

# Release Notes: Junos<sup>®</sup> OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices

28 January 2022

|                 |   |
|-----------------|---|
| <b>Contents</b> | <b>Introduction   4</b>   |
|                 | <b>New Features in Junos OS Evolved 20.1R2   4</b>                          |
|                 | <b>Junos OS Evolved Release Notes for PTX10003 and PTX10008 Devices   5</b> |
|                 | <b>What's New   6</b>   |
|                 | What's New in Junos OS Evolved Release 20.1R2   6                           |
|                 | What's New in Junos OS Evolved Release 20.1R1   17                          |
|                 | <b>What's Changed   45</b>  |
|                 | Changes in CLI Statements and Commands   46                                 |
|                 | What's Changed in Junos OS Evolved Release 20.1R2   46                      |
|                 | What's Changed in Junos OS Evolved Release 20.1R1   46                      |
|                 | <b>Known Limitations   48</b>   |
|                 | General Routing   48  |
|                 | Flow-Based and Packet-Based Processing   50                                 |
|                 | Infrastructure   50   |
|                 | Interfaces and Chassis   50   |
|                 | User Interface and Configuration   51                                       |
|                 | <b>Open Issues   51</b>   |
|                 | Class of Service (CoS)   52   |
|                 | Fault Management   52   |
|                 | General Routing   52  |
|                 | Infrastructure   57   |

|  |    |
|--|----|
| Interfaces and Chassis                             | 57 |
| Interfaces and Chassis                             | 58 |
| Network Management and Monitoring                  | 58 |
| Routing Protocols                                  | 58 |
| Resolved Issues                                    | 58 |
| Resolved Issues: 20.1R2                            | 59 |
| Resolved Issues: 20.1R1                            | 62 |
| Junos OS Evolved Release Notes for QFX5220 Devices | 64 |
| What's New   | 65 |
| What's New in Junos OS Evolved Release 20.1R2      | 66 |
| What's New in Junos OS Evolved Release 20.1R1      | 66 |
| What's Changed                                     | 71 |
| Changes in CLI Statements and Commands             | 71 |
| What's Changed in Junos OS Evolved Release 20.1R2  | 71 |
| What's Changed in Junos OS Evolved Release 20.1R1  | 71 |
| Known Limitations                                  | 72 |
| Interfaces and Chassis                             | 72 |
| Open Issues  | 73 |
| Routing Policy and Firewall Filters                | 73 |
| Routing Protocols                                  | 73 |
| System Management                                  | 73 |
| User Interface and Configuration                   | 74 |
| Resolved Issues                                    | 74 |
| Class of Service (CoS)                             | 75 |
| Infrastructure                                     | 75 |
| Interfaces and Chassis                             | 75 |
| Platform and Infrastructure                        | 75 |
| System Management                                  | 75 |
| User Interface and Configuration                   | 75 |
| Upgrade your Junos OS Evolved Software             | 76 |
| Finding More Information                           | 78 |
| Documentation Feedback                             | 78 |

Requesting Technical Support | 80

Self-Help Online Tools and Resources | 80

Creating a Service Request with JTAC | 81

Revision History | 81

# Introduction

Junos OS Evolved is the next-generation Junos OS. It has the same CLI, the same features, and, in some cases, even the same processes as on the previous versions of Junos OS. However, its infrastructure is entirely modernized.

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 20.1R2.

These release notes are cumulative and are updated for later releases.

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

## New Features in Junos OS Evolved 20.1R2

| Feature  | Release Note Section                   |
|--|--|
| Classification override support (PTX10008)   | <a href="#">“What’s New” on page 6</a> |
| Load balancing under enhanced-hash-key (PTX10008)  | <a href="#">“What’s New” on page 6</a> |
| Fabric resiliency (PTX10008)   | <a href="#">“What’s New” on page 6</a> |
| GRE tunnel decapsulation using flexible tunnel interfaces (FTI) (PTX10008)                               | <a href="#">“What’s New” on page 6</a> |
| Support for Pre-FEC BER and link degradation monitoring (PTX10008)                                       | <a href="#">“What’s New” on page 6</a> |
| QSFP28 100GE DWDM transceivers support (PTX10008)  | <a href="#">“What’s New” on page 6</a> |
| Support for DCU accounting and SCU accounting (PTX10008 routers)   | <a href="#">“What’s New” on page 6</a> |
| Support for port profile configuration on JNP10K-LC1201 line cards using Interfaces hierarchy (PTX10008) | <a href="#">“What’s New” on page 6</a> |
| Support for 2M routes (PTX10008)   | <a href="#">“What’s New” on page 6</a> |
| JTI support for kernel statistics (PTX10003)   | <a href="#">“What’s New” on page 6</a> |

| Feature   | Release Note Section                   |
|---|--|
| UDP (native) PFE sensor support for JTI (PTX10008)                      | <a href="#">“What’s New” on page 6</a> |
| 6VPE and 6PE for PTX10008   | <a href="#">“What’s New” on page 6</a> |
| Segment routing support (PTX10008)                                      | <a href="#">“What’s New” on page 6</a> |
| Local port mirroring support (PTX10008)                                 | <a href="#">“What’s New” on page 6</a> |
| Firewall filters support (PTX10008)                                     | <a href="#">“What’s New” on page 6</a> |
| Support for matching ip-options in IPv4 packet headers (PTX10008)       | <a href="#">“What’s New” on page 6</a> |
| Support for BGP LU FRR MPLS inter-AS link protection (PTX 10008)        | <a href="#">“What’s New” on page 6</a> |
| Support for FTI based decapsulation of IPv4 and IPv6 packets (PTX10008) | <a href="#">“What’s New” on page 6</a> |
| Support for passive monitoring (PTX10008)                               | <a href="#">“What’s New” on page 6</a> |
| Layer 3 VPN on PTX10008   | <a href="#">“What’s New” on page 6</a> |
| Layer 2 VPN support (PTX10008)  | <a href="#">“What’s New” on page 6</a> |

## Junos OS Evolved Release Notes for PTX10003 and PTX10008 Devices

### IN THIS SECTION

- [What’s New | 6](#)
- [What’s Changed | 45](#)
- [Known Limitations | 48](#)
- [Open Issues | 51](#)
- [Resolved Issues | 58](#)

These release notes accompany Junos OS Evolved Release 20.1R2 for PTX10003 and PTX10008 (with the JNP10008-SF3 SIB) Packet Transport Routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

## What's New

### IN THIS SECTION

- [What's New in Junos OS Evolved Release 20.1R2 | 6](#)
- [What's New in Junos OS Evolved Release 20.1R1 | 17](#)

Learn about new features introduced in Junos OS Evolved Release 20.1R2 for the PTX10003 and PTX10008.

### What's New in Junos OS Evolved Release 20.1R2

#### IN THIS SECTION

- [Class of Service \(CoS\) | 7](#)
- [Flow Based Packet Based Processing | 7](#)
- [High Availability and Resiliency | 8](#)
- [Interfaces and Chassis | 8](#)
- [IPv4 | 10](#)
- [Junos Telemetry Interface | 10](#)
- [MPLS | 11](#)
- [Network Management and Monitoring | 13](#)
- [Routing Policy and Firewall Filters | 14](#)
- [Routing Protocols | 14](#)
- [Services Applications | 16](#)
- [System Management | 16](#)
- [VPNs | 16](#)

Learn about new features introduced in Junos OS Evolved Release 20.1R2 for PTX10003 and PTX10008.

### **Class of Service (CoS)**

- **Classification override support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, we've extended support for classification override configured under a forwarding policy to PTX10008 devices.

[See [CoS Features and Limitations on PTX Series Routers](#) and [Overriding the Input Classification](#).]

### **Flow Based Packet Based Processing**

- **Load balancing under enhanced-hash-key (PTX10008)**—Starting with Junos OS Evolved Release 20.1R2, load balancing is introduced under the **[edit forwarding-options enhanced-hash-key]** hierarchy. The bit length of hash material is only limited by performance requirements. Many fields are included by default in hash computations (such as Layer 3, Layer 4, MPLS payloads, and inner tunnel packet fields), with an option to exclude some fields if problems are seen. Only a few fields, such as **incoming-interface-index**, are *not* included by default.

Fields in the hash computation do not depend on the way a packet is forwarded (for example, whether it is Layer 2 bridging, cross-connect, or Layer 3 forwarding, including MPLS forwarding). The CLI hierarchy **[edit forwarding-options enhanced-hash-key family any]** allows configuration of family-independent statements.

Load balancing includes:

- GRE key inclusion for transit IPv4 and IPv6 traffic
- IP Layer 3 fields
- IP Layer 4 fields
- IPv6 flow label inclusion
- MPLS labels
- MPLS port data
- MPLS pseudo-wire traffic
- TEID inclusion in GTP packets
- RSVP-TE load balancing in proportion to LSP bandwidth

The set of **enhanced-hash-key family** configuration attributes includes:

|                            |                             |
|----------------------------|-----------------------------|
| inet/no-source-port        | inet6/gtp-tunnel-identifier |
| inet/no-destination-port   | inet6/no-flow-label         |
| inet/type-of-service       | mpls/label-1-exp            |
| inet/gtp-tunnel-identifier | mpls/no-labels              |

|                           |                              |
|---------------------------|------------------------------|
| inet6/no-source-port      | mpls/no-payload              |
| inet6/no-destination-port | any/no-tunnel-payload        |
| inet6/type-of-service     | any/incoming-interface-index |

### High Availability and Resiliency

- **Fabric resiliency (PTX10008)**—PTX10008 routers with Junos OS Evolved Release 20.1R2 support fabric resiliency. The fabric resiliency includes handling of failures and faults related to the switch fabric, including errors in the fabric connectivity between the line card and switch fabric. Fault handling could include detecting and logging the error, raising alarms, and self-healing.

When errors related to the switch fabric occur, the software takes the following default actions:

- Creates a syslog entry for each error (irrespective of the error severity).
- When an error occurs during boot time which prevents the switch fabric to come online successfully, the software attempts to restart the impacted switch fabric. If the error persists, a maximum of three restart attempts are made. If all the attempts fail, the software moves the impacted switch fabric to **fault** state.
- When an error occurs during run time, the software handles the fault based on the severity of the error. In case of a major error during run time, the software logs the error and raises a major alarm. When a fatal error occurs, the software attempts to restart the impacted switch fabric. If the fatal error is not resolved after three attempts, software moves the impacted switch fabric to **offline** state.

[See [show system errors active](#).]

### Interfaces and Chassis

- **GRE tunnel decapsulation using flexible tunnel interfaces (FTI) (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, you can configure GRE tunnels on flexible tunnel interfaces. When you enable the **tunnel-termination** statement at the **[edit interfaces fti0 unit unit-number family (inet | inet6)]** CLI hierarchy, tunnels are terminated on the WAN interface before any other actions, such as sampling, port mirroring, or filtering, are applied.

GRE tunnel encapsulation is not supported.

- **Support for Pre-FEC BER and link degradation monitoring (PTX10008 routers)**—Starting in Junos OS Evolved Release 20.1R2, you can monitor the quality of physical links on Ethernet Interfaces and take corrective action when the link quality degrades beyond a certain value and the pre-forward error correction (pre-FEC) bit error rate (BER). To enable your device to monitor the links, use the **link-degrade-monitor** statement at the **[edit interfaces interface-name]** hierarchy level. This feature monitors the BER of the link and initiates corrective action when the BER value crosses a user-configured threshold.

[See [Link Degrade Monitoring Overview](#).]



- **Support for QSFP28 100GbE DWDM transceivers (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, the JNP10K-LC1201 line card supports the QSFP28 100GbE DWDM optical transceiver module (QSFP-100GE-DWDM2). This transceiver is suited for data center interconnect, which requires high-fiber-capacity for links up to 80 km.

The Junos OS Evolved features supported by the QSFP28 100GbE DWDM transceiver include the following:

- View the QSFP28 optics inventory information. [See [show chassis hardware and show chassis pic fpc-slot slot pic-slot slot.](#)]
  - View the diagnostics data, warnings, and alarms for interfaces. [See [show interfaces diagnostics optics.](#)]
  - Configure the optics loopback mode. [See [optics-options.](#)]
- **Support for DCU accounting and SCU accounting (PTX10008 routers)**—Starting in Junos OS Evolved Release 20.1R2, PTX10008 routers support destination class usage (DCU) accounting and source class usage (SCU) accounting. With SCU accounting, you can track traffic that originates from specific prefixes on the provider core. You can use DCU accounting to track traffic that originates from the customer edge and is destined for specific prefixes on the provider core router.

[See [Understanding Source Class Usage and Destination Class Usage Options.](#)]

- **Support for port profile configuration on JNP10K-LC1201 line cards using interfaces hierarchy (PTX10008 routers)**—Starting in Junos OS Evolved Release 20.1R2, you can configure the port speed on the JNP10K-LC1201 line card by using the port profile configuration statements in the **[edit interfaces]** hierarchy. To streamline the configuration, we've migrated the port profile configuration statements from the **[edit chassis]** hierarchy to the **[edit interfaces]** hierarchy for the JNP10K-LC1201 line card.

On the JNP10K-LC1201 line card:

- To configure the port speed of a physical port in a PIC, use the **speed** statement at the **[edit interfaces interface-name]** hierarchy. If you do not configure the speed, the default speed of the port is assigned to the non channelized interface. When you change the speed of a port, all the affected channelized or non channelized interfaces are deleted and new interfaces are created with the new speed.
- To configure the port speed of a physical port with a specified number of channelized interfaces, use the **speed** statement and specify the number of channelized interfaces by using the **number-of-sub-ports** statement at the **[edit interfaces interface-name]** hierarchy. All channelized interfaces have the same speed. You cannot configure an individual speed for each channelized interface. When you change the number of channelized interfaces, new channelized interfaces are created or existing channelized interfaces are deleted.
- To control the number of interfaces (channelized or non channelized) created on a physical port, use the **unused** statement at the **[edit interfaces interface-name]** hierarchy. If you configure a physical port as unused, no interfaces (channelized or non channelized) are created for that port. By default, all the physical ports in an interface are active. When you change the number of active ports, all affected channelized or non channelized interfaces are deleted and new interfaces are created.

In Junos OS Evolved releases before Release 20.1R2, you configure the port speed on the JNP10K-LC1201 line card by using the port profile configuration statements at the `[edit chassis]` hierarchy level. We no longer support the following statements:

- **pic-mode** statement at the `[edit chassis fpc fpc-slot-number pic pic-number]` hierarchy.
- **speed** statement at the `[edit chassis fpc fpc-slot-number pic pic-number port port-number]` hierarchy.

We no longer support port profile configuration for the JNP10K-LC1201 line card at this hierarchy level.

## IPv4

- **Support for two million routes (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, PTX10008 line cards support two million IPv4 routes (/32 routes).

In Junos OS Evolved releases before Release 20.2R1, PTX10008 line cards are able to scale to one million IPv4 routes.

**NOTE:** The scale for IPv6 routes still remains at 440,000 routes.

## Junos Telemetry Interface

- **JTI support for physical and logical interface sensors with gRPC services, gNMI services, and UDP (PTX10008)**—Junos OS Evolved 20.1R2 supports physical and logical interface sensors on PTX10008 routers. You can stream statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services, gRPC Network Management Interface (gNMI) services, or through UDP.

The following sensors are supported:

- `/junos/system/linecard/interface/`
- `/junos/system/linecard/interface/logical/usage/`

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

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

- **UDP (native) Packet Forwarding Engine sensor support for JTI (PTX10008)**—Junos OS Evolved Release 20.1R2 supports UDP (native) streaming of Packet Forwarding Engine sensors using Junos telemetry interface. UDP streaming for these sensors is not supported over the management interface.

The sensors supported are:

- Packet Forwarding Engine statistics. This sensor exports statistics for counters and provides visibility into Packet Forwarding Engine error and drop statistics (sensor name is `/junos/system/linecard/packet/usage/`).

- NPU memory utilization statistics. This sensor exports statistics such as Cuckoo Hash table statistics, Egress Descriptor Field (EDF) public words, and fast-look-up filter (FLT) block statistics (sensor name is `/junos/system/linecard/npu/memory`).
- NPU utilization statistics. This sensor exports Packet Engine utilization and packet load (sensor name is `/junos/system/linecard/npu/utilization/`).

