Release Notes

Published
2026-09-01

Junos OS Evolved Release 26.2R1

Table of Contents

Introduction

Junos OS Evolved runs on the following Juniper Networks® hardware: ACX Series, PTX Series and QFX Series. Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 26.2R1.

You can find release notes for all Junos OS Evolved releases at Junos OS Evolved Documentation | Juniper Networks

Note: The recommended release for QFX5000 Series switches is Junos OS Evolved 25.2X100-D21.

Junos OS Evolved Release Notes for ACX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

Class of Service (CoS)

Support for port-level excess rate and excess priority on CoS schedulers (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—You can allocate excess bandwidth on ports using `excess-rate` *`percent`* and `excess-priority` in port-level schedulers. After servicing the transmit-rate, also known as the committed information rate (CIR), queues with higher excess priority receive excess bandwidth, or excess information rate (EIR), up to their excess rate.

Configure `transmit-rate` *`percent`*, `excess-rate` *`percent`*, and `excess-priority` (`high`, `medium-high`, `medium-low`, `low`) in a scheduler at the `[edit class-of-service schedulers` *`scheduler-name`*`]` hierarchy level and apply it to an interface through a scheduler map. The configuration applies only to port-level schedulers and excludes low-latency, strict-high, low-medium, and low-high queues; hierarchical class of service (HCoS) is unsupported. If only `excess-priority` is set, the excess rate defaults to 100%.

[See Schedulers Overview for ACX Series Routers, excess-rate (Schedulers), and excess-priority (Schedulers).]

Support for IEEE 802.1p classification on IRB interfaces (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—You can apply Layer 2 IEEE 802.1p classification on IRB interfaces for consistent class of service (CoS) across routed and switched traffic. The IRB Layer 2 classifier propagates to underlying Layer 2 interfaces, temporarily overrides their classifiers, and restores the classifiers when you remove the IRB classifier. You can configure either Layer 2 or Layer 3 classification on an IRB interface, but not

both. To apply an IEEE 802.1p classifier to an IRB interface, set `classifiers ieee-802.1 classifier-name` at the `[edit class-of-service interfaces irb unit unit-number]` hierarchy level.

[See Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview.]

EVPN

EVPN-VXLAN cross-module forwarding consistency and hash verification using the CLI (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)—You can detect and diagnose EVPN-VXLAN forwarding inconsistencies across control, kernel, and forwarding modules to reduce traffic loss or duplication. Verify per-VLAN synchronization and errors with:

- `request system evpn forwarding vlans [detail | instance name | vlan name | module {mac-table|mac-ip-table}] [no-confirm]`
- `request system evpn hash-values vlans [instance name] [vlan name]`

Compute comparison digests and validate Type 2 route and next-hop consistency. These checks highlight missing programming elements or deleted bridge domains.

[See EVPN-VXLAN Forwarding Consistency and Hash Verification, request system evpn forwarding vlans, and request system evpn hash-values vlans.]

High Availability (HA) and Resiliency

High availability support for Hit-full MACsec recovery (ACX7332 and ACX7348)—We now support high availability for hit-full MACsec recovery. When high availability (HA) switchover occurs with MACsec enabled, a brief traffic interruption is expected in the hit-full interface. During the switchover, active sessions are torn down and re-established. As a result, the interface experiences a behavior similar to a MACsec down/up event after the switchover.

Junos Telemetry

Support for genstate YANG data models (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource

paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Decoupled rendering for telemetry data streaming (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming. Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

OpenConfig YANG model support for BGP (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)—Junos telemetry now supports the OpenConfig YANG model for BGP, providing a vendor-neutral framework for configuration and monitoring BGP. Retrieve BGP-specific operational and state data by subscribing to sensors under resource path hierarchies such as `inline communities`, `community-set`, `match-community-set`, `match-ext-community-set`, `peer-group`, `neighbor`, `policy-definition`, `administrative distance`, and `as-path-set`. OpenConfig-defined default behavior for BGP policies is implemented, with `community-level` policy options supporting `any-match` and `invert-match` configurations.

Use the Junos YANG Data Model Explorer to view supported sensor paths and platforms. The supported OpenConfig version for the openconfig-network-instance model is 4.4.1.

The following table lists the supported YANG models and their versions:

BGP YANG Model	Version
openconfig-routing-policy.yang	3.4.2
openconfig-bgp-common-multiprotocol.yang	9.9.1
openconfig-bgp-common-structure.yang	9.9.1
openconfig-bgp-common.yang	9.9.1
openconfig-bgp-errors.yang	6.1.0
openconfig-bgp-global.yang	9.9.1
openconfig-bgp-neighbor.yang	9.9.1
openconfig-bgp-peer-group.yang	9.9.1
openconfig-bgp-policy.yang	8.1.0
openconfig-bgp-types.yang	6.1.0
openconfig-bgp.yang	9.9.1
openconfig-types.yang	1.0.0

[See Junos YANG Data Model Explorer.]

Update and extend OpenConfig BGP telemetry (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)—Support has been added for OpenConfig BGP telemetry in these two ways:

- BGP telemetry has been updated to the latest OpenConfig models.
- Added support for the following paths:

```
/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/config/set-med-action
```

/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/state/set-med-action

/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/state/set-med

[See Junos YANG Data Model Explorer.]

Update to BGP link bandwidth community value (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)—Support BGP to allow regular expressions 0 and 2^32 in the link bandwidth community value. The OpenConfig path is

/routing-policy/defined-sets/bgp-defined-sets/ext-community-sets/ext-community-set/

See the YANG Data Model Explorer.

[See Junos YANG Data Model Explorer.]

Addition of send-community-type and telemetry support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)—Extend core functionality to add `send-community-type` and telemetry support.

The following paths are supported:

/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/afi-safis/afi-safi/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/afi-safis/afi-safi/state/send-community-type

The OpenConfig community-type doesn't define IPv6-ext-community types, so they are not controlled through the OC configuration and they are always advertised.

[See Junos YANG Data Model Explorer.]

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the [`edit event-options destinations` *`destination-name`* `archive-sites` *`url`*] hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

Layer 2 Features

Support for unmatched VLAN configuration (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—You can configure an interface to accept packets that are untagged or have a VLAN ID other than the VLAN IDs configured on the other interfaces on a port. The configured interface then supports bridging, VPLS, and CCC encapsulation.

If you configure a native VLAN ID on the interface, that configuration takes precedence over the unmatched configuration; the interface with the native VLAN ID configured would then accept the untagged packets.

[See Binding VLAN IDs to Logical Interfaces and vlan-id (VLAN ID to Be Bound to a Logical Interface).]

MPLS

Enhancements to MPLS automatic bandwidth for RSVP LSPs with per-LSP polling profiles and adjustable thresholds (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10003, PTX10004, and PTX10008)—You can dynamically allocate bandwidth for RSVP label-switched paths (LSPs) based on traffic to improve utilization and service assurance. Define multiple polling profiles with the `polling-profile` statement under the [`edit protocols mpls statistics auto-bandwidth`] hierarchy and associate each profile with specific LSPs to tailor polling intervals. Configure threshold settings through the `adjust-threshold-options` statement to set precise overflow and underflow triggers. Use enhanced CLI options to create polling profiles, bind them to LSPs, and set thresholds to improve monitoring and troubleshooting in complex networks.

[See Polling Profiles for Automatic Bandwidth Allocation for RSVP LSPs.]

Load balancing with entropy labels for SR-MPLS and SR-TE using IS-IS signaling (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—Use entropy labels for Segment Routing with MPLS (SR-MPLS) and segment routing-traffic engineering (SR-TE) to improve equal-cost multipath (ECMP) and link aggregation group (LAG) load balancing across paths. At ingress, the device inserts an Entropy Label Indicator (ELI) and Entropy Label (EL) at the bottom of the label stack. Transit devices calculate a hash based on the EL for per-flow distribution while forwarding packets on segment routing (SR) labels. Enable entropy label insertion is done globally or per path. By default, Entropy Label Capability (ELC) is enabled with the `load-balance-label-capability` statement at the [`edit forwarding-options`] hierarchy level. You can disable the capability by configuring the `no-load-balance-label-capability` statement at the [`edit forwarding-options`] hierarchy level.

You can enable entropy labels by configuring the `entropy-label` statement at the [`edit protocols isis source-packet-routing`] and at the [`edit protocols source-packet-routing source-routing-path`] hierarchy levels. After you configure the `entropy-label` statement, the IS-IS routes and the prefixes for SR-TE are installed with an ELI if the endpoint is entropy label capable.

[See Basic LSP Configuration.]

OpenConfig

OpenConfig gNMI heartbeat for on-change subscriptions (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—Use OpenConfig heartbeat for gNMI ON_CHANGE subscriptions to receive a full set of sensor data at each heartbeat interval, even when values do not change. The heartbeat provides periodic state validation and mitigates undetected data loss.

[See Configuring gNMI Heartbeat Interval for ON_CHANGE Subscriptions and Junos YANG Data Model Explorer.]

Routing Options

Route re-resolution optimization for L3VPN and recursive multipath with new diagnostic commands (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, AND PTX12008, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)—You can now use optimized route re-resolution in the resolver module to reduce CPU utilization during route churn. This optimization extends re-resolution beyond inet.0 and inet6.0 to `<instance>.inet.0`, `<instance>.inet6.0`, and L3VPN tables (bgp.l3vpn.0, bgp.l3vpn-inet6.0). The resolver also supports optimized re-resolution for recursive resolution cases and for resolution over multipath prefixes.

To monitor and diagnose resolution, you can use the following operational commands:

1. `show route resolution active-helpers` - Displays filtered routes used for the resolution, before the active routes were displayed
2. `show route resolution active-helpers statistics` - Helps identify the reason for which routes are re-resolved, which is helpful in debugging.
3. `show route resolution re-resolution-history` - Displays the routes that are re-resolved.
4. `show route resolution loop-history` - Helps identify if any resolution loop occurred.

[See show route resolution.]

Routing Protocols

Support for suppressed MPLS S=0 clone routes and associated next-hop entries (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—The rpd routing process ceases transmission of the MPLS S=0 clone routes and their associated next-hop entries, by default. This default configuration applies to:

- LDP
- RSVP
- BGP
- Segment Routing with IS-IS (SR-IS-IS)
- Segment Routing with OSPF
- Segment routing--traffic engineering (SR-TE)
- static LSPs
- Express segments
- Route express services

This default behavior of rpd prevents these routes from being pushed to the Packet Forwarding Engine, optimizing resource usage. Use the `show route table mpls.0` and `show route forwarding-table family mpls` commands to check for MPLS S=0 clone routes in the routing and forwarding tables.

[See show route forwarding-table and show route table.]

Software Installation and Upgrade

Support for TCP connectivity on management and WAN interfaces for non-SI applications (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E)—You can enable TCP connectivity for statically linked or other non-Socket Intercept (non-SI) applications over management and WAN interfaces. Start each non-SI application with the `NO_SI_INTERCEPT` environment variable to enable TCP traffic processing.

[See Running Applications Without the Socket Intercept Library.]

Source Packet Routing in Networking (SPRING) or Segment Routing

BGP RIB sharding for SRv6 routes (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)—We've extended support for Segment Routing over IPv6 (SRv6) with BGP routing information base (RIB) sharding to enhance route processing and utilize SRv6 next hops. RIB sharding splits BGP service resolution over multiple shards to improve BGP convergence and performance. Configure SRv6 under IS-IS or as static SRv6-Traffic Engineering (SRv6-TE) with BGP RIB sharding. We do not support BGP route resolution when RIB sharding is configured.

[See rib-sharding.]

Support for IS-IS IPv6 SRLG TLV 139 and interface-level SRLG configuration (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—This feature enables routers to advertise and process IPv6 Shared Risk Link Group (SRLG) information as defined in RFC 6119, *IPv6 Traffic Engineering in IS-IS*, ensuring accurate awareness of shared-risk failures. By extending SRLG support to IPv6 and integrating it directly into the IS-IS configuration, you can deploy SRv6 networks without relying on MPLS. This feature enhances Topology-Independent Loop-Free Alternates (TI-LFA) fast reroute path calculation. It enables faster and more reliable protection against link, node, and shared-risk failures for modern IP networks.

To enable per-interface SRLG configuration and advertisement in IS-IS for SRLG-aware TI-LFA, configure `srlg` *`srlg-name`* at the `[edit protocols isis interface interface-name]` hierarchy level.

[See Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS and OSPF.]

SR-MPLS-TE local computation with OSPF underlay (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, AND PTX12008, PTX10001-36MR, PTX10001-36MR-K, PTX10002-60MR, PTX10003, PTX10004, PTX10008, and PTX10016)—We support Segment Routing over MPLS with Traffic

Engineering (SR-MPLS-TE) local computation with OSPF as the underlay interior gateway protocol (IGP) for the following features:

- Endpoint IP address configured as anycast IP address
- Segment-list computation using anycast constraints
- Metric-type delay
- Protected and unprotected path selection

[See Segment Routing LSP Configuration and Configuring Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS.]

Storm Control

Storm control support (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—Storm control is applicable at the physical level and the logical level. You can configure different policer rates per traffic type by using the `set forwarding-options storm-control-profiles` command and these new options: `broadcast`, `registered-multicast`, `unknown-unicast`, and `unregistered-multicast`. By using the new options, you can monitor the type of traffic that you would like to monitor. If you configure only the `bandwidth-level` or the `bandwidth-percentage` option, the default burst size is 1500 kbits (Kb).

[See storm-control-profiles.]

User Access and Authentication

TACACS+ Over TLS for Secure Device Administration (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—This release introduces support for Transport Layer Security (TLS) for TACACS+ communication between network devices and TACACS+ servers. By leveraging TLS, this enhancement provides strong encryption, integrity protection, and mutual authentication, addressing the security limitations of legacy TACACS+ mechanisms. The previous weak protection methods are now deprecated, ensuring that sensitive AAA traffic is securely transmitted in modern deployments.

TACACS+ over TLS Support (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD,

QFX5241E-64OD, QFX5700, and QFX5700E)—This release introduces support for **TACACS+ over TLS**, enhancing the security of device administration by replacing the legacy MD5-based obfuscation mechanism with modern TLS encryption. With this update, communication between Juniper devices and TACACS+ servers is protected by TLS, ensuring **confidentiality, integrity, and mutual authentication** using X.509 certificates.

Key highlights:

- **Secure transport** for TACACS+ authentication, authorization, and accounting traffic.
- **Mutual TLS (mTLS)** support, enabling certificate-based validation of both client and server.
- **Failover support** between TACACS+ over TLS, TACACS+ over TCP, RADIUS, and local authentication.
- **Routing instance compatibility**, including default, non-default, and management instances.
- **CLI configuration enhancements** to define trusted CA groups and client certificates for TLS connections.

This feature allows TACACS+ deployments to meet modern compliance requirements (e.g., FIPS-140-3) and strengthens overall AAA security in Junos and Evo platforms.

Additional Features

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

Export timing data to collectors (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, PTX10004, and PTX10008)—Use Junos telemetry to export timing attributes for Precision Time Protocol (PTP) and Synchronous Ethernet (SyncE) to external telemetry data collectors. The device supports periodic streaming and on-change notifications. Subscribe to the following sensor paths to view PTP and Synchronous Ethernet sensor information:

- For Precision Time Protocol (PTP): `/state/protocols/ptp/`
- For Synchronous Ethernet: `/state/protocols/synce/`

For a complete list of sensor paths supported by the device, see Junos YANG Data Model Explorer.

Support for MPLS statistics on P2MP LSPs (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—We support packet and byte counter statistics for RSVP, LDP, and static P2MP LSPs at ingress and transit. Enable ingress statistics with the `set system packet-forwarding-options mpls lsp-ingress-stats-enable` command. Use the `set system packet-forwarding-options mpls lsp-transit-stats-enable` command to enable transit statistics. Enabling or disabling LSP statistics causes the Packet Forwarding Engine to restart. To view these statistics, issue the `show mpls lsp statistics p2mp` or `show ldp traffic-statistics p2mp` command.

[See show mpls lsp and show ldp traffic-statistics.]

Support for traffic statistics information for IRB interface with custom profile configuration (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—IRB statistics will be enabled when custom counter engine profile is created for IRB ingress and egress IFL. Statistics can be viewed by issuing the command `run show interfaces IRB extensive`.

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

Support for physical interface damping (ACX7020-AC, ACX7020-DC, ACX7024, ACX7100-32C, ACX7100-48L, ACX7332, and ACX7348)—Interface damping suppresses advertisements on physical interfaces during periodic flaps lasting 5 seconds or longer with 1-second up/down cycles. Interface damping reduces churn for upper-layer protocols and prevents unnecessary reconvergence. Configure `damping` at the `edit interfaces` *`interface-name`* hierarchy level. Monitor damping parameters and link status with the `show interfaces extensive` command.

[See damping (Interfaces), show interfaces extensive, and Damping Interfaces.]

Multicast-only fast reroute (MoFRR) with Protocol Independent Multicast (PIM) signaling (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—Support does not extend to Multipoint LDP-based MoFRR. Only local link failure is supported.

[See Understanding Multicast-Only Fast Reroute and Example: Configuring Multicast-Only Fast Reroute in a PIM Domain.]

NG-MVPN RSVP-TE P2MP link protection support (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—NG-MVPN RSVP-TE point-to-multipoint (P2MP) link protection support with a guaranteed convergence time of 50 ms.

[See Configuring Link Protection for Point-to-Multipoint LSPs.]

PTP feature support (ACX7332 and ACX7348)—The ACX7332 and ACX7348 support the following PTP features:

- Boundary clock and ordinary clock
- Unicast PTP over IPv4 and IPv6, including mixed-mode operation, with single, double, or no VLAN tagging
- PTP over IRB (with single or double VLAN tagging) for inter-VLAN or subnet communication
- TimeTransmitter-only, timeReceiver-only, and manual timeReceiver configurations
- Unicast negotiation with single or multiple TLVs and subnet timeTransmitter stream with automatic timeReceiver (IPv4/IPv6)
- Alternate best master clock algorithm (BMCA) and ITU-T G.8275.2 compliance (including Annex A and protocol parameters)
- G.8275.1 enhanced hybrid profile (PTP + Synchronous Ethernet) and enhanced servo

- Scalability up to 4 timeReceiver clocks and 512 timeTransmitter streams
- G.8273.4 Class B performance with enhanced G.8275.2 servo profile
- LAG for PTP over IPv4 and IPv6 with configurable interface asymmetry
- PTP recovery following GRES

[See Precision Time Protocol (PTP) Overview.]

sFlow ingress support for MPLS (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—You can configure sFlow on interfaces with the MPLS family to sample and report MPLS traffic. The system supports services such as LDP, RSVP-TE, SR-MPLS, VPLS, VPWS, L3-VPN, and FRR. However, the system does not support UHP and GRE encapsulation.

[See sFlow Technology Overview.]

EVPN Preferred Next-hop for All Active Multihoming (ACX7020-AC, ACX7020-DC, ACX7024, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)—The EVPN Preferred Nexthop feature enables an active/backup forwarding model for unicast data traffic on multi-homed PE's with all-active by selecting a preferred overlay nexthop IP address.

With this feature enabled, the network can achieve fast convergence when the primary PE goes down, by triggering BGP Prefix Independent Convergence [BGP-PIC] functionality in the PFE.

[See preferred-nexthop.]

What's Changed

IN THIS SECTION

Learn about what's changed 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 OS Evolved Problem Report Search application.

EVPN

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the `[edit routing-instances name protocols evpn]` hierarchy level.

  [See Preference-Based DF Election.]

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

General Routing

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for autoconfiguration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router?s link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid

unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Restrict PPPoE client lockout output to underlying logical interfaces with short-cycle protection**—The "show pope lockout" command shows PPPoE clients information only for underlying logical interfaces with `short-cycle-protection` enabled.

  [See show pppoe lockout.]

- **Traffic Engineering Database Import Identifier Update**—The help string for the `set protocols mpls traffic-engineering database import identifier` command has been updated to BGP-TE identifier for standard IGP instance. This change clarifies that the identifier configuration affects only the standard IGP instance. The BGP-LS advertisements in the `lsdist.0` table will now use the configured value instead of defaulting to Instance Identifier 0. This update ensures accurate Instance Identifier configurations in BGP-LS advertisements, enhancing network visibility and management. The routers running software later than Junos OS release 24.4.x version should verify that the correct Instance Identifier is being advertised. [See import (MPLS Traffic Engineering Database.]

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `edit services port-extender satellite <name> device-model` hierarchy, you can use the "?" to see a list of supported platforms. We have also added the `skip-lo0-config` statement under `edit service port-extender jnu` hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the `edit service port-extender satellite <name>` hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **ARP DDoS protection configuration restriction for unicast and broadcast policers (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—For ARP DDoS protection on devices running Junos OS Evolved in the ACX family, only the aggregate policer option is supported and effective under the system ddos-protection protocols arp hierarchy. The ucast and bcast ARP policer options are not supported on this Platform and are now treated as no-operations: if you configure these options, the system ignores the corresponding configuration blocks and reports a warning message "Warning: configuration block ignored: unsupported platform" indicating that the configuration block is ignored for the platform.

- **Pixel-Tracing Command Updates**—The pixel-tracing command suite has been updated to include the following options: `flush`, `disable`, and `enable`. These options allow you to control and manage pixel-tracing processes for routing tasks. Note that the `rib-sharding` option is removed from the CLI commands to prevent system crashes when sharding configuration is present in the system. Ensure that your configurations do not rely on unmentioned options to avoid potential disruptions in your workflow.

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- **Command to view interface-set information (ACX Series)**—Use the `show subscribers interfaces interface-set` command to view interface-set information.

  [See show subscribers interfaces.]

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Prevent multicast route installation on uninterested receivers**—Configure the `resolve-request-holdcount` statement under the `[edit routing-instances name routing-options multicast]` hierarchy to prevent the kernel from caching initial multicast packets during route resolution and avoid unnecessary route installation on devices without interested receivers. [See {https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/resolve-request-holdcount-edit-routing-instances-routing-options-multicast.html} resolve-request-holdcount].

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

- The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.
- The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.
- The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

[See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **BGP route-refresh warning messages added**—New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the

number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

[See show bgp diagnostics warnings.]

- **Inbound Convergence Status Added to BGP Neighbor Display (Junos OS and Junos OS Evolved)**—The `show bgp neighbor` command now displays the Inbound Convergence status and cause.

