overview of IPsec.] [See Configuring OSPF Authentication.] [See Configuring IPsec Security Associations.] • DHCP stateless relay MIB support. [See Enterprise-Specific MIBs for Junos OS Evolved.]
Protection against DDoS attacks	<ul style="list-style-type: none"> • Supports DDoS protection, which is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.] and protocols (DDoS) (ACX Series, PTX Series, and QFX Series).]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Platform and infrastructure	<ul style="list-style-type: none">• Platform resiliency support for hardware components of each FRU. <p>If a failure is detected on a hardware component, Junos OS Evolved:</p> <ul style="list-style-type: none">• Logs the message to give clear indication of failure details, including time stamp, module name, component name, and failure details.• Raises or clears alarms, if applicable.• Performs local action, such as self-healing and taking the component out of service.

Table 2: QFX5130-48C and QFX5130-48CM Feature Support (*Continued*)

Feature	Description
Precision Time Protocol	<ul style="list-style-type: none"> • Transparent Clock support <p>Transparent clocks improve synchronization between the timeTransmitter and timeReceiver clocks and ensure that the timeTransmitter and timeReceiver clocks are not impacted by the effects of packet delay variation.</p> <ul style="list-style-type: none"> • With or without VLAN encapsulation • With PTP over IPv4 • With PTP unicast or multicast • On LAG and MC-LAG • On all physical, IRB, and aggregated Ethernet interfaces • IEEE 1588 PTP Ordinary Clock/Boundary Clock Enterprise and media profiles. Includes support for: <ul style="list-style-type: none"> • Enterprise profile using PTP ordinary clock and PTP boundaryclock applications • Enterprise profile using PTP over IPv4 multicast transport • SMPTE profile, AES67, and AES67+SMPTE combined profiles • Media profiles using PTP ordinaryclock and PTP boundary clock applications • Media profiles using PTP over IPv4 multicast transport <p>[See PTP Profiles.]</p>

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> Support for PTP over IRB. QFX5130-48C supports Precision Time Protocol (PTP) over integrated routing and bridging (IRB) interfaces. <p>PTP boundary clock (BC) applications often need multiple PTP streams to share a local IP address for broadcast media. These packets are forwarded through Layer2 (L2) switching. Unlike traditional PTP configurations on physical interfaces, no physical interface logical units (IFLs) are created for each PTP physical interface. You can achieve this configuration through integrated routing and bridging (IRB) interfaces.</p> <p>[See PTP over IRB for Broadcast Profiles.]</p>
Routing options	<ul style="list-style-type: none"> Support for Unified Forwarding Table (UFT). <p>[See Understanding the Unified Forwarding Table.]</p>

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • Support for redistribution of IPv4 routes with IPv6 next hop into BGP. [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.] • Support for collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding. [See Telemetry Sensor Explorer.] • Support for maximum reference bandwidth increased to 4 TB for IGP protocols. [See reference-bandwidth (Protocols IS-IS).] [See reference-bandwidth (Protocols OSPF).] • Check for AS matches in BGP policy AS paths without regular expressions. [See Improve the Performance of AS Path Lookup in BGP Policy.] • Support for stripping or replacing BGP private AS. [See Autonomous Systems for BGP Sessions.] • BMP local RIB monitoring support for all RIBs with sharding. [See BGP Monitoring Protocol.] [See loc-rib.] [See rib-list.] • Support for bootstrapping route-validation database from a local file. [See validation (Origin Validation for BGP).]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support (*Continued*)

Feature	Description
Routing policy and firewall filters	<ul style="list-style-type: none"> • Sharding support for conditional route manager. [See Routing Policy Match Conditions.] [See rib-sharding.] [See show policy conditions.] • Support for fast lookup of origin and neighbor autonomous systems (ASs). [See policy-options.] [See policy-statement.] • Firewall filter support on Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions.] • Support for profiles to improve the firewall filter scale. [See Planning the Number of Firewall Filters to Create.] • EVPN-VXLAN firewall filtering and policing. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
System management	<ul style="list-style-type: none">• Secure boot and secure BIOS support. [See Secure Boot.]• CLI-based hash and ECMP resilient hashing support. [See enhanced-hash-key.] [See ecmp-resilient-hash.]• Support for dynamic load balancing (DLB). [See enhanced-hash-key.]• Support for configuring firewall filters and interfaces programmatically using JET APIs. [See Overview of JET APIs.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> • ZTP support. [See Zero Touch Provisioning.] • Secure zero-touch provisioning—You can use RFC-8572-based secure zero-touch provisioning (SZTP) to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before the remote network device is accessed for initiating ZTP. To enable mutual authentication, you need a unique digital voucher, which is generated based on the DevID (Digital Device ID or Cryptographic Digital Identity) of the network device. The DevID is embedded inside the Trusted Platform Module (TPM) 2.0 chip on the network device. Juniper Networks issues a digital voucher to customers for each eligible network device. You can switch between using SZTP and ZTP on secure platforms. The default behavior on this device is ZTP. To override the default behavior of your secure device, issue the request <code>system zeroize ztp-option secure-enable</code> command. [See Secure Zero Touch Provisioning, Generate Secure ZTP Vouchers, and Switching between Secure Zero Touch Provisioning and Zero Touch Provisioning.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
Services applications	<ul style="list-style-type: none">• Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agent.]• Inband Flow Analyzer (IFA) 2.0 transit node support. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
VPNs	<ul style="list-style-type: none"> • Support for EVPN Type 5 routes. [See Understanding EVPN Pure Type-5 Routes.] • Support for assisted replication (AR) integrated with optimized intersubnet multicast (OISM) in an EVPN-VXLAN edge-routed bridging (ERB) fabric. [See Assisted Replication Multicast Optimization in EVPN Networks.] [See Optimized Inter-Subnet Multicast in EVPN Networks.] • EVPN-VXLAN support with MAC-VRF routing instances. [See EVPN User Guide.] • Support for EVPN-VXLAN fabric with an IPv6 underlay. [See EVPN-VXLAN with an IPv6 Underlay.] [See Example: Configure an IPv6 Underlay for Layer 2 VXLAN Gateway Leaf Devices.] • Support for symmetric IRB with EVPN Type 2 routes. [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics.] [See irb-symmetric-routing.] • Support for MLDv1, MLDv2, and MLD snooping with OISM and AR in EVPN-VXLAN fabrics. [See Optimized Intersubnet Multicast in EVPN Networks.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> Support for determining IRB interface state changes based on local and remote connectivity states in EVPN fabrics. [See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics.] [See interface-state.] [See network-isolation.] Overlay and CE-IP ping and traceroute support for EVPN-VXLAN. [See Understanding Overlay ping and traceroute Packet Support.] Support for blocking asymmetric EVPN Type 5 routes. [See EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN.] [See ip-prefix-routes.] Support for DHCP relay in EVPN-VXLAN. [See DHCP Relay Agent over EVPN-VXLAN.] Support for coexistence of EVPN Type 2 and Type 5 routes . [See EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN.] Support for Interconnecting EVPN-VXLAN in a data center to an EVPN-VXLAN control plane in a WAN using a gateway model. [See Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network.]