[See [sensor \(Junos Telemetry Interface\)](#).]

## MPLS

- **6PE, 6VPE, and virtual router routing instance support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, you can configure the following features on PTX10008 routers:

- IPv6-over-MPLS using:
  - IPv6 VPN Provider Edge routers (6VPE)
  - IPv6 Provider Edge routers (6PE)
- Virtual router routing instance for:
  - BGP
  - IS-IS
  - OSPF and OSPFv3

[See [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks](#).]

- **Segment routing support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, you can configure the following Source Packet Routing in Networking (SPRING) or segment routing features on the PTX10008 router:

- MPLS (segment routing using IS-IS):
  - Ping and traceroute for single IS-IS node or prefix segment
- BGP-link state (BGP-LS):
  - Segment routing extensions for IS-IS
  - Segment routing extensions for OSPF
- BGP:
  - Binding SID for segment routing-traffic engineering (SR-TE)
  - Binding SID for SR-TE [*draft-previ-idr-segment-routing-te-policy*]
  - Programmable routing protocol process APIs for SR-TE policy provisioning
  - Static SR-TE policy with mandatory color specification
  - Static SR-TE policy without color specification

- IS-IS:
  - Adjacency SID
  - Advertising maximum link bandwidth and administrative color without RSVP-TE configuration
  - Anycast and prefix SIDs
  - Configurable segment routing global block (SRGB)
  - Node and link SIDs
  - Segment Routing Mapping Server (SRMS) and client
  - Topology-independent loop-free alternate (TI-LFA):
    - Link and node protection for IPv4 addressing (not required for IPv6 prefixes)
    - Link and node protection for IPv4 addressing (required for IPv6 prefixes)
    - Protection for SRMS prefixes
- OSPF:
  - Advertising maximum link bandwidth and administrative color without RSVP-TE configuration
  - Anycast SID
  - Configurable SRGB
  - Inter-area support
  - Node and link SID
  - Prefix SID
  - Segment Routing Mapping Server (SRMS) and client
  - Static adjacency SID
  - TI-LFA:
    - Link and node protection
    - Protection for SRMS prefixes
- MPLS ping and traceroute for single OSPF node or prefix segment
- IGP adjacency SID hold time
- Path Computation Element Protocol (PCEP) for segment routing LSPs
- BGP IPv4 labeled-unicast resolution over:
  - BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF
  - Non-colored IPv4 SR-TE with segment routing using IS-IS and OSPF
  - Static colored IPv4 SR-TE with segment routing using IS-IS and OSPF

- BGP Layer 3 VPN over:
  - Colored SR-TE tunnels and IPv4 protocol next hops
  - Non-colored SR-TE tunnels and IPv4 protocol next hops
- BGP-triggered dynamic SR-TE colored tunnels
- Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs
- First-hop label support for SID instead of an IP address
- Path specification using router IP addresses (segment routing segment list path ERO support using IP address as next hop and loose mode)
- SR-TE color mode:
  - 00—Route resolution fallback to IGP path
  - 01—Route resolution fallback to color only null routes
- Static LSPs with member-link next-hops for aggregated Ethernet bundles (also known as adjacent SID per LAG bundle or aggregated Ethernet member link)

[See [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

### ***Network Management and Monitoring***

- **Local port mirroring support (PTX10008)**—Starting with Junos OS Evolved Release 20.1R2, you can use port mirroring to copy packets entering or exiting a port or entering a VLAN and to send the copies to a local interface for local monitoring. You can send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

The following features are supported:

- Interface filter on ingress
- Interface filter on egress
- Forwarding-table filter (FTF) on ingress
- Unicast traffic on ingress and egress
- Multicast traffic on ingress and egress
- Families **inet** and **inet6**
- Firewall filter action **port-mirroring** in both ingress and egress directions
- Aggregated Ethernet interfaces a both ingress and egress
- Sampling range of 1-16M

- Packet clipping (**maximum-packet-length**)
- **run-length** attribute to specify the number of samples following trigger events

Use the following CLI hierarchies to configure port mirroring:

- **[edit interfaces]**
- **[edit forwarding-options port-mirroring]**
- **[edit firewall filter]**

You can configure **family inet** and **family inet6** in the **[edit interfaces]** and the **[edit forwarding-options port-mirroring]** hierarchies for this feature. The feature applies to global port mirroring only.

[See [Understanding Port Mirroring](#).]

### ***Routing Policy and Firewall Filters***

- **Firewall filters support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, the following firewall features are supported:
  - Forwarding IPv4 and IPv6
  - Firewall filter
  - Load balancing
  - MPLS fast reroute
  - Host path
  - Egress peer engineering

Non-replication mode, source and destination class usage (SCU, DCU), segmented filters, output filter binding on routing tables, and flexible offset filters are not supported in this release.

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

- **Support for matching ip-options in IPv4 packet headers (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, the PTX10008 supports the **ip-options any** match condition at the **[edit firewall family inet filter name term name from]** hierarchy level.

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

### ***Routing Protocols***

- **Support for BGP- LU FRR MPLS inter-AS link protection (PTX 10008)**—Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. For an MPLS inter-AS environment, you can enable link protection when labeled-unicast is used to send traffic between autonomous systems (ASs). The ingress router then chooses an alternate link through another interface to send traffic to its destination. Starting in Junos Evolved Release 20.1R2, to configure link protection on an interface, you can include the **protection** statement at the **[edit protocols bgp group group-name family inet labeled-unicast]** hierarchy level.

Junos OS Evolved now supports the following baseline BGP features on the PTX10008:

- Large communities
- Link bandwidth community aggregation
- Load balancing in proportion to available bandwidth community
- **local-as** option
- BGP Monitoring Protocol (BMP) version 3
- MTU discovery
- BGP multipath
- Multiple paths to IPv6 addresses
- PIC core and inter-AS link protection for BGP-LU
- Prefix prioritization
- Route reflection
- **tcp-mss** configuration
- Route leaking between routing instances, including default (inet) routing instance through RIB-groups and instance import/export

You can use the Junos Telemetry Interface to export telemetry data from supported interface hardware. Line-card sensor data, such as interface events, are sent directly to configured collection points without requiring polling.

We now support telemetry sensors for the following system resources:

- JTI: OXST: Openconfig BGP supporting configuration model with network instance XPATH (v2.1.1)
- OpenConfig: BGP configuration to become network-instance compliant (v4.0.1)
- CPU memory
- BGP peers (gRPC streaming only)
- Memory utilization for routing protocol tasks (gRPC streaming only)
- Network processing unit (NPU) memory and memory utilization
- Optical interfaces
- Inline flow sampling process (UDP streaming only)
- Chassis components
- Aggregated Ethernet interfaces configured with LACP (gRPC streaming only)
- ARP (gRPC streaming only)
- Ethernet interfaces configured with LLDP (gRPC streaming only)

- RSVP interface events (gRPC streaming only)
- Network Discovery Protocol table state (gRPC streaming only)
- Routing Engine internal interfaces (gRPC streaming only)

[See [Junos Telemetry Interface Feature Guide](#).]

### **Services Applications**

- **Support for FTI-based de-encapsulation of IPv4 and IPv6 packets (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, you can configure IP-IP de-encapsulation or tunnel termination on a flexible tunnel interface (FTI) by configuring tunnel termination at the **[edit interfaces fti0 unit unit tunnel encapsulation]** hierarchy level. Tunnel termination is completed in a single pass of packet processing, thus providing performance improvement over the existing process, which requires a recirculation path to support de-encapsulation.

### **System Management**

- **Support for passive monitoring (PTX10008)**—Starting with Junos OS Evolved Release 20.1R2, you can configure passive monitoring on any interface on the PTX10008 routers. After you enable passive monitoring, the device accepts and monitors traffic on the interface and forwards those packets to monitoring tools such as IDS servers and packet analyzers, or other devices such as routers or end-node hosts.

[See [Understanding Passive Monitoring](#).]

### **VPNs**

- **Layer 3 VPN on PTX10008**—Starting in Junos OS Evolved Release 20.1R2, you can configure MPLS-based Layer 3 VPNs on PTX10008 routers. The supported features include:
  - MPLS over Layer 3 VLAN-tagged sub interfaces
  - Per-next hop label allocation
  - Mapping of label-switched interface (LSI) logical interface label to the VPN routing and forwarding (VRF) routing table using the **vrf-table-label** statement
  - ICMP tunneling and MPLS traceroute
  - Disabling time-to-live (TTL) decrementing using **no-propagate-ttl**

[See [Layer 3 VPNs Feature Guide for Routing Devices](#).]

- **Layer 2 VPN support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R2, you can configure the following Layer 2 VPN on the PTX10008 router:
  - Transport of Layer 2 frames over MPLS (LDP signaling)
  - Layer 2 VPNs over tunnels (BGP signaling)
  - Simple Ethernet and VLAN based cross connect (also known as connections)



- Local and remote switching
- Ethernet and VLAN-CCC
- Single tagged CCC logical interfaces
- Control word
- Regular and aggregated Ethernet interfaces
- Layer 2 protocol pass-through
- Layer 2 circuit backup interface and backup neighbor
- Layer 2 circuit statistics and COS
- VCCV with type 2 and type 3

[See [Layer 2 VPNs and VPLS Feature Guide for Routing Devices](#) and [Translational Cross-Connect \(TCC\) Overview](#).]

## What's New in Junos OS Evolved Release 20.1R1

### IN THIS SECTION

- [Hardware](#) | 18
- [Authentication, Authorization, and Accounting](#) | 18
- [Class of Service](#) | 19
- [Dynamic Host Configuration Protocol](#) | 21
- [General Routing](#) | 21
- [High Availability and Resiliency](#) | 22
- [Interfaces and Chassis](#) | 23
- [Junos OS XML API and Scripting](#) | 26
- [Junos Telemetry Interface](#) | 27
- [Layer 2 Features](#) | 30
- [Layer 3 Features](#) | 31
- [Management](#) | 33
- [MPLS](#) | 34
- [Multicast](#) | 37
- [Network Security](#) | 37
- [Network Management and Monitoring](#) | 37
- [Port Security](#) | 39
- [Routing Policy and Firewall Filters](#) | 39

- Routing Protocols | 40
- Services Applications | 41
- Software Defined Networking | 43
- Software Installation and Upgrade | 43
- Spanning Tree Protocols | 44
- System Logging | 44
- System Management | 44

Learn about new features introduced in Junos OS Evolved Release 20.1R1 for PTX10003 and PTX10008 routers.

### **Hardware**

- **PTX10008 Packet Transport Router**—Starting with Junos OS Evolved Release 20.1R1, the modular chassis can now support more than 115 Tbps capacity with the addition of these new components:
  - *PTX10K-LC1201-36CD* 14.4 Tbps capacity, 400-Gigabit Ethernet line card—Provides 36 ports of QSFP56-DD. Each port defaults to 400 Gigabit Ethernet, but can be individually channelized using breakout cables to speeds of 200 Gbps, 100 Gbps, 50 Gbps, 25 Gbps, or 10 Gbps.
  - *JNP0K-RE1-E and JNP10K-RE1-E128* Routing and Control Boards (RCBs)—Provides 64 Gigabytes or 128 Gigabytes of memory running on an 8-core 2.1 GHz processor.
  - *JNP10008-SF3* Switch Interface Board (SIB)—Provides support for 14.4 Tbps line cards.

The new components are designed to work with the existing 5.5 KW power supplies, the JNP10008-FAN2 fan tray, and the JNP10008-FTC2 fan tray controller. The new components cannot be mixed with earlier RCBs, line cards, original fan trays, original fan tray controllers, or with 3 KW power supplies. [See [PTX System Overview](#).]

### **Authentication, Authorization, and Accounting**

- **Support for RADIUS and TACACS+ features (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 devices support the following AAA features for both RADIUS and TACACS+:
  - Local authentication with local authorization
  - AAA default routing instance and nondefault management routing instance (mgmt\_junos)
  - Template user authentication and authorization
  - Enhanced accounting to monitor attributes of logged-in users such as access method, remote port, and access privileges
  - Reachable over Packet Forwarding Engine interfaces

Additionally, VRF routing instances for AAA are supported only for TACACS+.

Junos OS Evolved Release 20.1R1 does not support the following features for TACACS+ and RADIUS:

- Password change policy enforcement
- Regular expressions to allow or deny CLI operational commands
- Local authentication with remote authorization
- Periodic refresh of the remote authorization profile that is received from the AAA server

[See [Junos OS User Authentication Overview](#), [RADIUS Authentication](#), [TACACS+ Authentication](#), [Configuring RADIUS System Accounting](#), and [tacplus-options](#).]

- **Link Layer Discovery Protocol (LLDP) (PTX10003)**—LLDP is supported on PTX10003 routers, including on em0 interfaces. Disabling of LLDP time, length, and value (TLV) messages is also supported.

[See [Device Discovery Using LLDP and LLDP-MED on Switches](#).]

- **Link Layer Discovery Protocol (LLDP) (PTX10008)**—LLDP is supported on PTX10008 routers, including on WAN and management interfaces. In addition, the LLDP MIB and Junos Telemetry Interface (JTI) for LLDP is also supported. Disabling of LLDP time, length, and value (TLV) messages is also supported.

The following are not supported in this release:

- LLDP-MED
- Port VLAN ID, VLAN name, and Chassis ID TLVs

[See [Device Discovery Using LLDP and LLDP-MED on Switches](#).]

- **Support for AAA features (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support the following AAA features:

- SSH and SSHv2
- Telnet

### ***Class of Service***

- **Support for CoS/QoS features for PTX10008 platforms**—Starting in Junos OS Evolved Release 20.1R1, you can configure Class of Service (CoS) / Quality of Service (QoS) features for 14.4 Tbps line cards or lite fixed chassis for PTX10008 platforms that include classifiers (packet, behavior aggregate [BA], fixed, and multifield [MF]), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control and drop-profiles.

The following CoS features are supported in this release:

- 16 forwarding classes
- 8 forwarding classes
- Classification: DSCP IPv4 and IPv6 for L3 logical interfaces
- Classification: IEEE 802.1p for L3 logical interfaces

- Classification: MPLS EXP for family bridge / ethernet-switching
- Classification: MPLS EXP for L3 logical interfaces
- Congestion Avoidance: 32 RED profile, with 64 fill levels per profile
- Congestion Avoidance: Weighted random early detection (WRED) and tail-drop profiles
- Congestion Avoidance: WRED Drop and Tail Drop Differentiation
- Host Outbound: Forwarding Class Selection
- Host Outbound: Rewrite
- Rewrite: 802.1p for L3 (together with IP)
- Rewrite: CoS supports 64 rewrite tables across L2/L3 rewrites. MPLS EXP, IP DSCP, and Ethernet 802.1p rewrites are supported. Rewrite of both L2 and L3 header CoS fields can occur simultaneously.
- Rewrite: IPv4 TOS/ DSCP and IPv6 DSCP
- Rewrite: MPLS EXP Push, MPLS EXP SWAP
- Scheduling and Shaping: Buffer size configuration
- Scheduling and Shaping: CoS scheduler configurability
- Scheduling and Shaping: eight queues, per queue priority scheduling, strict priority, and WRR
- Scheduling and Shaping: Physical Interface Scheduler
- Scheduling and Shaping: Physical Interface Shaping
- Scheduling and Shaping: Queue Shaping
- Scheduling and Shaping: Strict Priority Scheduling
- RSVP-TE: Class Based Forwarding (CBF)

[See [Understanding Class of Service.](#)]

- **Support for CoS features on 400-Gigabit Ethernet interfaces (PTX10003)**—Starting with Junos OS Evolved 20.1R1, the 400-Gbps interfaces on PTX10003 devices support standard classifier, rewrite, and scheduler functionality as listed in the following link.

[See [CoS Features and Limitations on PTX Series Routers.](#)]

- **Support for enabling a queue's buffer space to be 100 percent of the interface's buffer space (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, PTX10003 devices support the CLI option that enables you to set a queue's buffer space to be up to 100 percent of the interface's buffer space. This option allows the queue to grow as large as 100 percent of the interface's buffer if and only if it is the only active queue for the interface. You can enable this option by setting the **buffer-size-shared** statement at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.