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **System Stability Improvement**—When BGP fails to bind to its listening port, the routing protocol process (RPD) now stays running and logs a warning instead of terminating after the maximum retry count, preventing a service outage. The `show bgp diagnostics overview` output includes a new `BGP listening port status of v4 and v6` field that reports `Active` when BGP is successfully listening on the configured port and `Failed` when the port is unavailable or occupied by another process. Error logging is also enhanced with the `BGP_LISTEN_STATUS` message: `BGP listen failed on port [X] after 64 attempts with error-%d for address family identifier (IPv4 or IPv6). Administrative clear needed.` This message is generated after `BGP_MAX_LISTEN_START_ATTEMPTS` failed attempts to bind to the listening port.

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.

User Interface and Configuration

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.

  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.

- `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
- `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
- `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:
  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of `remove` operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents **statement not found** warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

IN THIS SECTION

Learn about known 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 OS Evolved Problem Report Search application.

General Routing

- During the transition from backup 1 to backup 2 in AFTS, backup 1 is disabled, and SPLL information from backup 2 is delayed. The system initially remains in the Phase Aligned state (Active, Ready) with backup 2. The delayed SPLL update causes a shift to the Acquiring state (Active, Not Ready) as synchronization with backup 2 begins. The system successfully re-aligns and returns to the Phase Aligned state (Active, Ready) with backup 2. PR1909821

Open Issues

IN THIS SECTION

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

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

EVPN

- On Junos OS Evolved platforms, if the router-id is not explicitly configured, the device dynamically selects a router-id during bootup. In this scenario, the device advertises EVPN (Ethernet Virtual Private Network) Type1 AD/ESI (AutoDiscovery/Ethernet Segment Identifier) and Type4 ES RT (Ethernet Segment Route Target) routes with a Route Distinguisher (RD) of 0:0. As a result, Route Targets (RTs) advertised by different Provider Edge (PE) routers end up having identical prefixes. Consequently, only one PE routers RT is advertised and learned. This behavior prevents EVPN all-active redundancy resulting in the loss of multihoming. PR1924472

General Routing

- On Junos Evolved ACX7024 and ACX7024X platforms, the system crashes, and a kernel crash is generated due to a random CPU error. PR1895749
- On ACX7100-32C platform, when the device is rebooted, 100G links do not not come up with TX LOS. PR1869296
- On ACX Series platforms running Junos OS Evolved, when renaming a Flexible Tunnel Interface (FTI) causes all Internet Protocol version 6 (IPv6) traffic on the FTI tunnel to be dropped.

  PR1896144
- After verifying DHCP relay active-leasequery subscriber still not fully bound.

  PR1877878
- On all Junos OS Evolved ACX7000 platforms, in an MPLS (Multiprotocol Label Switching) FRR (Fast Reroute) scenario, traffic switches to the backup path after a primary link failure but does not revert to the primary path once it recovers. This issue is caused by improper handling of a software-unsupported nexthop during MBB (Make-Before-Break) updates, where a stale nexthop entry remains in the session database instead of being cleared. This stale entry is associated with a session ID marked as down, and if the primary link flaps, it reuses this problematic session ID, causing the primary path to inherit the incorrect down state. Additionally, when MBB creates a new nexthop with a different session ID, the stale entry for the unsupported nexthop persists, leading to inconsistencies in the session database. As a result, the device continues forwarding traffic over the backup path even though the primary path is operational, causing suboptimal routing.

  PR1920646
- On Junos OS Evolved ACX7000 line of Cloud Metro Routers, when storm control is configured on a LAG (Link Aggregation Group) interface and one or more LAG member interfaces are deactivated, storm control stops working on all remaining active members. This results in an immediate increase in broadcast, unknown unicast, and multicast (BUM) traffic on the egress interface.

PR1855369

- On ACX7100 and ACX7348 platforms running Junos EVO 24.4R1, a commit error may occur when attempting to delete an interface (IFD) that has the hierarchical scheduler enabled and configured. The hierarchical scheduler is a required dependency for the output traffic control profile applied on interface sets (IFLSET). During commit validation, this dependency is enforced, preventing interface deletion until associated hierarchical scheduler configurations are deactivated. PR1833976
- Need to increase the committed-burst size from 10k to higher value - 20k or more. Otherwise, small percentage (3% - 5%) of traffic drop is seen at the policer level. (See #49 of 1947985) PR1956850

Infrastructure

- Logs are flooded with audit: PROCTITLE causes the message file to fill up quickly.

  PR1882686

MPLS

- MVPN route is programmed with same NH for primary and backup when LDP link-protection is enabled PR1955399

Routing Protocols

- On all Junos OS Evolved platforms, after Border Gateway Protocol (BGP) sessions configured with `family route-target` flaps, delayed route deletion causes the loss of the Route Target Filter (RTF), preventing the node from advertising L3VPN (Layer 3 Virtual Private Network) and direct routes (e.g., loopbacks and interface routes) to the BGP peer, leading to VPN route loss and service disruption.

  PR1849568

Services Applications

- bbe-smg-upd leaks in 16k LAC subscribers login and logout test PR1955734

Resolved Issues

IN THIS SECTION

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

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

Authentication and Access Control

- On all Junos OS Evolved platforms, if the username is more than sixteen characters, the username is truncated to sixteen characters in the accounting logs displayed for that user.

  PR1786580

General Routing

- On Junos OS Evolved ACX7332 and ACX7348 platforms, after Routing Engines(RE) switchover the Precision Time Protocol (PTP) delay message is sent from different REs when Link Aggregation Group (LAG) members are configured on different Flexible PIC Concentrators (FPCs) leading to PTP timing discrepancy. PR1892686

- On all Junos and Junos Evolved platforms supporting `chained-composite-next-hop` for BGP-LU (Border Gateway Protocol - Labeled Unicast), when having L2Circuit (Layer 2 Circuit) links configured, with indirect next-hop resolving on BGP-LU prefix, if `chained-composite-next-hop` enabled for BGP-LU and with `preserve-nexthop-hierarchy` configured, L2Circuit traffic will get dropped due to failure in installing the forwarding next-hop with labels for L2Circuit service.

  PR1931875

- On Junos OS Evolved platforms with TPM2.0, after executing the zeroize command, TPM (Trusted Platform Module) keys are not regenerated on reboot causing the TPM dependent features like sZTP, gRPC to fail.

  PR1899669

- An Improper Restriction of Excessive Authentication Attempts vulnerability in Juniper Networks Junos OS Evolved allows an unauthenticated, network-based attacker to cause a limited denial-of-

service (DoS) to the management plane. Refer to CEC Juniper Community for more information. PR1709385

- On all Junos OS Evolved platforms, the Packet Forwarding Engine Daemon (PICD) repeatedly creates new SNMP trap objects without removing old ones, which causes memory leaks in the distributord process over time. This is tracked under Technical Support Bulletin (TSB) 107540.

  PR1914708

- On Junos OS Evolved platforms, the link may continuously flap on Small Form-factor Pluggable(SFP) ports when auto-negotiation is enabled by default and a link status change occurs. When the link comes up, auto-negotiation is reconfigured, retriggering the auto-negotiation state machine and causing the link to go down again. PR1873534

- On Junos OS Evolved ACX and PTX platforms with Dual RE, a configuration commit will not sync to the standby RE (Routing Engine). If a switchover happens during this time, traffic will be impacted because both REs are not in sync. This is a rare case. PR1890518

- On all Junos OS Evolved platforms, high volumes of ARP (Address Resolution Protocol) sessions cause control message from one packet to persist and be incorrectly applied to the next packet during a processing loop. This internal data corruption leads to inconsistent ARP handling, traffic loss, and a crash of fibd (Forwarding Information Base Daemon) process. PR1865957

- On Junos OS Evolved ACX Series platforms, including QFX Series Switches, when the command `set system ddos-protection protocol` is configured or when the ddos-protection configuration is deleted, or the command `show ddos-protection protocol` is executed, the following logs appear: DDOS:POLICE: DdosPolicer does not exist DDOS:POLICE: Ddos group does not exist.

  PR1922793

- [ACX] HWD_ALARM_CLEAR_NOTICE: ClearFault: Fault(Location: /Chassis[0]/Psm[0] Device: psm0 Error: psm_cml_pec_fault) cleared PR1914062

- On all Junos OS Evolved platforms, the shared memory value in the CPU memory log is stored in bytes, but the memory leak detection logic interprets is as kilobytes. This unit mismatch leads to incorrect per-application memory usage calculations, resulting in memory leak error messages for various applications. PR1892348

- On Junos OS Evolved ACX platforms that support BNG CUPS (Broadband Network Gateway Control and User Plane Separation), L2TP (Layer 2 Tunneling Protocol) subscriber sessions will fail to come up. When this happens, the subscriber session does not bind correctly and no data traffic will pass for the affected subscribers.

  PR1916030

- On all Junos OS Evolved ACX7000 platforms, in an MPLS (Multi-Protocol Label Switching) scenario, due to the incorrect label pushing on PHP (Penultimate Hop Popping) pop operation, the MPLS traffic on the egress node is impacted.

  PR1920723

- On Junos OS Evolved ACX7509, ACX7348, ACX7332, and ACX7100-32C platforms with 1G fibre Optics, when the Inter-Packet Gap (IPG) setting intermittently falls below the IEEE minimum requirement of 8 bytes, and traffic passes through the PHY, certain receiving devices are sensitive to this reduced IPG, which leads to CRC (Cyclic Redundancy Check) errors and packet loss on the peer device connected to 1G fibre optics connection. PR1931077

- On all Junos Evolved platforms, on invalid interface fetching, display error is seen on

  `show arp interface` *`interface-name`* command.

  PR1934771

- On Junos Evolved ACX7000 series platforms, when FBF(Filter Based Forwarding) is configured with all L3 interfaces, forwarding of FBF filtered IP packets will not work properly when next-IP/next-IP6 in firewall action is defined with network-prefix IP(IPv4 address other than /32 and IPv6 address other than /128). In Filter based Forwarding FBF filtered IP packets are forwarded to predefined next hop which is defined in firewall action next-IP. When next-IP is configured with network-prefix IP, if ARP(Address Resolution Protocol)/NDP(Neighbor Discovery Protocol) to next hop becomes unreachable and if it becomes reachable, handling of software next-hop to hold next-hop resolution is not proper for network-prefix IP. Since next-hop resolution is not handled for network-prefix IP, FBF policy filtered IP packets will not be forwarded over FBF path.

  PR1947793

- On Junos OS Evolved ACX platforms, when HCOS (Hierarchical Class of Service) and class of service are enabled for VPLS (Virtual Private LAN Service), deactivating and reactivating the output-traffic-control-profile or making HCOS interface configuration changes might cause L2 (Layer 2) multicast traffic to be duplicated toward the access interface. PR1931648

- On all ACX-Series platforms running Junos Evolved Operation System (OS), the device stops sending and receiving control plane traffic and a packetio core-dump is generated while handling incorrect packets length received from ASIC. PR1930141

- On Junos OS Evolved ACX7100-32C, ACX7100-48L, ACX7100-32C-K and ACX7100-48L-K platforms when QSFP-100G-FR optics which has serial number starting with 2J4 is plugged in, it is observed that the port goes down with I2C(Inter-Integrated Circuit) access failure alarm and traffic is dropped. PR1935793

- On Junos OS Evolved ACX platforms, "broadsync_lock_failure" alarms are seen intermittently without any external trigger. These alarms are seen as a result of internal firmware misconfiguration and

affects PTP (Precision Time Protocol) functionality. The alarms are cleared after few seconds (~10 secs) once the broadsync lock is restored. PR1896521

- On ACX7000 devices running Junos Evolved, when Inline Sampled Flow (Sflow) is configured, sflow is unable to properly set the sampling rate per port or flow in hardware, instead it uses the global sampling rate configuration.

  PR1930078

- The dhcpd daemon may restart every 15 minutes, generating periodic syslog messages. This behavior occurs because the platform does not use DHCP/TFTP services for line card booting, but still inherits a protective mechanism designed for chassis systems with CPUbased FPCs.

  PR1936239

- On all Junos OS Evolved platforms with Border Gateway Protocol (BGP) signalled Segment Routing (SRv6) dynamic tunnel and connected to a third party vendor device, when an update packet with SRv6 Service Sub-TLV (SSTLV) length as zero is received from a peer, BGP SRv6 dynamic tunnel resolution fails. As tunnel resolution fails, traffic intended for this SRv6 path will be dropped.

  PR1926805

- ACX7348 running GRES/NSR redundancy may print an error message due to incorrect Broadcom Logical Interface logic: evo-pfemand [Error] Interface : BrcmPlusIfAlloc: status_t BrcmPlusLifIdxAllocator::setIdx PR1942169

- Due to FEC resource exhaustion, MPLS tunnel service programming fails on evo-pfemand on ACX7000 platforms. New MPLS tunnels and some forwarding entries cannot be programmed into the ASIC. This may impact traffic and cause instability in MPLS forwarding, even if the reported route scale is low.

  PR1942200

- On Junos OS Evolved ACX Series platforms configured in an EVPN-MPLS (Ethernet Virtual Private Network - Multiprotocol Label Switching) all-active multihoming environment using ESI-LAG (Ethernet Segment Identifier - Link Aggregation Group), when a BGP (Border Gateway Protocol) / BFD (Bidirectional Forwarding Detection) flap occurs between EVPN peers, BUM (Broadcast, Unknown-unicast, Multicast) traffic is sent back towards the same CE (Customer Edge) from which it was received due to missing ESI label encapsulation. As a result, BUM traffic is impacted and is not forwarded correctly. PR1953985

- On the Junos OS Evolved ACX7024X platform having EVPN-VxLAN (Ethernet VPN - Virtual Extensible LAN) or EVPN-MPLS (Ethernet VPN - Multiprotocol Label Switching) with VLAN-aware service-type configuration, traffic disruption will be observed when traffic is sent with an incorrect label. In this scenario, traffic may be dropped by the egress device or mistakenly forwarded to the

wrong bridge domain. This issue happens when the ACX7024X is peered with a multi-homing device that uses per bridge-domain label allocation. PR1917465

- PR1959318
- IRB (Integrated Routing and Bridging) next-hop is not updated with the latest QoS (Quality-of-service) attributes after output-traffic-control-profile is deactivated/activated on ESI (Ethernet Segment Identifier) LAG. This results in traffic blackholing after deactivating/activating the output-traffic-control-profile on an ESI LAG interface in EVPN scenarios. PR1944234
- On all Junos OS Evolved platforms, optics modules intermittently reports extremely high false temperature values in 'show interfaces diagnostics optics' command. This can trigger false high-temperature alarms, leading to unnecessary optics shutdown.

  PR1937870
- A license process crash is observed on Junos OS Evolved platforms when AE (Aggregated Ethernet) interface configuration is deactivated and activated or modified. PR1926569

Infrastructure

- On Junos OS Evolved platforms, when fibd reaches the process limit on the maximum number of allowed open file descriptors, fibd is unable to open new file descriptors and track new reader processes. Given this, processes such as Jade, SSHd and Stunnel are not found by fibd, triggering the core files for these processes affecting device management and Tx/Rx packet process. The Soft Limit value within the "cat /proc/<PID>/limits" output command will show the maximum number of file descriptors that fibd can have opened. PR1860358
- This PR address the following error while performing commit on the PTX EVO platform: # commit error: Failed to get RE list This is caused by a kernel software defect. PR1909995

Interfaces and Chassis

- On ACX Series routers running Junos OS Evolved, using both SP style (Service Provider style) and EP style (Enterprise style) Layer 2 configurations on the same system isn't supported. If these configurations are mixed, trunk interfaces might not come up.

  PR1926153

Layer 2 Features

- An rpd core occurs on the backup Routing Engine after a switchover due to an assertion failure during the context of L2VPN VPLS auto-site processing. The failure is due to a mismatch between the auto-site claim-id/site-id and the local site-id associated with interfaces in the VPLS IFL repository for a given instance. This mismatch results in stale interface entries being referenced during auto-site processing.

PR1885690

Routing Protocols

- NGMVPN multicast traffic is dropped post changing the route distinguisher and route target on L3VPN instances. This issue is intermittent and occurs when RD and RT are changed in all routing instances while traffic is active. In some cases, the issue is observed only in specific instances. When the issue occurs, traffic drops for the inclusive tunnel, affecting all flows on the inclusive tunnel. Recovery may not occur unless external events, such as a route change, take place.

  PR1909265

- Telemetry sensors on ACX EVO to pull /network-instances/network-instance/protocols/protocol/isis/interfaces only report the first ISIS interface PR1925611

Subscriber Access Management

- On all Junos OS Evolved platforms, the error message **UI_SCHEMA_SEQUENCE_ERROR** is observed when device is restarted. There is no traffic impact due to this issue.

  PR1813456

User Interface and Configuration

- When using the `load override` command, committing a scaled configuration on top of an already committed scaled configuration may take longer to complete. During this time, other sessions cannot enter configure mode, causing those sessions to appear unresponsive. Use the `load update` command instead of the `load override` command to load the configuration before committing.

  PR1906880

- On all Junos OS Evolved ACX Series platforms, which do not verify the syntax and integrity of the saved rescue configuration file when performing a configuration check, rescue configuration check fails.

  PR1907238

Junos OS Evolved Release Notes for PTX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

Hardware

New PTX12008 router (PTX Series) (PTX12008)—The Juniper Networks® PTX12008 Router is a high capacity 22-rack unit (22-U) modular chassis that runs Junos OS Evolved. It has eight Flexible PIC Concentrator (FPC) slots. It provides distributed forwarding, and up to 43.2 terabit per second (Tbps) of bandwidth per FPC slot and up to 345.6 Tbps of bandwidth per chassis. The PTX12008 delivers significant port density and scale while optimally balancing high thermal tolerance and low power consumption.

The router is available in both base and premium configurations, with support for both AC/HVAC/HVDC and DC power supply modules (PSMs), and front-to-back airflow. The PTX12008 supports fully redundant power supply, cooling, and Routing Engines (control and management plane).

You can use the PTX12008 in core, peering, metro aggregation, infrastructure edge, AI data center networks, DC spine, and Data Center Interconnect or Data Center Edge deployments.

See PTX12008 Router Hardware Guide

Software Features supported on PTX12008 routers

Feature	Description
Chassis	Resiliency support for Packet Forwarding Engine. Support for powering on, powering off, or restarting Packet Forwarding Engine. [See Resiliency, request chassis fpc, action (chassis error).] PIC online and offline is supported using the command `request chassis pic online/offline` [See request chassis pic.] **Support for SIB management and resiliency**—We support device initialization, management, monitoring and resiliency for Switch Interface Board (SIB). Use `show chassis sib` command to view the SIB status. Use the `request chassis sib slot` *`slot number`* `online` and `request chassis sib slot` *`slot number`* `offline` commands to bring the SIB online and offline. Chassis management features include temperature sensors, voltage sensors, and LED management. Use `show system firmware` to display the firmware version and `request system firmware` to upgrade the FRU firmware. [See Platform Resiliency and Installing and Upgrading Firmware.] **Support for fabric management**—Support for fabric management features including initialization, configuration and management of the Switch Interface Board (SIB), ASIC and additional fabric components. Fault detection and auto-heal are supported. Use `set chassis (sib | fpc) slot number power (on | off)` to turn the FPC or SIB on or off. Use `set chassis fabric event reachability-fault degraded error-threshold degradation percentage` command to configure the percentage of fabric degradation at which recovery action is taken. Use the `show chassis fabric (fpcs | sibs) extended` to display the state of the electric switch fabric link between the SIBs and FPCs. [See Fabric Plane Management, power, reachability-fault, show chassis fabric fpcs, and show chassis fabric sibs.] **Platform resiliency support**—We support platform resiliency features for RCB, PSM, Fan, CPU, BIOS, Memory, FPGA, ASIC, Optics, BITS and additional hardware components. Resiliency represents the system's ability to anticipate, withstand, and rapidly recover from disruptions while maintaining critical functionality. This capability monitors the health status of various device components and handles faults by taking necessary actions. [See Platform Resiliency.] **Fabric resiliency support**—You can detect, classify, and remediate SIB–FPC fabric faults in Junos EVO platforms. The software monitors fabric components, logs minor and major errors and raises alarms. The software attempts automatic recovery of fatal conditions by restarting SIBs or FPCs, up to three attempts, and performs plane auto-healing on CCL_LINK_FAULT_ERROR. You can troubleshoot with `show system alarms`, `show system errors active detail`, `show chassis fabric topology`. You can manually recover by issuing `request chassis sib slot` *`slot number`* `restart` or `request chassis fpc slot` *`slot number`* `restart`.

(Continued)

Feature	Description
	[See Fabric Resiliency and Degradation.] **Chassis management support**—Chassis management features include FRU management, power budgeting, and cooling. Maintain PSM health and N+1 redundancy while regulating fans and alarms. Use the CLI command `set chassis fru-poweron-sequence` to configure poweron sequence for an FPC, the `set chassis fpc` *`slot`* `max-power` *`power number`* command to override the default power number for FPC, the `set chassis no-power-budget` command to disable power management. Monitor environments and temperature thresholds and view shelf capacity using the commands `show chassis environment`, `show chassis environment fpc`, `show chassis temperature-thresholds`, `show chassis hardware`. [See Chassis Management.] **Software support for RCB**—Software support for the following features for Routing and Control Board (RCB): • Non-Volatile Memory Express (NVME) storage • Single and dual RCB • FRU health and fault monitoring • Powering Up • Power, voltage, humidity, and temperature sensors monitoring • Firmware upgrade • CLI commands for RCB configuration and show command outputs • Control board Ethernet switch port information • Secure boot • RCB Resiliency [See Chassis Management.]

(Continued)

Feature	Description

Class of service (CoS)	• 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.
	
	[See CoS Features and Limitations on PTX Series Routers.]
	
	• Support for priority-based flow control (PFC) watchdog, which detects and mitigates PFC pause storms received for PFC-enabled queues.
	
	[See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved and PFC Watchdog.]
	
	• 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).
	
	[See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic and CoS Explicit Congestion Notification.]
	
	• Support for VOQ queue-depth monitoring to report peak queue length for a given physical interface for each Packet Forwarding Engine.
	
	[See VOQ Queue-depth Monitoring.]
	
	• Support for export of physical interface queue statistics to an outside collector. Each physical interface has eight queues. The following counters are exported as part of this sensor for all configured physical interfaces:
	
	  • Transmitted packets and transmitted bytes
	
	  • Red drop packets and bytes
	
	  • Tail drop packets and bytes
	
	This feature includes zero suppression support. It does not include support for summed-up counters on aggregated Ethernet (ae) interfaces.
	