Table 2: QFX5130-48C and QFX5130-48CM Feature Support (*Continued*)

Feature	Description
	<ul style="list-style-type: none"> Support for OISM in an EVPN-VXLAN fabric. [See Optimized Inter-Subnet Multicast in EVPN Networks.] Support for service-provider-style interface configuration on EVPN-VXLAN Layer 3 gateways. [See Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network.] Overlapping VLAN support in EVPN-VXLAN fabrics on edge-routed bridging (ERB) overlay leaf devices. [See Overlapping VLAN Support Using VLAN Translation in EVPN-VXLAN Networks.] [See vlan-rewrite.]

- QFX5230-64CD switch (QFX Series)**—The QFX5230-64CD switch offers high-density 400GbE access ports in an SFP-DD form factor optimized for high-end spine and super-spine layers of the IP fabric multitier architecture in a 2-RU fixed form factor. The QFX5230-64CD switch provides a throughput of 25.6 Tbps and supports 64 400GbE access ports, along with support for 200-Gbps, 100-Gbps, 40-Gbps, 25-Gbps, and 10-Gbps speeds using breakout cables.

Table 3: QFX5230-64CD Feature Support

Feature	Description
CoS	<ul style="list-style-type: none"> • Support for CoS features on Layer 2 and Layer 3 interfaces. Both IPv4 and IPv6 unicast routing are supported. Other supported CoS features include: <ul style="list-style-type: none"> • 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. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2). • Weighted random early detection (WRED) and explicit congestion notification (ECN) • Telemetry support for CoS queue statistics exported using the sensor <code>/junos/system/linecard/qmon-sw/</code> . <p>[See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]</p>

Table 3: QFX5230-64CD Feature Support (*Continued*)

Feature	Description
EVPN	<ul style="list-style-type: none"> • Support for firewall filtering and policing on an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) network. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • Support for sFlow on an EVPN-VXLAN network. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers on an EVPN-VXLAN network. [See Port Mirroring and Analyzers.]
Forwarding and sampling	<ul style="list-style-type: none"> • Support for dynamic load balancing (DLB) and resilient hashing for equal-cost multipath (ECMP) routes. DLB and resilient hashing are not supported on a link aggregation group (LAG). [See Dynamic Load Balancing, Use of Resilient Hashing to Minimize Flow Remapping, and ecmp-resilient-hash.]

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Interfaces	<ul style="list-style-type: none"> • QFX5230-64CD has 64 QSFP56-DD ports and two SFP+ ports. The QSFP56-DD ports support the following speeds: <ul style="list-style-type: none"> • 400 Gbps • 200 Gbps • 100 Gbps • 50 Gbps • 40 Gbps <p>The QSFP ports also support the following speeds (with breakout cables):</p> <ul style="list-style-type: none"> • 50 Gbps • 25 Gbps • 10 Gbps <p>The SFP+ ports support 10-Gbps speed.</p> <p>QFX5230-64CD supports the following channelizations:</p> <ul style="list-style-type: none"> • 1x400GbE, 4x100GbE, 2x100GbE, and 2x50GbE on QSFP-DD ports • 2x50GbE, 1x50GbE, and 4x25GbE on QSFP28-DD ports • 1x40GbE and 4x10GbE on QSFP+ ports <p>[See Port Speed on QFX5230-64CD Switches.]</p>

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 2 unicast forwarding and VRRP. [See Understanding VRRP.] • Support for IGMP snooping, including: <ul style="list-style-type: none"> • IGMP snooping with 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) IGMP snooping • IGMP snooping with integrated routing and bridging (IRB) <p>[See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]</p>

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See DHCP Relay Agent.] • Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. We support both IPv4 and IPv6 unicast routing . [See Generic Routing Encapsulation (GRE).] • Support for Layer 3 multicast forwarding includes: <ul style="list-style-type: none"> • PIM first hop router (FHR) rendezvous point (RP) functionality • MSDP • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (aex) 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 (SM) • PIM dense mode (DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces: <ul style="list-style-type: none"> • IPv4 and IPv6 multicast • IGMP v1/v2/v3

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • Multicast Listener Discovery (MLD) v1/v2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes <p>[See Multicast Routing Protocols and PIM Overview.]</p>
Network management and monitoring	<ul style="list-style-type: none"> • Support for sFlow. <p>[See Overview of sFlow Technology.]</p> <ul style="list-style-type: none"> • Support for port mirroring and analyzers. The QFX5230-64CD switches supports a maximum of eight port mirroring sessions. <p>[See Understanding Port Mirroring and Analyzers.]</p>
Optics	<ul style="list-style-type: none"> • Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available. • We support performance monitoring for the QDD-400G-ZR optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics by using the <code>show interfaces transport pm</code> command and can manage optical transport link efficiently. <p>See show interfaces transport pm.</p>

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Platform and infrastructure	<ul style="list-style-type: none"> Platform resiliency support for hardware components of each FRU in the QFX5230-64CD switch. If a failure is detected on a hardware component, Junos OS Evolved: <ul style="list-style-type: none"> Logs the message to give clear indication of failure details, including time stamp, module name, and component name. Raises and clears alarms, if applicable. Raises SNMP trap. Makes the LED glow to indicate FRU fault, if an LED is present. Performs local actions such as self-healing or taking the component out of service. Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. <p>[See Overview of JET APIs.]</p>
Protection against DDoS attacks	<ul style="list-style-type: none"> Supports configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from DDoS attacks. By default, DDoS protection is enabled for many protocols on the QFX5230-64CD switches. <p>[See Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers, show ddos-protection statistics, and show ddos-protection version.]</p>

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Precision Time Protocol (PTP)	<ul style="list-style-type: none"> • Support for Precision Time Protocol transparent clock. [See PTP Transparent Clocks.] • Support for Precision Time Protocol enterprise and media profiles. [See PTP Profiles.] • Support for configuring Precision Time Protocol (PTP) over integrated routing and bridging (IRB) interfaces on AES67, SMPTE, AES67+SMPTE, and Enterprise PTP profiles. [See PTP over IRB for Broadcast Profiles.]
Routing policy and firewall filters	<ul style="list-style-type: none"> • Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Firewall Filters.]

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Services applications	<ul style="list-style-type: none">• Support for generic routing encapsulation (GRE) features:<ul style="list-style-type: none">• 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) <p>[See Generic Routing Encapsulation (GRE).]</p>

Table 3: QFX5230-64CD Feature Support (Continued)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> • Support for secure BIOS and secure boot implementation based on the UEFI 2.4 standard. [See Secure Boot.] • You can dynamically detect the port speed of WAN interfaces and use this information to create ZTP client ports with the same speed. ZTP automatically cycles through the WAN ports until it receives DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client. See Zero Touch Provisioning Using DHCPv6 Options. • You can use RFC-8572-based secure zero-touch provisioning (SZTP) to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before the remote network device is accessed for initiating ZTP. To enable mutual authentication, you need a unique digital voucher, which is generated based on the DevID (Digital Device ID or Cryptographic Digital Identity) of the network device. The DevID is embedded inside the Trusted Platform Module (TPM) 2.0 chip on the network device. Juniper Networks issues a digital voucher to customers for each eligible network device. You can switch between using SZTP and ZTP on secure platforms. The default behavior on this device is ZTP. To override the default behavior of your secure device, issue the request system zeroize ztp-option secure-enable command. See Secure Zero Touch Provisioning, Generate Secure ZTP Vouchers, and Switching between

Table 3: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
	<p>Secure Zero Touch Provisioning and Zero Touch Provisioning.</p> <ul style="list-style-type: none"> The QFX5230-64CD devices support the following firmware upgrade commands: <ul style="list-style-type: none"> request system firmware upgrade fpc slot 0 bcm-pfe request system firmware upgrade fpc slot 0 dp11 request system firmware upgrade fpc slot 0 dp11-cfg request system firmware upgrade fpc slot 0 opticscpld<0 1 2> request system firmware upgrade psm slot <0 1> request system firmware upgrade re bios request system firmware upgrade re fancpld request system firmware upgrade re fpga request system firmware upgrade re i210 request system firmware upgrade re ssd <disk1 disk2> request system firmware upgrade re xmcfgpa <p>[See request system firmware upgrade.]</p>

- **QFX5240 switches (QFX Series)**—The QFX5240-64OD and QFX5240-64QD are 800-Gigabit Ethernet (GbE) data center switches that run Junos OS Evolved. . These switches offer 64 800GbE OSFP and QSFP-DD ports. Using breakout cables, you can configure 64 800GbE ports, 128 400GbE ports, and 256 100GbE ports on the QSFP5240-OD and QSFP5240-QD.