[See [buffer-size \(Schedulers\)](#).]

- **Classification override support (PTX10003-80C and PTX10003-160C)**—Starting in Junos OS Evolved Release 20.1R1, we have extended support for classification override configured under a forwarding policy to PTX10003-80C and PTX10003-160C devices.

[See [CoS Features and Limitations on PTX Series Routers](#) and [Overriding the Input Classification](#).]

### ***Dynamic Host Configuration Protocol***

- **DHCPv4 and DHCPv6 Relay Agent (PTX10008)**—DHCPv4 and DHCPv6 relay agent is supported on PTX10008 routers.

The following DHCP Relay Agent features are supported in this release:

- DHCP Relay: L3 interfaces
- DHCP Relay: Option 82 for Layer 2 VLANs
- DHCP Relay: Option 82 for Layer 3 interfaces
- Extended DHCP Relay Agent
- Virtual router aware DHCP (VR-aware DHCP)

[See [Extended DHCP Relay Agent](#).]

- **Dynamic Host Control Protocol (DHCP) Client (PTX10003)**—DHCPv4 and DHCPv6 for dynamic configuration of IP addresses on client interfaces, including management interfaces, is supported on PTX10003 routers.

[See [DHCP for Routing Devices](#).]

### ***General Routing***

- **View ARP table entries and IPv6 neighbor information (PTX10003-80C and PTX10003-160C)**—Starting in Junos OS Evolved Release 20.1R1, on PTX10003 routers, you can view the following:
  - Address Resolution Protocol (ARP) details such as the next-hop reference count and the next-hop current state, using the commands **show arp reference-count** and **show arp state**.
  - IPv6 neighbor cache information, using the following commands: **show ipv6 neighbors flags**, **show ipv6 neighbors host**, **show ipv6 neighbors interface *interface-name***, **show ipv6 neighbors reference-count**, **show ipv6 neighbors vpn *vpn-name***.

[See [show arp](#) and [show ipv6 neighbors](#).]

- **Support for running applications signed by users (PTX10003, PTX10008, QFX5220)**—Starting with Junos OS Evolved Release 20.1R1, root users can generate signing keys which can be used to sign Linux applications and allow them to run on your device.

Signing keys are stored in the **system-keystore** by using the following CLI command: **request security system-keystore import key-name *key-name* private-key *key-path* x509-cert *certificate-path***. Once a

key is stored in the **system-keystore**, it can be used to sign a file with the following CLI command: **request security integrity measure file *file-name* key *key-name***.

[See [Junos OS Evolved Overview](#)]

### **High Availability and Resiliency**

- **Packet Forwarding Engine resiliency (PTX10008 router)**—Starting in Junos OS Evolved Release 20.1R1, in the PTX10008 router, the Packet Forwarding Engine software supports resiliency. The software detects, reports and takes action on Packet Forwarding Engine faults. Actions are taken based on the default configuration or a user-configuration available for the errors. You can use the **show system errors active** and **show system errors active detail** commands to view the error details.

[See [show system errors active](#).]

- **Fabric link autoheal (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10003 router supports fabric link autoheal, a mechanism that attempts recovery of faulty fabric links from a link error condition. Autoheal involves bringing down the faulty fabric link and then training it. The system attempts to recover a faulty link from a maximum of three link error instances (per link) within the span of 24 hours or until the FPC or the switch fabric is rebooted. This means, for instance, if the same link encounters a fourth error within 24 hours from the first error instance, autoheal is not attempted on that link. You can use the existing **show chassis fabric errors autoheal** command to view the details the autoheal actions performed by the system.

**NOTE:** The FPC and switch fabric are logical entities in the PTX10003.

[See [show system errors active](#) and [show chassis fabric errors](#).]

- **Resiliency application supports application-level restart (PTX10003-80C and PTX10003-160C)**—Starting in Junos OS Evolved Release 20.1R1, the resiliency application on PTX10003 routers supports application-level restart, in case of failures. This feature ensures that a failure of the resiliency application such as crash or termination is handled gracefully without causing the entire platform to reboot.

[See [PTX10003 System Overview](#)]

- **Error management at the switch fabric level (PTX10003-80C and PTX10003-160C)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10003 router supports the configuration of error scope, category, threshold, and corrective actions at the switch fabric level for the fabric errors. The PTX10003, being a fixed configuration chassis, does not have any physical switch fabric card that can be installed. Instead, it provides two logical SIBs. You can use the existing command **show chassis hardware** to view information about the logical SIBs in the system. The router also supports multilevel fault actions for repeated faults—that is, to move the switch fabric (logical SIB) to the offline state when it is rebooted consecutively for three times within a time span of 900 seconds. The software supports a new error scope **switch** and an error category **internal**. All the errors from the fabric modules and a few errors from Packet Forwarding Engine are mapped to the internal category.

[See [error](#).]

- **BFD sessions for LAGs (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, the following are supported with interface IP addresses:
  - Micro BFD sessions
  - Child links of aggregated Ethernet or LAG bundle

[See [Configuring Micro BFD Sessions for LAG](#).]

### *Interfaces and Chassis*

- **Power budgeting and support for overriding the default power budget (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10008 router supports power budgeting, management, and environment monitoring. You can also override the default power budget for a line card, by using the CLI command **set chassis fpc slot max-power power-number**. The PTX10008 supports power management at the switch fabric level. The switch fabric has a higher priority over the line cards with regard to power allocation. You can configure the ambient temperatures of 25°C and 40°C on this router. The PTX10008 also supports the following power-related alarms:
  - **Insufficient Power - FRU(s) went offline**—Raised in case FRUs are powered off because of insufficient power.
  - **No Redundant Power**—Raised in case the power redundancy requirements are not met.

[See [Managing Power](#).]

- **Support for AC and DC PSMs (PTX10008)**—Starting in Release 20.1R1, Junos OS Evolved supports the following power supply modules (PSMs) on the PTX10008 router:
  - **JNP10K-PWR-AC2**—This power supply supports AC, high voltage AC (HVAC), or high voltage DC (HVDC) systems in either a 20-A or a 30-A mode. It feeds AC input and provides a DC output of 5000 W with a single feed and 5500 W with dual feeds. For AC systems, the operating input voltage is 180 to 305 VAC and for DC systems, the operating input voltage is 190 to 410 VDC.
  - **JNP10K-PWR-DC2**—This DC power supply provides two power supplies (PS\_0 and PS\_1) in a single housing that accepts either 60 A or 80 A using four redundant input power feeds. PS\_0 and PS\_1 together provide a combined output of up to 5500 W.

[See [PTX10000 Power System](#).]

- **Cooling management on PTX10008**—Starting in Junos OS Evolved Release 20.1R1, the PTX10008 routers support cooling management with the fan tray JNP10008-FAN2 and the fan tray controller JNP10008-FTC2. The PTX10008 supports two fan trays. The fan tray contains 22 fans. Under normal operating conditions, the fans run at less than full speed. However, if a fan is faulty, the remaining working fans in that fan tray run at the maximum speed. If one fan tray is removed, all the fans in the other fan tray run at the maximum speed.

**NOTE:** You must use the fan tray JNP10008-FAN2 only along with the fan tray controller JNP10008-FTC2.

[See [PTX10008 Cooling System](#).]

- **60-A power source for PTX10003 DC power supply (PTX10003-80C and PTX10003-160C)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10003 DC power supply supports a 60-A power source. If you choose the 60-A setting, the power supply limits its output power capacity to 2700 W at an input voltage of 48 V and linearly increases the output power if the input voltage increases. The power supply provides a maximum output of 3000W at input voltage greater than 55V, assuming 92 percent efficiency. If the voltage drops below 48 V but is above 40 V, the software raises a minor alarm **Input Under Voltage Warning** and reduces the output power capacity to 2200 W. The PSM is powered off if the voltage drops below 40 V.

[See [PTX10003 Power System](#)]

- **Dedicated virtual routing and forwarding (VRF) instance for management Ethernet interface (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can confine the management interface on your device to a dedicated VRF management instance, `mgmt_junos`. To configure the `mgmt_junos` routing instance, include the **management-instance** statement at the **[edit system]** hierarchy level. After you enable `mgmt_junos`, the management traffic does not share the default routing table (`default.inet.0`) with other control or protocol traffic in the system. This improves the security of management traffic and makes it easier to troubleshoot issues on your device.

[See [Management Interface in a Nondefault Instance](#) and [management-instance](#).]

- **Support for 400-Gbps speed using PTX10K-LC1201-36CD optics (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can use the QSFP56-DD optic transceiver to configure 400-Gbps speed on the PTX10K-LC1201-36CD line card of PTX10008 routers. The complete list of transceivers is available in the [Hardware Compatibility Tool](#) (HCT).
- **Support for rate selectability on PTX10K-LC1201-36CD line cards (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, we introduce a new fixed-configuration, rate-selectable line card, PTX10K-LC1201-36CD, with 36 built-in ports. The default port speed is 400 Gbps for all ports. Each PTX10K-LC1201-36CD line card provides a maximum bandwidth of 14.4 Tbps.

On the PTX10K-LC1201-36CD, you can choose to configure the following port speeds on all 36 ports:

- 4x10 Gbps, 4x25 Gbps, and 2x50 Gbps
- 40 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps

[See [Introduction to Rate Selectability](#).]

- **Support for resilient hashing and consistent hashing (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can configure resilient hashing on PTX10008 routers to minimize flow remapping across link aggregation groups (LAGs) or equal cost paths. Resilient hashing works in conjunction with



the default static hashing algorithm. When you configure resilient hashing on LAGs, the configuration is applicable to a specific aggregated Ethernet Interface.

Also, you can configure consistent hashing across ECMP groups by using BGP multipath on PTX10008 routers. You can configure a maximum of 64 member links for resilient hashing on LAG and consistent hashing on BGP.

[See [Resilient Hashing on LAGs and ECMP groups](#) and [Load Balancing for a BGP Session](#)].

- **Support for Link Aggregation Control Protocol Features (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can configure LACP to bundle several physical interfaces to form one logical aggregated Ethernet interface. By default, LACP is disabled on aggregated Ethernet interfaces. You can configure LACP in active or passive mode. If LACP is configured, it is in passive mode, by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode. LACP monitors the links to check whether both ends of the bundle are connected to the correct group. You can also configure the following supported LACP features on PTX10008 routers:
  - Configuration Synchronization for MC-LAG groups
  - Distributed LACP (AFT-based PPMAN)
  - Load balancing
  - MC-LAG Links or Interfaces with limited LACP capability to be up
  - Telemetry Support (Periodic streaming and On-change)

[See [Configuring Aggregated Ethernet Interfaces and LACP](#).]

- **Support for LACP features (PTX10003)**—In Junos OS Evolved Release 20.1R1, the PTX10003 supports LACP in the default periodic packet management (PPM) mode. The router supports the following LACP features:
  - LACP load balancing
  - OpenConfig
  - Telemetry
  - **Sync-reset** with minimum-link
  - **LACP force-up**
  - **LACP hold up**
- **Support for Layer 2 bridging (PTX10003)**—In Junos OS Evolved Release 20.1R1, PTX10003 routers support Layer 2 bridging. You create a bridge domain by adding a set of Layer 2 logical interfaces (on your device) to represent a broadcast domain. All the member ports of the bridge domain participate in Layer 2 learning and forwarding. You can configure one or more bridge domains to perform Layer 2 bridging. You can optionally disable learning on a bridge domain. You can configure the Layer 2 interfaces either by configuring the access and trunk port of the **ethernet-switching** family or by **vlan-tagging**.

[See [Configuring Layer 2 Bridging Interfaces](#).]

- **Support for resilient hashing (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10003 supports resilient hashing by link aggregation groups (LAGs) and equal-cost multipath (ECMP) sets. Working together with the default static hashing algorithm, resilient hashing enhances LAGs and ECMP sets by minimizing destination remapping when a new member is added to or deleted from the LAG or ECMP set. When a LAG member change affects the traffic flow, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table. Both consistent hash and resilient hash (on LAG) support only 64-way ECMP routing.

[See [Understanding Consistent Load Balancing Through Resilient Hashing on ECMP Groups](#).]

### *Junos OS XML API and Scripting*

- **The jcs:load-configuration template supports loading the rescue configuration (PTX Series and QFX Series)**—Starting in Junos OS Evolved Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to "rescue" to replace the active configuration with the rescue configuration.

[See [Changing the Configuration Using SLAX and XSLT Scripts](#) and [jcs:load-configuration Template](#).]

- **Support for automation scripting features (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers include support for the following automation scripting features:
  - SLAX and XSLT commit, event, op, and SNMP scripts
  - Python commit, event, op, and SNMP scripts
  - Event policies and event policy features including:
    - Event correlation using **attributes-match**
    - Event policy action to change the configuration using Junos OS configuration mode commands
    - Ability to override the system log priority of the triggering event
    - Ability to limit permissions for executing event policy actions
  - Dedicated **lib** directory for user script libraries
  - Global variable provided to Junos OS automation scripts
  - Loading scripts from flash
  - Refreshing scripts from a configured or specified source location
  - Configuring checksums to verify the integrity of scripts
  - Script dampening
  - Junos PyEZ library on box
  - Allowing operational commands in op scripts

- Uploading files generated by event scripts or other event policy actions
- Remote RPC support for event scripts
- libslax distribution, including the libslax default extension libraries: **bit**, **curl**, **db**, **os**, and **xutil**

[See [Junos Automation Scripts Overview](#).]

- **Support for ECMP path trace tool using APIs (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, you can use ECMP tracer JET APIs to monitor traffic flow, trace the ECMP path used by packets entering and exiting a CLOS network, and get flow counters without changing the configuration on your device. You provide match criteria to the ECMP path trace tool, `ecmptacerd`, through JET API requests. The `ecmptacerd` process then installs firewall filters for the inputs on all the interfaces on the device except the private interfaces. The counters associated with these filters increment as packets matching the firewall filter criteria arrive and leave. These APIs provide similar functionality to the `ecmp-tracer` on-box shell utility. Do not use the JET client and the `ecmp-tracer` on-box shell utility simultaneously.

[See [ecmp](#), [Juniper Extension Toolkit Developer Guide](#), and [Juniper Engineering Network website](#).]

### *Junos Telemetry Interface*

- **IPFIX process statistics on JTI (PTX10003)**—Junos OS Evolved Release 20.1R1 supports IPFIX process statistics using Junos telemetry interface (JTI) and remote procedure calls (gRPC) services. Use the inline J-flow sensor to monitor Inline J-flow service statistics on the router and export these to a collector at configurable intervals. The minimum interval is 30 seconds.

Use the resource path `/junos/system/linecard/services/inline-jflow` in a gRPC subscription to export these statistics to a collector.

[See [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **JTI support for Packet Forwarding Engine statistics for (PTX10008)**—Junos OS Evolved Release 20.1R1 supports the Junos telemetry interface (JTI) export of Packet Forwarding Engine statistics using remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. Using gRPC or gNMI and JTI, data is streamed from devices to outside collectors at configurable intervals.

Export the following statistics:

- Packet Forwarding Engine statistics (resource path `/junos/system/linecard/packet/usage`)
- Network processing unit (NPU) memory statistics (resource path `/junos/system/linecard/npu/memory/`)
- NPU utilization statistics (resource path `/junos/system/linecard/npu/utilization/`)

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

To provision the sensor to export data through gNMI services, use the **Subscribe** RPC. The **Subscribe** RPC and subscription parameters are defined in the `gnmi.proto` file. Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

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

- **Chassis, environmental, and power sensor support using JTI (PTX10008)**—Junos OS Evolved Release 20.1R1 supports Junos telemetry interface (JTI) sensors that provide the operational state of Routing Engines, power supplies and fan trays, FPCs, SIBs, and PICs on PTX10008 routers. JTI enables the export of statistics from these sensors to outside collectors at configurable intervals using remote procedure call (gRPC) services.

To provision a sensor to export data through gRPC services, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module.