	[See sensor (Junos Telemetry Interface) and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]
	
	• 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

(Continued)

Feature	Description
	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. [See hierarchical-scheduler and Hierarchical Class of Service in ACX Series Routers.] • Support for policy maps. The device supports the following types of packet marking for policy maps: INET-Precedence, DSCP, IEEE 802.1p, and IEEE 802.1ad. [See policy-map (Class of Service) and Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps.] • Support for on-chip queue buffer for PFC-enabled queues. The device views and installs a priority-based flow control (PFC)-enabled queue with a buffer-size less than 450 microseconds as a PFC-enabled on-chip queue. When a queue is in PFC on-chip mode, the entire virtual output queue (VOQ) buffer is always on-chip and is not scaled based on bandwidth usage. [See buffer-size (Schedulers).] • Support for on-chip buffer (OCB) with explicit congestion notification (ECN). Keeping the queue on-chip eliminates the need to access slower external memory. This approach is ideal for low-latency, high-throughput applications. [See CoS Explicit Congestion Notification (ECN).] • Per-queue accounting of explicit congestion notification (ECN) packets. You can track ECN-marked traffic with ECN-CE packet and byte counters in queue statistics and telemetry, helping you monitor congestion and validate ECN behavior. [See ECN Packets per Queue and show interfaces queue.]

(Continued)

Feature	Description
EVPN	• Support for EVPN-VXLAN Layer 2 (L2) gateways and Layer 3 (L3) gateways with EVPN Type 5 routes. Support includes the following: • EVPN-VXLAN L2 and L3 gateway operations in edge-routed bridging (ERB) and centrally routed bridging (CRB) overlay fabrics • MAC-VRF EVPN instances with VLAN-based, VLAN-bundle, and VLAN-aware bundle service types • Pure EVPN Type 5 route VRF model • IPv4 and IPv6 IRB traffic • Q-in-Q VNI mapping • Overlapping VLAN IDs across MAC-VRF instances • ECMP underlay reachability • Active-active multihoming • ARP and NDP suppression • Multicast IRB without IGMP or MLD snooping • CoS support on EVPN-VXLAN tunnel interfaces [See EVPN User Guide.] • Support for static VXLAN Layer 2 gateway deployments [See Static VXLAN, remote-vtep-list, and static-remote-vtep-list.] • Support for overlay and CE-IP ping and traceroute operations in EVPN-VXLAN overlays. [See Understanding Overlay ping and Overlay traceroute Packet Support.] • Support for lightweight PE-CE loop detection in EVPN-VXLAN fabrics. [See EVPN-VXLAN Lightweight Leaf to Server Loop Detection.] • Support for ECN copy on EVPN-VXLAN tunnels. [See CoS Support on EVPN VXLANs.]

(Continued)

Feature	Description
	• Support for EVPN-VXLAN fabrics with an IPv6 underlay. [See EVPN-VXLAN with an IPv6 Underlay and Example: Configure an IPv6 Underlay for Layer 2 VXLAN Gateway Leaf Devices.] • Support for EVPN-VXLAN Type 2 and Type 5 stitching. [See interconnect.] • Support for interconnecting EVPN-VXLAN data centers with EVPN-MPLS in WAN using gateway nodes.[See Overview of EVPN-VXLAN Interconnect through EVPN-MPLS WAN Using Gateways.] • Support for DSCP copy on EVPN-VXLAN tunnels. Routers originating EVPN-VXLAN tunnels copy DSCP bits from inner packet headers to outer packet headers during encapsulation. [See CoS Support on EVPN VXLANs.]
Firewall filter	• Support for firewall filters • 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 firewall filter actions such as `count`, `accept`, `drop`, `reject`, `sample`, and `police` on packets classified into SCU/DCU classes. You can configure both source-class and destination-class based filters at the same time. [See firewall filter and Understanding the Use of Policers in Firewall Filters, and Guidelines for Configuring SCU.]
High availability (HA) and resiliency	Support for graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) [See Understanding Graceful Routing Switchover and Understanding Nonstop Active Routing.]

(Continued)

Feature	Description
Interfaces	Support for Neighbor Discovery Protocol (NDP) proxy and duplicate address detection (DAD) proxy for multiple interfaces for interface-restricted mode and interface-unrestricted mode. [See NDP Proxy and DAD Proxy.] Hold up/hold down timers and interface dampening is supported. Commands: `set interface interface name> hold-time`, `set interface <interface-name> damping` [See Damping Interfaces and damping (Interfaces).] Support for Digital Optical Monitoring using the command `show interfaces diagnostics optics interface-name`. [See show interfaces diagnostics optics.] Support for performance monitoring of Optics related parameters using the command `show interfaces transport pm optics all interface-name`. [See show interfaces transport pm.] Support for media access control (MAC) accounting for source and destination MAC addresses learnt (both statically and dynamically) over Layer 3 physical interfaces and aggregated Ethernet interfaces. [See mac-learn-enable and show interfaces mac-database .]

(Continued)

Feature	Description
IP and infrastructure	Support for IP and infrastructure features: • Telemetry support for Packet Forwarding Engine and device monitoring sensors. Use Junos telemetry to stream Packet Forwarding Engine (PFE), network processing unit (NPU), and device monitoring statistics to an external data collector. Subscribe to telemetry sensor paths to obtain state information, performance metrics, and device statistics. Supported sensors include PFE statistics, NPU memory utilization, CPU utilization, Adaptive Forwarding Table (AFT), and device monitoring sensors. Stream device-level data, such as interface status, forwarding status, and trap statistics. Use telemetry data to analyze NPU resources, route and next-hop utilization, and packet and drop metrics. Native and OpenConfig telemetry sensors are supported. View the supported sensor paths in the Junos YANG Data Model Explorer. [See Junos YANG Data Model Explorer.] • Support for distributed denial-of-service (DDoS) protection and higher DDoS bandwidth for Layer 2 and Layer 3 protocols. [See protocols (DDoS) (ACX Series, PTX Series, and QFX Series) and show ddos-protection protocols isis.] • Support for load balancing. The following load balancing features are supported: • Up to 128-way equal-cost multipath (ECMP) • Weighted Equal Cost Multiple Path (WECMP) • Load balancing on aggregated Ethernet (AE) interfaces • Load balancing on mixed-rate aggregated Ethernet interfaces • Resilient hashing • Enhanced hash key • Native equal-cost LAG with adaptive load balancing (ALB) Fast reroute (FRR) is supported to minimize traffic loss. [See show forwarding-options load-balance, enhanced-hash-key, Adaptive Load Balancing, and Configuring Per-Packet Load Balancing.] • Support for classification override for unicast traffic.

(Continued)

Feature	Description
	[See CoS Features and Limitations on PTX Series Routers, Overriding the Input Classification, and classification-override.] • Support for CoS-based forwarding (CBF), unicast reverse-path forwarding (unicast RPF) check for both IPv4 and IPv6 addresses, source class usage (SCU) and destination class usage (DCU), accounting and IPv6 transit packet accounting. [See CoS-Based Forwarding, Example: Configuring Unicast RPF (On a Router), Understanding Source Class Usage and Destination Class Usage Options, and Configuring IPv4 and IPv6 Accounting.] • Support for passive monitoring, including support for passive monitoring on MPLS-encapsulated packets. [See passive Monitoring, and passive-monitor-mode .]

(Continued)

Feature	Description
Junos telemetry	Export telemetry data to external collectors. Use Junos telemetry to stream device performance and management data from PTX12008 routers to external data collectors. This information facilitates enhanced performance management and monitoring. To configure a telemetry subscription to deliver data to your collector, see Telemetry Data Subscriptions over gNMI and Establish a Dial-in Telemetry Connection. Subscribe to the following resource paths for telemetry data: • /junos/system/cmerror/configuration/ • /junos/system/cmerror/counters/ • /junos/system/linecard/environment/ • /junos/system/linecard/optics/ • /junos/system/linecard/interface/ • /interfaces/interface/state/ • /interfaces/interface/ethernet/state/ • /interfaces/interface/ethernet/state/counters/ • /interfaces/interface/ethernet/state/counters/distribution/ • /state/interfaces/interface/ethernet/fec/ • /state/interfaces/interface/ethernet/counters/ • /components/component/state/ • /components/component/properties/ • /components/component/subcomponents/ • /components/component/port/ • /components/component/power-supply/ • /components/component/fan/ • /components/component/integrated-circuit/pipeline-counters/errors/

(Continued)

Feature	Description
	• /components/component/integrated-circuit/backplane-facing-capacity/state/ • /components/component/integrated-circuit/pipeline-counters/packet/fabric-block/ • /components/component/integrated-circuit/pipeline-counters/errors/fabric-block/ • /components/component/transceiver/state/ • /components/component/transceiver/physical-channels/ • /components/component/optical-channel/config/ • /components/component/optical-channel/state/ • /system/alarms/alarm/ • /state/chassis/modules/module/health/ • /components/component/linecard/config/ • /components/component/linecard/state/ • /components/component/chassis/utilization/resources/ • /components/component/healthz/ • /components/component/state/temperature/ View the complete list of supported sensors in the Junos YANG Data Model Explorer. [See Establish a Dial-in Telemetry Connection and Junos YANG Data Model Explorer.]
	Interface-level telemetry for IPv4 and IPv6 transit traffic. Junos telemetry streams statistics and rates for interface-level (IFL) IPv4 and IPv6 transit traffic on logical and aggregated Ethernet interfaces. Prerequisite: Configure the device to enable route-based accounting and telemetry streaming using the following CLI commands: `set forwarding-options family inet route-accounting set forwarding-options family inet6 route-accounting set services analytics zero-suppression no-zero-suppression` The following native resource paths are supported:

(Continued)

Feature	Description
	• /junos/system/linecard/interface/logical/family/ipv4/usage/ • /junos/system/linecard/interface/logical/family/ipv6/usage/ The following counters for ingress or egress traffic are streamed (per IFL): **IPv4 counters** • /interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/in-octets • /interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/in-pkts • /interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/out-octets • /interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/out-pkts **IPv6 counters** • /interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/in-octets • /interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/in-pkts • /interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/out-octets • /interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/out-pkts Use these counters to monitor IPv4 and IPv6 transit traffic at a granular interface level. For a complete list of sensor paths supported by the device, see Junos YANG Data Model Explorer.

(Continued)

Feature	Description

Layer 2 features	Support for the following Layer 2 features:
	• Layer 2 basic learning, bridging, and flooding:
	  • Enterprise-style bridging (support both trunk and access mode)
	  • Service provider-style bridging (also known as subinterface 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

(Continued)

Feature	Description

Note: You cannot configure enterprise-style Layer 2 logical interfaces under the flexible-ethernet-services encapsulation.

- Virtual switch
- Persistent MAC learning (sticky MAC)
- Service provider bridging for multiple logical interfaces on the same physical interface that are part of the same bridge domain and Ethernet bridge encapsulation

[See Layer 2 Bridging, Address Learning, and Forwarding User Guide.]

- Support for IRB interfaces that enables routing of Layer 3 traffic between a bridge domain and another routed interface.

  Supported IRB features include:

  - 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 addresses and statistics
  - IRB Layer 3 multicast support with flooding only
  - Address family support for IPv4 and IPv6, and support for IPv4 MTUs and IPv6 MTUs with different MTU values
  - IRB interface in VRF routing instances
  - Directed subnet broadcast support with IRB

    [See Integrated Routing and Bridging.]

  Support for VRRP on IRB. ISSU, proxy ARP, and MC-LAG are not supported with VRRP.

  [See Understanding VRRP.]

  - Support for IRB in VRF instance

    [See Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port and show interfaces irb.]

(Continued)

Feature	Description
	• Support for interface MAC limiting [See Configuring MAC Limiting and packet-action.] • Support for Q-in-Q tunneling and VLAN translation [See Configuring Q-in-Q Tunneling and Q-in-Q Tunneling and VLAN Translation.] • Support for LLDP, xSTP, and BPDU block/filter LLDP is supported, including on em0 interfaces. Disabling of LLDP time, length, and value (TLV) messages is also supported. Support for RSTP, VSTP, MSTP, STP, root protection for STP, concurrent configuration of RSTP and VSTP. virtual switch, and BPDU protection for spanning-tree protocols. [See Configuring STP.] Support for BPDU protection for EVPN-VXLAN. [See Understanding BPDU Protection for EVPN-VXLAN.]

(Continued)

Feature	Description

Layer 3 features	Support for the following Layer 3 features:
	• You can configure the following Layer 3 forwarding features on the router:
	  • IPv4 and IPv6, MPLS, LAG, ECMP, MTU checks, ICMP, OSPF, IS-IS, ARP, NDP, BGP, BFD, LACP, LDP, RSVP, LLDP, VRF-lite, TTL expiry, IP options, IP fragmentation, and DDoS
	• 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.]
	• BFD support, including:
	  • Distributed BFD and BFD-triggered local repair (BFD authentication is not supported.)
	  • Independent micro-BFD sessions enabled on a per-member link basis for a LAG bundle
	  • Inline BFD
	  [See Understanding BFD.]
	• Support for BGP flowspec signaling. The following match conditions are not supported:
	  • ICMP codes alone [`inet/inet6`]
	  • Source/destination prefix with offset for `inet6`
	  • Flow label for inet6 fragment [for `inet6`]
	  Junos OS Evolved running on the PTX12008 router doesn't support the traffic marking action.
	  [See Understanding BGP Flow Routes for Traffic Filtering.]

(Continued)

Feature	Description

Layer 3 tunnels	Support for Layer 3 tunnels:
	• Support for FTI tunnels, which are interface-based tunnels configured by using FTIs:
	  • FTI-based encapsulation and de-encapsulation of IPv4 and IPv6 packets.
	  • UDP tunnel encapsulation on FTIs.
	  • GRE tunnel encapsulation on FTI by using a loopback-based interface.
	  • FTI support for GRE tunnel de-encapsulation.
	  [See Tunnel and Encryption Services Interfaces User Guide for Routing Devices], encapsulation (interfaces-fti), Components of Filter-Based Tunneling Across IPv4 Networks, tunnel-end-point, Understanding Filter-Based Tunneling Across IPv4 Networks, Firewall Filter Terminating Actions.]
	• Support for configuring MPLS protocols over FTI tunnels, thereby transporting MPLS packets over IP networks that do not support MPLS.
	  [See Configuring Flexible Tunnel Interfaces.]
	• Support for IP-over-IP encapsulation to facilitate IP overlay construction over an IP transport network. The following are not supported:
	  • Dynamic tunnel de-encapsulation operation
	  • Next-hop-based statistics for dynamic tunnels
	  • IP fragmentation at tunnel start point and path MTU discovery for IPv4/IPv6
	  [See Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation]
	• Filter-based GRE encapsulation and de-encapsulation and filter-based MPLS-in-UDP de-encapsulation.
	  [See Configuring a Firewall Filter to De-encapsulate GRE traffic]
	• Support for output filter-based GRE.
	  [See Configuring a Filter to De-Encapsulate GRE Traffic, and decapsulate (Firewall Filter).]
	• Redistribution of IPv4 routes with IPv6 next hop into BGP.
	  [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]

(Continued)

Feature	Description
MACsec	Support for Media Access Control security (MACsec) bounded delay protection [See Configuring Bounded Delay Protection.] Support for MACsec on physical and logical interfaces. Support for dynamic power management of MACsec block in a port group. [See Understanding Media Access control Security (MACsec.]

(Continued)

Feature	Description
MPLS	Support for MPLS features including: • Static and dynamic LSP support • Support for both LDP and RSVP • MPLS Traffic Engineering • FRR (link protection) • FRR (node protection) • Penultimate-hop popping (PHP) • MPLS with ECMP/WECMP • MPLS over LAG • MPLS OAM LSP ping • MPLS traceroute • MTU and fragmentation • LSP statistics using telemetry • No decrement/No propagate TTL • LDP over RSVP • Entropy labels • MPLS over tagged subinterfaces • Explicit NULL • Auto-bandwidth • Container LSPs [See MPLS Applications User Guide .]

(Continued)

Feature	Description

	• Support for MPLS Layer 3 VPN to enable a single core network to transport packets from multiple customers, including:
	  • IPv4 L3VPN with label allocation per policy next hop, per prefix, or per VRF instance.
	  • IPv6 L3VPN over IPv4 core (6VPE) with label allocation policy per next hop, per prefix, per VRF
	  • IPv6 over IPv4 core (6PE) (with and without explicit null)
	  • Ultimate-hop popping of LSP
	  • Ping and traceroute support
	  • IPv4 and IPv6 BGP Labeled Unicast (BGP-LU)
	  • Inter-AS (support for options A, B, and C)
	  • Carrier supporting Carrier (CsC) support
	  [See Layer 3 VPNs Feature Guide for Routing Devices.]
	• Support for segment routing over MPLS (SR-MPLS), including:
	  • IS-IS, OSPF, and static segment routing with support for adjacency SID, node SID and prefix SID
	  [See Understanding Source Packet Routing in Networking (SPRING).]
	  • Operation, Administration, and Maintenance (OAM)
	  [See Operations and Maintenance (SR-MPLS).]
	  • Segment Routing interoperability with RSVP and LDP
	  • Segment routing Topology Independent Loop-Free Alternate (TI-LFA)
	  [See Topology-Independent Loop-Free Alternates (TI-LFA)]
	  • Segment routing telemetry statistics per SID for ingress and egress MPLS traffic
	  [See Overview of Segment Routing Statistics and Junos YANG Data Model Explorer.]

(Continued)

Feature	Description
	• Segment routing–traffic engineering (SR-TE): • BGP SR-TE, static SR-TE, and dynamic SR-TE [See Color-Based Mapping of VPN Services for SR-MPLS Segment Routing LSPs • SR-TE telemetry statistics per tunnel and per path [See Enabling Streaming of Telemetry Sensor Information for SR-TE policies • Periodic and on-change streaming of SR-TE colored policy telemetry data, per-segment-list and per-prefix statistics for colored tunnel [See Enabling Streaming of Telemetry Sensor Information for SR-TE policies.] • L3VPN, 6VPE, and 6PE services over colored and uncolored SR-TE LSPs [See Color-Based Mapping of VPN Services for SR-MPLS Segment Routing LSPs • Weighted ECMP (WECMP) and hardware ECMP • Support for scaled-up static and BGP segment routing policies, where each policy contains an average of eight segment routing paths with 16 labels per path. [See source-packet-routing, Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface, and Understanding OpenConfig and gRPC on Junos Telemetry Interface.] • Support for Layer 2 VPN features, including: • Transport of Layer 2 frames over MPLS (LDP signaling) • Layer 2 VPNs over tunnels (BGP signaling) • Simple Ethernet and VLAN-based cross-connect • Local and remote switching • Ethernet and VLAN CCC • Single-tagged CCC logical interfaces

(Continued)

Feature	Description
	• 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, and VCCV with type 2 and type 3 • Support for VLAN tag manipulation: pop, push, and swap [See Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough.] • Support for virtual circuit connection verification (VCCV) protocol [See Configuring BFD for VCCV for Layer 2 Circuits, BFD Support for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS, MPLS Pseudowires Configuration, show ldp database,and show route instance.] Support for inner VLAN transparency [See Understanding VLAN Manipulation (Normalization and VLAN Mapping) on Ethernet Services, Layer 2 VPNs and VPLS User Guide for Routing Devices and TCC Overview.] • Support for flow-aware transport for pseudowires (FAT-PW) labels on ingress routers, with parsing that includes all the payload fields in the hash calculation. These flow labels are supported: • Layer 2 circuit, LDP-signaled pseudowires • L2VPN, BGP-signaled pseudowires • L2VPN with FEC129 (BGP autodiscovery) [See flow-label-transmit, flow-label-receive, and show l2circuit connections.] • Support for VLAN ID lists for Layer 2 circuits [See vlan-id-list (Ethernet VLAN Circuit) and Configuring VLAN Identifiers for VLANs and VPLS Routing Instances.] • Support for redistribution of IPv4 routes with IPv6 next hop into BGP per RFC 5549 [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]

(Continued)

Feature	Description
	• Support for color-based mapping of VPN services over SR-TE [See Color-Based Mapping of VPN Services.] [See Layer 3 VPNs Feature Guide for Routing Devices, Carrier-of-Carrier VPNs, Layer 2 VPNs and VPLS Feature Guide for Routing Devices, and Translational Cross-Connect (TCC)]

(Continued)

Feature	Description
Multicast	Support for multicast features: • IPv4 and IPv6 multicast support including MSDP, support for PIM sparse mode (PIM-SM) as the first-hop router (FHR) or last-hop router (LHR), and for anycast, static, or local rendezvous point (RP). • Support for multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows. MoFRR is supported for PIM sparse mode (SM) and source-specific multicast (SSM) modes only. Support does not extend to multipoint LDP-based MoFRR. [See Understanding Multicast-Only Fast Reroute.] • Support for bidirectional PIM for multicast traffic. [See pim-snooping.] • Support for RSVP-based and multicast LDP-based point-to-multipoint (P2MP) LSPs with graceful restart. [See Point-to-Multipoint LSPs Overview.] • Support for MPLS features P2MP ping and P2MP LSPs traceroute. MPLS ping and traceroute provide the mechanism to detect data-plane failure and isolate faults in the MPLS network. The traceroute or ping is initiated to validate LSP paths on P2MP. [See MPLS Applications User Guide .] • Optimized fast branch updates. The method of making fast-branch updates to a multicast replication tree has been refined. Now, any membership changes in the tree trigger fast make-before-break (FMBB) re-optimization of the tree and ensure that there is no traffic loss. [See Multicast Shortest-Path Tree.] • Multicast support for next-generation MVPN (NG-MVPN) including 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 such as auto-rendezvous point (RP), bootstrap router (BSR), and embedded RP. • Support for automatic ingress LSP policing in P2MP LSPs

(Continued)

Feature	Description
	[See Point-to-Multipoint LSP Configuration, MPLS Applications User Guide, Multicast Shortest-Path Tree Multiprotocol BGP MVPNs Overview, Understanding Next-Generation MVPN Concepts, Understanding Next-Generation MVPN Control Plane, and Configuring Automatic Policers.]
Network Management	• Support for NETCONF Call Home [See NETCONF Call Home.] • Support for syslog over TCP and TLS • Support for configuring a source IP address and routing instance for legacy remote procedure call (gRPC) service dial-out connections [See Configure gRPC service for Dial-out over TCP Connections.] • Support for loading set-formatted and XML-based configuration files when your device provisions for zero-touch provisioning (ZTP). [See Zero Touch Provisioning.] • Support for aggregating telemetry data directly on the device to provide unified statistics for some of the distributed telemetry data originating from the Packet forwarding Engine.