Table 4: QFX5240 Feature Support

Feature	Description
Chassis	<ul style="list-style-type: none"> Support for inbuilt Routing Engine, Control Board (CB), power supply units, fan trays, Flexible PIC Concentrators (FPCs), and PICs. <p>[See QFX5240 Switch Hardware Guide.]</p>
CoS	<ul style="list-style-type: none"> Support for CoS features on Layer 2 and Layer 3 interfaces. Supported CoS features include: <ul style="list-style-type: none"> IPv4 and IPv6 unicast routing. Classification and rewrite rules (DSCP, IEEE-802.1p) Port scheduling Shared buffer Priority-based flow control (PFC) based on IEEE-802.1p. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2). Weighted random early detection (WRED) and explicit congestion notification (ECN) Telemetry support for CoS queue statistics exported using the sensor <code>/junos/system/linecard/qmon-sw/</code>. <p>[See Understanding How Class of Service Manages Congestion.]</p>
Forwarding and sampling	<ul style="list-style-type: none"> Support for dynamic load balancing (DLB) (for port speeds over 50 Gbps) and resilient hashing for ECMP routes. DLB and resilient hashing are not supported on a link aggregation group (LAG) or when a LAG is one of the egress ECMP members. <p>[See Dynamic Load Balancing, Use of Resilient Hashing to Minimize Flow Remapping, and ecmp-resilient-hash.]</p>

Table 4: QFX5240 Feature Support *(Continued)*

Feature	Description
Interfaces	<ul style="list-style-type: none"> • The QFX5240 switches have 64x800GbE OSFP ports on the QFX5240-64OD and 64x800GbE QSFP-DD ports on the QFX5240-64QD. The last two ports (64 and 65) on both the QFX5240 variants are 2x10GbE SFP ports. <p>The ports on the QFX5240-64OD and QFX5240-64QD support the following speeds:</p> <ul style="list-style-type: none"> • 1x800 Gbps • 2x400 Gbps • 4x200 Gbps • 8x100 Gbps <p>NOTE: On the QFX5240 switches, the runts (under Input errors) and fragment frames (under MAC statistics) counters do not increment in the output of the <code>show interfaces extensive</code> command. These counters are not supported due to a hardware limitation.</p> <p>[See Port Speed on QFX5240 Switches.]</p>

Table 4: QFX5240 Feature Support *(Continued)*

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 2 unicast forwarding and VRRP. [See Understanding VRRP.] • Support for IGMP snooping includes: <ul style="list-style-type: none"> • IGMP snooping with 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) IGMP snooping • IGMP snooping with integrated routing and bridging (IRB) <p>[See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]</p>

Table 4: QFX5240 Feature Support *(Continued)*

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. We support both IPv4 and IPv6 unicast routing. [See Generic Routing Encapsulation (GRE).] • Support for Layer 3 multicast forwarding includes: <ul style="list-style-type: none"> • PIM first hop router (FHR) rendezvous point (RP) functionality • Multicast Source Discovery Protocol (MSDP) • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (aex) 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 (SM) • PIM dense mode (DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces: <ul style="list-style-type: none"> • IPv4 and IPv6 multicast • IGMP v1/v2/v3 • Multicast Listener Discovery (MLD) v1/v2 • 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. The switches support DHCPv4 and DHCPv6. [See DHCP Relay Agent.]

Table 4: QFX5240 Feature Support *(Continued)*

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> Support for sFlow. [See Overview of sFlow Technology.] Support for analyzers and port mirroring. The QFX5240-64OD and QFX5240-64QD switches support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]
Optics	<ul style="list-style-type: none"> Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
Platform and infrastructure	<ul style="list-style-type: none"> Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.]
Protection against DDoS attacks	<ul style="list-style-type: none"> Supports 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 QFX5240-64OD and QFX5240-64QD 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	<ul style="list-style-type: none"> Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]

Table 4: QFX5240 Feature Support *(Continued)*

Feature	Description
Services applications	<ul style="list-style-type: none">• Support for generic routing encapsulation (GRE) features:<ul style="list-style-type: none">• 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) <p>[See Generic Routing Encapsulation (GRE).]</p>

Table 4: QFX5240 Feature Support (Continued)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> Support for firmware upgrade commands: <ul style="list-style-type: none"> <code>request system firmware upgrade re bios</code> <code>request system firmware upgrade re i210</code> <code>request system firmware upgrade re ssd disk1</code> <code>request system firmware upgrade re ssd disk2</code> <code>request system firmware upgrade cb fancpld</code> <code>request system firmware upgrade cb fpga</code> <code>request system firmware upgrade cb port-fpga</code> <code>request system firmware upgrade fpc slot 0 bcm-pfe</code> <code>request system firmware upgrade fpc slot 0 dp11</code> <p>[See request system firmware upgrade.]</p> Support for USB booting. <p>NOTE: On QFX5240 switches, only UEFI boot media (UEFI USB, UEFI NVME, UEFI network, and so on) is supported. You must select USB (UEFI USB) manually from the BIOS menu or use the <code>request node reboot re0 usb</code> command to boot from USB.</p>

- QFX5700 switches (QFX Series)**—Supported transceivers, optical interfaces, and DAC cables -- Select your product in the Hardware Compatibility Tool (<https://apps.juniper.net/hct/product/>) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

Authentication and Access Control

- Support for outbound SSH through HTTP proxy servers (PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5230-64CD)**—You can establish outbound SSH connections through an HTTP proxy server to enable secure remote management of devices, even when firewalls block outbound SSH connections.

Use the `outbound-ssh client name proxy-server` configuration statement to configure proxy server details.

[See [outbound-ssh](#).]

Additional Features Optimized for AI-ML Fabrics

For more information about features optimized for AI-ML fabrics, see the [AI-ML Data Center Feature Guide](#).

- **BGP Support for Global Load Balancing in DC Fabric (QFX5240)**— In a DC fabric, hashing is unable to ensure even load distribution over all ECMP links, which might result in congestions on certain links and underutilization on other links. Dynamic load balancing helps to avoid congested links to mitigate local congestion. However, dynamic load balancing cannot address some congestions. For example, AI ML traffic that has elephant flows and lacks entropy causes congestions in the fabric. In this case, global load balancing (GLB) helps to mitigate these congestions. Global load balancing is hashing a route with multiple ECMP links onto several links for load balancing.

In a CLOS network the congestions on the first two next hops impacts the load balancing decisions of the local node and the previous hop nodes triggering global load balancing. If the route has only one next-next-hop, a simple path quality profile is created. If the route has more than one next next-hop node then a simple path quality profile is created for each next next-hop node.