[See [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#) and [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Support for gNMI services with JTI for Routing Engine and Packet Forwarding Engine sensors (PTX10003)**—Junos OS Evolved Release 20.1R1 introduced support to export the statistics using JTI and remote procedure call (gRPC) services. In Junos OS Evolved Release 20.1R1, you can also export statistics for these sensors using gRPC Network Management Interface (gNMI) services.
  - Routing Engine internal interfaces (physical interface state and statistics) (resource path `/interfaces/`)
  - Firewall filter counters and policer counters (resource path `/junos/system/linecard/firewall/`)
  - Packet Forwarding Engine (resource path `/junos/system/linecard/packet/usage/`)
  - NPU memory utilization statistics (resource path `/junos/system/linecard/npu/memory`)
  - NPU utilization statistics (resource path `/junos/system/linecard/npu/utilization/`)

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

To provision the sensor to export data through gNMI services, use the **Subscribe** RPC. The **Subscribe** RPC and subscription parameters are defined in the `gnmi.proto` file. Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

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

- **Supporting gRPC services with JTI for Routing Engine sensors (PTX10008)**—Junos OS Evolved 20.1R1 supports the following Routing Engine sensors on the PTX10008. You can export statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services.
  - Routing Engine logical interface statistics (resource path `/interfaces/interface/`). Applicable for management interfaces (re0:mgmt-0 and re1:mgmt-0).
  - Address Resolution Protocol (ARP) (resource path `/arp-information/`)
  - Network Discovery Protocol (NDP) (resource path `/nd6-information/`)
  - Operational state of AE with LACP (resource path `/lACP/`)
  - Operational state of LLDP (neighbors information) (resource path `/lldp/`)
  - Sensor for NDP router-advertisement statistics (resource path `/ipv6-ra/`)

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

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

- **JTI support for kernel statistics (PTX10003)**—Junos OS Evolved Release 20.1R1 supports Junos telemetry interface (JTI) sensors for Address Resolution Protocol (ARP/IPv4 routes), Network Discovery Protocol (NDP/IPv6 routes), and management interface statistics on PTX10003 routers. Statistics are delivered with gRPC Network Management Interface (gNMI) to stream statistics at regular intervals to an outside collector.

To stream statistics, include the resource paths in a subscription:

- Sensor for ARP statistics for IPv4 routes (resource path `/arp-information/`)
- Sensor for NDP table state information for IPv6 routes (resource path `/nd6-information/`)
- Sensor for management interface statistics (re0:mgmt-0) (resource path `/interfaces/interface/`)

To provision a sensor to export data through gNMI, use the **Subscribe** RPC defined in the [gnmi.proto](#) file to specify request parameters.

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

- **IPFIX process statistics on JTI (PTX10008)**—Junos OS Evolved Release 20.1R1 supports IPFIX process statistics using Junos telemetry interface (JTI) and Remote Procedure Calls (gRPC) services. Use the inline Jflow sensor to monitor Inline Jflow service statistics on the router and export these to a collector at configurable intervals. The minimum interval is 30 seconds.

Use the resource path `/junos/system/linecard/services/inline-jflow` in a gRPC subscription to export these statistics to a collector.

[See [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **Packet Forwarding Engine sensor enhanced on JTI (PTX10003)**—Junos OS Evolved Release 20.1R1 supports enhanced Packet Forwarding Engine sensor enhancements for firewall filters and policers to stream statistics using Junos telemetry interface (JTI) and Remote Procedure Calls (gRPC) services or Network Management Information (gNMI) services.

Use the resource path `/junos/system/linecard/npu/memory/` in a gRPC or gNMI subscription to export these statistics to a collector.

[See [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **Sensor level statistics support on Junos Telemetry Interface (JTI) (QFX Series, PTX Series)**—Starting with Junos OS Evolved Release 20.1R1, you can issue the Junos operational mode command **show network-agent statistics** to provide more information on a per-sensor level for statistics being streamed to an outside collector by means of remote procedure calls (gRPC) and JTI. Only sensors exported with gRPC are supported. The command does not support UDP-based sensors.

[See [show network-agent statistics](#).]

- **Routing Engine sensor support with gNMI services and JTI (PTX10008)**—Junos OS Evolved Release 20.1R1 supports Routing Engine statistics using Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services.

Include the following resource paths in a gNMI subscription to periodically stream statistics to an outside collector:

- Address Resolution Protocol (ARP) (resource path `/arp-information/`)
- Network Discovery Protocol (NDP) (resource path `/nd6-information/`)
- Physical interface traffic (resource path `/interfaces/interface/`) on management interfaces (re0:mgmt-0)

[See [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

## Layer 2 Features

- **Support for Q-in-Q tunneling (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, Q-in-Q tunneling allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using this feature, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.

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

- **Proxy ARP (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008, both restricted and unrestricted ARP are supported. With restricted ARP, the device responds to the ARP requests in which the physical networks of the source and target are not the same and the device has an active route to the target address in the ARP request. The device does not reply if the target address is on the same subnet and the same interface as the ARP requestor. With unrestricted ARP, the device responds to any ARP request, on the condition that the device has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route. By default, proxy ARP is unrestricted.

Proxy ARP is supported on et and ae interfaces and interface routes only.

To enable unrestricted proxy ARP, enable the **unrestricted** statement at the `[edit interfaces interface-name proxy-arp]` CLI hierarchy.

To enable restricted proxy ARP, enable the **restricted** statement at the `[edit interfaces interface-name proxy-arp]` CLI hierarchy.

To enable default mode (unrestricted), enable the **proxy-arp** statement at the `[edit interfaces interface-name]` CLI hierarchy.

[See [Restricted and Unrestricted Proxy ARP Overview](#).]

- **Support for basic Layer 2 features (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, PTX10003 routers supports the following basic Layer 2 features:

- Layer 2 bridging with trunk and access modes
- MAC learning and aging
- Handling BUM (broadcast, unknown unicast and multicast) traffic, including split horizon
- Spanning tree functionalities
- MAC limiting
- Static MAC addresses

[See [Understanding Layer 2 Bridge Domains](#), [Understanding Layer 2 Learning and Forwarding](#).]

- **Layer 2 features on PTX10003**—Starting in Junos OS Evolved Release 20.1R1, you can configure the following Layer 2 features on the PTX10003 router:
  - Circuit cross-connect (CCC) with and without control word
  - Transport protocols in CCC
  - Layer 2 connection
  - Layer 2 circuit
  - Layer 2 VPN configuration with **ethernet-ccc** and **vlan-ccc** encapsulation types

The following Layer 2 features are not supported:

- Translational cross-connect (TCC)
- Layer 2 configuration with **flexible-ethernet-services** encapsulation type

[See [Layer 2 VPNs and VPLS Feature Guide for Routing Devices](#), and [Configuring Circuit Cross-Connect \(CCC\) and Translational Cross-Connect \(TCC\)](#).]

### **Layer 3 Features**

- **Support for Layer 3 Features (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can configure the following Layer 3 forwarding features on the PTX10008 router:
  - IPv4
  - IPv6
  - IPv6 accounting, includes interface statistics for explicit null, PTX10008 hop popping (UHP), 6PE, and BGP labeled unicast (BGP LU).
  - IPv6 subnet support on loopback interfaces
  - IPv4 and IPv6 forwarding – limited longest prefix match (LPM) entries for OSPFv3, ISIS, BGP, static routes
  - IPv4 and IPv6 traceroute
  - IPv4 and IPv6 ping

- IPv4 and IPv6 multicast and unicast traffic statistics
- Junos Telemetry Interface (JTI), including routing process daemon (RPD) and RPD Task Memory
- Interface ranges
- Flexible VLAN-tagged Layer 3 logical interfaces
- Single VLAN-tagged Layer 3 logical interfaces
- Virtual router routing instances for unicast protocols
- MPLS
- LAG
- ECMP
- Path MTU discovery
- MTU checks
- Routing protocol next-hop acknowledgements and time intervals
- ICMP
- OSPF, including support for:
  - LDP synchronization
  - Multitopology
  - OSPF database protection for OSPF and OSPFv3
  - OSPF export and import policies for network-summary LSAs
  - OSPFv2
  - OSPFv3
  - Policy based support for LFA in OSPF
  - Remote LFA support in OSPF (IGP and LDP)
  - Shortcuts
  - SNMP MIB support for OSPFv3
  - Stub network overload and external route leaking

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

- IS-IS, including support for:
  - IPv6 and IPV6 unicast topology
  - JTI adjacency, link state database (LSDB) streaming, state, interface, adjacency statistics, and LSDB export
  - LDP synchronization



- LFA and LFA policy based support
- Multitopology
- Remote LFA support for LDP in IS-IS
- Shortcuts
- Stub network overload

[See [IS-IS User Guide](#).]

- Route-filter walkup
- ARP
- Static ARP
- NDP
- BGP
- BFD
- LACP
- LDP
- RSVP
- LLDP
- VRF-lite
- TTL expiry
- IP options
- IP Fragmentation
- DDOS

### **Management**

- **Programmable Routing Protocol API support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, the Programmable Routing Protocol (prpd) is supported on PTX10008. Through JET, the prpd provides public APIs to program routing systems, making it possible for users to directly access the APIs to customize, create, and modify the behavior of their network. The supported APIs include:
  - BGP route monitoring
  - BGP route APIs used to program the virtual Route Server (vRS)
  - BGP route APIs used to program inet and inet6 routes
  - BGP route APIs used to program l3vpn-inet and l3vpn-inet6 in Route Reflector
  - BGP Labeled Unicast (BGP-LU) route monitoring

- RIB route monitoring
- RIB MPLS routes
- RIB static routes
- RIB APIs that support the preference, tag, and color attributes
- Interface notifications

[See [Juniper Extension Toolkit Developer Guide](#) and [Juniper Engineering Network website](#).]

- **Support for management features (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support the following management features:
  - DNS
  - FTP
  - IPv6 support for SSH, Telnet, path MTU, CoS, VRRPv3, Radius, TACACS+, AAA, NTP, syslog, on network interfaces
  - SFTP
  - Configuring the terminal width with the [set cli screen-width](#) command
  - Support for forwarding structured system log messages to a remote system log server
  - System logging (syslog) over IPv4 and IPv6
  - System logging (syslog) over IPv4 and IPv6: Management vRF support

## **MPLS**

- **sFlow support (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, sFlow supports the sampling and detecting of MPLS traffic on the PTX10003 router.

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

- **MPLS Support for PTX10008**—In Junos OS Evolved Release 20.1R1, the following MPLS features are supported:
  - CLI support for monitoring MPLS label usage
  - Inline MPLS and IPv6 lookup for explicit null
  - 32K transit LSPs
  - Explicit null support for MPLS LSPs
  - MPLS Label Block Configuration
  - MPLS over untagged Layer 3 interfaces
  - MPLS OAM - LSP ping
  - JTI: OCST: MPLS operational state streaming (v2.2.0)

- 2K ingress LSP support
- 2K egress LSP support
- Entropy Label support
- MPLS: JTI: Junos Telemetry Interface MPLS self-ping, TE++, and misc augmentation
- LDP, including:
  - Configurable label withdraw delay
  - Egress policy
  - Explicit null
  - Graceful restart signalling
  - IGP synchronization
  - Ingress policy
  - IPv6 for LDP transport session
  - Strict targeted hellos
  - Track IGP metric
  - Tunneling (LDP over RSVP)
- RSVP++
- RSVP-TE, including:
  - Bypass LSP static configuration
  - Ingress LSP statistics in a file
  - RSVP-TE Hitless-MBB with no artificial delays
  - 32K transit LSPs
  - Auto bandwidth
  - Class based forwarding (CBF) with 16 classes
  - CBF with next-hop resolution
  - Convergence and scalability
  - Graceful restart signalling
  - JTI interface statistics and LSP event export
  - LSP next-hop policy
  - LSP self ping
  - MPLS fast reroute (FRR)
  - MTU signaling

- Optimize adaptive teardown
- Node/link protection
- Refresh reduction
- Soft preemption
- Shared Risk Link Group (SRLG)
- Static LSPs with IPv4 next-hop, IPv6 next-hop, and IPv6 next-hop with next-table support for bypass
- Traffic engineering, including:
  - TE++: Dynamic ingress LSP splitting
  - Traffic engineering extensions (OSPF-TE and ISIS-TE)
  - Traffic engineering options: bgp, bgp-igp, bgp-igp-both-ribs, mpls-forwarding

[See [MPLS Applications User Guide](#).]

- **no-decrement-ttl and no-propagate-ttl statements support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can configure the following configuration statements at the **[edit protocols mpls]** hierarchy level on the PTX10008 routers:
  - **no-decrement-ttl**—Decrements the IP time-to-live (TTL) by 1 before encapsulating the IP packet within an MPLS packet.
  - **no-propagate-ttl**—Disables normal TTL decrementing for all RSVP-signaled or LDP-signaled LSPs.

[See [no-decrement-ttl](#) and [no-propagate-ttl](#).]

- **RSVP and LDP LSPs on PTX10003**—Starting in Junos OS Evolved Release 20.1R1, you can configure RSVP and LDP point-to-multipoint label-switched paths (LSPs) on the ingress, transit, egress, PHP and bud interfaces of PTX10003 routers.

The following features are supported:

- RSVP-based point-to-multipoint LSPs
- Multicast LDP-based point-to-multipoint LSPs
- Point-to-multipoint LSP statistics

The following features are not supported:

- Point-to-multipoint LSP correlation
- Fast make-before-break (FMBB) protected Packet Forwarding Engine (PFE) support
- Point-to-multipoint LSP traceroute
- Virtual private LAN service (VPLS)
- Ethernet VPN (EVPN)

[See [Point-to-Multipoint LSPs Overview](#).]

- **Higher scale of transit LSPs with link protection (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support a higher scale of transit RSVP-TE signaled MPLS label-switched paths (LSPs) that are enabled with link protection.

### **Multicast**

- **MoFRR for PIM SM and SSM modes (PTX10003)**—Starting with Junos OS Evolved Release 20.1R1, PTX10003 routers support Multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows. MoFRR minimizes multicast packet loss in Protocol Independent Multicast (PIM) domains when there are link failures. When a failure is detected on the primary path, the repair is made locally by changing the interface on which packets are accepted to the secondary interface for the backup path, which greatly improves subsequent convergence times.

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](#).]

### **Network Security**

- **Control plane protection against DDoS attacks (PTX10008)**—Distributed denial-of-service (DDoS) attacks typically use a large number of network control packets to trigger exceptions in the network, consuming resources and bringing down network operations. Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers have control plane DDoS protection enabled by default for many Layer 2 and Layer 3 protocols. Control plane DDoS protection uses firewall filters and policers to discard or rate-limit control plane traffic so that malicious traffic does not interfere with device operations. You can configure **ddos-protection** options at the **[edit system]** hierarchy level to disable DDoS protection or change default policer parameters for a protocol group or for supported packet types in a protocol group.

[See [Control Plane Distributed Denial-of-Service \(DDoS\) Protection Overview](#).]

### **Network Management and Monitoring**

- **Local port mirroring and analyzer support (PTX10003)**—Starting with Junos OS Evolved Release 20.1R1, you can use port mirroring and port analyzers to copy packets entering or exiting a port or entering a VLAN and to send the copies to a local interface for local monitoring. You can send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

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

- **Support for adding custom YANG data models to the Junos OS schema (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support loading custom YANG data models on the device, which enables you to add RPCs or configuration hierarchies that are customized for your operations. The ability to add data models to a device is beneficial when you want to create

device-agnostic RPCs and configuration models that can be used on different devices from one or more vendors.

[See [Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS](#).]

- **Chef and Puppet provided as third-party applications (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support Chef and Puppet, which are included in the software image.

[See [Chef for Junos OS Getting Started Guide](#) and [Puppet for Junos OS Administration Guide](#).]

- **Support for running a Puppet agent as a Docker container (PTX10003-80C, PTX10003-160C, QFX5220-32CD, and QFX5220-128C)**—Starting in Release 20.1R1, devices running Junos OS Evolved support running a Puppet agent as a Docker container. As an alternative to using the Puppet agent that is integrated into the Junos OS Evolved software image, you can use the Puppet agent Docker container provided by Juniper Networks. Using the container enables you to use standard Docker tools to manage the container and mount or unmount the Puppet agent as needed.

[See [Installing Puppet for Junos OS](#).]