(Continued)

Feature	Description

Port Mirroring	• Local port mirroring support including:
	  • Interface filter on ingress and egress
	  • Forwarding table filter (FTF) on ingress
	  • Families `inet` and `inet6`
	  • Aggregated Ethernet interfaces at both ingress and egress
	  [See Understanding Port Mirroring and Analyzers.]
	• Remote port mirroring with ToS or DSCP settings.
	  [See instance (Port Mirroring).]
	• Support for additional family `any` in port mirroring. You can configure family `any` for local port mirroring and remote port mirroring. You can use the `family any` configuration option to configure the families any, ccc, ethernet-switching, or mpls.
	  Note:
	  You use the family any configuration option to process all 4 families.
	  The following configuration statements are no longer part of the port mirroring configuration on PTX Series devices:
	  • `next-hop` for `family any`
	  • `family vpls`
	  • `no-filter-check`
	  • `hosted-service`
	  • `server-profile`
	  [See port-mirroring.]
	• Support for EVPN-VXLAN filter and port mirroring based on VNI match conditions. This feature supports redirecting traffic to a global port-mirroring instance.
	  [See Firewall Filter Match Conditions and Actions (PTX Series Routers).]

(Continued)

Feature	Description
	• Support for on-device packet capture • [See Example: Configuring Remote Port Mirroring on PTX Routers, port-mirroring, See Example: Configure Port Mirroring with Family any and a Firewall Filter, Configuring Port Mirroring and Analyzers, and On-Device Packet Capture.]

(Continued)

Feature	Description
Protocols	Support for Virtual Router Redundancy Protocol (VRRP), Link Aggregation Group (LAG), Link Aggregation Control Protocol (LACP), and Link Fault Management (LFM): • VRRP The following features are not supported for VRRP on Junos OS Evolved: • ISSU • Proxy ARP • MC-LAG • Distribution support on aggregated Ethernet interfaces • IRB • Inline delegation [See Understanding VRRP.] • LAG (aggregated Ethernet) [See Ethernet Link Aggregation] • LACP link protection—We support the 1:1 and N:N link-protection on LACP. [See link-protection.] • Support for link fault management (LFM)—We support IEEE 802.3ah OAM LFM to monitor point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The following LFM features are supported: • Link discovery with active and passive modes • Detecting loss of continuity • Remote loopback • Loopback tracking • Action profile

(Continued)

Feature	Description
	• GRES and non-graceful Routing Engine switchover [See Introduction to OAM Link Fault Management (LFM]
Services Applications	Inline monitoring services support for packet mirroring with metadata and Juniper Resiliency Interface support. [See Inline Monitoring Services Configuration and Juniper Resiliency Interface.] Inline active flow monitoring support, including IPFIX and version 9 data templates for IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPV6 traffic, support for multiple collectors, and support from the mgmt_junos VRF.. [See Understand Inline Active Flow Monitoring.]

(Continued)

Feature	Description
sFlow technology	• Support for sFlow monitoring technology for high-speed switched or routed networks. • Support for sFlow on interfaces configured with the `mpls family`. • Support for ingress and egress sFlow functionalities for transit nodes for IPv4-in-IPv4, IPv6-in-IPv4, and regular IPv4/IPv6 traffic. The IP-in-IP encapsulated packet can transit through the transit-only device without any change, be de-encapsulated and forwarded, or be de-encapsulated and encapsulated and forwarded, based on the next-hop configuration. The packet can traverse multiple VRF instances while getting forwarded. The functionality also caters to sFlow capabilities for ECMP traffic. • sFlow technology support for exporting extended tunnel egress structure for traffic entering IPv4 or IPv6 GRE tunnels. These attributes provide information about the GRE tunnel into which a packet entering the device is encapsulated. The feature is supported only when sFlow is enabled in the ingress direction wherein firewall-based GRE happens on IPv4 or IPv6 packets. You can use the feature in the following traffic scenarios: • Incoming IPv4 or IPv6 traffic that undergoes IPv4 GRE • Incoming IPv4 or IPv6 traffic that undergoes IPv6 GRE • Support for 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. • Ingress and egress sFlow support for Layer2 traffic. Use sFlow to monitor ingress or egress data and control traffic forwarded through L2 logical interfaces associated with a bridge domain. This feature supports both enterprise-style and service provider-style L2 configurations. [See sFlow Technology Overview.]

(Continued)

Feature	Description
Software install and upgrade	Support for zero-touch provisioning (ZTP). [See Zero Touch Provisioning.] Support for secure zero-touch provisioning (SZTP). [See Secure Zero Touch Provisioning.] Firmware upgrade support. [Seerequest system firmware upgrade (Junos OS Evolved).] Secure boot and common BIOS support--The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms.[See Junos OS Evolved Overview and request system firmware upgrade (Junos OS Evolved).]

[See PTX12008 Router Hardware Guide.]

PTX10002-60MR router (PTX Series) (PTX10002-60MR)—PTX10002-60MR is a closed-door, 2 RU, high-density platform with the below port specifications:

- 48 QSFP28 (Quad Small Form-factor Pluggable 28) ports are high-density, hot-pluggable interfaces designed for 100 Gbps port speeds. They utilize 4 lanes, each carrying up to 28 Gbps (often 25 Gbps), and are backward compatible with 40Gbps QSFP+ modules.
- 12 QSFP112 DD ports - High-density, 8-lane interface designed for 800G networking, supporting 100G PAM4 per lane. These ports offer backward compatibility with 400G, 200G, and 100G modules.
- The chassis is designed with 1+1 PSU (power supply unit) redundancy and supports both AC and DC SKUs. The PSUs and Fans will plug in from the rear side.

To install the PTX10002-60MR router and perform initial configuration, routine maintenance, and troubleshooting, see the juniper.net/documentation/us/en/hardware/ptx10002-60mr. See Feature Explorer for the complete list of features for any platform.

Table 1: PTX10002-60MR Feature Support

Feature	Description
Chassis	• Chassis management support for PSU and fans. • Low power mode is not supported. • Support for front-to-back (AFO) front airflow. • Support for Main Board Sensor management and LED management. • Packet Forwarding Engine resiliency. We provide resiliency feature support for the Packet Forwarding Engine, which enables the system to detect, report, and take action on Packet Forwarding Engine faults. Actions are taken based on the default configuration or a user configuration available for the errors. [See Infrastructure resiliency: PFE and hardware resiliency.] • Resiliency feature support for platform components such as PSU, fan, and interfaces. The resiliency feature tracks component health, notifies users about errors, and performs corrective actions to restore normal operation based on error severity. [See Platform Resiliency.] • Routing Engine and Control Board resiliency support. Provides fault-handling capabilities and indicates errors using LEDs. • Supported transceivers, optical interfaces, and DAC cables. Select your product in the Hardware Compatibility Toolto view the 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.

(Continued)

Feature	Description
Class of service	• 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, buffer management, and drop profiles. [See CoS Features and Limitations on PTX Series Routers | Junos OS | Juniper Networks.] • Support for priority-based flow control (PFC) watchdog, which detects and mitigates PFC pause storms received for PFC-enabled queues. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved and PFC Watchdog.] • Hierarchical CoS support: The router supports up to four levels of scheduling on an interface (physical interfaces, logical interface sets, logical interfaces, and queues). [See Hierarchical Class of Service Overview.] • Support for priority-based flow control (PFC) at Layer 3 for untagged traffic and explicit congestion notification (ECN). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic and CoS Explicit Congestion Notification.]
Dynamic Host Configuration Protocol (DHCP)	Support for DHCPv4 relay agent and DHCPv6 relay agent, including: • DHCP relay: Layer 3 (L3) interfaces • DHCP relay: Option 82 for Layer 2 VLANs • DHCP relay: Option 82 for L3 interfaces [DHCP Relay Agent.]
EVPN	• Virtual private wire service (VPWS) with EVPN signaling mechanisms and flexible cross-connect support. [See Overview of VPWS with EVPN Signaling Mechanisms] • Support for EVPN-MPLS L2 and L3 features. [See EVPN Overview.] • Support for EVPN-VXLAN Layer 2 (L2) gateways and Layer 3 (L3) gateways with EVPN Type 5 routes. [See EVPN User Guide.] • Support for ping and traceroute for EVPN-VXLAN. [See Understanding Overlay ping and traceroute Packet Support.] • Support for Static VXLAN (L2 gateway). [See Static VXLAN.]

(Continued)

Feature	Description
Firewall filter and policer	• Firewall filter support. [See Firewall filter support.] • Policer support. [See Traffic Policers User Guide.]
High availability	• Support for VRRP. [See Understanding VRRP.]
Infrastructure	Support for the following IP and Infrastructure features: • Junos telemetry for exporting operational data from Junos devices to external collectors for real-time monitoring, analytics, and troubleshooting. [See Junos YANG Data Model Explorer.] • Support for distributed denial-of-service (DDoS), IS-IS classification, and higher DDoS bandwidth for Layer 2 and Layer 3 protocols. [See show ddos-protection protocols isisand protocols (DDoS).] • Support for load balancing. [See enhanced-hash-key.]
Layer 2 features	• Multichassis link aggregation groups (MC-LAGs) support for active/standby Mode. [See Understanding Multichassis Link Aggregation Groups.] • Support for the Layer 2 basic learning, bridging, and flooding features. [See Layer 2 Bridging, Address Learning, and Forwarding User Guide.] • Support for integrated routing and bridging: [See Integrated Routing and Bridging.] • Support for interface MAC limit action. [See Configuring MAC Limiting and packet-action.] Support for the following protocols: • LAG (aggregated Ethernet) • LACP • LLDP

(Continued)

Feature	Description
Layer 3 features	BFD support, including: • Distributed BFD and BFD-triggered local repair. • Independent micro-BFD sessions enabled on a per-member link basis for a LAG bundle. • Inline BFD. [See Understanding BFD.]
MACsec	• Media Access Control Security (MACsec) support on physical interfaces and logical interfaces. [See Understanding Media Access control Security (MACsec).]

(Continued)

Feature	Description
MPLS	Support for MPLS fast reroute (FRR). [See Fast Reroute Overview.] Support for MPLS features, including: • Static and dynamic LSP support • Support for both LDP and RSVP • MPLS Traffic Engineering • Penultimate-hop Popping [PHP] • MPLS with ECMP/WECMP • MPLS over LAG • MPLS OAM LSP ping • MPLS traceroute • MTU and fragmentation • LSP Statistics through Telemetry • No Decrement/No Propagate TTL • LDP over RSVP • Entropy labels • MPLS over tagged subinterfaces • Explicit NULL • Auto-Bandwidth • Container LSPs Enhanced scaling for the following MPLS features: • RSVP transit LSPs with link and node protection

(Continued)

Feature	Description
	• RSVP ingress and egress LSPs with ultimate-hop popping (UHP) and penultimate-hop popping (PHP) • LDP-over-RSVP LSPs • Packet Forwarding Engine statistics • FRR and make before break (MBB) • Weighted ECMP • Ping and traceroute • Clone route • Transit statistics [See MPLS Applications User Guide.]
Multicast	• Support for MPLS features point-to-point (P2MP) ping and P2MP LSPs traceroute. [See MPLS Applications User Guide.] • Optimized fast branch updates. The method of making fast branch updates to a multicast replication tree has been refined. Any membership changes in the tree trigger fast make-before-break (FMBB) re-optimization of the tree and ensure that there is no traffic loss. [See Multicast Shortest-Path Tree.] • Support for RSVP-based and LDP-based 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. [See Point-to-Multipoint LSP Configuration.] • Support for MVPN. [See Multiprotocol BGP MVPNs Overview, Understanding Next-Generation MVPN Concepts, and Understanding Next-Generation MVPN Control Plane.] • MVPN BIER with MPLS encapsulation. The Bit Index Explicit Replication (BIER) architecture simplifies 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.] • Support for distributed and inline modes for connectivity fault management (CFM). [See Configure Connectivity Fault Management (CFM) | Junos OS | Juniper Networks.]

(Continued)

Feature	Description
Network management and monitoring	• Support for the sFlow technology, which is a monitoring technology for high-speed switched or routed networks. The sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. [See Understanding How to Use sFlow Technology for Network Monitoring.] • Local and Remote Port Mirroring Support. [See Understanding Port Mirroring and Analyzers.] • NDP proxy and DAD proxy support. [See NDP Proxy and DAD Proxy.
Segment Routing	• Support for SRv6 network programming in IS-IS. [See How to Enable SRv6 Network Programming in IS-IS Networks.] • Support for SRv6 network programming and Layer 3 Services over SRv6 in BGP. [See Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP.] • OAM ping support for SRv6 network programming. • Support for OAM ping/traceroute for micro-SIDs. • SRv6 support for static SR-TE policy. [See Understanding SR-TE Policy for SRv6 Tunnel.]
Services applications	• Inline active flow monitoring support, including IPFIX and version 9 data templates for IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPv6 traffic, support for multiple collectors, and support from the mgmt_junos VRF. [See Understand Inline Active Flow Monitoring.] • Inline monitoring services support for packet mirroring with metadata and Juniper Resiliency Interface (JRI) support. [See Inline Monitoring Services Configuration andJuniper Resiliency Interface.] • Support for additional RPCs for the gNOI certificate management (cert) service. Junos OS Evolved supports the following gRPC Network Operations Interface (gNOI) cert service RPCs: CanGenerateCSR()— Query if the target device can generate a certificate signing request (CSR) with the specified key type, key size, and certificate type. RevokeCertificates()—Revoke certificates on the target device. [See gNOI Certificate Management (Cert) Service.]

(Continued)

Feature	Description
Software installation and upgrade	• Support for zero-touch provisioning (ZTP) and secure zero-touch provisioning (SZTP). [See Zero Touch Provisioning.] • Secure boot and common BIOS support. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms. [See Junos OS Evolved Overview.] • Firmware upgrade support. [See request system firmware upgrade (Junos OS Evolved).]
Timing	• Support for Synchronous Ethernet timing, Synchronous Ethernet over LAG, and Timing SNMP and MIB (SYNCE). • Support for G.8275.1 profiles over both LAG and physical interfaces. • G.8275.1 hybrid mode with clear separation of PTPoE and Synchronous Ethernet clock domains. one-step timestamping mode, and PTPoE support. [See Precision Time Protocol (PTP) Overview and PTP over Ethernet Overview.]
Tunneling	Support for the following Packet Forwarding Engine tunnel features: • Support for flexible tunnel interface (FTI) - based encapsulation and de-encapsulation of IPv4 and IPv6 packets. • Support for configuring GRE, IP-in-IP (IPIP), and UDP tunnel encapsulation on FTIs. [See encapsulation (interfaces-fti).] • Support for configuring MPLS protocols over FTI tunnels, thereby transporting MPLS packets over IP networks that do not support MPLS. GRE and UDP tunnels support the MPLS protocol for both IPv4 and IPv6 traffic. [See Flexible Tunnel Interfaces Overview.] • Support for GRE tunnel encapsulation using loopback-based interface on FTIs. [See encapsulation (interfaces-fti).] • Support for GRE tunnel de-encapsulation on flexible tunnel interfaces (FTIs). [See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.]
VPLS	Support for VPLS. [See Introduction to VPLS.]

Support for PTX10K-LC1301-36DD line card with JNP10K-RE3-E RCB (PTX10004 and PTX10008) (PTX10004 and PTX10008)—We support PTX10K-LC1301-36DD line cards with JNP10K-RE3-E RCBs on PTX10004 and PTX10008 routers.

Class of Service (CoS)

Support for latency-based ECN with a configurable global threshold and per-queue or per-scheduler enablement (PTX10002-36QDD, PTX10002-60MR, PTX10004, PTX10008, and PTX12008)—You can use latency-based explicit congestion notification (ECN) to mark congestion based on instantaneous per-packet latency instead of buffer occupancy, reducing loss and optimizing utilization. Enable optional hardware latency tracking with `set class-of-service options latency-tracking`. Define a global ECN latency threshold with `set class-of-service options latency-ecn threshold` *`value`*, where the threshold value ranges from 1ns-100ms. Activate marking per queue or scheduler using `set class-of-service forwarding-classes class` *`class-name`* `queue-num` *`queue-number`* `latency-ecn` or `set class-of-service schedulers` *`scheduler-name`* `latency-ecn`. Latency-based ECN cannot coexist with buffer-based ECN on the same port or queue.

[See Latency-Based ECN.]

EVPN

EVPN-VXLAN cross-module forwarding consistency and hash verification using the CLI (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)—You can detect and diagnose EVPN-VXLAN forwarding inconsistencies across control, kernel, and forwarding modules to reduce traffic loss or duplication. Verify per-VLAN synchronization and errors with:

- `request system evpn forwarding vlans [detail | instance` *`name`* `| vlan` *`name`* `| module {mac-table|mac-ip-table}] [no-confirm]`
- `request system evpn hash-values vlans [instance` *`name`*`] [vlan` *`name`*`]`

Compute comparison digests and validate Type 2 route and next-hop consistency. These checks highlight missing programming elements or deleted bridge domains.

[See EVPN-VXLAN Forwarding Consistency and Hash Verification, request system evpn forwarding vlans, and request system evpn hash-values vlans.]

Interfaces

MAC address accounting for Gigabit interfaces (PTX10002-60MR)—We support MAC address accounting for the following interfaces:

- 10GbE/25GbE/40GbE

- 100G/2x100G
- 400G
- 800G

To enable MAC accounting, use the existing `mac-learn-enable`command under the *[edit interfaces interface-name* `ether-options` *ethernet-switch-profile]* or *[edit interfaces aex aggregated-ether-options ethernet-switch-profile]* hierarchy level.

[See show interfaces mac-database.]

Transit Traffic Statistics and Telemetry (PTX10002-60MR)—We provide IPv4 and IPv6 transit traffic statistics at the logical interface (IFL) level on PTX Series routers running Junos OS Evolved.

Track ingress and egress packets and bytes, along with real-time rates (pps and bps), for scalar (et) and aggregated Ethernet (AE) interfaces. Stream counters use Junos telemetry sensors for enhanced visibility and analytics. This feature supports include native IPv4 and IPv6 traffic but excludes encapsulated traffic types.

[See actual-transit-statistics (Dynamic Profiles).]

Port Speed and Channelization (PTX10002-60MR)—QSFP28 ports support speeds of 10G and 100G by default, with selective support for 4x25G, 4x10G, and 40G on specific port groups. QSFP112-DD ports support speeds up to 800G, including 8x100G channelization, without constraints.

[See Port Speed on PTX Series Routers.]

Line-rate performance with 100-byte packets (PTX10002-36QDD, PTX10004, and PTX10008)—We support line-rate performance with 100-byte packet sizes across all supported speed groups. The implementation optimizes packet processing for smaller packet sizes while maintaining system stability.

[See line-rate.]

Enhanced LAG scalability (PTX10002-36QDD, PTX10004, PTX10008, PTX10016, and PTX12008)—Increase Link Aggregation Group (LAG) scalability by configuring up to 256 member links per LAG. Use CLI and validation scripts to efficiently manage large-scale aggregated interfaces. The `show interfaces` command accurately displays LAG details at this scale, maintaining parity with smaller aggregated Ethernet (AE) configurations. This enhancement supports high-density deployments and optimizes interface utilization in large network environments.

[See Ethernet Link Aggregation.]

Interfaces and Chassis

Application selection enhancement (PTX10004 and PTX10008)—ZR or ZR-M optics advertises supported applications for a particular speed. You can select any of the applications that best meets your requirements. You can also switch between the applications.

[See Application Selection.]

Junos Telemetry

Support for genstate YANG data models (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Telemetry for BGP diagnostics via YANG model and sensors (PTX10004, PTX10008, and PTX10016)—You can stream BGP diagnostics telemetry using a dedicated YANG module to accelerate detection and remediation with less manual effort. Subscribe to sensors for global, group, neighbor, NLRI, and local RIB data through management interfaces such as NETCONF. Examples include session state, flap reasons, import and export evaluation triggers, and label and threshold metrics. You do not need to change the CLI. Certain recommendation and route fields emit static values.

[See Junos YANG Data Model Explorer.]

OpenConfig QoS queue statistics keyed by queue name rather than queue ID (PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)—You can stream OpenConfig QoS interface

queue statistics using a push-based telemetry sensor that exports counters by queue name rather than queue ID. Subscribe to <codeph>/qos/interfaces/interface/output/queues/queue[name=qname]/state</codeph> via gRPC to receive transmit and drop packet/octet counters aggregated on-box across forwarding cards. This streaming telemetry enables clearer monitoring and correlation of QoS performance without polling. Configure queue names to avoid system-generated identifiers; names beginning with "__" are restricted. You do not need new CLI commands.

[See Junos YANG Data Model Explorer.]

Periodic and event-based gNMI telemetry for OpenConfig storage error counters (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)—You can use this feature to monitor storage health and predict failures by streaming OpenConfig storage counters over gNMI. Perform proactive maintenance and reduce the risk of update or boot issues from storage degradation. Subscribe to <codeph>/components/component/storage/state/counters</codeph> for reallocated-sectors, offline-uncorrectable-sectors-count, and percentage-used when available using periodic or event-based threshold mode. To enable, configure gRPC/gNMI under <codeph>system services extension-service</codeph>; optionally load local certificates under <codeph>security certificates local</codeph>.

[See Junos YANG Data Model Explorer.]

Graceful restart for IS-IS and BGP via OpenConfig without native CLI dependency (PTX10001-36MR, PTX10008, and PTX10016)—Use OpenConfig to manage IS-IS and BGP graceful restart, reducing reliance on a native global option and simplifying vendor-neutral automation.

To configure, enable graceful restart and set a 30–300 second T3 timer with */network-instances/.../isis/global/graceful-restart/config/enabled* and *.../config/restart-time.*

Note: You do not need additional CLI commands. IS-IS graceful restart operates only in the default instance**.** Avoid partial protocol GR enabling. Native configurations take precedence in mixed mode.

[See Junos YANG Data Model Explorer.]