To enable global load balancing, include the `global-load-balancing` statement at the `[edit protocols bgp]` hierarchy level. We have disabled this statement by default.

[See [global-load-balancing](#).]

- **Configurable FlowSet table in DLB flowlet mode (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—Dynamic load balancing (DLB) uses the FlowSet table to determine the egress interface of flows. The table holds 32,768 entries, distributed among 128 DLB equal-cost multipath (ECMP) groups. By default, each ECMP group receives 256 entries. You can modify this distribution to accommodate more flows per ECMP group, thereby enhancing flow distribution. [See [Configure Flowset Table Size in DLB Flowlet Mode](#).]
- **Reactive path rebalancing (QFX5240-64OD and QFX5240-64QD)**—Use the enhanced flowlet mode in dynamic load balancing (DLB) to configure an inactivity interval for traffic on an outgoing interface. If the outgoing link quality deteriorates over time without exceeding the inactivity timer, reassign the traffic to a better quality link within the flowlet mode. This approach overcomes the limitations of the classic flowlet mode and ensures optimal traffic distribution. [See [Reactive Path Rebalancing](#).]
- **SNMP support for PFC, ECN, and CoS ingress packet drop accounting (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—We have introduced SNMP support that helps to account for the packets that are dropped because of ingress port congestion. You can view and export the error counters data for ECN, ingress drops, and PFC using the following commands:
 - `Show snmp mib walk ifJnxTable`

- Show snmp mib walk jnxCosPfcPriorityTable

[See [show snmp mib](#) and [SNMP MIBs Supported by Junos OS and Junos OS Evolved](#).]

- **Extended sFlow functionality (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—We have extended the sFlow monitoring functionality to support the export of sFlow sample packets through the mgmt_junos interface and non-default VRF WAN ports.

The management Ethernet interface provides the out-of-band management network by default. Deploying the mgmt_junos VRF instance ensures management traffic uses private IPv4 and IPv6 routing tables. The new routing-instance option at the [edit protocol sflow collector] hierarchy specifies the routing instance name.

sFlow can now export sample packets through non-default VRF WAN ports, which allows it to sample traffic on configured ports based on sample rate and port information.

The sFlow system comprises an agent embedded in the device and up to four external collectors. With these updates, collectors can be spread across different VRFs, and the software forwarding infrastructure daemon (SFID) determines the correct next-hop address for collector IPs, ensuring proper routing.

The show sflow collector detail command now displays the additional field “Routing Instance Name” to indicate the VRF name on which collector is reachable and “Routing Instance Id” that is corresponding to that VRF.

[See [collector](#), [show sflow collector](#), and [System Logging and Routing Instances](#).]

- **Remote port mirroring to IPv4/IPv6 address (GRE encapsulation) with DSCP, source-address, and rate-limiting parameters (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can configure DSCP, source-address, and rate-limiting parameters in your configuration for remote port mirroring to IPv4 or IPv6 addresses. You use remote port mirroring to copy packets entering a port or VLAN and send the copies to the IPv4 or IPv6 address of a device running an analyzer application on a remote network (sometimes referred to as “extended remote port mirroring”). The mirrored packets are GRE-encapsulated.

You configure source-address or source-ipv6-address, dscp, and forwarding-class options—either in the analyzer configuration or the port-mirroring configuration—under these hierarchies, respectively:

- [edit forwarding-options analyzer instance *instance-name* output]
- [edit forwarding-options port-mirroring instance *instance-name* family inet|inet6 output]

You configure the forwarding class and the shaping-rate option under the class-of-service hierarchy, as follows:

- set class-of-service forwarding-classes class *class-name* queue-num *queue-number*
- set class-of-service interfaces *interface-name* scheduler-map *map-name*

- `set class-of-service scheduler-maps map-name forwarding-class class-name scheduler scheduler-name`
- `set class-of-service schedulers scheduler-name shaping-rate rate`

[See [Port Mirroring and Analyzers](#).]

Chassis

- **Optics EM policy support (QFX5230-64CD)**—We've extended the Junos Environment Monitoring (EM) policy to include optics temperature sensors for QFX5230-64CD switches. The policy includes the following features:
 - The Optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed.
 - Junos OS Evolved automatically triggers shutdown of 100GbE and 400GbE optics when the system breaches the Fire Shutdown threshold.
 - The Optics EM policy is enabled by default on all 100GbE and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables.

You can use the `set chassis fpc fpc_slot pic pic_slot port port_no no-temperature-monitoring` command to explicitly disable the Optics EM policy on specific WAN ports. Use the `show chassis environment` command to view the optics temperature.

[See [temperature-sensor](#).]

- **Support for multiple speeds (QFX5130E-32CD)**—You can configure 10-Gbps, 25-Gbps, 40-Gbps, or 100-Gbps port speed on QFX5130E-32CD switches.

[See [speed \(Ethernet\)](#).]

- **Chassis management on new FEB (QFX5700E)**—We provide support for chassis management features on QFX5700E-FEB, a new Forwarding Engine Board (FEB) for QFX5700 switches. You can use the following commands to manage power and view QFX5700E-FEB information:
 - `set chassis feb slot slot-number power (off | on)` command to turn on or turn off the FEB
 - `show chassis feb` and `show chassis environment feb` commands to display FEB status information
 - `request chassis feb slot slot-number (offline | online | restart)` to take the specified FEB offline, bring it online, or restart the FEB

[See [power](#), [show chassis feb](#), [request chassis feb](#), and [show chassis environment](#).]

- **Resiliency support (QFX5130-48C, QFX5240-64OD, and QFX5240-64QD)**—We support resiliency for platform components on the QFX5130-48C, QFX5240-64OD, and QFX5240-64QD devices.

[See [Resiliency](#).]

Class of Service

- **Support for dynamic ECN (QFX5240-64OD and QFX5240-64QD)**—QFX5240 switches support dynamic explicit congestion notification (ECN). Regular ECN sends an alert when a configured traffic threshold is exceeded and takes mitigating action based on the configured profile. Dynamic ECN automatically adjusts these threshold settings based on real-time conditions such as queue size and traffic patterns. This enhances responsiveness and management during periods of network congestion.

[See [CoS Explicit Congestion Notification](#) and [Example: Configuring Static and Dynamic ECN](#).]

- **Per-queue accounting of ECN packets (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—Counters on explicit congestion notification (ECN)-enabled queues increment when the queues experience congestion or receive packets that encountered congestion on another device. You can view these per-queue ECN accounting statistics using the `show interfaces queue` command.

[See [CoS Explicit Congestion Notification](#) and [show interfaces queue](#).]

- **Support for PFC watchdog (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Use the priority-based flow control (PFC) watchdog to PFC-enabled ports for pause frame storms and automatically resolve them. The watchdog detects the storms, mitigates the condition triggering the storm, and restores the network.

[See [PFC Watchdog](#).]

Forwarding Options

- **Support for faster forwarding mode (QFX5130, QFX5220, QFX5230-64CD, QFX5240, and QFX5700)**—You can configure all interfaces on your switch to use cut-through forwarding mode for faster packet forwarding and optimized packet processing. By default, the switch uses store-and-forward mode to forward packets.

[See [Configuring Forwarding Mode on Switches](#).]