- **Support for running a Chef client as a Docker container (PTX10003-80C, PTX10003-160C, QFX5220-32CD, and QFX5220-128C)**—Starting in Junos OS Evolved Release 20.1R1, devices running Junos OS Evolved support running a Chef client as a Docker container. As an alternative to using the Chef client that is integrated into the Junos OS Evolved software image, you can use the Chef client Docker container provided by Juniper Networks. Using the container enables you to use standard Docker tools to manage the container and mount or unmount the Chef client as needed.

[See [Deploying Chef for Junos OS](#).]

- **sFlow support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, Juniper Networks supports sFlow on the PTX10008 router. The sFlow technology 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. The following sFlow features are supported:

- Adaptive sampling
- Adaptive sampling fallback, which decreases the sampling load when the traffic load decreases after adaptive sampling has taken place.
- Ability to configure forwarding class and DSCP values per collector.

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

- **Open-R supported (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, third-party application Open-R can operate on PTX10008 routers.
- **sFlow support for MPLS (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can configure sFlow on interfaces configured with the MPLS family to sample and report MPLS traffic on the PTX10008 routers.

See [\[Overview of sFlow Technology\]](#).

- **Support for DCU accounting and SCU accounting (PTX10003 routers)**—Starting in Junos OS Evolved Release 20.1R1, the destination class usage (DCU) accounting and source class usage (SCU) accounting are supported on PTX10003 routers. You can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets, which are defined as source classes and destination classes. SCU counts packets sent to customers by performing lookups on the source IP address and the destination IP address. SCU accounting enables you to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. DCU counts packets from customers by performing lookups of the IP destination address. DCU accounting enables you to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.

### **Port Security**

- **Media Access Control Security (MACsec) enhancements (PTX10003)**—The following MACsec enhancements are supported on PTX10003 routers:
  - Pre-shared key (PSK) hitless rollover keychain
  - Fallback PSK
  - Configuration option for the destination EAPoL address

MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is standardized in IEEE 802.1AE.

[See [Understanding Media Access Control Security \(MACsec\)](#).]

### **Routing Policy and Firewall Filters**

- **Firewall filters support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, you can configure IPv4 and IPv6 firewall filters on your PTX10008. Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on the router from a source address to a destination address. You configure firewall filters at the **[edit firewall]** hierarchy level.

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

- **MPLS LSP policer support (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, you can apply MPLS LSP policing on your PTX10003. Policing allows you to control the amount of traffic forwarded through a particular LSP and ensures that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. The router supports automatic policing only on regular LSPs. Automatic policing allows you to provide strict service guarantees for network traffic. Such guarantees are especially useful in the context of Differentiated Services for traffic-engineered LSPs, providing better emulation for ATM wires over an MPLS network.

[See [Configuring MPLS Firewall Filters and Policers on Routers](#).]

## Routing Protocols

- **BGP PIC edge support for inet and MPLS VPNs (PTX10003)**—Prefix-Independent Convergence (PIC) edge support is extended to BGP with multiple routes in the global tables such as **inet** and **inet6 unicast**, and **inet** and **inet6 labeled unicast**. When the BGP PIC feature is enabled on a router, BGP installs in the forwarding table the second best path in addition to the calculated best path to a destination. When an IGP loses reachability to a prefix, the router uses this backup path to reduce traffic loss until the global convergence through BGP is resolved, thereby drastically reducing the outage duration.

You can also install a Layer 3 VPN route in the forwarding table as an alternate path, enabling fast failover when a provider edge (PE) router fails or you lose connectivity to a PE router. This already installed path is used until global convergence through the IGP is resolved.

To enable BGP PIC edge in an MPLS VPN, include the **protect-core** statement at the **[edit routing-instances routing-instance-name routing-options]** hierarchy level. Both IS-IS LDP and OSPF LDP are supported. When BGP PIC edge is enabled, the **show route extensive command** now displays the weight assigned to the indirect hop.

[See [Configuring BGP PIC Edge for MPLS Layer 3 VPNs](#) and [Use Case for BGP PIC for Inet.](#)]

- **Redistribution of IPv4 routes with IPv6 next hop into BGP (PTX10008)**—Devices running Junos OS Evolved Release 20.1R1, can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, IPv4 traffic is tunneled from CPE devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers, which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop.

To configure a dynamic IPv4-over-IPv6 tunnel, include the **dynamic-tunnels** statement at the **[edit routing-options]** hierarchy level.

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

- **Support for inline BFD (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, PTX10003 routers support inline BFD.

[See [Understanding BFD for Static Routes for Faster Network Failure Detection.](#)]

- **Support for 128-way ECMP for MPLS transit case (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, 128-way equal-cost multipath routing (ECMP) is supported. The following features do not support 128-way ECMP:
  - consistent hashing
  - link protection (MPLS)
  - class-based forwarding
- **BGP flow specification (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX10008 devices running



14.4 Tbps line cards. Propagating firewall filter information as part of BGP enables you to propagate firewall filters against denial-of-service (DOS) attacks dynamically across autonomous systems.

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]

The following actions are not supported:

- Traffic Marking

To configure flow routes statically, configure the match conditions and actions under the **[edit routing-options]** hierarchy level.

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

- **Distributed BFD and BFD-triggered local repair (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support distributed BFD and BFD-triggered local repair. BFD authentication is not supported.

[See [BFD-Triggered Local Repair for Rapid Convergence](#) .]

- **Support for NTP (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support NTP. The following features are supported:
  - NTP over WAN ports with IPV4 or IPV6 configuration.
  - Authentication Over WAN ports
  - Support for MD5, SHA1, and SHA256 authentication
  - Support for the following NTP modes:
    - Client/server
    - Peer mode
    - Broadcast mode

### ***Services Applications***

- **Support for real-time performance monitoring (RPM) on PTX10003**—Starting with Junos OS Evolved Release 20.1R1, PTX10003 routers support RPM. However, in this release, Junos OS Evolved supports only the following RPM features:
  - User configuration parsing for RPM
  - Probe generation and reception (client) as well as reflection (server) for the following RPM probe types:

- icmp-ping probes
- icmp-timestamp probes
- udp-ping probes
- udp-timestamp probes
- Probe history management as well as reporting through syslog.

[See [Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.](#)]

- **Support for MPLS-based J-Flow (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10003 routers supports MPLS-based inline J-Flow for ingress sampling. MPLS-based flow records are based on Version 10 (IPFIX) and Version 9.

[See [Configuring Inline Active Flow Monitoring on PTX Series Routers.](#)]

- **MPLS-based J-Flow support (PTX10008)**—Starting with Junos OS Evolved Release 20.1R1, the PTX10008 routers supports MPLS-based inline J-Flow for ingress sampling. MPLS-based flow records are based on Version 10 (IPFIX) and Version 9.

[See [Configuring Inline Active Flow Monitoring on PTX Series Routers.](#)]

- **Support for J-Flow monitoring (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, the PTX10008 router with the JNP10K-LC1201 line card supports inline J-Flow monitoring. J-Flow templates support both Version 9 and Version 10 (IPFIX) formats.

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

- **Support for real-time performance monitoring (RPM) on PTX10008**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 routers support RPM. However, in this release, Junos OS Evolved supports only the following RPM features:
  - User configuration parsing for RPM
  - Probe generation and reception (client) as well as reflection (server) for the following RPM probe types:
    - icmp-ping probes
    - icmp-timestamp probes
    - udp-ping probes
    - udp-timestamp probes
  - Probe history management and reporting through syslog

Consider the following caveats while configuring RPM:

- Configure **probe-server** at the **[edit services rpm]** hierarchy level mandatorily for ICMP and UDP probes.
- Specify **address** (or **inet6-address**) keyword for an IPv4 (or IPv6) address for **target** at the **[edit services rpm probe owner1 test test1 rpm-scale]** hierarchy level. The **target** type is deduced.
- Need not specify **std-dev-egress**, **std-dev-ingress**, and **std-dev-rtt** thresholds for the probes at the **[edit services rpm probe owner test test-name]** hierarchy level.
- Configure **offload-type (none|pfe-timestamp)** instead of **hardware-timestamp** to enable timestamping.

[See [Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.](#)]

### **Software Defined Networking**

- **PCE-initiated point-to-multipoint LSPs support (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, the Path Computation Element Protocol (PCEP) support on PTX10003 routers is extended to include Path Computation Element (PCE)-initiated point-to-multipoint LSPs.

With this feature, a PCE can initiate and provision a point-to-multipoint LSP dynamically without the need for local LSP configuration on the Path Computation Client (PCC). This enables the PCE to control the timing and sequence of the point-to-multipoint path computations within and across the PCEP sessions, thereby creating a dynamic network that is centrally controlled and deployed.

[See [Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs.](#)]

- **PCEP support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, the Path Computation Element Protocol (PCEP) is supported on PTX10008 routers.

[See [PCEP Overview.](#)]

### **Software Installation and Upgrade**

- **Zero touch provisioning dynamically detects port speed of WAN interfaces (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, Zero Touch Provisioning (ZTP) dynamically detects the port speed of WAN interfaces and uses this information to create ZTP server ports with the same speed. ZTP cannot proceed unless the WAN and ZTP server interface speeds match. To make sure the interface speeds match, ZTP configures the interfaces in the various port speeds supported by the platform. This happens automatically and does not require any configuration by the user.

ZTP automatically cycles through the WAN ports that can support different speeds until it receives DHCP arguments from the port connected to the DHCP server, and then connects your device to the Dynamic Host Configuration Protocol (DHCP) server to perform the bootstrap process.

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

- **Zero touch provisioning on management interface (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, PTX10008 devices support automation of the device configuration and software upgrade over the management interface of Routing Engine 0 (RE0).

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

### **Spanning Tree Protocols**

- **Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocol (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP) support (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, You can use STP, RSTP, MSTP, and VSTP to provide Layer 2 loop prevention.

[See [Spanning-Tree Protocol Overview](#).]

### **System Logging**

- **Support for system log notifications over MQTT when configuring JET applications (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you can configure JET applications by using the CLI and then enable the applications to subscribe to system log events. The Junos OS Evolved eventd process sends system log notifications to the event broker. The broker then sends the notifications formatted as JSON messages to the JET application over the Message Queuing Telemetry Transport (MQTT) protocol.

To enable a JET application to subscribe to system log events, configure the **extension-service** statement at the **[edit system extensions]** hierarchy level. You can manage the JET application by using commands at the **[edit system extensions]** and **[edit system services extension-service]** hierarchy levels.

[See [JET Architecture](#), [request extension-service \(start | stop\)](#), [show extension-service status](#), [extensions](#), and [extension-service \(System Services gRPC\)](#).]

### **System Management**

- **Commit Synchronize redesigned (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, **commit synchronize** is enabled by default. If you issue **commit** from the master Routing Engine, you see that the backup Routing Engine is automatically synchronized. You can see the details of various tasks performed during commit of both the Routing Engines in the **commit** output.

[See [Synchronizing Configurations Across Routing Engines](#).]

- **SNMP and RMON Support (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, the following SNMP and RMON features are supported on the PTX10008 router:
  - SNMP version 1, version 2, and version 3
  - RMON events, alarms, and history—Junos OS supports monitoring routers from remote devices. These values are measured against thresholds and trigger events when the thresholds are crossed. You configure remote monitoring (RMON) alarm and event entries to monitor the value of a MIB object.
  - Juniper Networks enterprise-specific SNMP utility MIB
  - Subagent statistics using **show snmp statistics subagent**
  - Enabling SNMP management routing instances using the **routing-instance-access** statement at the **[edit snmp]** hierarchy level

[See [Junos OS SNMP FAQs](#).]

- **Passive monitoring support (PTX10003)**—Starting in Junos OS Evolved Release 20.1R1, you can enable passive monitoring on your PTX10003 so that any network interface can be configured as a passive monitoring interface. Passive monitoring provides filtering capabilities for monitoring ingress and egress traffic at the Internet point of presence (PoP) where security networks are attached. When you enable passive monitoring, the device accepts and monitors traffic on the interface and forwards the traffic to monitoring tools such as IDS servers and packet analyzers, or other devices such as routers or end-node hosts. To enable this feature, include the **passive-monitor-mode** statement at the **[edit interface]** hierarchy level.

[See [Understanding Passive Monitoring](#).]

- **Run third-party applications in Docker or LXC containers (PTX10008)**—Starting with Junos OS Evolved Release 20.1R1, you can run applications inside containers. You can run third-party applications in two types of containers: Docker containers and LXC containers.

[See [Junos OS Evolved Overview](#).]

## What's Changed

### IN THIS SECTION

- [Changes in CLI Statements and Commands | 46](#)
- [What's Changed in Junos OS Evolved Release 20.1R2 | 46](#)
- [What's Changed in Junos OS Evolved Release 20.1R1 | 46](#)

Learn about what changed in Junos OS Evolved Release 20.1R2 for the PTX10003 and PTX10008.

## Changes in CLI Statements and Commands

- To see the list of CLI statements and commands changed in Junos OS Evolved, see *How Junos OS Evolved Differs from Junos OS* in the [Introducing Junos OS Evolved Guide](#).

## What's Changed in Junos OS Evolved Release 20.1R2

### *Class of Service (CoS)*

- **PTX10001-36MR and PTX10008 routers support a maximum of two drop profile pairs (PTX Series)**—Pair one drop probability must be less than or equal to 25%. Point two drop probability value must be greater than point one drop probability value. Pair two fill level must be greater than or equal to 1.2 times the pair one fill level.

[See [CoS Features and Limitations on PTX Series Routers](#).]

### *General Routing*

- **Support for full inheritance paths of configuration groups to be built into the database by default (PTX10003 PTX10008, and QFX5220)**—Starting with Junos OS Evolved Release 20.1R2, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

[See [commit \(System\)](#).]

### *Network Management and Monitoring*

- **Enhancement to the show snmp mib command**—In Junos OS Evolved Release 20.1R2, we've introduced a new option, **hex**, to display the SNMP object values in hexadecimal format. In earlier releases, the **show snmp mib** command displays the SNMP object values in ASCII and decimal format only.

[See [show snmp mib](#).]

### *Routing Protocols*

- **Recommendation to include the local-address statement when configuring IBGP and multihop EBGP**—When a device peers with a remote device's loopback interface address, use the **local-address** statement at the **[edit protocols bgp group internal-peers]** hierarchy to specify the source information in BGP update messages. Although a BGP session can be established when only one of the paired routing devices has local-address configured, we strongly recommend that you configure **local-address** on both paired routing devices for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

[See [local-address \(Protocols BGP\)](#) and [BGP Peering Sessions](#).]

## What's Changed in Junos OS Evolved Release 20.1R1

### *Interfaces and Chassis*

- **Modified Ethernet interface configuration (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you'll no longer see the **gigether-options** statement at the **[edit interfaces interface-name]** hierarchy level. To configure Ethernet interfaces and link aggregation groups (LAGs), use the **set interfaces interface-name ether-options** command instead.

The following options are added to the **ether-options** statement:

- **fec**
- **loopback-remote**

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

### *Junos Telemetry Interface (JTI)*

- **LLDP ON\_CHANGE statistics support with JTI (PTX Series, QFX Series)**—Enhanced telemetry ON\_CHANGE event support provides the following LLDP attributes: - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes. - ON\_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered

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

### *MPLS*

- **Introduction of name option (PTX10003-160C and PTX10008)**—We've introduced the **name** option for the **show path-computation-client lsp** command and the **show path-computation-client lsp extensive** command. You can filter a particular LSP from a large number of LSPs by entering the specific name of the LSP.

See [show path-computation-client lsp](#).

- **Root XML tag change for show rsvp pop-and-forward | display xml command (PTX10003, PTX10008, and QFX5220)**—We've changed the root XML tag for the **show rsvp pop-and-forward | display xml** command to **rsvp-pop-and-fwd-information** to make it consistent with the XML tag convention. In earlier releases, the command output displays **rsvp-pop-and-fwd-info** XML tag. Update the scripts with the **rsvp-pop-and-fwd-info** XML tag to reflect the new **rsvp-pop-and-fwd-information** XML tag.

[See [Junos XML API Explorer - Operational Tags](#).]

### System Logging

- **Removal of `show system queues` command (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you'll no longer be able to use the `show system queues` command because the information in the command output is not relevant.