Interface Service API enables gRPC queries for interface state and state-change notifications for physical and LAG interfaces (PTX10002-36QDD, PTX10008, and PTX12008)—You can query and monitor physical and LAG interface state by using the Interface Serivce API RPCs described in the jnx_interfaces_service.proto proto definition file. Use the *InterfaceGetState* RPC with wildcards and the *DetailLevel* RPC to obtain admin/operational status, type, speed, MTU, and LAG membership. Use *InterfaceNotifyState* to stream add, change, and delete events for automation and faster troubleshooting. Enable the service with:

- *set system services extension-service request-response grpc ssl port*

- *set security certificates local mycert load-key-file*

- *set system services extension-service request-response grpc ssl address 0.0.0.0*

- *set system services extension-service request-response grpc ssl local-certificate mycert*

To download Proto files, see the GitHub Repository.

[See Understanding gRPC Services for Managing Network Devices and Junos YANG Data Model Explorer.]

Unified naming and parent hierarchy for OpenConfig components in telemetry (PTX10001-36MR, PTX10001-36MR-K, PTX10002-36-OD, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—Use this feature to obtain consistent, hierarchical telemetry component names and accurate parent-child relationships under /components/component and /components/component/state/parent, simplifying data consumption and automation. A central nomenclature library and APIs return standardized names (for example, CHASSIS0:FPC0:PIC0), identify roots and parents, and map types to openconfig-platform. To consume telemetry, enable gRPC/gNMI and SSL if required. Update telemetry consumers and tests that reference legacy names.

[See Important Guidelines and Naming Conventions for Sensor Paths and Junos YANG Data Model Explorer.]

Decoupled rendering for telemetry data streaming (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming. Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

OpenConfig YANG model support for BGP (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)—Junos telemetry now supports the OpenConfig YANG model for BGP, providing a vendor-neutral framework for configuration and monitoring BGP. Retrieve BGP-specific operational and state data by subscribing to sensors under resource path hierarchies such as `inline communities`, `community-set`, `match-community-set`, `match-ext-community-set`, `peer-group`, `neighbor`, `policy-definition`, `administrative distance`, and `as-path-set`. OpenConfig-defined default behavior for BGP policies is implemented, with `community-level` policy options supporting `any-match` and `invert-match` configurations.

Use the Junos YANG Data Model Explorer to view supported sensor paths and platforms. The supported OpenConfig version for the openconfig-network-instance model is 4.4.1.

The following table lists the supported YANG models and their versions:

BGP YANG Model	Version
openconfig-routing-policy.yang	3.4.2
openconfig-bgp-common-multiprotocol.yang	9.9.1
openconfig-bgp-common-structure.yang	9.9.1
openconfig-bgp-common.yang	9.9.1
openconfig-bgp-errors.yang	6.1.0
openconfig-bgp-global.yang	9.9.1
openconfig-bgp-neighbor.yang	9.9.1
openconfig-bgp-peer-group.yang	9.9.1
openconfig-bgp-policy.yang	8.1.0
openconfig-bgp-types.yang	6.1.0
openconfig-bgp.yang	9.9.1
openconfig-types.yang	1.0.0

[See Junos YANG Data Model Explorer.]

Update and extend OpenConfig BGP telemetry (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)—Support has been added for OpenConfig BGP telemetry in these two ways:

- BGP telemetry has been updated to the latest OpenConfig models.

- Added support for the following paths:

/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/config/set-med-action

/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/state/set-med-action

/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/state/set-med

[See Junos YANG Data Model Explorer.]

any-match option for BGP telemetry (PTX10001-36MR, PTX10008, and PTX10016)—Support for OpenConfig match-set options in routing policy (the `any-match`option is added to the Junos CLI). The following paths are implemented:

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-community-set/config/match-set-options

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/state/match-set-options

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/config/match-set-options

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-community-set/state/community-set

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-community-set/config/community-set

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/state/ext-community-set

/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/config/ext-community-set

See the YANG Data Model Explorer.

[See Junos YANG Data Model Explorer.]

Update to BGP link bandwidth community value (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)—Support BGP to allow regular expressions 0 and 2^32 in the link bandwidth community value. The OpenConfig path is

/routing-policy/defined-sets/bgp-defined-sets/ext-community-sets/ext-community-set/

See the YANG Data Model Explorer.

[See Junos YANG Data Model Explorer.]

Configuration of administrative distance for BGP (PTX10008)—You can configure preferences (administrative distance) separately (internal or external) for BGP learned routes and corresponding paths. See the YANG Data Model Explorer.

[See Junos YANG Data Model Explorer.]

Enhancement of BGP XPath policy (PTX10001-36MR, PTX10008, and PTX10016)—Add OpenConfig BGP XPath enhancements for default policy and policy definitions. See the YANG Data Model Explorer.

[See Junos YANG Data Model Explorer.]

Addition of send-community-type and telemetry support (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10004, PTX10008, and PTX10016)—Extend core functionality to add `send-community-type` and telemetry support.

The following paths are supported:

`/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/config/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/config/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/config/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/config/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/afi-safis/afi-safi/config/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/state/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/state/send-community-type`

`/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/afi-safis/afi-safi/state/send-community-type`

The OpenConfig community-type doesn't define IPv6-ext-community types, so they are not controlled through the OC configuration and they are always advertised.

[See Junos YANG Data Model Explorer.]

DDoS rate limiting for telemetry traffic (PTX10003)—Junos PTX10003 routers now support DDoS rate limiting for telemetry traffic to enhance host path protection and provide improved control over telemetry traffic volume. Telemetry traffic, which uses TCP to stream network data for real-time analysis, was previously classified under the default *unclassified* DDoS protocol. Telemetry traffic is now mapped to a dedicated DDoS protocol, enabling targeted rate limiting, bandwidth control, and traffic monitoring.

The telemetry server IP address (optional, depending on configuration) and TCP port (mandatory) are used as match terms by the DDoS filter for telemetry traffic. DDoS rate limiting is implemented on a packets-per-second basis and is independent of packet size. The telemetry DDoS policer applies to ingress traffic and rate-limits TCP acknowledgment (ACK) packets from the collector to the device.

This feature applies only to TCP-based telemetry traffic and does not support UDP telemetry traffic. You can subscribe to the `/junos/system/linecard/ddos/` resource path to stream DDoS statistics.

[See Junos YANG Data Model Explorer, DDoS Rate Limiting for Telemetry Traffic, protocols (DDoS) (ACX Series, PTX Series, and QFX Series), and Configuring Control Plane DDoS Protection.]

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the [`edit event-options destinations` *`destination-name`* `archive-sites` *`url`*] hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

MPLS

Enhancements to MPLS automatic bandwidth for RSVP LSPs with per-LSP polling profiles and adjustable thresholds (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10003, PTX10004, and PTX10008)—You can dynamically allocate bandwidth for RSVP label-switched paths (LSPs) based on traffic to improve utilization and service assurance. Define multiple polling profiles with the `polling-profile` statement under the [`edit protocols mpls statistics auto-bandwidth`] hierarchy and associate each profile with specific LSPs to tailor polling intervals. Configure threshold settings through the `adjust-threshold-options` statement to set precise overflow and underflow triggers. Use enhanced CLI options to create polling profiles, bind them to LSPs, and set thresholds to improve monitoring and troubleshooting in complex networks.

[See Polling Profiles for Automatic Bandwidth Allocation for RSVP LSPs.]

Network Management and Monitoring

Support for NDP and DAD Proxy (PTX10008)—Use NDP Proxy and Duplicate Address Detection (DAD) Proxy to manage IPv6 neighbor discovery and detect duplicate address in supported deployments. You can enable restricted mode or unrestricted mode for NDP Proxy and DAD Proxy. To enable NDP Proxy, use the `ndp-proxy interface-restricted` or `ndp-proxy interface-unrestricted` configuration statement. To enable DAD Proxy, use the `dad-proxy interface-restricted` or `dad-proxy interface-unrestricted` configuration statement.

[See NDP Proxy and DAD Proxy.]

Disable system-level ICMPv6 redirect generation (PTX10001-36MR, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, and PTX10016)—Disable ICMPv6 redirect generation to prevent the device from sending ICMPv6 redirect messages. When an IPv6 packet enters and exits through the same interface, the device generates an ICMPv6 redirect message to indicate that a better path. Use the `set system no-redirects-ipv6` configuration statement to disable ICMPv6 redirect generation across the device. To re-enable ICMPv6 redirects, delete the `no-redirects-ipv6` statement from the configuration.

PiTrace binary tracing for IS-IS, BGP, RSVP, and MPLS (PTX10001-36MR, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—You can enable always-on PiTrace to collect compact event histories with minimal overhead. This efficient binary tracing infrastructure enhances root-cause analysis without compromising scalability. PiTrace actively traces key IS-IS events, monitors BGP for CPU-intensive operations, and gathers operational data for RSVP and MPLS. The router supports binary and legacy text logging to function simultaneously.

To configure tracing levels and retention, use the commands `set system pixel-tracing level {any|critical|debug|info|none}` and `set system pixel-tracing files {1..1000}`. Clear the logs by executing the command `request system pixel-tracing routing flush`.

[See https://www.juniper.net/documentation/test/us/en/software/junos/network-mgmt/topics/concept/Pitrace-Binary-Trace-Overview.html.]

OpenConfig

OpenConfig gNMI heartbeat for on-change subscriptions (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—Use OpenConfig heartbeat for gNMI ON_CHANGE subscriptions to receive a full set of sensor data at each heartbeat

interval, even when values do not change. The heartbeat provides periodic state validation and mitigates undetected data loss.

[See Configuring gNMI Heartbeat Interval for ON_CHANGE Subscriptions and Junos YANG Data Model Explorer.]

Precision Time Protocol (PTP)

PTP passive port monitoring with SNMP MIB support and timeReceiver performance monitoring support (JNP10K-LC1301 LINE CARDS ON PTX10004 AND PTX10008)—You can enhance router performance monitoring through PTP passive port monitoring and Simple Network Management Protocol (SNMP) Management Information Base (MIB) integration. Passive port monitoring alerts you to potential fiber asymmetries or clock failures, improving network reliability by identifying issues when phase thresholds are exceeded. Use the show ptp passive-port-monitor-status command to display performance statistics.

SNMP exposes device information through MIB objects, which represent structured data points describing PTP operation, status and performance. The SNMP get, get-next, and walk management capabilities are enabled for timing features.

Use the timeReceiver performance monitoring feature to collect detailed metrics for active PTP timeReceiver ports using IEEE 1588-2019 standards, including timeTransmitter clock-timeReceiver clock delay and to synchronize message transmission. Use the show ptp performance-monitor status command to view the status of timeReceiver port performance monitoring.

[See PTP Passive Port Performance Monitoring, SNMP MIB for Timing on Routing Platforms, and PTP TimeReceiver Port Performance Monitoring.]

Routing Options

Route re-resolution optimization for L3VPN and recursive multipath with new diagnostic commands (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, AND PTX12008, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, and PTX10016)—You can now use optimized route re-resolution in the resolver module to reduce CPU utilization during route churn. This optimization extends re-resolution beyond inet.0 and inet6.0 to `<instance>.inet.0`, `<instance>.inet6.0`, and L3VPN tables (bgp.l3vpn.0, bgp.l3vpn-inet6.0). The resolver also supports optimized re-resolution for recursive resolution cases and for resolution over multipath prefixes.

To monitor and diagnose resolution, you can use the following operational commands:

1. `show route resolution active-helpers` - Displays filtered routes used for the resolution, before the active routes were displayed
2. `show route resolution active-helpers statistics` - Helps identify the reason for which routes are re-resolved, which is helpful in debugging.

3. `show route resolution re-resolution-history` - Displays the routes that are re-resolved.

4. `show route resolution loop-history` - Helps identify if any resolution loop occurred.

[See show route resolution.]

Routing Policy and Firewall Filters

MPLS payload TCP flag and egress filter support (PTX10002-36QDD, PTX10004, and PTX10008)—We now support TCP flag match on MPLS payload traffic and use MPLS filter for ingress and egress directions to enhance filtering capabilities. Logical operators `AND` and `OR` are supported.

[See Firewall Filter Match Conditions for MPLS Traffic, Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic, and Firewall Filter Match Conditions Based on Bit-Field Values.]

Routing Protocols

Support for handling unknown optional transitive BGP attributes (PTX10008)—You can drop unknown optional transitive BGP attributes or withdraw routes carrying these attributes to prevent the propagation of malformed or harmful attributes across BGP peers. Apply the configurations at the BGP global, group, or neighbor level for comprehensive control. You can specify exceptions for attributes that should not be dropped or cause route withdrawal.

`Note:` If you configure multiple actions for the same attribute, then you can perform only the withdrawal action.

To display the unknown optional transitive BGP attributes, use the enhanced `show BGP neighbor` and `show log messages` commands.

[See BGP Overview and show bgp neighbor.]

Services Applications

Use FQDN targets in RPM probes (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—Configure an RPM (real-time performance monitoring) probe test with a fully qualified domain name (FQDN) target (instead of a IPv4 or IPv6 address) to adapt to endpoint address changes. The device resolves the FQDN through the Domain Name System (DNS) at the start of each test, then sends probes to the resolved address at each interval. Telemetry first displays the FQDN until resolution, then shows the destination address to help you track changes. If DNS resolution fails, the probe results retain the FQDN as the target address. Configure the `fqdn` or the `fqdnv6` option at the `[edit services monitoring rpm owner test test-name target]` hierarchy level.

[See target.]

Software Installation and Upgrade

Block configuration commits during a Junos OS Evolved software upgrade or rollback (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, QFX5250-64OE-DOT2L, QFX5700, and QFX5700E)—For supported products, Junos OS Evolved now blocks configuration commits:

1. When a Junos OS Evolved software upgrade or rollback is in progress.
2. After the upgrade or rollback command has finished, but before you reboot the system.

[See commit, request system software add (Junos OS Evolved), and request system software rollback (Junos OS Evolved).]

Support for TCP connectivity on management and WAN interfaces for non-SI applications (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E)—You can enable TCP connectivity for statically linked or other non-Socket Intercept (non-SI) applications over management and WAN interfaces. Start each non-SI application with the `NO_SI_INTERCEPT` environment variable to enable TCP traffic processing.

[See Running Applications Without the Socket Intercept Library.]

Firmware upgrade support for CMIS 4.0 optics modules (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—We now support upgrading the firmware for CMIS 4.0 optics modules. For our optics modules, we provide a tar file that contains the firmware binary, checksum details of the firmware binary, and meta files with the JPN and firmware revision mapping details. For third-party optics, you must acquire the tar file from your supplier and load it onto the device. You upgrade the firmware for CMIS 4.0 optics using the `request system firmware optics update` command. If the transceiver is advertising two CDB instances, you can use the `request system firmware optics switch` command to choose between the two instances. To view the firmware status, use the `show system firmware optics status` command or the `show chassis pic` command. To view the CDB advertisements, use the `show system firmware optics cdb-advertisement` command.

[See request system firmware optics update (Junos OS Evolved), request system firmware optics switch (Junos OS Evolved), show system firmware (Junos OS Evolved), and Upgrade Firmware on Junos OS Evolved Devices.]

Source Packet Routing in Networking (SPRING) or Segment Routing

BGP RIB sharding for SRv6 routes (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10002-36QDD, PTX10004, PTX10008, and PTX10016)—We've extended support for Segment Routing over IPv6 (SRv6) with BGP routing information base (RIB) sharding to enhance route processing and utilize SRv6 next hops. RIB sharding splits BGP service resolution over multiple shards to improve BGP convergence and performance. Configure SRv6 under IS-IS or as static SRv6-Traffic Engineering (SRv6-TE) with BGP RIB sharding. We do not support BGP route resolution when RIB sharding is configured.

[See rib-sharding.]

Support for IS-IS IPv6 SRLG TLV 139 and interface-level SRLG configuration (ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, and PTX12008)—This feature enables routers to advertise and process IPv6 Shared Risk Link Group (SRLG) information as defined in RFC 6119, *IPv6 Traffic Engineering in IS-IS*, ensuring accurate awareness of shared-risk failures. By extending SRLG support to IPv6 and integrating it directly into the IS-IS configuration, you can deploy SRv6 networks without relying on MPLS. This feature enhances Topology-Independent Loop-Free Alternates (TI-LFA) fast reroute path calculation. It enables faster and more reliable protection against link, node, and shared-risk failures for modern IP networks.

To enable per-interface SRLG configuration and advertisement in IS-IS for SRLG-aware TI-LFA, configure `srlg` *`srlg-name`* at the `[edit protocols isis interface interface-name]` hierarchy level.

[See Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS and OSPF.]

SR-MPLS-TE local computation with OSPF underlay (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, AND PTX12008, PTX10001-36MR, PTX10001-36MR-K, PTX10002-60MR, PTX10003, PTX10004, PTX10008, and PTX10016)—We support Segment Routing over MPLS with Traffic Engineering (SR-MPLS-TE) local computation with OSPF as the underlay interior gateway protocol (IGP) for the following features:

- Endpoint IP address configured as anycast IP address
- Segment-list computation using anycast constraints
- Metric-type delay
- Protected and unprotected path selection

[See Segment Routing LSP Configuration and Configuring Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS.]

SRv6-TE micro-SID support for transport and L3VPN (PTX10002-36QDD, PTX10004, and PTX10008)—We extend micro-SID support for Segment Routing over IPv6 with Traffic Engineering (SRv6-TE). We support segment routing–traffic engineering (SR-TE) micro-SID only with default block configurations across the network domain. If block configurations exist, they must be uniform throughout the network. The Packet Forwarding Engine supports bit shifting for both <block>:<uN>:<uA> and <block>:<uA> routes. Configure the full SID as advertised in the IS-IS interior gateway protocol (IGP), either <block>:<uN> or <block>:<uN>:<uA>.

We support:

- `micro-srv6-sid` statement at the `[edit protocols source-packet-routing segment-list` *`name hop-name`*`]` hierarchy level to configure micro-SIDs in an SR-TE SRv6 segment list.
- `strict-adjacency` statement at the `[edit protocols source-packet-routing segment-list` *`name hop-name`*`]` hierarchy level to strictly follow the micro-adjacency SID.

You can configure the segment list containing micro-SIDs using the existing `srv6` statement, similar to the traditional SRv6 configuration. For SRv6-TE segment-list configuration, use the `srv6-sid` statement. For micro-SID configuration, use the `micro-srv6-sid` statement.

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

User Access and Authentication

TACACS+ Over TLS for Secure Device Administration (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—This release introduces support for Transport Layer Security (TLS) for TACACS+ communication between network devices and TACACS+ servers. By leveraging TLS, this enhancement provides strong encryption, integrity protection, and mutual authentication, addressing the security limitations of legacy TACACS+ mechanisms. The previous weak protection methods are now deprecated, ensuring that sensitive AAA traffic is securely transmitted in modern deployments.

TACACS+ over TLS Support (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—This release introduces support for **TACACS+ over TLS**, enhancing the security of device administration by replacing the legacy MD5-based obfuscation mechanism with modern TLS encryption. With this update, communication between Juniper devices

and TACACS+ servers is protected by TLS, ensuring **confidentiality, integrity, and mutual authentication** using X.509 certificates.

Key highlights:

- **Secure transport** for TACACS+ authentication, authorization, and accounting traffic.
- **Mutual TLS (mTLS)** support, enabling certificate-based validation of both client and server.
- **Failover support** between TACACS+ over TLS, TACACS+ over TCP, RADIUS, and local authentication.
- **Routing instance compatibility**, including default, non-default, and management instances.
- **CLI configuration enhancements** to define trusted CA groups and client certificates for TLS connections.

This feature allows TACACS+ deployments to meet modern compliance requirements (e.g., FIPS-140-3) and strengthens overall AAA security in Junos and Evo platforms.

Additional Features

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

Support for multiple bridge domains and local switching in VPLS (PTX10004 and PTX10008)—We support the following VPLS features:

- Multiple bridge domain within a single VPLS routing instance

[See Introduction to Configuring VPLS.]

- Local switching with hierarchical VPLS

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

Export timing data to collectors (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, PTX10004, and PTX10008)—Use Junos telemetry to export timing attributes for Precision Time Protocol (PTP) and Synchronous Ethernet (SyncE) to external telemetry data collectors. The device supports periodic streaming and on-change notifications. Subscribe to the following sensor paths to view PTP and Synchronous Ethernet sensor information:

- For Precision Time Protocol (PTP): `/state/protocols/ptp/`
- For Synchronous Ethernet: `/state/protocols/synce/`

For a complete list of sensor paths supported by the device, see Junos YANG Data Model Explorer.

Interface-level telemetry for IPv4 and IPv6 transit traffic (PTX10008)—We support Transit traffic rates in bits per second (bps) and packets per second (pps) at interface-level for both IPv4 and IPv6 traffic. Use the following commands to enable transit statistics accounting:

- `set forwarding-options family inet route-accounting (for IPv4 traffic)`
- `set forwarding-options family inet6 route-accounting (for IPv6 traffic)`

Junos telemetry now streams statistics and rates for IPv4 and IPv6 transit traffic on logical and aggregated Ethernet interfaces.

The following native resource paths are supported:

- `/junos/system/linecard/interface/logical/family/ipv4/usage/`
- `/junos/system/linecard/interface/logical/family/ipv6/usage/`

The following counters for ingress or egress traffic are streamed (per logical interface):

IPv4 counters

- `/interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/in-octets`
- `/interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/in-pkts`
- `/interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/out-octets`
- `/interfaces/interface/subinterfaces/subinterface/ipv4/state/counters/out-pkts`

IPv6 counters

- `/interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/in-octets`
- `/interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/in-pkts`
- `/interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/out-octets`
- `/interfaces/interface/subinterfaces/subinterface/ipv6/state/counters/out-pkts`

Use these counters to monitor IPv4 and IPv6 transit traffic at a granular interface level. Use the `show interfaces` *`interface`* `statistics` command to display the traffic rates.

[See Configuring IPv4 and IPv6 Accounting and Junos YANG Data Model Explorer.]

BITS and GPS support (JNP10K-LC1202 LINE CARDS ON PTX10004 AND PTX10008)—We have added the following support:

- Building Integrated Timing Supply (BITS) is a centralized, highly stable timing distribution system used to synchronize communication and data networks. It is used to distributes accurate reference frequencies to all networking devices within a single building, campus, or data center.
- Global positioning system (GPS) serves as a highly accurate external timing source by delivering precise clock signals derived from satellite-based clocks, which is used as a primary reference for network synchronization.

[See BITS and GPS Support.]