High Availability

- **Unified ISSU support for 802.1X protocol (QFX5230-64CD)**—802.1X clients no longer experience traffic disruptions during unified ISSU. You can use this feature to maintain continuous network connectivity and minimize traffic loss. The feature is enabled by default.

Use the `show dot1x sync-pending-sessions` command to monitor authentication sessions pending synchronization after unified ISSU.

[See [Understanding Unified ISSU](#).]

- **EVPN support during unified ISSU on leaf devices (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Spine-and-leaf topologies with external BGP (EBGP) connections support EVPN data traffic forwarding on the spine device without any traffic loss while the leaf device performs a unified ISSU.

[See [Performing a Unified ISSU](#)]

Interfaces

- **Support for performance monitoring and TCA (QFX5220-32CD)**—We support performance monitoring for the QDD-400G-ZR-M optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics by using the `show interfaces transport pm` command and can manage optical transport link efficiently. See [show interfaces transport pm](#).
- **Support for performance monitoring and TCA (QFX5130-32CD)**—We support performance monitoring for the QDD-400G-ZR-M optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics by using the `show interfaces transport pm` command and can manage optical transport link efficiently. See [show interfaces transport pm](#).

Junos Telemetry Interface

- **Support for Health Monitoring telemetry data for standby nodes and FPCs (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—We've expanded telemetry data to support Health Monitoring telemetry beyond the primary node to include standby nodes and Flexible PIC Concentrators (FPCs). You can stream statistics that include load average, process parameters, and component CPU utilization using either Juniper's proprietary remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) transport from the device to the collector.

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

- **Real-time hardware resource monitoring (QFX5130, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, and QFX5700)**—You can monitor hardware resource utilization statistics in real time for efficient capacity planning. The Junos telemetry interface (JTI) streams NPU resource utilization statistics through gRPC and the gNMI transport from a device to an external collector.

[See <https://www.juniper.net/documentation/us/en/software/junos/interfaces-telemetry/topics/concept/junos-telemetry-interface-grpc-sensors.html>, <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/sensor-edit-services-analytics.html>, and [Junos YANG Data Model Explorer](#).]

- **Telemetry and SNMP support for network counters and buffer usage (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—This feature introduces telemetry streaming and SNMP support for explicit congestion notification (ECN), priority-based flow control (PFC) counters, and ingress buffer usage. Monitor the MAC Priority Flow Control Statistics, ECN Marked Packets, and Ingress Resource Errors counters using fields in the `/state/interfaces/interface/counters/` sensor. Track the Egress Priority Group Peak Buffer Occupancy statistics with the `/junos/system/linecard/qmon-sw/` sensor. The feature introduces SNMP support for these statistics using the respective fields.
 - In the `/state/interfaces/interface/counters/` sensor:
 - MAC Priority Flow Control Statistics:
 - `/state/interfaces/interface[name=' ']/counters/pfc[priority=' ']/priority`
 - `/state/interfaces/interface[name=' ']/counters/pfc[priority=' ']/in-pkts`
 - `/state/interfaces/interface[name=' ']/counters/pfc[priority=' ']/out-pkts`
 - ECN Marked Packets:
 - `/state/interfaces/interface[name=' ']/counters/errors/out-ecn-ce-marked-pkts`
 - Ingress Resource Errors:
 - `/state/interfaces/interface[name=' ']/counters/errors/in-resource-drops`
 - In the `/junos/system/linecard/qmon-sw/` sensor:
 - Egress Priority Group Peak Buffer Occupancy:
 - `/cos/interfaces/interface[name=' ']/priority-groups/priority-group[pg=' ']/peakBufferOccupancy`

SNMP support is introduced with the following fields:

- MAC Priority Flow Control Statistics:
 - `JnxCosPfcPriorityTable` (`jnxCosPfcPriorityRequestsTx` & `JnxCosPfcPriorityRequestsRx`)
- ECN Marked Packets:
 - `IfJnxTable` (`ifJnxOutEcnMarkedPackets`)
- Ingress Resource Errors:
 - `IfJnxTable` (`ifJnxInQDrops`)

[See [Junos YANG Data Model Explorer](https://www.juniper.net/documentation/us/en/software/junos/netconf/topics/concept/yang-junos-modules-overview.html) and <https://www.juniper.net/documentation/us/en/software/junos/netconf/topics/concept/yang-junos-modules-overview.html>.]

- **Configuring IP source address for legacy gRPC dial-out connections (QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)**—Set a source IP address and routing instance for legacy gRPC service dial-out connections. Previously, the outgoing interface IP address was used as a source address without an option to configure a source IP address. This feature supports FLEX deployments, enabling dial-out from a specified IP address or interface (such as a loopback0 address). Use the routing-instance and local-address statements at the [edit services analytics export-profile *profile-name*] hierarchy level.

[See <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/local-address-edit-services-analytics-export-profile.html>, <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/export-profile-edit-services-analytics.html>, <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/routing-instance-edit-services-analytics-export-profile.html>, and <https://www.juniper.net/documentation/us/en/software/junos/interfaces-telemetry/topics/topic-map/telemetry-grpc-dialout-ta.html>]

- **IPv4 and IPv6 traffic statistics telemetry (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—This feature supports streaming of IPv4 and IPv6 transit statistics using the native resource paths `/state/interfaces/interface[name='']/counters/ipv4/` and `/state/interfaces/interface[name='']/counters/ipv6/`. The exported fields include in-pkts, in-bytes, out-pkts, and out-bytes. To enable transit statistics for the physical port, configure route accounting by including the route-accounting statement at the [edit forwarding-options family *family-name*] hierarchy level.

[See [Junos YANG Data Model Explorer](#) and <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/route-accounting-edit-forwarding-options.html>.]

- **OpenConfig compatibility for Junos OS Evolved telemetry sensors (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—We've introduced Junos telemetry interface (JTI) sensor support for renaming existing noncompatible interface sensor leaves to OpenConfig-compatible sensor leaves. We've removed a few nonstandard leaves for OpenConfig compatibility. Current sensors support target-defined streaming, but new sensors do not support ON-CHANGE. [For a complete list of all other sensors available in the above sensor path, see [Junos YANG Data Model Explorer](#).]
- **Telemetry streaming support (QFX Series)**—Junos OS Evolved supports telemetry on QFX Series switches using the following sensors:

```
lastchange
init-time
parent-ae-name
high-speed
counters/in-pkts
counters/in-octets
```

```

counters/in-unicast-pkts
counters/in-multicast-pkts
counters/in-broadcast-pkts
counters/in-pause-pkts
counters/out-pkts
counters/out-octets
counters/out-unicast-pkts
counters/out-multicast-pkts
counters/out-broadcast-pkts
counters/out-pause-pkts
counters/in-errors
counters/in-discards
counters/carrier-transitions,
counters/out-errors
counters/out-discards

```

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

Multicast

- **Enhanced L3 multicast operational commands (ACX7100-32C, PTX10004, and QFX5130-32CD)**—The show instance command is now extended to all routing instances for the following commands. Earlier, only specific PIM-enabled routing instances were displayed.

- show pim join instance all
- show pim rps instance all
- show pim statistics instance all
- show multicast route instance all
- show multicast statistics instance all

Additionally, the show pim statistics output will display V2 Sparse Join and V2 Sparse Prune counters.

The show igmp statistics output will also display the V1/V2/V3 Membership Query field.

[See [show pim statistics](#), [show multicast statistics](#), and [show igmp statistics](#).]