### System Management

- **Change of default wait time between Routing Engine switchovers (PTX10008)**—Starting in Junos OS Evolved Release 20.1R1, we've increased the default wait time between Routing Engine switchovers from 120 seconds to 360 seconds when using the `request chassis routing-engine master switch`.

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

## Known Limitations

### IN THIS SECTION

- General Routing | 48
- Flow-Based and Packet-Based Processing | 50
- Infrastructure | 50
- Interfaces and Chassis | 50
- User Interface and Configuration | 51

Learn about limitations in this release for the PTX10003 and PTX10008. For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online [Juniper Problem Report Search](#) application.

### General Routing

- On PTX10008 routers, control packets statistics are not accurate for aggregated Ethernet subinterfaces (logical units). [PR1460631](#)
- On PTX10008 routers, the subsystem within the Packet Forwarding Engine continues to monitor the ASIC for new interrupts even for the ASIC for that all the interfaces are disabled. On an ASIC which has all the interfaces disabled due to a fatal error, all the new interrupts that are reported after the fatal event should be ignored. [PR1470391](#)
- The `error` option in the `clear system` command is not supported. [PR1472827](#)



- On PTX10008 routers, the counter per aggregated Ethernet child member for egress statistics are not accurate for aggregated Ethernet subinterfaces (logical units) and they should be ignored. [PR1475078](#)
- While the FPC is offlining, there a 2-5 seconds window when the pic-status command fails due to communication failure with hwdfpc. This is expected and has no functional impact. [PR1484344](#)
- On PTX10008 routers, counter per AE child member for egress statistics are not accurate for AE sub-interface (logical units) and they should be ignored. [PR1490130](#)
- UDP sensor streaming is not currently supporting FPC restart and apps that use Junos OS Evolved restart (for example. evo-aftmand-bt, evo-cda-bt). [PR1492096](#)
- Learnt MAC's on PTX10003-80C and PTX10003-160C are transported as Mac learning packets that is, MLP's towards Packet Forwarding Engine CPU. When there is burst of learning events, we can land into scenarios where these packets use more bandwidth than towards's lcpu interfaces supports. We debugged this with ASIC team and they have confirmed that these errors appear when more traffic is enqueued towards LCPU than it can support. Once MAC's are learnt these errors stop. Also there is no functional impact as such of these errors because TOE ucode would retry with MAC learning packets, if these are dropped anywhere in packet path. [PR1493132](#)
- If **vrf-table-label** configuration is enabled (LSI label), then SLU does the tunnel termination (early lookup). Once SLU terminates the MPLS tunnel, downstream blocks only get inner headers and this is an ASIC limitation. If the **vrf-table-label** configuration is disabled (L3VPN label) then label lookup happens in DLU. In this scenario, the applied MPLS firewall filters applied work. [PR1493504](#)
- UDP sensor streaming is not currently supporting FPC restart and apps that use Junos OS Evolved restart (eg. evo-aftmand-bt, evo-cda-bt). [PR1495108](#)
- On PTX10008, leaf under /interfaces/interface/ethernet/state/ is not supported. [PR1496275](#)
- Support for ECMP path trace tool using APIs (PTX10003) >Starting in Junos OS Evolved Release 20.1R1 or 19.4R2, you can use ECMP tracer JET APIs to monitor traffic flow, trace the ECMP path used by packets entering and exiting a CLOS network, and get flow counters without changing the configuration on your device. You provide match criteria to the ECMP path trace tool, ecmptracerd, through JET API requests. The ecmptracerd process then installs firewall filters for the inputs on all the interfaces on the device except the private interfaces. The counters associated with these filters increment as packets matching the firewall filter criteria arrive and leave. These APIs provide similar functionality to the ecmp-tracer on-box shell utility. Do not use the JET client and the ecmp-tracer on-box shell utility simultaneously. [See [ecmp](#), [Juniper Extension Toolkit Developer Guide](#), and [Juniper Engineering Network website](#).] [PR1496613](#)
- PCIe aer uncorrectable errors might be seen on the console when sibs are powered off during system boot, CLI offline or when new Routing Engine connects to the sibs after mastership switchover. These are expected and are not indicative of underlying functional issues. [PR1501647](#)
- This behaviour occurs because order of execution of **clear ethernet switching** and traffic stop. Had traffic been stopped before and we had done **clear ethernet switching** we would not have seen this issue. What happens is that after **clear ethernet switching**, there are MLP's in-flight because traffic was continuously running. Now since CPU was busy clearing MAC's it would taken more than 20 seconds

to acknowledge these MAC's. Once we go beyond the 20 seconds time TOE ucode would have invalidated these entries (i.e. cleared up the SLU hash table entry). Now when Packet Forwarding Engine tries to acknowledge the entries towards the ucode it doesn't find entry in the hardware and returns a failure to L2alm. L2alm retries the Mac add 5 times and returns an error back to L2ald when Packet Forwarding Engine continues to return failure. L2ald would clean up these MAC's once it receives the message from L2alm. Until such time the MAC's show up in **show ethernet switching**. This does not affect functionality because if traffic is started the MAC's would get learnt regularly. [PR1507812](#)

- This is a harmless and expected interrupt/minor alarm. If this is noted that on link down then it can be ignored. This is cleared automatically. This alarm cannot be disabled as in valid scenarios we still need this alarm to be enabled so that it can be reported in case of problem. [PR1509779](#)
- bandwidth-percent is not supported under policer configuration **set firewall policer p if-exceeding bandwidth-percent**. [PR1513149](#)

## Flow-Based and Packet-Based Processing

- Two cases are listed in this PR 1. VRF-Label-Table configuration 2. UHP configuration. In both of these cases, the SLU module in the forwarding path will strip off the MPLS labels. After the labels are stripped off, this packet is seen as an IPv4 or IPv6 packet, hence it does not get sampled as an MPLS packet. This is the behavior in BT ASIC. [PR1495502](#)

## Infrastructure

- This is a corner case that could occur in Junos OS Evolved Release 20.1R1 for dual Routing Engine PTX10008. If this happens, the way to recover the device is to reboot the Routing Engines. [PR1496366](#)

## Interfaces and Chassis

- 1. Physical interface is bounced when **vlan-tagging** configuration changes. 2. PTX10003-80C and PTX10003-160C-AFT code re-initiates the ASIC WANIO blocks, which results in the physical interface flapping. This limitation is acceptable because a VLAN tagging change is more or less a network topology change (because the tagging has to be changed on the far-end peer interface also). Hence a link flap during this event is not harmful. [PR1382130](#)
- When Picd App is restarted, sometimes, 10GE Link status remains Down. Link can be recovered by disabling/enabling remote end. [PR1488146](#)
- When the loopback configuration is applied to a LAG (aggregated Ethernet) interface the interface goes down. [PR1497591](#)
- PTX10008 is physical FPC with a dedicated CPU on the FPC. PTX10003 is virtual FPC with no CPU. Hence for PTX10008 the vty connects to the right CPU and fails if incorrect FPC number is given. PTX10003 - the vty is a dummy cmd to the FPC which is virtual. We just run a script to connect to the

CLI-Packet Forwarding Engine. Hence this behavior. This has to be lived with - there's no impact. And there is one CLI-Packet Forwarding Engine for the entire set of FPC. Hence FPC granularity not needed.  
[PR1499413](#)

- Leaf sensor for logical is not supported on Junos OS Evolved.[PR1511474](#)

## User Interface and Configuration

- This is a corner case with a workaround. The system is not clogged with stale users, but merely the messages coming when user enters configuration mode. This will not have any functional impact on the system. The issue will not be seen with the case of graceful reboot or a clean exit. Command **request system logout user *username* terminal pts/0** can be used to remove the stale user as a workaround.  
[PR1409851](#)

## Open Issues

### IN THIS SECTION

- [Class of Service \(CoS\) | 52](#)
- [Fault Management | 52](#)
- [General Routing | 52](#)
- [Infrastructure | 57](#)
- [Interfaces and Chassis | 57](#)
- [Interfaces and Chassis | 58](#)
- [Network Management and Monitoring | 58](#)
- [Routing Protocols | 58](#)

Learn about open issues in this release for the PTX10003 and PTX10008. For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online [Juniper Problem Report Search](#) application.

## Class of Service (CoS)

- cosd core file seen @ CosRewriteRulesFCLossSniplet::populateOperObject, when core profile configuration rollback is done to the base configuration. [PR1475637](#)
- CoS ieee-802.1p default classifier is not working as expected, after deleting the customer ieee-802.1p classifier/rewrite. [PR1478563](#)
- COS "default" classifiers and rewrite rules are not working for both **dscp** and **dscp-ipv6**. [PR1479326](#)
- 32 DSCP classifier/rewrite functionality is not working after the cosd restart ( the PTX10008 platform supports only 16 DSCP classifier/rewrite). [PR1479676](#)
- 32 DSCP classifier/rewrite functionality is not working after the cosd restart ( the PTX10008 platform supports only 16 DSCP classifier/rewrite). [PR1479676](#)
- Support import - Classifier/Rewrite. [PR1483505](#)
- Before this PR fix is committed "Below are the valid interface names supported by set class-of-service interface configuration: et-\* or et-\*/\*/ or ae\* or fti\* or any valid interface name (without wildcard eg.et-1/\*/\*). Cosd might core when the wildcard interface is used such as et-1/\*/\*" Post this PR fix Junos OS Evolved based products support below interface name wildcards for **set class-of-service interface** configuration. 1) Specific interface (for example: et-0/0/0 or et-0/0/0:0 or ae0 or fti0) 2) All channels in a port (eg. et-0/0/0:\*) 3) All Ports in a Pic (eg. et-0/0/\*.\* or et-0/0/\*) 4) All Pica/Ports in a FPC(Slot) (for example: et-0/\*/\*.\* or et-0/\*/\*) 5) All ET interfaces: et-\*/\*/ or et-\*//\*/\*.\* or et-\* For AE, ae\* and For FTI, fti\* wild cards are supported. Note: Preference is be provided for a more specific configuration over a wildcard. for example: et-1/\*/\* over et-\*/\*/ - et-1/1/\* over et-1/\*/\* - ae4 over ae\* - fti0 over fti\* and so on. [PR1500722](#)

## Fault Management

- CRC errors on an interface might result in Cmerror and ASIC errors in logs. [PR1499291](#)

## General Routing

- Static link protections are not supported on PTX10003. [PR1382695](#)
- In a scale scenario of million routes and thousands of LSPs, **show mpls lsp statistics** is slow due to high volume of route installation and statistics queries. [PR1416363](#)
- **irp\_intr\_smp\_trap** error messages are seen on PTX10003 console while bouncing the IGP protocols and restarting routing the daemon. [PR1422881](#)
- There are differences in the XML tags for the CLI **show version** command as compared to the Junos OS CLI. [PR1429640](#)
- There are differences in XML tags for the Junos OS **show interfaces**, **show interfaces terse**, and **show interfaces statistics** CLI commands compared to Junos OS. [PR1433459](#)

- Differences in XML tags for **show system virtual-memory**. [PR1438110](#)
- For the **inet6** family, DSCP action is not supported for BGP flowspec filters. If such a configuration is tried, the **inet6** family filtering does not happen for BGP flowspec rules. [PR1439366](#)
- **AftCore: AftValidate: Validation failed for token** errors are seen while validating MPLS LSP counters. [PR1447452](#)
- During fabric link bringup/fabspoked-fchip restart/sib offline or sib online, /re0/fabspoked-fchip might not respond to CLI commands (**show chassis sibs**). This is indicated by following output on CLI: **error: communication failure with /re0/fabspoked-fchip/**

- The producer app is either down or unresponsive.

Run **show system processes node node | grep app** to check if app is running.

Run **show system application app app** to check the state of the app

- If there is a change in mastership recently, please wait for the switchover to complete for the app to be online.

Run **request chassis routing-engine master switch check** to check the status of switchover.

If the command is reissued after training/detraining reaches a steady state, valid output shows up in CLI.

#### [PR1459430](#)

- The SSH service is configured with certain key-exchange algorithms, the SSH server will be unresponsive. For example: set system services ssh key-exchange dh-group1-sha1. [PR1467467](#)
- For PTX10008 weighted ECMP MPLS is not supported. [PR1472889](#)
- Traffic loss of up to 400 ms is seen during an MPLS FRR test with the following scale: 600,000 IPv4 routes, 40,000 IPv6 routes, 19,000 transit LSPs, 1600 ingress LSPs, 10,000 BGP-LU. [PR1472908](#)
- J-flow sampling traffic is not rate-limited in host pipe. [PR1473844](#)
- On a PTX10008, ungraceful line-card or PIC removal is not supported with traffic running. Offline and online commands should be used. [PR1475694](#)
- Use Traffic statistics instead of Packet Forwarding Engine statistics to understand the flow of traffic. [PR1478244](#)
- We might see a FPC evo-aftmand-zx process generate a core file when OSPF is configured and one of the interface in the configuration undergoes a major change. A major interface configuration change would involve deletion and re-installation of the interface and its subcomponents in the forwarding path. To recover from this crash, the FPC needs to be rebooted. [PR1478811](#)
- BGP started creating route pointing to List NH having INDIRECT members inside after comcast N+1 feature. But resolver treats LIST NH as a leaf node and does not do loop detection inside List NH members (Day 1 behavior). This causes wedges in Packet Forwarding Engine as loop detection does not happen. [PR1485478](#)

- The local repair time for fast reroute is 50 ms range. If the system has scaling configuration or heavily loaded for processing, it local repair time may observe longer than 50 ms. In this case it gets 65 ms local repair time. [PR1489139](#)
- MPLS LSP Label is not getting imposed on for BGP V6Prefix in Traceroute response with Ultimate-hop-popping. This issue will not be addressed in 20.1R2 and 20.2R1. [PR1489796](#)
- Routing Engine switchover when fabric links are in transition state is not supported in Junos OS Evolved Release 20.1R2, and can leave the system in an inconsistent state. Such fabric link transitions are triggered by the following events: SIB Offline - SIB Online - FPC Offline - FPC Online. We recommend that you wait till all fabric link transitions are complete before triggering Routing Engine switchover. [PR1490182](#)
- On PTX10008 routers, the subsystem within the Packet Forwarding Engine reports a series of new interrupts for the ASIC for which the PCIe address bus or PCIe link is down. On an ASIC that has issues with the PCIe address bus or PCIe link, all the new interrupts that are reported are false alarms and should be ignored. [PR1491128](#)
- When all SIBs are taken offline, use chassis restart to bring them back online. Failure to restart can cause a traffic black-hole condition (packets get dropped silently). [PR1491770](#)
- The scenario is MPLS segment routing SWAP case. The impact here is that the next hop reported in the export record is incorrect, while all other fields exported to the collector are correct. [PR1492616](#)
- 'Fabric drops' counter shows 0 in the **show pfe statistics traffic** CLI command output. [PR1494226](#)
- Scheduler ingress Packet Forwarding Engine VOQ drop counters do not match egress queue drop counters (diff > 100,000). [PR1494785](#)
- JNP10K-LC1201 - Interfaces do not show up after the line card is taken offline or brought back online, experienced OIR, or is restarted and brought online. [PR1494793](#)
- All the IPv6 addresses learned from DHCP server continue to have the /64 prefix. There will not be any reachability issue with this. [PR1494823](#)
- Offline/Online FPC feature is not fully supported on PTX10003 series. We can still do offline, but for bringing FPC or the associated ports back online requires a reboot of system. It is recommended we do not use FPC offline/online this feature is completely supported. It was also seen in some cases with scale configurations and heavy churn in system, off-lining an FPC were disabling packet forwarding on other FPCs as well. To recover, system reboot the system. [PR1495118](#)
- The copying of files to the RCB over WAN ports is slow. This is observed across all Junos OS Evolved platforms. [PR1496895](#)
- There is a discrepancy of 22 bytes for the same exact packets between the firewall filter in Junos OS Evolved Release and in Junos. From RPM perspective, probes are sent and received correctly, as shown in the description. Issue is in the byte counts reported by the fw filter. [PR1503145](#)
- Routing engine removal and insertion will rarely result in hwdre process restart(crash), with no functionality impact. [PR1503158](#)