Support for flood policer over EVPN-MPLS and EVPN-VxLAN (PTX120008)—Support is added for flood policer for EVPN-MPLS and EVPN-VxLAN

[See Traffic Policer Types.]

Support for displaying the load-balancing hash result (PTX10008)—Support is added for displaying the load-balancing hash result using the `show forwarding-options load-balance` CLI command.

[See show forwarding-options load-balance.]

Support for 1x200GbE, 2x200GbE, 3x200GbE, and 4x200GbE interfaces (PTX10008)—Supported on PTX10008 with PTX10K-LC1301-36DD line card

Support for QDD-2X400G-VR4-P optics (PTX10008)—Supported on PTX10008 with PTX10K-LC1301-36DD line card

Monitor performance of OSFP-800G-ZR optics (PTX12008)—Monitor the performance of OSFP-800G coherent ZR optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Use the `show interfaces transport pm` command to view current and historical performance data.

[See optics-options, Coherent Optics Performance Monitoring, and show interfaces transport pm.]

Support for appsel number configuration. Support for display of Q-factor and Q-margin in 400ZR and 400G OpenZR+ optics (PTX10004 and PTX10008)—Supported on PTX10004 and PTX10008 with PTX10K-LC1301-36DD line card

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

Monitor network performance and configure user-configurable PM intervals (PTX10004 and PTX10008)—Supported on PTX10004 and PTX10008 with the PTX10K-LC1301-36DD line card. Configure performance monitoring (PM) intervals for long-haul coherent ZR optics across supported platforms to gain more granular and precise visibility. Adjust intervals from the default 15 minutes to shorter durations. Use the `set chassis optics pm interval-length (10s|30s|1min|5min|15min)` command to adjust the intervals according to your specific monitoring needs.

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

SNMP MIB walk and traps support for coherent ZR optics (PTX10004 and PTX10008)—Use SNMP MIB walk and traps to efficiently monitor and manage coherent ZR optics, including 100ZR, 400ZR, 400ZR-M, and 400ZR-M-HP. Use this feature to retrieve OID information sequentially and receive notifications when alarms are triggered or cleared.

[See Enterprise-Specific MIBs Supported by Junos OS Evolved, show snmp mib, and SNMP MIB Explorer.]

SNMP support for coherent ZR optics performance monitoring and threshold alerts (PTX10004 and PTX10008)—You can monitor the performance of coherent ZR optics (100ZR, 400ZR, 400ZR-M, and 400ZR-M-HP) and receive threshold crossing alerts using SNMP. Retrieve real-time, historical, and statistical data for various performance parameters through SNMP Get requests.

[See Enterprise-Specific MIBs Supported by Junos OS Evolved, show snmp mib, and SNMP MIB Explorer.]

EVPN-VXLAN fabric with an IPv6 underlay (PTX10008)—You can use EVPN-VXLAN overlays over an IPv6 underlay to deliver IPv4 and IPv6 services.

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

Seamless Stitching of EVPN-VXLAN to EVPN-MPLS Gateways Across a Data Center Interconnect (PTX10008)—We support EVPN-VXLAN to EVPN-MPLS DCI stithing on the listed hardware platforms.

[See Configure EVPN-VXLAN Data Center Stitching Through Interconnected EVPN-MPLS WAN Gateways.]

MAC Pinning in EVPN (PTX10001-36MR, PTX10002-36QDD, PTX10002-60MR, PTX10004, PTX10008, and PTX10016)—We support mac-pinning in EVPN networks.

[See EVPN MAC Pinning Overview.]

What's Changed

IN THIS SECTION

Learn about what's changed 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 OS Evolved Problem Report Search application.

EVPN

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

Forwarding and Sampling

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

General Routing

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for autoconfiguration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router?s link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Restrict PPPoE client lockout output to underlying logical interfaces with short-cycle protection**—The "show pope lockout" command shows PPPoE clients information only for underlying logical interfaces with `short-cycle-protection` enabled.

  [See show pppoe lockout.]

- **Traffic Engineering Database Import Identifier Update**—The help string for the `set protocols mpls traffic-engineering database import identifier` command has been updated to BGP-TE identifier for standard IGP instance. This change clarifies that the identifier configuration affects only the standard IGP instance. The BGP-LS advertisements in the `lsdist.0` table will now use the configured value instead of defaulting to Instance Identifier 0. This update ensures accurate Instance Identifier configurations in BGP-LS advertisements, enhancing network visibility and management. The routers running software later than Junos OS release 24.4.x version should verify that the correct Instance Identifier is being advertised. [See import (MPLS Traffic Engineering Database.]

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `edit services port-extender satellite <name> device-model` hierarchy, you can use the "?" to see a list of supported platforms. We have also added the `skip-lo0-config` statement under `edit service port-extender jnu` hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the `edit service port-extender satellite <name>` hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **IPv6 Underlay Support for EVPN-VXLAN (PTX10008)**—The `vtep-source-interface inet6` configuration statement is allowed only when the `[edit chassis interoperability]` configuration is set to either `express5-enhanced` or `express6-enhanced`.

[See vtep-source-interface and interoperability.]

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- **Pixel-Tracing Command Updates**—The pixel-tracing command suite has been updated to include the following options: `flush`, `disable`, and `enable`. These options allow you to control and manage pixel-tracing processes for routing tasks. Note that the `rib-sharding` option is removed from the CLI commands to prevent system crashes when sharding configuration is present in the system. Ensure that your configurations do not rely on unmentioned options to avoid potential disruptions in your workflow.

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
```

```
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)              - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)       -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Prevent multicast route installation on uninterested receivers**—Configure the `resolve-request-holdcount` statement under the `[edit routing-instances name routing-options multicast]` hierarchy to prevent the kernel from caching initial multicast packets during route resolution and avoid unnecessary route installation on devices without interested receivers. [See {https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/resolve-request-holdcount-edit-routing-instances-routing-options-multicast.html} resolve-request-holdcount].

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading,

copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

[See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **BGP route-refresh warning messages added**—New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

  [See show bgp diagnostics warnings.]

- **Inbound Convergence Status Added to BGP Neighbor Display (Junos OS and Junos OS Evolved)**—The `show bgp neighbor` command now displays the Inbound Convergence status and cause.

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **System Stability Improvement**—When BGP fails to bind to its listening port, the routing protocol process (RPD) now stays running and logs a warning instead of terminating after the maximum retry count, preventing a service outage. The `show bgp diagnostics overview` output includes a new `BGP listening port status of v4 and v6` field that reports `Active` when BGP is successfully listening on the configured port and `Failed` when the port is unavailable or occupied by another process. Error logging is also enhanced with the `BGP_LISTEN_STATUS` message: `BGP listen failed on port [X] after 64 attempts with error-%d for address family identifier (IPv4 or IPv6). Administrative clear needed.` This message is generated after `BGP_MAX_LISTEN_START_ATTEMPTS` failed attempts to bind to the listening port.

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.

User Interface and Configuration

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the

corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

- `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
- `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
- `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
- `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
- `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:
  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of `remove` operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents **statement not found** warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

IN THIS SECTION

Learn about known 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 OS Evolved Problem Report Search application.

General Routing

- The default forward error correction (FEC) mode for the QSFP28-END-100G-4x25G-DAC-B0-5M optical module is FEC 74. However, FEC 74 is not supported on the BX chipset, which constitutes a hardware constraint. To ensure proper operation of the optical module, the interface and its connected peer must be configured to use either FEC 91 or FEC None, as these are the forward error correction modes supported by the BX chipset. PR1934424

- When a loopback is applied to channelized (4x10) on all channels simultaneously, only the first link becomes active. This behavior occurs only during a combined loopback and does not appear when applied individually. This issue is specific to loopback, and does not impact traffic. PR1925465

- Immediately after the PICD restart, wait for a few seconds (>15) before applying a new PICD configuration.

  PR1924485

- When AH is in progress and GRES happens, there is a possibility that AH may 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

Open Issues

IN THIS SECTION

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

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

EVPN

- On Junos OS Evolved platforms, if the router-id is not explicitly configured, the device dynamically selects a router-id during bootup. In this scenario, the device advertises EVPN (Ethernet Virtual Private Network) Type1 AD/ESI (AutoDiscovery/Ethernet Segment Identifier) and Type4 ES RT (Ethernet Segment Route Target) routes with a Route Distinguisher (RD) of 0:0. As a result, Route Targets (RTs) advertised by different Provider Edge (PE) routers end up having identical prefixes. Consequently, only one PE routers RT is advertised and learned. This behavior prevents EVPN all-active redundancy resulting in the loss of multihoming. PR1924472

Forwarding and Sampling

- On Junos Evolved PTX devices, next-hop churn, such as adding and deleting routes, can lead to the exhaustion of available pfetokens. These tokens are claimed but not released, which causes low pfetoken availability. This issue results in traffic drops and unexpected system behaviour.

PR1833330

- For scaled configuration having 256 member links in an aggregated Ethernet (AE), if users disable all the child members of the AE interface and then immediately (within 30 seconds to 1 minute) issue the `show interfaces` command, they might see absolute statistics values instead of relative values. Clearing statistics is a software reset and is only cosmetic; the actual or absolute traffic statistics remain unchanged.

  Example: Before disabling, the Carrier Transition (CT) count is 10. Performing a clear operation resets the software counter to 0 (snapshotted CT = 10). Relative statistics on the next `show interfaces` command would be calculated as (current value - snapshot value) = 10 - 10 = 0 (relative values). After disabling and immediately issuing `show interfaces`, users might observe the absolute CT value of 11 (as the snapshot is deleted, and the CT value increments by 1 after the interface flap or disable command). PR1947492

General Routing

- For PTX10001-36MR, the software driver reads the voltage threshold erroneously causing the **Host 0 Voltage Threshold Crossed** alarm to present on the device. PR1592258
- PTX10008: Junos OS Evolved CB1, SIB0 and FPC0 (JNP10K-LC1201) went to fault after system power outage in lab.

  PR1749793
- On Junos OS Evolved based platforms, while using ping, traceroute, or other utility that requires host name resolution, an error is raised indicating hostname resolution has failed. PR1822994
- INLINEKA:CFM-SLM-RESPONDER No entry error might be seen with CFM configuration having more than one remote peer MEP in inline mode. There is no functional impact due to this error. PR1865524
- PTX10003-160C: While verifying the multicast traffic, zero packet loss does not occur as expected.

  PR1917000
- On Junos OS Evolved running on PTX10008 and PTX10016 platforms, Switch Interface Board (SIB) initialization failures occur after system reboot. Traffic loss occurs due to affected SIBs.

  PR1900242
- When users perform a Routing Engine switchover with an FPC in an offline state in the chassis, after the switchover, the FPC becomes operational with some links in a fault state.

  Recovery: Take the FPC offline and then bring it online to recover the link faults or bring the FPC online before performing the Routing Engine switchover.

PR1936673

- For ZX/Express-4/5/6 platforms, hosting VLAN-CCC interfaces, if *NO* port protocol (e.g., LACP) is configured on the port, native VLAN ID can be enabled on the port so that *untagged* control packets (e.g., LACP) can be switched. This behavior aligns with Layer 2 transparency. Tagged control packets are also switched. However, if a port protocol is configured, native VLAN ID must *NOT* be configured to ensure untagged control packets are punted to the host path and tagged control packets are switched. PR1935588

- In Junos OS Evolved, ARP resolution requests are throttled on the FPC at the logical interface level. In Junos OS Evolved, ARP resolution requests are throttled on the Flexible PIC Concentrator (FPC) per logical interface. If a resolution request on a logical interface expires, a throttle timer starts on that logical interface, and no other resolution request can be generated on that logical interface while the throttle timer is active. If a resolution request on a logical interface expires, a throttle timer is initiated for that logical interface, preventing additional resolution requests during the timer's duration. In this issue, the RE netstack sends packets to the FPC with a hint to the resolution request for an already resolved IP address, triggering the throttle timer. In this issue, the Routing Engine (RE) netstack transmits packets to the FPC with a hint for a resolution request targeting an already resolved IP address, which activates the throttle timer. As a result, a resolution request for a second IP address on the same logical interface cannot be generated for a certain period, causing the phone-home application's 10-second timeout. Consequently, a resolution request for a second IP address on the same logical interface is delayed, leading to the phone-home application's 10-second timeout.

  PR1883158

- Problem: On systems using a JNP10K-RE3 Routing Engine with an LC1301 line card, a "PLL LOCK Failure" major alarm may be reported after a reboot or power cycle. The alarm may clear automatically after some time or remain active indefinitely. Although the alarm indicates a PLL lock failure, it does not affect system functionality with the JNP10K-RE3 and LC1301 hardware combination. Symptoms: A major alarm similar to the following may be displayed: Major FPC <x> zl30642a PLL LOCK Failure The alarm is also reported in the output of: show system errors active detail with an error similar to: Error Name : PLL lock failure Conditions: This issue may occur following a system reboot or power cycle on systems using a JNP10K-RE3 (McLaren) Routing Engine with an LC1301 (Aegon) line card. Workaround: There is no workaround. The alarm does not impact any feature supported with the JNP10K-RE3 and LC1301 hardware combination. PR1936269

- 800G ports in Whistler 100 use MPM54524 as the point-of-load. This device is not reporting "power-bad" flag, hence SW is unable to detect power-bad status and raise alarm. PR1965458

Infrastructure

- Logs are flooded with audit: PROCTITLE causes the message file to fill up quickly.

  PR1882686

Routing Policy and Firewall Filters

- On all Junos OS Evolved platforms, configuring the reserved prefix 0.0.0.0/8 in management interface firewall filters or host-inbound policies will cause unintended filtering, resulting in loss of management traffic, including SSH access. PR1910579

Routing Protocols

- On all Junos OS Evolved platforms, after Border Gateway Protocol (BGP) sessions configured with `family route-target` flaps, delayed route deletion causes the loss of the Route Target Filter (RTF), preventing the node from advertising L3VPN (Layer 3 Virtual Private Network) and direct routes (e.g., loopbacks and interface routes) to the BGP peer, leading to VPN route loss and service disruption.

  PR1849568

- On all Junos OS and Junos OS Evolved platforms, telemetry data collection after an interface flap event, on an interface with PIM (Protocol Independent Multicast) sparse-dense mode configured causes rpd (Routing Protocol Daemon) process crash and core-dump generation. This results in protocol restarts and service disruption. PR1937530

Resolved Issues

IN THIS SECTION

Learn about resolved issues in this release for PTX Series routers.

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

Authentication and Access Control

- On all Junos OS Evolved platforms, if the username is more than sixteen characters, the username is truncated to sixteen characters in the accounting logs displayed for that user.

  PR1786580

Class of Service (CoS)

- On all Junos OS Evolved platforms, when a forwarding class is attached to a logical interface and the forwarding-class queue number is modified, the logical interface configuration reflects this updated queue number. However, the traffic continues to be forwarded through the previous queue.

  PR1927601

- On PTX10000 line of Routers, minor alarm **Anomaly errors reported for application cosd [context=default] on node re0** is raised post second Routing Engine switchover.

  PR1938824

- `set class-of-service host-outbound-traffic` DSCP code point was disabled on PTX Series Routers in Junos OS Evolved Release 25.4R1. Later enabled in Junos OS Evolved Release 26.2R1 and backported to Junos OS Evolved Release 25.4R2.

  PR1930528

Forwarding and Sampling

- When Layer 2 circuit (l2circuit) or Layer 2 VPN (l2vpn) is configured, Layer 2 control packets with destination MAC address 01:00:0C:CC:CC:CD (used by Spanning Tree Protocol) are discarded and not forwarded to customer edge devices. Spanning Tree Protocol does not establish across the Layer 2 circuit or Layer 2 VPN connection. Packet discard counters increment under Packet Forwarding Engine statistics.

  PR1914535

General Routing

- Routing protocols sees forwarding impact on Junos Evolved PTX10003 platforms when a nexthop transitions between forwardable and hold/reject states, causing the nexthop token to fail reacquisition and leading to unresolved forwarding plane dependencies. PR1923943

- On all Junos OS Evolved PTX Series platforms with multiple VLANs on an Aggregate Ethernet (AE) interface and an Firewall filter with family ethernet switching and bandwidth policer attached to both

the AE interface and its associated VLAN interfaces, a modification or deletion of the VLANs will lead to a FPC crash. All the traffic over the affected FPC will be dropped. PR1934506

- On Junos OS Evolved PTX platforms with optics EM policy enabled, when the temperature of an optics module exceeds the fire threshold, an over-temperature shutdown alarm is raised and then unexpectedly cleared. Despite a log message indicating optics shutdown, the link remains operational, resulting in inconsistent optics behavior.

  PR1875956

- An Improper Restriction of Excessive Authentication Attempts vulnerability in Juniper Networks Junos OS Evolved allows an unauthenticated, network-based attacker to cause a limited denial-of-service (DoS) to the management plane. Refer to CEC Juniper Community for more information. PR1709385

- On all Junos OS Evolved platforms, the Packet Forwarding Engine Daemon (PICD) repeatedly creates new SNMP trap objects without removing old ones, which causes memory leaks in the distributord process over time. This is tracked under Technical Support Bulletin (TSB) 107540.

  PR1914708

- On all Junos OS Evolved platforms, when FIPS is enabled, the JSD server fails to start if a local SSL certificate is configured on telemetry; as a result, telemetry services using SSL certificates are unavailable.

  PR1902691

- The JavaScript Object Notation (JSON) encoding of leafs of type "ieeefloat32"(https://github.com/openconfig/public/blob/master/release/mode ls/types/openconfig-types.yang#L127) is not correct, causing gRPC Network Management Interface (gNMI) data outputs unreadable. PR1923848

- A cli-pfe show command may continue to gather data in the background after the user issues a Ctrl-C. Eventually the background command will complete. However, the CPU usage for the cli-pfe process will continue to be high while it is still running. PR1915225

- On all Junos OS Evolved platforms, Time-zone info changes to default UTC after upgrade is done with restart-upgrade PR1803511

- On all Junos and Junos OS Evolved platforms, after a system reboot, sending a very large number of RPC (Remote Procedure Call) calls per second (20K or more) will prevent gNMI (gRPC Network Management Interface), gNOI (gRPC Network Operations Interface) and P4RT connections from being established. This will result in nginx session failures. There will be no traffic impact due to this issue. PR1915308

- On all Junos Evolved PTX series platforms, when acting as egress PE router, it drops IPv6 packets with MPLS header from any ingress PE device adds padding to small IPv6 packets going over MPLS

L3VPN and arriving on egress PTX router. If the padding is removed, they are forwarded as expected. PR1926159

- On Junos OS Evolved platforms with telemetry streaming configured, the sysd (System Daemon) process crashes on the backup RE (Routing Engine), and error logs are observed when sysd memory usage continues to grow over time, and the memory limit is exceeded. This issue is observed when the 64-bit sysd consumes more memory than the configured system limit when the device has been running for a long time (typically more than 100 days) in the dual RE scenario. PR1929823
- The "PICD_XCVR_REMOVE" and "PICD_XCVR_INSERT" syslog messages, which are generated when a QSFP is removed or inserted, may not be reported PR1859581
- On Junos platforms with Micro-BFD and LACP (Link Aggregation Control Protocol) configured, if Micro-BFD (Bidirectional Forwarding Detection) is deactivated on the peer device and the member interfaces on that peer are flapped, the AE (Aggregated Ethernet) interface remains down even after the member links come back up. Due to this, the AE will be down and traffic loss occurs temporarily. PR1912709
- On all Junos OS Evolved PTX platforms, IRB (Integrated Routing and Bridging) routed traffic is dropped when both source and destination hosts are reachable through the same Service Provider (SP)-style Layer 2 IFL (Logical Interface) and associated IRB interface. PR1928605
- On the PTX10002, when the system is rebooted with the power-optimized setting enabled, a "PLL lock error" may be detected, and an "FPC 0 Major Errors" alarm may be observed. PR1931372
- On all Junos and Junos evolved platforms, the Routing Protocol Process Daemon (RPD) core is observed due to aggressive flapping of aggregated interface (ae). This causes the routing protocol process daemon to restart resulting in control plane going down and re-converging. PR1601226
- On Junos OS Evolved PTX Series platforms with incoming MPLS packets and outgoing MPLS over UDP configured over a Flexible Tunnel Interface (FTI), MPLS over UDP traffic for the same label would have the same source port. This would affect load balancing mechanism and can lead to inefficient utilization of links. PR1921577
- On PTX platforms, in deployments containing PTX Series platforms, Seamless Bidirectional Forwarding Detection (SBFD) sessions associated with Segment Routing (SR) or Segment Routing Traffic Engineering (SR-TE) Label Switched Paths (LSPs) may remain permanently down when Intermediate System to Intermediate System (ISIS) Explicit Null is configured. This issue occurs due to incorrect checksum handling during SBFD packet reflection processing between MX and PTX devices. As a result, SBFD sessions fail to establish successfully, causing the associated Label Switched Paths (LSPs) to remain down and resulting in traffic impact.

  PR1941741

- On PTX Series platforms running Junos OS Evolved, the Packet Forwarding Engine might become unresponsive when scaled firewall filters using TCAM-based match conditions are configured, which can delay the installation of new firewall filters PR1938941
- On PTX10002-36QDD platform and LC1301 line card, under high-traffic conditions with active counter processing on a loopback interface, an ASIC error may be triggered, causing the Packet Forwarding Engine (PFE) to enter a disabled state. This results in all interfaces going down and leads to complete traffic loss. PR1949690
- On PTX10008 and PTX10004 platforms running Junos Evolved affected releases with LC1301 line cards, Ethernet interfaces remain down when the 2x100G interface speed is configured before Quad Small Formfactor Pluggable (QSFP) transceivers are inserted and the optical signal is up. This condition affects dataplane connectivity and can result in loss of traffic on the impacted interfaces. PR1945978
- On Junos Evolved PTX platforms, FPC (Flexible PIC Concentrator) major errors can be seen after FPC restart/offline or multiple concurrent PFE (Packet Forwarding Engine) restart/offline in a highly-scaled scenario that can result in the PFEs failing to come back online. This is due to offline timeout, FPC go into fault state causing services to impact. This is a rare case scenario. PR1937796
- On certain PTX platforms running Junos EVO, VPLS (Virtual Private LAN Service) instances configured with multiple Bridge Domains fail to forward control-plane (CP-class) traffic received on the Logical Service Interface (LSI). Under this condition, control-plane protocols carried over the VPLS instance, such as ARP (Address Resolution Protocol), routing, multicast, and other control protocols, fail to operate correctly because packets are dropped during Layer 2 forwarding, while data-plane traffic forwarding remains unaffected. PR1945292
- On Junos EVO PTX platforms, IPv6 hosts lose connectivity in EVPN (MAC-VRF) CRB (Centralized Routing and Bridging) fabric because Router Solicitation (RS) packets are incorrectly dropped by the Layer 2 address-learning module (l2alm), preventing Router Advertisement (RA) exchange and causing IPv6 neighbor discovery to fail. PR1944894
- The dhcpd daemon may restart every 15 minutes, generating periodic syslog messages. This behavior occurs because the platform does not use DHCP/TFTP services for line card booting, but still inherits a protective mechanism designed for chassis systems with CPUbased FPCs.