Network Management and Monitoring

- **Higher adaptive sampling rate for sFlow (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—The adaptive sampling rate increases from 9500 to 30000 packets per second (PPS). This enhancement allows for more efficient data sampling and improved network performance.

[See [adaptive-sample-rate](#) and [Adaptive Sampling Overview](#).]

Platform and Infrastructure

- **NIST purge method for media sanitization (QFX5130-32CD, QFX5130-48CM, QFX5220, and QFX5230-64CD)**—We've extended support for NIST media sanitization for SATA hard disk drives to include:
 - Cryptographic scramble and block erase priorities for the purge method.
 - Enhanced secure erase priority for the clear method.

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 disk 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](#).]

Precision Time Protocol (PTP)

- **Transparent clock support (QFX5130E-32CD)**—QFX5130E-32CD switches support the Precision Time Protocol (PTP) transparent clock feature. Use the transparent clocks 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](#).]

- **Enterprise profile support on PTP ordinary and boundary clocks (QFX5130E-32CD)**—The PTP ordinary and boundary clocks support enterprise profile functionality on QFX5130E-32CD devices. You can configure the enterprise profile to meet specific network requirements, ensuring seamless integration and optimal operation.

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

Routing Policy and layer2-policer Firewall Filters

- **Change default processing behavior of firewall filters with filter based forwarding actions (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**

A firewall filter term that has a filter based forwarding action is always processed first, regardless of its order of placement in the firewall filter configuration. You use `force-fbf-terms` to change this default behavior. When you apply this configuration, the firewall filter terms in a firewall filter are always

processed in the order of their placement in the configuration, irrespective of whether a firewall filter term has a filter based forwarding action or not.

[See [force-fbf-terms](#).]

- **Selectively enable or disable dynamic load balancing (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can selectively enable or disable dynamic load balancing based on `rdma-opcode` match or any match available in firewall filters using the new `dynamic-load-balance` configuration statement. You can modify port load and port queue metrics from their default values so that when selective load balancing is enabled, the metrics are used to determine an optimal link. Use the new `egress-quantization` configuration statement to configure the desired ratio of port load metric to port queue metric based on the traffic pattern.

[See [rdma-opcode](#), [dynamic-load-balance-selective](#), and [egress-quantization](#).]

- **Support for forwarding matched packets to a specific VLAN (QFX5130-32CD, QFX5130-48C, and QFX5700)**—To activate this action profile on these platforms, you have to apply the `set system packet-forwarding-options firewall profiles actions ethernet-switching profile1` configuration. You can configure the `vlan vlanID` action in port and VLAN firewall filter rules.

[See [Firewall Filter Match Conditions and Actions \(QFX and EX Series Switches\)](#).]

Routing Protocols

- **Minimum ECMP (QFX5130-32CD, QFX5130-48C, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—We support conditional advertising and withdrawal of BGP routes based on certain constraints such as bandwidth and minimum available next-hop ECMP. When a BGP receiver learns the same route from multiple BGP peers, BGP updates the active BGP path and the routing information base (RIB), also known as the routing table. The BGP export policy determines whether to advertise the BGP route to these next hops based on the number of ECMP BGP peers it receives the prefix from. A BGP route that has multiple ECMP BGP peers creates better resiliency in case of link failures. You can configure a BGP export policy to withdraw a BGP route unless it receives the BGP route prefix from a minimum number of ECMP BGP peers.

[See [link](#).]

Software Installation and Upgrade

- **Switching between SZTP and ZTP on secure platforms (QFX5130-48C and QFX5130-48CM)**— Use RFC-8572-based secure zero-touch provisioning (SZTP) to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before initiating ZTP.

To enable mutual authentication, the system generates a unique digital voucher based on the Digital Device ID or Cryptographic Digital Identity (DevID) of the network device. The DevID is embedded

inside Trusted Platform Module (TPM) 2.0 chip on the network device. We issue a digital voucher to customers for each eligible network device.

You can switch between using SZTP and zero touch provisioning (ZTP) on secure platforms. The default behavior on this device is ZTP. To override the default behavior of your secure device, you can issue the request `system zeroize ztp-option secure-enable` command.

[See [Secure Zero Touch Provisioning](#), [Generate Secure ZTP Vouchers](#), and [Switching between Secure Zero Touch Provisioning and Zero Touch Provisioning](#).]

- **Overlay load balancing in an EVPN-VXLAN network (QFX5130-32CD and QFX5700)**—You can provision QFX5130-32CD and QFX5700 switches to function as leaf or spine devices in an EVPN-VXLAN network to support load balancing among different virtual tunnel endpoints (VTEPs). Overlay load balancing is enabled by default. We support overlay load balancing:
 - With centrally-routed bridging (CRB) overlays and edge-routed bridging (ERB) overlays
 - When a leaf device is multihomed to multiple spine devices
 - When a host is multihomed to multiple leaf devices

You configure each multihomed aggregated Ethernet interface, logical interface, or physical interface with an Ethernet segment identifier (ESI). You can use a maximum of 256 ESIs with overlay load balancing.

[See [show ethernet-switching vxlan-tunnel-end-point svlbnh](#).]

- **ZTP with DHCPv4/DHCPv6 support (QFX5240-64OD and QFX5240-64QD)** — Use a DHCPv4 or DHCPv6 client and zero-touch provisioning (ZTP) to set up a device. During the bootstrap process, the device requests image and configuration file information from the DHCP server using the DHCPv4/DHCPv6 client. The device is provisioned with ZTP using either DHCPv4 or DHCPv6 based on the response received from the DHCP client.

[See [Zero Touch Provisioning](#).]

Additional Features

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

- **BGP autodiscovery underlay in EVPN-VXLAN** (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)

[See [BGP Auto-Discovered Neighbors](#).]

- **BGP link bandwidth community** (QFX5130-48CM) BGP can communicate link speeds to remote peers, enabling better optimization of traffic distribution for load balancing.

[See [auto-sense](#), and [group \(Protocols BGP\)](#).]

- **EVPN-VXLAN L2 gateway and ARP suppression** (QFX5230-64CD). For Layer 2 (L2) gateway support, we've added the overlay configuration statement under system packet-forwarding-options. Use the overlay statement options to set the size of the overlay interfaces and next hops to ensure optimal resource allocation based on network requirements.

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#), [EVPN Proxy ARP and ARP Suppression](#), and [overlay \(Packet Forwarding Options\)](#).]

- **Inband Flow Analyzer (IFA) 2.0 transit node support** (QFX5240-64OD and QFX5240-64QD)

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

- **L2PT with Q-in-Q over VXLAN tunnels in EVPN-VXLAN bridged overlay networks** (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700). Support includes additional protocol options to tunnel the STP family of protocols:

- STP, Multiple STP (MSTP), and Rapid STP (RSTP)—stp option
- Per-VLAN STP (PVSTP) and PVSTP Plus (PVSTP+)—pvstp option
- VLAN STP (VSTP)—vstp option

[See [Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks, Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network](#), and [I2pt \(Destination Tunnels\)](#).]

- **Minimum ECMP**(QFX5130-48CM). We support conditional advertising and withdrawal of BGP routes based on certain constraints such as bandwidth and minimum available next-hop ECMP. You can configure a BGP export policy to withdraw a BGP route unless it receives the BGP route prefix from a minimum number of ECMP BGP peers.
- **Non-revertive preference-based EVPN DF Election** (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)

[See [EVPN Multihoming Designated Forwarder Election, preference \(DF Election\), df-election-type](#).]