- In PTX10008, for Junos OS Evolved Release 20.1R2, AE interface flap causes unexpected behavior. Few of the outgoing traffic streams carrying wrong encap information cause them to be dropped in the next hop router. This issue does not happen every time, its very sporadic. Issue seen when router having around 800,000 IP routes, 2,000 ingress LSPs and around 500 L3VPNs and LSPs have link protection enabled. [PR1503260](#)
- As per the current implementation, To display 'switchover status' under **show system switchover** command on backup only the readiness of 'configuration DB', 'object DB' and 'Applications' ready state' are considered. It does not check for the output of **request chassis routing-engine master switch check** command. Due to which in some of the cases output of both the commands could be different for 'switchover readiness' for a short duration of time. [PR1503269](#)
- L2 circuit ping that uses CC type 1, also referred to as "PWE3 Control Word with 0001b as first nibble" is not supported. [PR1503724](#)
- The "hwdr: HWD\_FIRMWARE\_VERSION\_READ\_ERROR\_NOTICE: Failed to get firmware version for fru: Sib" message is seen for each SIB after a graceful offline and online. This is expected behavior since the SIB version is read correctly during an online transition. Please check that the SIB versions are being displayed correctly using "show system firmware" after the SIBs are online. [PR1504156](#)
- During SIB offline or online, if **show chassis sib** CLI command is run continuously from multiple sessions as in this PR: **show chassis sib |refresh 1** - this could affect fabric process performance causing delays in SIB initialization and traffic recovery. [PR1504267](#)
- evo-aftmand crashes on vlan changes on ingress and egress ports. [PR1504339](#)
- When changing a physical interface from vlan-id-list to multiple single vlan ids or vlan-id-lists that fall within the original vlan range, we may see one or more of the new physical interfaces not notified. E.g. In the current commit we have the following configuration **set interfaces ae100 unit 1 vlan-id-list 1-100**. In the next commit we do the following **set interfaces ae100 unit 1 vlan-id 1 set interfaces ae100 unit 2 vlan-id 2**. We may see that one of the two physical interfaces is not notified. [PR1504374](#)
- DHCPv6 clients might fail to bind after the dhcp-relay configuration is deleted and then quickly re-added after a short time later in a given routing instance. [PR1505993](#)
- With all 8 output queues on a port congested, it may sometimes not be able to honour more than 70% guaranteed-rate on a single output queue. [PR1506855](#)
- In PTX10008: - 6 SIBs are required to carry line rate traffic - no fabric redundancy even when ingress traffic rate is such that 5 SIBs are sufficient to carry ingress traffic (e.g., traffic is less than 1280Gbps per BT's Data path), graceful SIB offline may result in small transient loss of traffic, till system re-routes traffic around fabric paths going offline. [PR1506866](#)
- When firewall app is restarted, it takes more than 90s to be online. Traffic may be impacted until the firewall app becomes 'Online Ready'. [PR1507433](#)
- When the ERO of a particular branch is updated and the update goes to PCC just for this branch then the issue is seen. [PR1508412](#)

- With one or more fabric cards down due to bias voltage failure, Routing Engine mastership switch could lead to non-deterministic fabric state in the system. [PR1510410](#)
- 400G interface takes a long time to come up sometimes. [PR1510587](#)
- In a scenario where there are multiple 'degraded\_fabric\_reachability\_to\_peer\_pfe' to different SIBs, on SIB restart there is inconsistency in errors that gets cleared. Errors are seen even though there is recovery of few planes from fabric degradation. [PR1510763](#)
- Alarms for SIB Link error under **show system alarms** are getting cleared after recovery of few planes corresponding to a restarted SIB even though there are active link faults to other SIB. [PR1510766](#)
- In a scaled setup with 4,000 logical interfaces and 8,000 policers, when there is a reactivate and activate of firewall and interface configuration, evo-aftmand process spins for 100% and aggregated Ethernet interfaces do not come up. [PR1511597](#)
- With fabric degradation detection enabled, if a fabric card goes down due to a fault at run time, relevant alarms or errors might not be raised and related fault action is not taken. [PR1511910](#)
- **fabric\_down\_condition\_on\_pfe** errors not cleared when fpc is offlined after fabricHub app restart. [PR1511918](#)
- Issue: Firewall crash, when a filter is deleted and added back, in single commit. Impact: Firewall core. Firewall recovers after the core. Workaround: Delete and add should not be done in single commit. If there is an attempt to program Filter (F1) . Delete F1 ; Commit . Add F1, Terms, ..; Commit [PR1512065](#)
- In scenarios with fabric degradation enabled, both errors **degraded\_fabric\_reachability\_to\_peer\_pfe**, **degraded\_fabric\_condition\_on\_pfe** are seen on Packet Forwarding Engine. However Packet Forwarding Engine is disabled and all the interfaces on the Packet Forwarding Engine go down due to above condition. [PR1512270](#)
- With fabric degradation enabled, once degraded\_fabric\_condition\_on\_pfe is cleared on Packet Forwarding Engine, Evo applications fail to set degraded\_fabric\_reachability\_to\_peer\_pfe on the same Packet Forwarding Engine thereby. [PR1512271](#)
- SIB <> FPC Link Errors seen prior to switchover do not get cleared when switchover is followed by SIB restart. [PR1512272](#)
- Under rare circumstances, when multiple fabric cards are offlined and online in quick succession, the device might crash and go for a reboot after reporting a kernel 'msmi error' or 'zookeeper session failed' error. [PR1512633](#)
- Fabric cards might get stuck in 'offlining' state some times if offlined/online in quick succession without allowing for the cards to get to steady state. [PR1512814](#)
- evo-cda-bt core file might be seen sometimes during ASIC initialization when FPC is in the process of coming up as "online". [PR1513067](#)
- On dual Routing Engine box, by default, the commit process synchronizes the configuration across both the Routing Engines. Rarely, commit synchronization process malfunction and thus could not synchronize the RE0 configuration to the RE1. Under the problem scenario, the following steps can synchronize the



RE0 configuration to the RE1. 1. Login to the RE1 shell 2. `rm /config/.mgdInitialized` 3. `systemctl restart mgd-init` 4. `systemctl restart mgd` 5. `systemctl restart config-sync` [PR1513142](#)

- If FPC reboots ungracefully (including using `node reboot fpc <slot>` ) then we see a traffic impact OR host path it self not functional, due to this FPC might not be useful. [PR1513451](#)
- Under rare circumstances, with fabric cards in fault state, graceful Routing Engine master switchover followed immediately by **request chassis sib slot X offline** command might cause the fabric cards in fault state to be restarted rather than being moved to 'offline' state. [PR1513772](#)
- After picd app restart or FPC restart, Inphi colored optics does not link up. [PR1514058](#)
- When there have been no errors ever in the system, the **show system errors active** command with or without fru filters would give empty output, that is, no counts are shown. This does not affect functionality since what is shown is correct. When FRUs are offlined, errors that might have been raised on them would get deleted from the DB. If in case all the errors that have been raised get deleted, we come to the same situation as mentioned earlier, and the **show system errors active** command with or without fru filters would give empty output. [PR1514105](#)
- During system comes up, if fpc offline is performed, evo-aftmand crash might be seen. [PR1514722](#)
- For fabspoked-pfe, 'Restart Supported' and 'Restart Node' flags are incorrect in the output of `show system applications app fabspoked-pfe detail`. Show command shows that app can be restarted without needing node reboot, but restarting the app leads to node reboot. [PR1515790](#)

## Infrastructure

- IPv6 ping does not support **do-not-fragment** in Junos OS Evolved Release 20.1R1. [PR1362365](#)
- Rebooting the PTX10003 router during a broadcast storm on the management port might cause a fault on the PTX10003 router. [PR1423216](#)
- Problem: When a TACACS user or RADIUS user logs in through Telnet, the username displayed at the login prompt **show cli authorization** output, **show system users** output, and accounting logs is the template username and not the actual username that is used to log in. This is a display issue where the username displayed is the template username of the logged-in user and there is no functional issue. The problem is specific to Telnet; and when the user logs in through SSH, actual logged-in username is displayed. [PR1482363](#)
- A SIB board that has BIAS voltage failure is not listed in the output of 'show chassis sibs'. Kernel logs indicate the presence of the SIB. Since none of the i2c devices on the SIB are accessible if BIAS voltage has failed, the SIB can not be identified and shows up as 'unsupported' in the Kernel logs. [PR1499398](#)

## Interfaces and Chassis

- `show chassis fpc` output will not show Temperature, CPU Utilizations and Memory for the fixed form chassis's PTX10003 having no physical FPCs. [PR1367671](#)

- For a PTX10003-160C setup with all optics plugged in and all 4 PSMs operational, if somehow we experience a triple fault in the PSM, then two FPCs should go down due to insufficient power. However, in the process, the picd app might be hung. This occurs only when the setup is a scaled condition with all or most of the optics plugged in. To recover from this condition, reboot the router. [PR1460185](#)
- CRC errors on an interface might result in Cmerror and ASIC errors in logs. [PR1499291](#)

## Interfaces and Chassis

- **show chassis fpc** output does not show temperature, CPU Utilizations and memory for the fixed form chassis PTX10003 having no physical FPCs. [PR1367671](#)
- For PTX10003-160C setup with all optics plugged in and all 4 PSMs operational, if somehow we incur a triple fault in PSM - 2 FPCs should go down due to insufficient power. However, in the process of that, the picd app might be hung. This only occurs when the setup is a scaled condition with all or most of the optics plugged in. To recover from this condition, reboot the router. [PR1460185](#)

## Network Management and Monitoring

- mplsL3VpnVrfCreationTime and mplsL3VpnVrfConfLastChanged mib objects will display timestamp as 0 in 20.1R2. [PR1499757](#)
- This issue can potentially happen when someone does ctrl+c when the command is running. Here in this case, it is **show system boot-messages**. When this issue happens, workaround is to find the pids of "journalctl -b" processes and kill them. [PR1513447](#)

## Routing Protocols

- Keychain authentication configuration is not supported on Junos OS Evolved platforms. [PR1490351](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.1R2 | 59](#)
- [Resolved Issues: 20.1R1 | 62](#)

Learn which issues were resolved in Junos OS Evolved Release 20.1R2 for the PTX10003 and PTX10008.

## Resolved Issues: 20.1R2

### *Class of Service (CoS)*

- Recommended FC configuration to aid CBF and Classifier override. [PR1454501](#)
- cosd anomalies are observed after committing the EBGp baseline configuration. [PR1494573](#)

### *Flow-based and Packet-based Processing*

- In 6vPE scenario route records are not reported in exported flow records. [PR1494706](#)

### *General Routing*

- The aggregated Ethernet Queue stats are all 0 for more than 3 minutes (where as rate/pps/bps stats are correct) after the **clear interfaces statistics all** commands and when one AE member link flaps (24x10G ae0). [PR1423134](#)
- The aggregated Ethernet LACP Distributed Mode:Tracking PR for display of FPC in Distribution address in run show ppm adjacencies [PR1458228](#)
- For a Routing Engine-driven local repair, about packet loss is seen for 500 ms with 3000 LSPs. [PR1459265](#)
- Sometimes a dhclient core file is generated during ZTP retry when the dhclient process is exiting. [PR1460906](#)
- "xmlproxyd" binding queue complete-deleted anomalies seen for the type "net::juniper::config::services::agent". [PR1470872](#)
- Block unified ISSU commands from the CLI for Junos OS Evolved Release 19.4R1. [PR1470943](#)
- picd top-level periodic is hogging CPU. [PR1475753](#)
- After OIR feed to the PSM, PSM capacity gets stuck at 0 W capacity with output error in show system alarm. The system does not recover from this failure even after the request system shutdown reboot command is issued. [PR1477975](#)
- Default election priority goes to backup state, when you configure disabled state. [PR1478717](#)
- Routing and Control Board: Power and temperature sensor data is missing in the output of the **show chassis environment cb** command after graceful removal and re-insertion. [PR1478951](#)
- Taking SIBs offline can cause packet forwarding wedges. [PR1479263](#)
- Routing Engine mastership switchover is allowed even when a critical application is not in Online state on the backup Routing Engine. [PR1480076](#)
- hwdre might crash, while FPCs are being inserted into the chassis or taken out of the chassis after performing back-to-back mastership switchovers. [PR1482453](#)
- In StrictPriority mode, with any form of rate-limiter/shaper configured, it may create some accuracy gaps in the scheduler across Queues. [PR1485478](#)

- Scheduler might not work correctly when configuring **transmit-rate exact**. [PR1485486](#)
- For traffic schedulers, excess rate configuration might not show accurate results in a few scenarios. [PR1485730](#)
- Minor alarm to be raised for CPUs running at low frequency. [PR1486559](#)
- The Component(s) field of **show agent sensors** is implementation-defined.
- Junos OS evolved **set system processes app failover** configuration not supported.
- With 80% traffic, single SIB yank (ungraceful OIR) leads to a traffic black-hole and major errors across other SIBs and some FPCs. [PR1489544](#)
- Ping to any neighboring routers fails to work whenever PTX10008 comes up with BGP sessions configured. [PR1489569](#)
- GRPC telemetry is failing for ingress LSPs. [PR1489605](#)
- Routing Engine switchover when fabric links are in transition state is not supported in Junos OS Evolved Release 20.1R2, and can leave the system in an inconsistent state. Such fabric link transitions are triggered by the following events: SIB Offline - SIB Online - FPC Offline - FPC Online. We recommend that you wait till all fabric link transitions are complete before triggering Routing Engine switchover. [PR1490182](#)
- Exp rewrite at egress interface is not working as expected; after rewrite, BE queue is carrying the traffic instead of AF. [PR1490814](#)
- QSFP28-100G-CU3M not detected on fresh boot. [PR1491116](#)
- Traffic statistics are not updated for bypass LSP even though the traffic is flowing through the bypass LSP. [PR1491467](#)
- DCDCEdge-VPNTunnelMulticastL3L2:evo-aftmand-bt generates a core file on FPC0 starting NodeStateRecycle::NotifyReady () --> IflibCtrHandler --> EallffHandler::OnAdd () --> EallffHandler::checkCollision () --> Eallff::makeKey () --> get\_iflIndex (), while adding or deleting CoS interface during the verification of DSCP and DSCPv6 classifiers on ieee802.1ad interface on an aggregated Ethernet interface bundle. [PR1491548](#)
- evoaft-jvisiond core file generated after GRES. [PR1492059](#)
- sflow ingress sampling reports wrong next hop in case IPv6 traffic is routed through lsp at ingress PE. [PR1492076](#)
- Behavior change in clients with multiple gRPC channels to the same target. [PR1492088](#)
- l2ald core seen in vlan member scale configuration is re-applied. [PR1492546](#)
- PTX10008: sflow ingress sampling reports nexthop as 0.0.0.0 in case IPv6 is label swapped through segment routing lsp. [PR1492616](#)
- sflow ingress sampling reports nexthop as 0.0.0.0 in case lrv6 is label swapped through segment routing lsp. [PR1492667](#)

- After flapping the aggregated Ethernet interface through child link flaps every 5 minutes for around 30 times, memory leak is seen. [PR1492667](#)
- After egress PIC offline/online or deactivate/activate triggers, the IPv4 and IPv6 CBF forwarding-table functionality is not working as expected; The AF and NC queues are not transmitting packets from the second egress interface. [PR1492707](#)
- cmddd crashes when the **request system shutdown reboot disk1** command is run. [PR1492955](#)
- For PTX10008 LC1201, Routing Engine queue statistics show a low value for tail-dropped bits per second. [PR1493787](#)
- The rpd process generates a core file @idr\_encode\_rr\_as\_seg\_idl\_0 after rpdagent restart. [PR1493963](#)
- Default firewall filter "host-bound-classification-ipv6-filter" reported to the jvision server. [PR1494467](#)
- Basic IPv6 and IPv4 CBF (Cos-based forwarding) is broken in the latest Junos OS Evolved Release 20.1R1 dailies 20.1-202003090136.0-EVO and 20.1-202003101244.0-EVO. After the CBF action, the packets are not transmitted using correct queues. [PR1495119](#)
- Software upgrade to PTX-fixed image is not blocked on PTX10008. [PR1495655](#)
- PTX10008: Longevity: evo-aftmand-bt generates a core file at JzphrHandleMplsLkupEntry::getFdbMplsEntKey.
- Consistent hash information is missing : **error: timeout communicating with Evo-Aft BT daemon daemon.** [PR1496077](#)
- Aggregated Ethernet logical interface disappear after switchover. [PR1496887](#)
- DCDCEdge-VPNTunnelMulticastL3L2: picd core in net::juniper::picd\_cmd::initialize\_optic\_port\_information\_DataIterator::receive\_request]. [PR1497285](#)
- Traceroute in a MPLS network may not show the complete path. [PR1498949](#)
- Any new unilist NH token does not get updated into the indirect NH token. [PR1501284](#)
- Routes are being installed in the Packet Forwarding Engine even when the interface is down or disabled. [PR1501321](#)
- evo-aftmand-bt core file is generated at FPC4 starting jzphr\_fdb\_action\_data\_get --> jzphr\_fdb\_halfwide\_action\_data\_get --> JexprHandleNh::embedNhBuild --> jprds\_nh\_ext\_format\_get --> \_\_GI\_raise, after deactivate-activate all forwarding-classes on LSP CBF, modify lsp path for cbf and then after configuring non-lsp-next-hop using static LSP for lsp cbf. [PR1502365](#)
- DHCPv6 clients cannot bind after jdncpd restart. [PR1503747](#)
- PTX10008: Each FPC is exporting data for all interfaces. [PR1505521](#)