  PR1936239
- On all Junos OS Evolved platforms with coherent optics (ZR/ZR-M/ZR-M-HP), when interface speed configuration is changed along with the optics application parameters in quick succession before the optics get initialized, the picd (port interface concentrator daemon) crashes, resulting in the interfaces going down. PR1931835
- On Junos Evolved PTX platforms, when Class-of-Service (CoS) forwarding-class is enabled, there's a mis-handle during route install. Due to ARP/NDP (Address Resolution Protocol (IPv4) / Neighbor

Discovery Protocol (IPv6)) timeout, the internal Next-hop changes are not updated properly, causing the route to point to invalid memory. It will affect the forwarding traffic to a destination that is tied to a Class-of-Service forwarding-class, with the classification-override feature enabled. As a result the route points to invalid Next-hop memory in hardware and traffic drops. PR1947926

- On PTX10004, PTX10008 and PTX10016 platforms running Junos OS Evolved, when a Switch Interface Board (SIB) experiences a power failure and is later replaced with a healthy SIB, the device fails to recognize the installed SIB. The affected SIB slot remains in an Empty state, resulting in reduced switch fabric capacity, loss of fabric redundancy, and potential traffic impact under load. PR1936725
- Reverting the **Null Route Detection** from the **default-on** to the **default-off**. This change ensures consistency with pre-Junos OS Evolved 25.4-EVO software behavior. PR1939686
- On the Junos OS Evolved PTX1000136MR platform, sFlow (Sampled Flow) samples were not generated because the sflowapp was not running on the device. PR1927361
- On all Junos OS Evolved platforms using the securityd application for MACsec, the securityd process crashes during restart scenarios such as FPC restart, PIC restart, or securityd restart causing temporary service disruption and Secure traffic drops or fail to establish. When PIC initialization is delayed and accessed prematurely, it results in a core dump of the securityd process, leading to disruption in MACsec services. PR1950757
- On the PTX10002-36QDD platform running Junos Evolved, any configuration that includes a Wide Area Network (WAN) interface with the keyword unused fails to commit, without impacting protocols or traffic on the WAN interface. PR1921097
- On Junos OS Evolved PTX platforms, the traffic will be dropped when IPv6 traffic is over IPv4 MPLS (Multiprotocol Label Switching.) tunnel under BGP (Border Gateway Protocol) 6VPE (IPv6 VPN Provider Edge) environment. PR1940918
- On PTX-series Junos EVO platforms, when no-graceful-switchover is configured and a forced RE switchover is performed, the CFM and LLDP subsystems fail to start on the newly active RE. Routing protocols and traffic forwarding are not affected. PR1940278
- On all Junos OS Evolved platforms, optics modules intermittently reports extremely high false temperature values in 'show interfaces diagnostics optics' command. This can trigger false high-temperature alarms, leading to unnecessary optics shutdown.

  PR1937870
- A license process crash is observed on Junos OS Evolved platforms when AE (Aggregated Ethernet) interface configuration is deactivated and activated or modified. PR1926569

Infrastructure

- On Junos OS Evolved platforms, when fibd reaches the process limit on the maximum number of allowed open file descriptors, fibd is unable to open new file descriptors and track new reader processes. Given this, processes such as Jade, SSHd and Stunnel are not found by fibd, triggering the core files for these processes affecting device management and Tx/Rx packet process. The Soft Limit value within the "cat /proc/<PID>/limits" output command will show the maximum number of file descriptors that fibd can have opened. PR1860358

- This PR address the following error while performing commit on the PTX EVO platform: # commit error: Failed to get RE list This is caused by a kernel software defect. PR1909995

- On Junos Evolved platforms, when Open Shortest Path First version 3 (OSPFv3) Internet Protocol Security (IPsec) SA is applied at the interface (IFD/IFL) level, Label Distribution Protocol (LDP) is able to discover the neighbor via User Datagram Protocol (UDP) Hellos, but fails to bring up the Transmission Control Protocol (TCP) session, preventing the LDP session from becoming operational. Disabling the OSPFv3 IPsec SA immediately restores normal LDP operation. During the service impact, the session will be down and traffic will not be send. Once the LDP session comes up, traffic can be routed through it. PR1919313

Interfaces and Chassis

- On Junos OS Evolved platforms with dual Routing Engine (RE) systems or Virtual Chassis (VC) setup, in an MPLS (Multiprotocol Label Switching) network along with Label Distribution Protocol (LDP) signaling protocol and Point-to-Multipoint (P2MP) is used and recursive Forwarding Equivalence Class (FEC) is configured such as `set protocols ldp p2mp recursive fec` on the Label Edge Router (LER)/ Provider Edge (PE) device, the multicast traffic statistics don't show as expected. Additionally, in the above setup, if an RE switchover is performed or occurs due to reasons such as component failure, kernel issue on primary RE or the primary RE fails (but not limited to) on the LER/PE, the multicast traffic drop is observed. However, the multicast traffic drop occurs only during the brief moment of RE switchover. Once the switchover is completed and the label is reprogrammed, the traffic recovers automatically, resulting in minimal traffic loss.

  PR1917186

- On all Junos OS Evolved platforms, when enabling GRES while not enabling NSR simultaneously, dependency errors for vrrpd are reported 15 minutes after enabling GRES and occur after switching over the primary role if VRRP is in use.

  PR1949135

Network Management and Monitoring

- System connections are not established while subscribing to OC path using SSL local certificates after upgrading to Junos OS Evolved 24.4R2-S3. This issue applies only to Junos OS Evolved 24.4R2-S3 and 24.4R2-S4.

PR1933395

- On all Junos OS Evolved platforms, for AE interfaces, the SNMP OID ifPhysAddress (iso.3.6.1.2.1.2.2.1.6) returns the MAC address of a bundled physical member interface instead of the configured logical MAC address on the AE interface, resulting in an inaccurate result.

  PR1945323

- On Junos OS Evolved platforms, SNMP(Simple Network Management Protocol) management systems do not receive related to optics or transceiver lane and module alarm traps. PR1932267

Routing Policy and Firewall Filters

- On Junos OS Evolved platforms, commit operations with ~38005000 lo0 filter binds take ~4m42s (vs <1m expected). If commit exceeds 5 minutes, mgd times out, which cause configuration load failure and service impact. PR1935328

Routing Protocols

- On all Junos OS Evolved, when Long-Lived Graceful Restart (LLGR) is configured and a Border Gateway Protocol (BGP) neighbor goes down and reconnects after the LLGR stale timer expires, the Routing Protocol Daemon (rpd) process crashes leading to service disruption.

  PR1915893

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

EVPN

EZ-LAG simplified CLI settings for native VLAN ID and access mode interfaces, and trunk mode interface configuration change (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—You can use the following new configurations to define the peer provider edge (PE) device interfaces to the servers in an easy EVPN LAG (EZ-LAG) simplified configuration:

- Specify a user-defined native VLAN ID for trunk mode interfaces using the `native-vlan-id` *`vlan-id`* statement at the `[edit services evpn server` *`server-name`*`]` hierarchy level. You must include the existing `vlan-id-list [`*`vlan-id-list ...`*`]` statement to generate a trunk mode interface configuration.
- Generate PE-to-server interface configurations in access mode instead of trunk mode using the `access-vlan-id` *`access-vlan-id`* statement instead of the `vlan-id-list` statement. You can set `access-vlan-id` at the `[edit services evpn server` *`server-name`*`]` or `[edit services evpn server` *`server-name`* `mac-vrf-instance` *`instance-id`*`]` hierarchy levelss.

Also, the EZ-LAG commit script now generates trunk interface configurations with the `flexible-vlan-tagging` option instead of the `vlan-tagging` option.

[See Add an Access Mode Interface Or a Native VLAN ID for Trunk Interfaces, server (Easy EVPN LAG Configuration), and mac-vrf-instance (Easy EVPN LAG Configuration).]

Enhanced OISM with M-VLAN IRB method for external multicast connectivity (also called L2 uplink with EVPN-enabled IRB interfaces) (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—With optimized intersubnet multicast (OISM) in an EVPN-VXLAN network, you can connect the OISM border leaf (BL) devices to an external PIM domain using one of a few supported methods. When you use enhanced OISM, these platforms now support the multicast VLAN (M-VLAN) integrated routing and bridging (IRB) method, which establishes a Layer 2 (L2) uplink on the BL devices to the external PIM router. The BL devices use the M-VLAN and associated IRB interfaces in the EVPN instance to exchange multicast control and data traffic with external multicast sources and receivers. In the EVPN instance, configure the `pim-evpn-gateway irb` statement with the M-VLAN IRB logical interface at either the `[edit protocols evpn oism]` or the `[edit routing-instances` *`name`* `protocols evpn oism]` hierarchy level.

[See Configure Border Leaf Device OISM Elements with M-VLAN IRB Method and pim-evpn-gateway.]

EVPN-VXLAN cross-module forwarding consistency and hash verification using the CLI (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)—You can detect and diagnose EVPN-VXLAN forwarding inconsistencies across control, kernel, and forwarding modules to reduce traffic loss or duplication. Verify per-VLAN synchronization and errors with:

- `request system evpn forwarding vlans [detail | instance` *`name`* `| vlan` *`name`* `| module {mac-table|mac-ip-table}] [no-confirm]`
- `request system evpn hash-values vlans [instance` *`name`*`] [vlan` *`name`*`]`

Compute comparison digests and validate Type 2 route and next-hop consistency. These checks highlight missing programming elements or deleted bridge domains.

[See EVPN-VXLAN Forwarding Consistency and Hash Verification, request system evpn forwarding vlans, and request system evpn hash-values vlans.]

Anycast multihoming for EVPN-VXLAN Type 2 symmetric IRB and BUM traffic (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)—Use anycast multihoming with egress link protection (ELP) to achieve subsecond EVPN-VXLAN convergence for link and node failures with Type 2 symmetric IRB and Type 3 enhanced broadcast, unknown unicast, and multicast (BUM) traffic handling. This approach shifts load balancing to the underlay, reduces control-plane churn, and enables global repair. To configure anycast multihoming, use the following commands:

- `set protocols evpn anycast-multihoming [mlag-interop-mode | router-mac | inet/inet6 address | loopback-interface]`
- `set forwarding-options evpn-vxlan reroute-address`
- `set interfaces` *`aeX`* `esi anycast`

- (For BUM convergence) `set interfaces aeX esi df-election-granularity per-esi`

You can use Auto-ESI through Type 1 LACP. Constraints include no single-home support and no Type 5 anycast support.

EVPN maintenance mode scheduling (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—You can schedule the EVPN maintenance mode CLI to minimize traffic loss and streamline upgrades across `erb-leaf`, `core-distribution`, `collapsed-core`, and `ip-clos` deployments. Use `set protocols evpn maintenance-mode (collapsed-core | core-distribution | erb-leaf | ip-clos) action-type (force | validate-and-set) schedule-start-at time [schedule-end-at time]` to trigger timed designated forwarder (DF) preference changes, network isolation (NISO), and Type 5 route withdrawal. Enable the NTP-based DF with `set protocols evpn df-election-ntp`. The `force` option overrides precheck results. Validate with `request evpn validate-maintenance-mode-pre-checks deployment-mode` and monitor with `show evpn maintenance-mode-status`. For Data Center Interconnect (DCI) gateways, delay interconnect Ethernet segment identifier (I-ESI) recovery using `set protocols evpn interconnect-esi-bringup-delay seconds`.

[See EVPN Maintenance Mode for Multihomed Leaf Isolation.]

Global AS override for auto-derived route targets in EVPN (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5230-64CD, and QFX5700)—EVPN now supports global autonomous system (AS) override for auto-derived route targets. You can configure independent AS control for import and export route target generation by using the `import-as` and `export-as` statements. Alternatively, you can configure unified AS control for both import and export route target generation by using the `as-num` statement. This feature provides additional control over the AS value used for EVPN auto-derived route target generation in multi-AS deployments.

[See Global AS Override for Auto-Derived Route Targets.]

Junos Telemetry

Support for genstate YANG data models (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

`/genstate/evpn-instance-information/evpn-instance`

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Decoupled rendering for telemetry data streaming (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming. Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the `[edit event-options destinations` *`destination-name`* `archive-sites` *`url`*`]` hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

Network Management and Monitoring

NDP and DAD proxy support with user-route proxy (QFX5130-32CD)—User route proxy enables the device to use protocol-installed NDP, host routes to determine IPv6 reachability and respond to neighbor solicitation requests with its own MAC address. This capability helps maintain IPv6

connectivity between hosts that are in the same subnet but located behind different Layer 3 gateways. To enable user route proxy globally, use the `set protocols neighbor-discovery ndp-proxy user-route-proxy` configuration statement.

[See NDP Proxy and DAD Proxy and user-route-proxy.]

Support for full-length SNMP engine ID (QFX5220)—Use the `unique-id` *`hex-string`* option in the `engine-id` configuration statement to enable a full-length SNMP engine ID. The unique ID uses a hexadecimal string of the length between 5 and 32 bytes.

[See engine-id, Configure SNMP in Junos OS, and Configure Engine ID on SNMPv3.]

OpenConfig

OpenConfig gNMI heartbeat for on-change subscriptions (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—Use OpenConfig heartbeat for gNMI ON_CHANGE subscriptions to receive a full set of sensor data at each heartbeat interval, even when values do not change. The heartbeat provides periodic state validation and mitigates undetected data loss.

[See Configuring gNMI Heartbeat Interval for ON_CHANGE Subscriptions and Junos YANG Data Model Explorer.]

Precision Time Protocol (PTP)

Transparent clock support (QFX5700 and QFX5700E)—QFX5700 and QFX5700E switches now support the Precision Time Protocol (PTP) transparent clock feature. 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.

Note:

QFX5700 and QFX5700E switches:

- do not support syntonized transparent clock feature.
- support transparent clock only when the Routing Engine in slot 0 is the timeTransmitter and the Forwarding Engine Board (FEB) in slot 0 is used to handle traffic.

[See PTP Transparent Clocks.]

Routing Protocols

Interface-scoped TCP MD5, TCP keychain, and TCP-AO authentication and NSR support for BGP autodiscovered neighbors (QFX5700)—We support interface-scoped authentication for BGP autodiscovered peers in single-hop connectivity and is best suited for deployments with symmetric routing, where traffic for a TCP flow traverses the same logical interface in both directions. Authentication keys are associated directly with logical interfaces instead of the IP addresses or prefixes as BGP unnumbered peering uses IPv6 link-local addresses. Authentication is applied based on the logical interface over which the BGP session is established, and is independent of the dynamically learned peer address. We also support nonstop active routing (NSR) support for interface scoped authentication for BGP autodiscovered neighbors.

[See dynamic-neighbor and Interface-Scoped TCP Authentication.]

Software Installation and Upgrade

Block configuration commits during a Junos OS Evolved software upgrade or rollback (PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5250-64OE, QFX5250-64OE-DOT2L, QFX5700, and QFX5700E)—For supported products, Junos OS Evolved now blocks configuration commits:

1. When a Junos OS Evolved software upgrade or rollback is in progress.
2. After the upgrade or rollback command has finished, but before you reboot the system.

[See commit, request system software add (Junos OS Evolved), and request system software rollback (Junos OS Evolved).]

Support for TCP connectivity on management and WAN interfaces for non-SI applications (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10001-36MR-K, PTX10002-36QDD, PTX10002-60MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E)—You can enable TCP connectivity for statically linked or other non-Socket Intercept (non-SI) applications over management and WAN interfaces. Start each non-SI application with the `NO_SI_INTERCEPT` environment variable to enable TCP traffic processing.

[See Running Applications Without the Socket Intercept Library.]

User Access and Authentication

TACACS+ Over TLS for Secure Device Administration (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—This release introduces support for Transport Layer Security (TLS) for TACACS+ communication between network devices and TACACS+ servers. By leveraging TLS, this enhancement provides strong encryption, integrity protection, and mutual authentication, addressing the security limitations of legacy TACACS+ mechanisms. The previous weak protection methods are now deprecated, ensuring that sensitive AAA traffic is securely transmitted in modern deployments.

TACACS+ over TLS Support (ACX7020-AC, ACX7020-AC-C, ACX7020-DC, ACX7020-DC-C, ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10002-36QDD, PTX10003, PTX10004, PTX10008, PTX10016, PTX12008, QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—This release introduces support for **TACACS+ over TLS**, enhancing the security of device administration by replacing the legacy MD5-based obfuscation mechanism with modern TLS encryption. With this update, communication between Juniper devices and TACACS+ servers is protected by TLS, ensuring **confidentiality, integrity, and mutual authentication** using X.509 certificates.

Key highlights:

- **Secure transport** for TACACS+ authentication, authorization, and accounting traffic.
- **Mutual TLS (mTLS)** support, enabling certificate-based validation of both client and server.
- **Failover support** between TACACS+ over TLS, TACACS+ over TCP, RADIUS, and local authentication.
- **Routing instance compatibility**, including default, non-default, and management instances.
- **CLI configuration enhancements** to define trusted CA groups and client certificates for TLS connections.

This feature allows TACACS+ deployments to meet modern compliance requirements (e.g., FIPS-140-3) and strengthens overall AAA security in Junos and Evo platforms.

Additional Features

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

OISM with M-VLAN IRB method for external multicast connectivity (also called L2 uplink with EVPN-enabled IRB interfaces) (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—With

optimized intersubnet multicast (OISM) in an EVPN-VXLAN network, you can connect the OISM border leaf (BL) devices to an external PIM domain using one of a few supported methods. When you use regular OISM, these platforms now support the multicast VLAN (M-VLAN) integrated routing and bridging (IRB) method, which establishes a Layer 2 (L2) uplink on the BL devices to the external PIM router. The BL devices use the M-VLAN and associated IRB interfaces in the EVPN instance to exchange multicast control and data traffic with external multicast sources and receivers. In the EVPN instance, configure the `pim-evpn-gateway irb` statement with the M-VLAN IRB logical interface at either the `[edit protocols evpn oism]` or the `[edit routing-instances name protocols evpn oism]` hierarchy level.

[See Configure Border Leaf Device OISM Elements with M-VLAN IRB Method and pim-evpn-gateway.]

What's Changed

IN THIS SECTION

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

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

EVPN

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

General Routing

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for autoconfiguration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router?s link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Restrict PPPoE client lockout output to underlying logical interfaces with short-cycle protection**—The "show pope lockout" command shows PPPoE clients information only for underlying logical interfaces with `short-cycle-protection` enabled.

  [See show pppoe lockout.]

- **Traffic Engineering Database Import Identifier Update**—The help string for the `set protocols mpls traffic-engineering database import identifier` command has been updated to BGP-TE identifier for standard IGP instance. This change clarifies that the identifier configuration affects only the standard IGP instance. The BGP-LS advertisements in the `lsdist.0` table will now use the configured value instead of defaulting to Instance Identifier 0. This update ensures accurate Instance Identifier configurations in BGP-LS advertisements, enhancing network visibility and management. The routers running software later than Junos OS release 24.4.x version should verify that the correct Instance Identifier is being advertised. [See import (MPLS Traffic Engineering Database.]

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `edit services port-extender satellite <name> device-model` hierarchy, you can use the "?" to see a list of supported platforms. We have also added the `skip-lo0-config` statement under `edit service port-extender jnu` hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the `edit service port-extender satellite <name>` hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

[See port-extender (Services).]

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Pixel-Tracing Command Updates**—The pixel-tracing command suite has been updated to include the following options: `flush`, `disable`, and `enable`. These options allow you to control and manage pixel-tracing processes for routing tasks. Note that the `rib-sharding` option is removed from the CLI commands to prevent system crashes when sharding configuration is present in the system. Ensure that your configurations do not rely on unmentioned options to avoid potential disruptions in your workflow.

- **IFA metadata stack enhancement (QFX Series)**—Additional class-of-service fields are now encoded in the Inband Flow Analyzer (IFA) metadata. Along with the existing fields, three egress-queue related counters are included: Egress Queue Tx Bytes, Egress Queue Current Bytes (Current Buffer Occupancy), and Egress Queue Maximum Bytes (Peak Buffer Occupancy).

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

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
```

```
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)              - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)       -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Prevent multicast route installation on uninterested receivers**—Configure the `resolve-request-holdcount` statement under the `[edit routing-instances name routing-options multicast]` hierarchy to prevent the kernel from caching initial multicast packets during route resolution and avoid unnecessary route installation on devices without interested receivers. [See {https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/resolve-request-holdcount-edit-routing-instances-routing-options-multicast.html} resolve-request-holdcount].

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading,

copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

[See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **BGP route-refresh warning messages added**—New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

  [See show bgp diagnostics warnings.]

- **Inbound Convergence Status Added to BGP Neighbor Display (Junos OS and Junos OS Evolved)**—The `show bgp neighbor` command now displays the Inbound Convergence status and cause.

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **System Stability Improvement**—When BGP fails to bind to its listening port, the routing protocol process (RPD) now stays running and logs a warning instead of terminating after the maximum retry count, preventing a service outage. The `show bgp diagnostics overview` output includes a new `BGP listening port status of v4 and v6` field that reports `Active` when BGP is successfully listening on the configured port and `Failed` when the port is unavailable or occupied by another process. Error logging is also enhanced with the `BGP_LISTEN_STATUS` message: `BGP listen failed on port [X] after 64 attempts with error-%d for address family identifier (IPv4 or IPv6). Administrative clear needed.` This message is generated after `BGP_MAX_LISTEN_START_ATTEMPTS` failed attempts to bind to the listening port.

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.

User Interface and Configuration

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the

corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

- `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
- `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
- `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
- `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
- `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:
  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of `remove` operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents **statement not found** warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for QFX Series switches.