- **Statically identify multihoming peer OISM leaf devices in an EVPN-VXLAN network running enhanced OISM** (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700). The statement to identify an optimized intersubnet multicast (OISM) device's multihoming peers changes with this release from multihoming-peer-gateways at the [edit protocols evpn] hierarchy level to static-multihoming-peer at the [edit protocols evpn] hierarchy level. The multihoming-peer-gateways statement isn't available anymore.

[See [Statically Identify Multihoming Peers With Enhanced OISM To Improve Convergence](#) and [static-multihoming-peer](#).]

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

- **Supported transceivers, optical interfaces, and DAC cables**—Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
- **Support for centralized and inline BFD** (QFX5130E-32CD)
[See [Understanding How BFD Detects Network Failures](#).]
- **Wake-on-LAN (WOL) targeted broadcast support for EVPN-VXLAN** (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)
[See [Targeted Broadcast](#) and [targeted-broadcast](#).]
- **WRED drop counter for ECN telemetry** (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)
[See [CoS Explicit Congestion Notification, Example: Configuring Static and Dynamic ECN](#), and [show interfaces queue](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 158](#)
- [Class of Service \(CoS\) | 158](#)
- [EVPN | 158](#)
- [General Routing | 158](#)
- [Interfaces and Chassis | 159](#)
- [Junos Telemetry Interface | 159](#)
- [Junos XML API and Scripting | 159](#)
- [Network Management and Monitoring | 159](#)
- [Routing Protocols | 160](#)
- [User Interface and Configuration | 160](#)

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

Authentication and Access Control

- Disabled CDN auto download (Junos OS Evolved)— The PKI process periodically, by default every 24 hours, polls the CDN server for the latest default trusted CA bundle and updates the list for any changes to the trusted CAs in the bundle. If there are any changes, PKI process loads them in the background. The auto download of CA certificates might generate core files. We've disabled the service of PKI query to CDN server periodically to download the latest trusted CA bundle.
- On Junos OS Evolved, password authentication for SCP based configuration archival is supported.

Class of Service (CoS)

- On QFX5000 Series switches running Junos OS Evolved, egress buffer thresholds are not applicable for lossless traffic. Only ingress priority group thresholds are used for admission control. We use ingress alpha (dynamic-threshold) for ECN threshold calculation on lossless queues. In DCQCN deployment, this means that most of time during congestion, ECN marking occurs before PFC generation, as desired. This fix also results in relatively more ECN-marked packets on lossless queues due to a more accurate threshold calculation.

EVPN

- **EVPN system log messages for CCC interface up and down events**—Devices will now log EVPN and EVPN-VPWS interface up and down event messages for interfaces configured with circuit cross-connect (CCC) encapsulation types. You can look for error messages with message types EVPN_INTF_CCC_DOWN and EVPN_INTF_CCC_UP in the device system log file (/var/log/syslog).

General Routing

- Change to the commit process—In prior Junos OS Evolved releases, if you use the commit prepare command and modify the configuration before activating the configuration using the commit activate command, the prepared commit cache becomes invalid due to the interim configuration change. As a result, you cannot perform a regular commit operation using the commit command. The CLI shows an error message: 'error: Commit activation is pending, either activate or clear commit prepare'. If you now try running the commit activate command, the CLI shows an error message: 'error: Prepared commit cache invalid, failed to activate'. You then must clear the prepared configuration using the clear system commit prepared command before performing a regular commit operation. From this

Junos and Junos OS Evolved release, when you modify a device configuration after 'commit prepare' and then issue a 'commit', the OS detects that the prepared cache is invalid and automatically clears the prepared cache before proceeding with regular 'commit' operation.

[See [Commit Preparation and Activation Overview](#).]

Interfaces and Chassis

- **Maximum Configurable MTU Size (QFX5130-32CD, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5700, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)**— You can configure a maximum MTU size of 9408 using the `set interfaces` command. For QFX5130-48C and QFX5130-48CM, maximum configurable MTU size is 9368.

Junos Telemetry Interface

- The `show agent sensors` command output for gRPC sensors is truncated on the Junos OS Evolved platform to align with the output format of the Junos OS platform.

Junos XML API and Scripting

- **Commit script input to identify software upgrades during boot time (ACX Series, PTX Series, and QFX Series)**—The `junos-context` node-set includes the `sw-upgrade-in-progress` tag. Commit scripts can test the `sw-upgrade-in-progress` tag value to determine if the commit is taking place during boot time and a software upgrade is in progress. The tag value is `yes` if the commit takes place during the first reboot after a software upgrade, software downgrade, or rollback. The tag value is `no` if the device is booting normally.

[See [Global Parameters and Variables in Junos OS Automation Scripts](#).]

Network Management and Monitoring

- In a firewall filter configured with a `port-mirror-instance` or `port-mirror` action, if `l2-mirror` action is also configured, then `port-mirroring` instance family should be any. In the absence of the `l2-mirror` action, `port-mirroring` instance family should be the firewall filter family.

- **Python 2 interpreter option deprecated for Juniper Extension Toolkit (JET) applications (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX10K-LC1202-36MR (line cards for PTX10016, PTX10008 and PTX10004), QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-64OD, QFX5240-QD, QFX5700, and QFX5700E)**—Python 2.7 is already not supported on Junos OS Evolved devices as of an earlier release. The python statement at the **edit system extensions extension-service application file <filename>** hierarchy level was used to interpret JET applications written in Python 2. This statement is now deprecated. To run daemonized on-device JET applications written in Python 3, use the python3 statement.

[See [file \(JET\)](#).]

Routing Protocols

- **Update to IGMP snooping membership command options**— The instance option is now visible when issuing the `show igmp snooping membership ?` command. Earlier, the instance option was available but not visible when `?` was issued to view all possible completions for the `show igmp snooping membership` command.

[See [show igmp snooping membership](#).]

- **MLD snooping proxy and l2-querier source-address (ACX7024, ACX7100-32C, PTX10001-36MR, QFX5120-32C, and QFX5130-32CD)**— The source-address configured for proxy and l2-querier under the `mld-snooping` hierarchy should be an IPv6 link-local address in the range of `fe80::/64`. The CLI help text has been updated to "Source IPv6 link local address to use for proxy/L2 querier". In earlier releases, the CLI help text read, "Source IP address to use for proxy/L2 querier."

[See [source-address](#).]

User Interface and Configuration

- **Compact format deprecated for JSON-formatted state data (ACX Series, PTX Series, and QFX Series)**—We've removed the `compact` option at the `[edit system export-format state-data json]` hierarchy level because Junos devices no longer support emitting JSON-formatted state data in compact format.
- **Access privileges for request support information command (ACX Series, PTX Series, and QFX Series)**—The `request support information` command is designed to generate system information for troubleshooting and debugging purposes. Users with the specific access privileges `maintenance`, `view`, and `view-configuration` can execute `request support information` command.

Known Limitations

IN THIS SECTION

- [General Routing](#) | 161

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

- The recommended procedure for FPC card removal is to do a graceful OIR by following the documented steps. However, if the FPC card is ungracefully jacked out or plugged in several times in a short interval, the software may show the FPC status as "FAULT". Below are the steps to recover from the fault state: 1. Jack out the FPC 2. Wait for 1 minute for FPC state to become 'Fault' in 'show chassis fpc pic-status' command 3. Jack In FPC 4. Wait for 1 minute for FPC state shows as Present in 'show chassis fpc pic-status' command 5. Run cli command 'request chassis fpc online slot fpc slot number ' to bring FPC in online state 6. Wait for 3 minutes for all interfaces to come up [PR1799333](#)
- When a VXLAN encapsulated IP packet, or an IP packet with UDP port matching the VXLAN UDP port, is received on a vlan-tagging enabled interface, the switch drops the frame. This issue is not seen if the incoming port is an untagged interface, or if the interface is actually doing VXLAN encaps/decap operations. In such cases, the device forwards the frame correctly. [PR1805922](#)
- On QFX5240 QFX5230 platforms with 23.4R2, when PFC watchdog feature configured with recovery action of "drop" and if PFC storm continuously detected and recovered , CRC error counter could increase with small number under "show interface extensive interface" output. These CRC errors could occur only for larger frames (mtu>1024) received at high rate. Packet drops will be seen on the interface due to PFC watchdog action of drop, which is expected. CRC errors will not be seen with PFC watchdog recovery action of "forward". [PR1807420](#)
- On QFX5700E, " pci 0000:xx:xx.x BAR xx: failed to assign" messages are seen during boot in the logs. These messages have no impact to functionality. [PR1807706](#)
- IPv4/IPv6 reserved multicast and L2 multicast traffic received over VXLAN access port will be flooded out of all ports of the VXLAN except vtep. [PR1811158](#)

- Storm control does not work on multicast packets on QFX5130, QFX5700 platform if broadcast packets are excluded from the storm control profile. To support IGMP, MLD snooping functionality, we've added flood in vlan rules in asic pipeline which causes mcast packets to flooded as broadcast packets due to BCM TD4 asic limitation.[PR1813514](#)
- On EVO QFX5K platforms , PFC XON and MRU values configured under congestion notification profile are properties of Ingress port priority-group. Junos CLI doesn't have any equivalent construct for ingress port priority-group, unlike queue which has forwarding-class construct in CLI. Codepoint (priority) to priority-group mapping happens in PFE based on CNP configuration using internal logic. As there are only 6 lossless priority groups per ingress port, there is possibility of more than one IEEE 802.1P / DSCP code points to get mapped to same PG (fate sharing). In this case if the code points mapped to same PG configured with different XON values, then the XON value associated with highest codepoint will be programmed in HW for the corresponding priority group. This is due to entries in CNP profile are processed and programmed in ascending order based on codepoint. If the code points mapped to same PG configured with different MRU values, then the MRU value associated with lowest codepoint will be used to calculate PFC headroom for that PG. [PR1822023](#)
- The PFC watchdog can be triggered when it is set with a very low detection timer value, like 4 ms, while continuously receiving PFC XOFF frames from the peer device. On the peer device, two different priorities have been configured for PFC. One priority has a very high PFC XON offset (greater than 10,000), while the other priority uses the default PFC XON offset (20). As part of the PFC feature, BCM supports a PFC refresh functionality. When a priority experiences congestion and the current buffer utilization exceeds the PFC XOFF threshold, a PFC XOFF frame is sent. If the buffer utilization does not fall back to the PFC XON threshold within the default PFC refresh time, the port will generate a new PFC XOFF refresh frame to the peer device. For a 100G port, the default refresh time is 262 microseconds. This is why multiple PFC XOFF frames may be observed before a PFC XON frame is sent. This behavior is expected for the priority with the higher XON offset. However, due to hardware design limitations, the PFC refresh timer operates on a per-port basis. Therefore, when the per-port PFC refresh timer expires, the port triggers PFC refresh XOFF frames for all priorities that are in the XOFF state at that time. The hardware cannot distinguish which priority's refresh timer has expired. As a result, even for a priority with the default XON offset, multiple refresh XOFF frames may be sent continuously due to the expiration of the port-level PFC refresh timer. This could cause the peer device to detect a PFC storm for that priority as well. Since this is a hardware limitation, it cannot be resolved. Aside from the continuous XOFF frames that may trigger PFC watchdog detection on the peer, there are no other functional impacts due to this design. The recommendation is that if a user sets a very high XON offset for any priority on a port, which could lead to PFC refresh timer expiry and continuous XOFFs, the peer device should be configured with a longer PFC watchdog detection timer. For instance, if a PFC XON offset of 10,000 is set for a priority, the peer device should have a PFC watchdog detection timer of at least 10 ms.[PR1833562](#)
- On restarting the picd app on QFX5700 , interfaces flap [PR1842469](#)
- EP-style L2 IFL statistics (ingress / egress) is not supported on EVO QFX platforms.[PR1843854](#)

Open Issues

IN THIS SECTION

- [General Routing](#) | 163

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 QFX5230-64CD platform, 400G DAC cable of 2.5m length and 4X100G DAC BO might not link up with some peer devices. This issue is not seen with all peer devices. If this happens, please replace the 2.5m cable with a 1m DAC cable or use supported 400G Optics instead.[PR1747315](#)
- On QFX5130-48C, when image upgrade is done using CLI, the log messages have kernel trace reported. There is no functionality impact due to the messages.[PR1755406](#)
- Whenever a QFX5130-48C device is reset using rear panel reset button by pressing the button for short/long duration, the port LEDs on ports 56/57 might glow in amber for a few seconds during boot up. The port LEDs amber status turn off and reflect correct port state once device is up.[PR1792619](#)
- QFX5230-64CD, QFX5130-48C/48CM: Help string for MTU settings on management interface are not correct. The range is showed as 256-9408 instead of 256-9216.[PR1813591](#)
- When user has custom dedicated buffer configuration using `set class-of-service dedicated-buffer` along with congestion-notification-profiles on multiple interfaces which requires a larger lossless headroom buffer space which is more than the default shared buffer space, during full class-of-service configuration deletion/deactivation syslog error message Invalid config : **Not enough Ingress Lossless headroom** is seen due to order of delete events. This error message is seen due to a transient system state, however this does not affect the system functionality after the completion of configuration commit. [PR1815246](#)
- QFX5700 MACsec: Minor packet drops (0.0000001%) observed when MACsec is enabled.[PR1816407](#)

- Link with the DAC (SFP56-50G-DAC-3M) comes up with 25G default speed configuration when switch is rebooted. Hence, when 50G speed configuration is applied , peer side sees the link flap.[PR1836697](#)

Upgrade Your Junos OS Evolved Software

Products impacted: ACX7024, ACX7024X-DC, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX10K-LC1202-36MR (line cards for PTX10016, PTX10008 and PTX10004), PTX10002-36QDD, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-64OD, QFX5240-QD, QFX5700, and QFX5700E.

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/>.

Documentation Updates

This section lists the errata and changes in Junos OS Evolved Release 24.4R1 documentation.

The Time Management Administration Guide is renamed to Timing and Synchronization Guide. See, [Timing and Synchronization Guide](#).

Licensing

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

The major components of the framework are:

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

For information about the list of supported products, see [Juniper Flex Program](#).

Finding More Information

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

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

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

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

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

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



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

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

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

Requesting Technical Support

IN THIS SECTION

- [Self-Help Online Tools and Resources | 167](#)
- [Creating a Service Request with JTAC | 167](#)

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

Revision History

28 January 2025—Revision 4, Junos OS Evolved Release 24.4R1

17 January 2025—Revision 3, Junos OS Evolved Release 24.4R1

6 January 2025—Revision 2, Junos OS Evolved Release 24.4R1

30 December 2024—Revision 1, Junos OS Evolved Release 24.4R1

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