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

General Routing

- On QFX5130 platforms, in Virtual Extensible LAN (VXLAN) topologies, symmetric equal-cost multipath (ECMP) hashing may not be enforced when resilient hashing is concurrently enabled. PR1907397

Open Issues

IN THIS SECTION

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

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

EVPN

- On all Junos OS Evolved platforms with EVPN-VXLAN multihoming scenarios, in EVPN multihomed (ESI) deployments without core isolation, temporary traffic duplication might occur during a spine reboot or events that force forwarding reconvergence.

  PR1879311

General Routing

- QFX5700 TD4 chip: 16x100G PIC: commit check added when 40G and 100G transceivers are put to share a vendor serdes (serializer/deserializer) core.

  PR1900911

- During USB boot with a fully populated system, error messages were observed in the boot logs indicating IRQ was not handled. This has no functional impact, as the driver is not used on the QFX5241-32/64 platforms. PR1868237

- With a back-to-back connection using 800G ZR* optics, if one end is configured with native speed (default) and other end is configured wit 8x100G mode, then the expectation is that link is down on both ends because of the configuration mismatch between the ends; but one of the two ends shows as Up. The same behaviour is presented by 400G ZR* optics also. PR1884559

- DNS resolution for traceroute does not work for a router using the management VRF with Junos OS Evolved Release 23.4R2. PR1858650

- 400G Optical Transceivers on QFX5130-32C and QFX5220-32C have a nominal voltage drop of upto 2% . Although this remains within specification, optical modules are rated for 3.3V +/-5%, meaning any reduction in operating voltage directly erodes noise margin. The link becomes more susceptible to transient noise events, which may manifest as simultaneous interface flaps. PR1940885

- On all Junos OS Evolved QFX5000 platforms, the transceiver-related telemetry sensor paths does not function at the leaf level, as data retrieval from the leaf level does not return any information. This results in a loss of telemetry data for leaf-level sensor paths.

  PR1926570

Interfaces and Chassis

- 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

- The traffic works fine with RE0 and FEB0 combination.So customers should use the same combination instead of RE1 and FEB1. PR1870507

- On QFX5230-64CD platform, 400G DAC cable of 2.5m length and 4X100G DAC BO may 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

- Junk values are shown for Carrier transition counter sometimes on initiating "clear interfaces statistics all" command and picd app restart. This issue is observed on QFX5220-32CD, QFX5230-64CD, QFX5240-64OD. PR1864375

Routing Protocols

- QFX5240: After configuring IGMP snooping, the group count is not as expected because groups are showing in IRB interfaces. PR1888392

Resolved Issues

IN THIS SECTION

Learn about resolved issues in this release for QFX Series switches.

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

EVPN

- Multiple logical interfaces (IFLs) in service provider (SP) style on the same port and part of the same bridge domain is an unsupported configuration on Junos OS Evolved-based QFX Series devices, such as the QFX5130 Switch. This limitation is due to hardware constraints, where the ASIC (responsible for VLAN membership and egress translation) is mis-programmed under this configuration. This restriction applies to both VXLAN and non-VXLAN configurations.

  PR1927899

General Routing

- On Junos OS Evolved platforms with TPM2.0, after executing the zeroize command, TPM (Trusted Platform Module) keys are not regenerated on reboot causing the TPM dependent features like sZTP, gRPC to fail.

  PR1899669

- On all Junos OS Evolved platforms, Time-zone info changes to default UTC after upgrade is done with restart-upgrade PR1803511

- ARP resolution failure when quick flap occurs on core-facing interface in scaled setup.

  PR1816344

- In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) deployment, a remote MAC move event can cause the remote MAC IP entry to be removed after it is initially installed, resulting in missing Address Resolution Protocol (ARP) state and incomplete hardware programming on the remote Virtual Tunnel Endpoint (VTEP), which can lead to traffic loss for the affected endpoint. PR1936648

- When Weighted ECMP using BGP link bandwidth is configured, the total ECMP member count displayed in the pfe debug command `show npu utilization info` shows only half the utilization. The actual entries programmed in the asic will be double the displayed count.

  PR1896618

- On Junos OS Evolved QFX5130 and QFX5700 platforms with EVPN-VxLAN (Ethernet Virtual Private Network - Virtual Extensible LAN) configured, when manually adding/deleting leaked routes across multiple VRFs (Virtual routing and forwarding), some VRFs will be skipped unintentionally. This will leave stale routes in the routing table, leading to incorrect next-hop entries. As a result, traffic related to the affected IRB (Integrated Routing and Bridging) ARP (Address Resolution Protocol)/NDP (Neighbor Discovery Protocol) next-hop will be dropped. PR1944812

- On Junos Evolved QFX5130/5700 switches participating as Spine or Leaf devices in Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) architectures, the Virtual Gateway Media-Access Control address (VMAC) configured or the default VRRP-based MAC as the Layer3 gateway MAC might be

deleted. This results in loss of connectivity to the Layer3 MAC. Some MACIP deletion events might trigger this condition.

PR1940786

- In Junos Evolved QFX5220, QFX5230 and QFX5240 platforms, when IGMP-snooping is enabled on IRB VLAN interface and OSPF is configured on same IRB VLAN interface, OSPF adjacency fails with peer neighbour. PR1947511
- On all Junos OS Evolved platforms, optics modules intermittently reports extremely high false temperature values in 'show interfaces diagnostics optics' command. This can trigger false high-temperature alarms, leading to unnecessary optics shutdown.

  PR1937870

Interfaces and Chassis

- On QFX5240, QFX5241, QFX5130 and QFX5220 platforms, configuring interface range using the wildcard et-0/0/* causes the system to incorrectly expand the range beyond the actual number of physical Ethernet ports. As a result, any subsequent configuration applied to the affected interface range fails to commit due to the presence of invalid interfaces. PR1853302
- On QFX5130-48CM platforms, when the Physical Interface Card (PIC) daemon (picd) restarts, interfaces et-0/0/0 and et-0/0/1 will go down and fail to recover, resulting in traffic loss on both ports. PR1898893

Subscriber Access Management

- On all Junos OS Evolved platforms, the error message **UI_SCHEMA_SEQUENCE_ERROR** is observed when device is restarted. There is no traffic impact due to this issue.

  PR1813456

Junos OS Evolved Release Notes for Third-Party White Box

IN THIS SECTION

What's New

IN THIS SECTION

Learn about new features introduced in this release for Third-Party White Box.

Hardware

UfiSpace-S9500-22XST Gateway Router (UFISPACE-S9500-22XST)—The S9500-22XST Gateway Router is a 1-U third-party white box from UfiSpace. The UfiSpace router supports Junos OS Evolved for Layer 2 Bit Stream Access (L2BSA) in a broadband network gateway (BNG) application.

L2BSA allows the access infrastructure provider to provide wholesale connectivity between subscribers and their respective network service providers. Subscriber traffic is bundled transparently through EVPN VPWS tunnels to the S9500-22XST Gateway Router and handed off to the respective service providers across an A10-NSP interface.

Service providers can manage their customer premises equipment by connecting their data communication network (DCN) to a Layer 3 VPN (L3VPN) service established between the S9500-22XST and the relevant leaf switches.

The S9500-22XST supports all necessary features specific to the L2BSA application, including but not limited to the following:

General Area	Summary
IPv4/IPv6/MPLS infrastructure	• Basic infrastructure to support L2 circuit, L2VPN, L3VPN, EVPN-VPWS, including regular router functions such as routing and forwarding, longest prefix match, fragmentation, exception handling, VLANs (single, double, untagged), ARP, ND, ICMP, LAG and LACP, LLDP, BFD, DDOS, and more • LSR, LER, SRLG, LDP, RSVP-TE, MPLS protection • RIP, OSPF, and IS-IS including respective traffic engineering extensions where applicable • BGP including multipath, prefix independent convergence, core and edge, with support for `inet`, `inet6`, `inet-vpn`, and `inet6-vpn` address families • Equal cost multipath (ECMP) • Unicast reverse path forwarding (loose and strict, including dual stack) • IP and MPLS fast reroute (FRR) • IP FRR, LFA, and R-LFA with MPLS next hop • E-Line, E-LAN, E-ACCESS, and E-Transit services
SR-MPLS	• Maximum four labels (two labels for transport and two labels for service) • IP FRR (on post convergence path) and Topology-Independent Loop-Free Alternate (TI-LFA) • Prefix/node and anycast segments • SR-to-LDP mapping
EVPN-VPWS	• EVPN-based E-Line service • Control word, VLAN maps, composite chain next hop • Single-homing and multihoming • Flexible cross connect (FXC) aware and unaware

(Continued)

General Area	Summary
L3VPN	• Full mesh, hub-spoke, and more • Inter-AS options A and B • Chained-composite mode • LDP-DoD
L2 circuit and L2VPN	• LDP (L2 circuit) and BGP (L2VPN) • Control word, VLAN maps, configurable TPIDs (Tag Protocol Identifiers) • Indirect and composite next hop models
Firewall and policing	• Support for `inet`, `inet6`, `ccc`, `mpls` (ingress), and `any` address families • Ingress policer • Multifield classification and rewrite including DSCP, PCP, and EXP

NOTE: You cannot use the S9500-22XST for any other application.

Class of Service (CoS)

Support for new CoS CLI configuration attributes and variables (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—We've added new class-of-service (CoS) CLI configuration attributes and variables for dynamic profiles:

- shaping-rate-granularity attribute at the `[edit dynamic-profiles` *`profile-name`* `class-of-service traffic-control-profiles` *`profile-name`*`]` hierarchy level—Set a shaping rate granularity of the default 2.6Mbps or a smaller rate of 128Kbps.
- number-of-queues attribute at the `[edit dynamic-profiles` *`profile-name`* `class-of-service traffic-control-profiles` *`profile-name`*`]` hierarchy level—Define the number of queues per subscriber at the traffic control profile (TCP) level.
- `temporal-absolute` buffer-size option at the `[edit dynamic-profiles` *`profile-name`* `class-of-service schedulers` *`scheduler-name`*`]` hierarchy level—Set the buffer size as an absolute value in bytes for a dynamic profile scheduler.

- restricted-queue attribute at the [edit dynamic-profiles *profile-name* class-of-service forwarding-classes class] hierarchy level—Map a forwarding class to an output queue number when number-of-queues is set to 4.

Junos Telemetry

Telemetry support for business subscriber accounting (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—Logical interfaces are provisioned as part of a business subscriber service and are used for accounting and telemetry. Junos telemetry now supports streaming business subscriber accounting telemetry data to data collectors. You can stream logical interface, class-of-service (CoS) queue, and hierarchical policer (h-policer) statistics over the Junos telemetry infrastructure. Use the new extended-telemetry CLI option under the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy to enable telemetry data collection per logical interface. Subscribe to the following new resource path to retrieve telemetry data:

/state/interfaces/interface/subinterfaces/subinterface/extended-counters

You can retrieve the following counters through a telemetry subscription:

- Logical interface-level statistics:
  - input-bytes
  - input-packets
  - output-bytes
  - output-packets
- CoS queue statistics:
  - queue-id (0–7)
  - bytes (0–7)
  - packets (0–7)
  - drop-packets (0–7)
- Hierarchical policer (H-policer) statistics:
  - policer-level (0–3)
  - bytes (0–3)
  - packets (0–3)
  - drop-packets (0–3)

The *interface-information* group is defined in the *extended-interface-telemetry.yang* data model.

For a complete list of sensor paths supported by the device, see Junos YANG Data Model Explorer.

[See Junos YANG Data Model Explorer and extended-telemetry.]

Export sensor data to external collectors using Junos telemetry (UFISPACE-S9500-22XST, UFISPACE-S9600-102XC, UFISPACE-S9600-32X, and UFISPACE-S9600-72XC)—Use Junos telemetry to stream device and operational statistics from UFISPACE-S9600-32X, UFISPACE-S9600-72XC, UFISPACE-S9600-102XC, and UFISPACE-S9500-22XST devices to external collectors. Monitor system health by subscribing to sensor paths for chassis temperature, CPU and memory utilization, fan speed, transceiver parameters, interface statistics, Timex statistics, and subscriber details. Obtain platform, routing, process, system resource, port, optical, H-policer, PSM fan, and event metrics by subscribing to the respective resource paths.

You can view all the supported sensor paths in the Junos YANG Data Model Explorer.

Multicast

Establish multicast group memberships manually, and install forwarding routes using control-plane signaling (UFISPACE-S9600-102XC, UFISPACE-S9600-32X, and UFISPACE-S9600-72XC)—You can manually establish multicast group memberships using PIM static joins to bypass IGMP or MLD join messages. Configure PIM static joins by enabling the `static` configuration statement at the `[edit protocols pim]` hierarchy level. To install multicast (S,G) forwarding routes using control-plane signaling, configure the `control-driven-sg` option at the `[edit protocols pim]` hierarchy level.

[See Configuring PIM Static Joins, static (Protocols PIM), and control-driven-sg.]

Subscriber Management and Services

Fragmentation support for large IPv4 packets (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—You can manage maximum transmission unit (MTU) exceptions across IPv4, IPv6, and PPP traffic by enabling packet fragmentation or generating ICMP error messages. For IPv4 packets, the device generates an `ICMP Packet Too Big` message when the Don't Fragment (DF) bit in the packet header is `1`. When the DF bit is `0`, you can enable fragmentation using the `set forwarding-options family inet fragmentation-enable` command. For IPv6 packets, the device can only generate an `ICMPv6 Packet Too Big` message because intermediate nodes cannot fragment IPv6 packets.

To avoid excessive CPU load caused by fragmentation and ICMP/ICMPv6 message generation, you can configure DDoS rate limiting for MTU-exceeded packets using the `set system ddos-protection protocols exceptions mtu-exceeded bandwidth rate` command.

[See forwarding-options (Routing Instance), bandwidth (DDoS), mtu (interfaces), and mtu (Dynamic and Static PPPoE).]

Support for PAO REST interface (UFISPACE-S9500-22XST)—Use the Pod Access Orchestrator (PAO) REST interface to manage A10NSP devices in the DT-A4 network using REST APIs. The device's REST agent processes requests from PAO and supports secure REST communication for management operations. Supported functionalities include HTTPS communication, JWT-based token authentication, NETCONF RPC requests, running configuration updates, ephemeral database updates, REST API counters, and communication through in-band and out-of-band management VRFs. Use this enhancement to automate device management and operational workflows in DT-A4 A10NSP deployments.

Support for L2BSA services management (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—You can use Layer 2 Bitstream Access (L2BSA) services to wholesale subscriber traffic to a network service provider. This feature supports dynamic VLAN creation for L2BSA subscribers and forwards subscriber traffic through Ethernet VPN Virtual Private Wire Services (EVPN-VPWS). You can manage L2BSA services in the following ways:

- Use the `set system services subscriber-management external-management vlans l2bsa-dynamic-profile` *`profile-name`* command to create dynamic VLAN profiles for L2BSA subscribers
- Use the `[edit dynamic-profiles]` hierarchy to dynamically configure and apply services such as CoS and firewall.
- Use the `show subscribers detail` and `show subscribers extensive` commands to view L2BSA subscriber statistics, including policer and queue statistics. The output displays the subscriber type as `L2BSA-EXTM`.

You must complete these configurations before any L2BSA subscriber logs in. You can use multiple routing instances to isolate L2BSA subscribers with different configurations.

[See external-management, dynamic-profiles, demux0 (Dynamic Interface), and show subscribers.]

Support for new VSA attribute and variables (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—The authentication, authorization, and accounting (AAA) service framework supports the vendor-specific attribute (VSA) 157 and variables that map to class-of-service (CoS) parameters for traffic shaping, queuing, and rewrite rules.

We've added the following variables to the new VSA attribute `Jnpr-Cos-Rewrite-Rules`:

- `$junos-cos-rewrite-rules-dscp`
- `$junos-cos-rewrite-rules-dscp-ipv6`
- `$junos-cos-rewrite-rules-ieee802-1`

These variables dynamically apply rewrite-rule parameters to map forwarding classes and loss priorities.

We've added the following variables to the existing VSA attribute `Jnpr-Cos-Parameter-Type`:

- `$junos-cos-shaping-rate-granularity`

- `$junos-cos-number-of-queues`

These variables dynamically apply traffic-shaping and scheduling parameters to manage output queues and bandwidth shaping.

[See Juniper Networks VSAs Supported by the AAA Service Framework and Predefined Variables That Correspond to RADIUS Attributes and VSAs.]

Additional Features

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

Support for configuring output firewall filters on aggregated Ethernet (AE) interface (UFISPACE-S9500-22XST)—Support is added for configuring output firewall filters (only 16 AE IFLs) on AE interface with policer configuration.

[See Configuring Firewall Filters.]

Firewall Filter and Policers Support (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—

- Support for logical interface policer type only for enhanced hierarchical policer for static IFLs.
- Support of a maximum of four levels of enhanced hierarchical policer for static IFLs by configuring CIR (Committed Information Rate), CBS (Committed Burst Size), and max-CIR (Peak Information Rate/ Shaping Rate) as 0 (zero) for unrequired levels.
- Support of optional max-CIR rate for a given policer level. Auto calculation of max-CIR would be aggregate of current-level CIR and all previous levels' CIR. In the following example configuration, `medium-high` and `medium-low` levels are optionally skipped by configuring `0`. Also, max-CIR of `low` is also optionally skipped. Note that all levels are mandatory in the config, but with `0` values a level can be ignored.

```
set firewall enhanced-hierarchical-policer hpol-test filter-specific

set firewall enhanced-hierarchical-policer hpol-test high committed-burst-size 15k

set firewall enhanced-hierarchical-policer hpol-test high committed-information-rate 1m

set firewall enhanced-hierarchical-policer hpol-test high max-committed-information-rate 1m

set firewall enhanced-hierarchical-policer hpol-test high then discard

set firewall enhanced-hierarchical-policer hpol-test medium-high committed-burst-size 0

set firewall enhanced-hierarchical-policer hpol-test medium-high committed-information-rate 0

set firewall enhanced-hierarchical-policer hpol-test medium-high max-committed-information-rate 0

set firewall enhanced-hierarchical-policer hpol-test medium-low committed-burst-size 0
```

```
set firewall enhanced-hierarchical-policer hpol-test medium-low committed-information-rate 0

set firewall enhanced-hierarchical-policer hpol-test medium-low max-committed-information-rate 0

set firewall enhanced-hierarchical-policer hpol-test low committed-burst-size 1500

set firewall enhanced-hierarchical-policer hpol-test low committed-information-rate 10m

set firewall enhanced-hierarchical-policer hpol-test low then discard
```

- Support of global firewall filter for family INET of non-default loopback interfaces (lo0.1, lo0.2..lo0.x) in ingress direction.
- Support of global firewall filter for family INET6 of non-default loopback interfaces (lo0.1, lo0.2..lo0.x) in ingress direction.
- Support of `forwarding-class` firewall filter match for family inet/inet6/mpls/ccc/ethernet-switching firewall filters in ingress direction for all IFLs.
- Support of `dscp` firewall filter match for inet matching the outer tunnel header's `dscp` value for L2TP terminated PPP packet flows in downstream ingress direction.

[See Routing Policies, Firewall Filters, and Traffic Policers User Guide.]

Firewall filter support for L3 split horizon group for management VRFs on Leaf Devices (UFISPACE-S9600-102XC and UFISPACE-S9600-72XC)—

- Support of `interface-group` match for inet/inet6 family filters in egress direction for static IFLs.
- Support of `interface-group` configuration per inet/inet6 family per IFL with value 0-1.
- Support for packet dropping and counting to support L3 split horizon.
- Support of global firewall filters for inet and inet6 families in egress direction for static IFLs.

[See Routing Policies, Firewall Filters, and Traffic Policers User Guide.]

Support for EVPN-virtual private wire service (UFISPACE-S9500-22XST) (UFISPACE-S9500-22XST)—
[See Overview of VPWS with EVPN Signaling Mechanisms.]

What's Changed

IN THIS SECTION

Learn about what's changed in this release for Third-Party White Box.

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

General Routing

- **Optics diagnostics vendor material number support (UfiSpace S9600-32X, S9600-72XC, S9600-102XC, and S9500-22XST)**—Command introduced in Junos OS Evolved Release 26.2 to display vendor material numbers for optics on UfiSpace-based systems. The show interfaces diagnostics optics vendor-extension interface-name command now provides vendor material information for the interface. Example: user@host> show interfaces diagnostics optics vendor-extension et-0/0/30 Physical interface: et-0/0/3 Vendor material number : T-UNISFP40340960 The vendor material number is applicable only on UfiSpace's vendor optics.

Known Limitations

There are no known limitations in hardware or software in this release for Third-Party White Box.

Open Issues

There are no known issues in hardware or software in this release for Third-Party White Box.

Resolved Issues

There are no resolved issues in this release for Third-Party White Box.

Migration, Upgrade, and Downgrade Instructions

For products impacted, see Feature Explorer.

Follow these steps to upgrade your Junos OS Evolved software:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: 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 Junos OS Evolved Dates & Milestones - Juniper Networks.

Documentation Updates

There are no documentation updates for this release.

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

- **Software Release Notification**—Technical Support Bulletin that describes the open issues and resolved issues in <Junos OS/Junos OS Evolved xxx>. 25.2R1-S2-EVO: Software Release Notification for JUNOS Evolved Software
- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.

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

  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.

  Compliance Advisor.

Requesting Technical Support

IN THIS SECTION

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit Product Warranty Policy - Support - Juniper Networks.
- 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: Support
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: Documentation
- Find solutions and answer questions using our Knowledge Base: CEC Juniper Community

- Download the latest versions of software and review release notes: Support
- Search technical bulletins for relevant hardware and software notifications: CEC Juniper Community
- Join and participate in the Juniper Networks Community Forum: Elevate - Elevate Community

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

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: Support
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: Documentation
- Find solutions and answer questions using our Knowledge Base: CEC Juniper Community
- Download the latest versions of software and review release notes: Support
- Search technical bulletins for relevant hardware and software notifications: CEC Juniper Community
- Join and participate in the Juniper Networks Community Forum: Elevate - Elevate Community

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

Creating a Service Request with JTAC

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

- Visit Juniper Support Portal: Case Management, Product Support & More
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see Contact - Support - Juniper Networks.

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

- Visit Juniper Support Portal: Case Management, Product Support & More
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see Contact - Support - Juniper Networks.

Revision History

15-Jul-26 - Revision 3, Junos OS Evolved Release 26.2R1.

08-Jul-26 - Revision 2, Junos OS Evolved Release 26.2R1.

29-Jun-26 - Revision 1, Junos OS Evolved Release 26.2R1.