### **Infrastructure**

- fibd daemon crashes and generates a core file. [PR1487437](#)

### **Interfaces and Chassis**

- IPv4 and IPv6 uRPF test cases failed on aggregated Ethernet interface after restarting ifmand. [PR1453461](#)
- ET interface linklocal self ping is not working in PTX10003-80C and PTX10003-160C routers. [PR1476971](#)
- After application restart - false optics alarms continue to persist although underlying optics does not have the failures. [PR1493230](#)
- cos agg.d.re generates a core file with LAG link protect configuration, @ AELpAttachmentManager::onMemberActive <- AELPManager::setLpActive. [PR1495651](#)
- ifmand crashed at IFAManager::findIfaInSameSubnet. [PR1496361](#)
- Direct Attach Copper (DAC) cables on PTX10003 might experience delay in coming up or may fail to come up. [PR1504622](#)

### **MPLS**

- In scenarios where multiple links in IS-IS are brought down and brought up again, LSPs might not come up. [PR1477780](#)

### **Platform and Infrastructure**

- A specific IPv4 packet can lead to FPC restart. (CVE-2020-1638) [PR1493176](#)

### **User Interface and Configuration**

- Initiate an empty commit from the master Routing Engine when a different Routing Engine is inserted in backup node. [PR1465291](#)
- The aggregated Ethernet bundle is not coming up after VLAN tagging is configured and the configuration is rolled back. [PR1491955](#)

## **Resolved Issues: 20.1R1**

### **Authentication and Access Control**

- TACACS authentication might not fall back to the remote user template for platforms based on Junos OS Evolved releases. [PR1463705](#)

### **Class of Service (CoS)**

- cosd core is seen at CosRewriteRulesFCLossSniplet::populateOperObject when core profile configuration rollback is done to the base configuration. [PR1475637](#)
- COS EXP classifier and rewrite with protocol option **mpls-inet-both-non-vpn** is not working as expected. [PR1479575](#)
- 32 dscp classifier/rewrite functionality does not work after the cosd restart. [PR1479676](#)

## General Routing

- LSP statistics CLI is slow in a scaled scenario during installation time. [PR1416363](#)
- Differences in XML tags for the **show pfe statistics traffic** command. [PR1430821](#)
- Output MTU Error statistics does not increase on **show pfe statistics traffic**. [PR1434860](#)
- **show route forwarding-table** commands are not supported in Junos OS Evolved releases. [PR1438744](#)
- evo-aftmand-zx generates a core file while loading a BGP configuration. [PR1439366](#)
- Operational and configuration CLIs changed. [PR1451455](#)
- Traffic hitting SRTE label or transit route might have multiple primary next hops, leading to silent dropping of packets. [PR1453446](#)
- With 3000 LSPs, about 500 ms of traffic loss is observed for Routing Engine driven local-repair. [PR1459265](#)
- Ports are in disabled state after continuous triggers are given on setup and incomplete objects in app-controller and binding queue. (This happened on MTU change.) [PR1460063](#)
- System software issue where PSM firmware version is not refreshed after hot-plugging an updated PSM into the chassis. [PR1460409](#)
- Sometimes dhclient generates a core file during ZTP retry when the dhclient process is exiting. [PR1460906](#)
- Support information is not available in the **show system errors** CLI hierarchy. [PR1461691](#)
- Policer drop statistics are not updated for the software policer at the Packet Forwarding Engine level. [PR1463997](#)
- UDP feature is not enabled for the Packet Forwarding Engine statistics on PTX10003-80C and PTX10003-160C. [PR1467333](#)
- Ports are in disabled State. Platform binding queue points to incomplete objects in evo-aftmand-zx with "Waiting usr inc : true" in one guid and also it can be recovered only by reboot. [PR1467546](#)
- Block unified ISSU commands from the CLI for Junos OS Evolved Release 19.4R1. [PR1470943](#)
- Ping to an IPv6 logical interface fails after multiple triggers of an aggregated Ethernet member. [PR1472033](#)
- The master-only configuration in the management Ethernet configuration does not work. [PR1473388](#)
- Enhancements to **request system software sync** command. [PR1473572](#)
- All images except current and rollback removed when /var is full on the backup Routing Engine. [PR1475655](#)
- NPU memory/utlization sensor - show agent sensor blank after the interface is deactivated or deleted. [PR1476886](#)

- USB upgrade fails from Junos OS Evolved Release 19.3R2.1 to Junos OS Evolved Release 19.4R1.8. [PR1480371](#)
- Hardware crashes when removing and replacing FPCs after performing back to back Routing Engine mastership switch-overs. [PR1482453](#)

### **Infrastructure**

- Ping does not work when we set record-route. [PR1474649](#)

### **Interfaces and Chassis**

- A 400GE interface might take a longer time (about 50s) to come up. [PR1450606](#)
- picd becomes unresponsive while powering off three PSMs in 16T setup. [PR1460185](#)
- SIB major error seen upon removal of three DC 80-A PSMs. [PR1461179](#)

### **Network Management and Monitoring**

- mib2d generates a core file in jnxHrStorageEntry\_get\_value (lastmatch=< optimized out>, pktp=< optimized out>, vbp=0x7f1978e35a80, object=0x7f181564cd60) at ../../src/storage\_mthd.cpp:357. [PR1475159](#)

### **User Interface and Configuration**

- CMDD application restart issue: command-handler subsystem is not responding. [PR1434409](#)
- Delete **show system ntp source-address** command before upgrading to Junos OS Evolved Release 19.3R1 from Junos OS Evolved Release 19.1R1. [PR1462728](#)
- Initiate an empty commit from the master Routing Engine when a different Routing Engine is inserted in backup node. [PR1465291](#)
- Configd generates a core file after commit, restart, and rollback 1 with a firewall filter configuration if the filter-term has more than 122 **from source-address** statements. [PR1477772](#)

## **Junos OS Evolved Release Notes for QFX5220 Devices**

### **IN THIS SECTION**

- [What's New | 65](#)
- [What's Changed | 71](#)
- [Known Limitations | 72](#)



- [Open Issues | 73](#)
- [Resolved Issues | 74](#)

These release notes accompany Junos OS Evolved Release 20.1R2 for QFX5220-32CD and QFX5220-128C Switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

## What's New

### IN THIS SECTION

- [What's New in Junos OS Evolved Release 20.1R2 | 66](#)
- [What's New in Junos OS Evolved Release 20.1R1 | 66](#)

Learn about new features introduced in Junos OS Evolved Release 20.1R2 for the QFX5220.

## What's New in Junos OS Evolved Release 20.1R2

### *Routing Protocols*

- **Recommendation to include the `local-address` statement when configuring IBGP and multihop EBGP**—When a device peers with a remote device's loopback interface address, use the **`local-address`** statement at the **[`edit protocols bgp group internal-peers`]** hierarchy to specify the source information in BGP update messages. Although a BGP session can be established when only one of the paired routing devices has `local-address` configured, we strongly recommend that you configure **`local-address`** on both paired routing devices for IBGP and multihop EBGP sessions. The **`local-address`** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

[See [local-address \(Protocols BGP\)](#) and [BGP Peering Sessions](#).]

## What's New in Junos OS Evolved Release 20.1R1

### *Forwarding and Sampling*

- **Customizing hashing parameters and shared-buffer alpha values for better load balancing (QFX5220)**—The QFX5220 switches achieve load balancing through the use of a hashing algorithm, which determines how to forward traffic over LAG bundles or to next-hop devices when ECMP is enabled. Starting with Junos OS Evolved Release 20.1R1, you can explicitly configure some hashing parameters to make hashing more efficient. The shared-buffer pool is a global memory space that all ports on the switch share dynamically as they need buffers. The switch uses the shared-buffer pool to absorb traffic bursts after the dedicated-buffer pool is exhausted. The shared-buffer pool threshold is dynamically calculated based on a factor called alpha. You can also specify the alpha, or dynamic threshold, value to determine the change threshold of shared buffer pools for both ingress and egress buffer partitions.

To specify hashing parameters:

```
user@switch# set forwarding-options enhanced-hash-key hash-parameters (ecmp | lag)
```

To specify a threshold value for a particular queue:

```
user@switch# set class-of-service shared-buffer (ingress | egress) buffer-partition buffer
dynamic-threshold value
```

[See [hash-parameters](#) and [buffer-partition](#).]

### *General Routing*

- **View ARP table entries and IPv6 neighbor information (QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, on QFX5220 switches, you can view the following:
  - Address Resolution Protocol (ARP) details such as the next-hop reference count and the next-hop current state, using the commands **`show arp reference-count`** and **`show arp state`**.

- IPv6 neighbor cache information, using the following commands: **show ipv6 neighbors flags**, **show ipv6 neighbors host**, **show ipv6 neighbors interface *interface-name***, **show ipv6 neighbors reference-count**, **show ipv6 neighbors vpn *vpn-name***.

[See [show arp](#) and [show ipv6 neighbors](#)]

- **Support for running applications signed by users (PTX10003, PTX10008, QFX5220)**—Starting with Junos OS Evolved Release 20.1R1, root users can generate signing keys which can be used to sign Linux applications and allow them to run on your device.

Signing keys are stored in the **system-keystore** by using the following CLI command: **request security system-keystore import key-name *key-name* private-key *key-path* x509-cert *certificate-path***. Once a key is stored in the **system-keystore**, it can be used to sign a file with the following CLI command: **request security integrity measure file *file-name* key *key-name***.

[See [Junos OS Evolved Overview](#)]

### **Interfaces and Chassis**

- **Support for dynamic load balancing (QFX5220)**—In Junos OS Evolved Release 20.1R1, QFX5220 switches support dynamic load balancing (DLB) for ECMP. DLB is an enhancement to static load balancing. DLB considers member bandwidth utilization along with packet content for member selection.

You can use the following DLB modes to load-balance traffic:

- Flowlet
- Assigned flow
- Per-packet

To configure DLB for ECMP, include the **ecmp-dlb** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy level.

**NOTE:** You cannot configure both DLB and resilient hashing at the same time.

[See [Understanding Dynamic Load Balancing](#) and [Configuring Dynamic Load Balancing](#).]

### **Junos OS XML API and Scripting**

- **The jcs:load-configuration template supports loading the rescue configuration (PTX Series and QFX Series)**—Starting in Junos OS Evolved Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

[See [Changing the Configuration Using SLAX and XSLT Scripts](#) and [jcs:load-configuration Template](#).]

- **Requests library support in Python scripts (QFX5220-32CD and QFX5220-128C)**—Starting in Junos OS Evolved Release 20.1R1, QFX5220 switches support the Requests library, an HTTP library for Python. Python scripts can import the **requests** module to send HTTP/1.1 requests. The Requests library provides additional methods for supporting initial deployments as well as for performing routine monitoring and configuration changes on devices running Junos OS Evolved.

[See [Using the Requests Library for Python on Devices Running Junos OS.](#)]

### **Multicast**

- **Support for multicast forwarding (QFX5220-32CD and QFX5220-128C)**—Starting with Junos OS Evolved Release 20.1R1, the following multicast features are supported:
  - IPv4 and IPv6 multicast
  - Internet Group Management Protocol (IGMP)
  - Multicast Listener Discovery (MLD)
  - Protocol Independent Multicast source-specific multicast (PIM SSM)
  - Protocol Independent Multicast sparse mode (PIM SM)

Support does not extend to features beyond those listed above. That is, in this release, IGMP snooping, MLD snooping, MVPN Multicast, PIM MoFRR, PIM First Hop router (FHR), Rendezvous point (RP), and Last Hop Router (LHR) are not supported. In addition, IRB interface is not supported as either a source or receiver, and make before break (MBB) is not supported for existing layer 3 Aggregated Ethernet (AE) and link aggregation group (LAG) receivers (member addition/deletions or up/downs).

[See [Multicast Routing Protocols](#) ].

### **Network Management and Monitoring**

- **Support for running a Puppet agent as a Docker container (PTX10003-80C, PTX10003-160C, QFX5220-32CD, and QFX5220-128C)**—Starting in Release 20.1R1, devices running Junos OS Evolved support running a Puppet agent as a Docker container. As an alternative to using the Puppet agent that is integrated into the Junos OS Evolved software image, you can use the Puppet agent Docker container provided by Juniper Networks. Using the container enables you to use standard Docker tools to manage the container and mount or unmount the Puppet agent as needed.

[See [Installing Puppet for Junos OS.](#)]

- **Support for running a Chef client as a Docker container (PTX10003-80C, PTX10003-160C, QFX5220-32CD, and QFX5220-128C)**—Starting in Release 20.1R1, devices running Junos OS Evolved support running a Chef client as a Docker container. As an alternative to using the Chef client that is integrated into the Junos OS Evolved software image, you can use the Chef client Docker container provided by Juniper Networks. Using the container enables you to use standard Docker tools to manage the container and mount or unmount the Chef client as needed.

[See [Deploying Chef for Junos OS.](#)]

### *Routing Policy and Firewall Filters*

- **Support for 2000 egress firewall filters (QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you can configure up to 2000 VLAN firewall filters on the QFX5220 switch. This feature is supported only in the egress direction (traffic exiting the VLAN). To configure, include the **egress-scale** option under the **eracl-profile** statement at the **[edit system packet-forwarding-options firewall]** hierarchy level.

[See [eracl-profile](#) and [Configuring Firewall Filters](#).]

### *Software Installation and Upgrade*

- **Zero touch provisioning using WAN interfaces (QFX5220-32CD)**—Starting in Junos OS Evolved Release 20.1R1, Zero Touch Provisioning (ZTP) uses either WAN interfaces or the management interface to automatically download and install the appropriate software and the configuration file on your device during the bootstrap process.

ZTP starts with the management interface, and if it does not receive the minimal set of DHCP parameters, it will then try any of the interfaces that has WAN links that are up. If none of the interfaces provide a minimal set of DHCP parameters, ZTP retries the process again.

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

- **Support for asynchronous notifications, SIOCETHTOOL ioctl, and multipath routes (QFX Series)**—Junos OS Evolved Release 20.1R1 supports the following features:

- Use the **set system netlink-async-mode** configuration to enable NETLINK\_ROUTE asynchronous notifications. This feature is disabled by default. Use **show nsld mode** to show the current netlink asynchronous mode.
- **SIOCETHTOOL ioctl**, which can be used by other applications.
- Multipath next-hop route information through netlink route attributes.

[See [Running Linux Tools on Junos OS Evolved](#).]

- **Support for curl binary on Junos OS Evolved (QFX Series)**—The curl binary is a command-line utility, used from the shell, that you can use to perform operations over several transport protocols including FILE, FTP, HTTP, and SMTP. The curl binary can improve your workflow by making it much easier to gather information about a device. Junos OS Evolved Release 20.1R1 supports curl version 7.50.1 and libcurl version 7.50.1.

### Spanning Tree Protocols

- **Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocol (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP) support (QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, You can use STP, RSTP, MSTP, and VSTP to provide Layer 2 loop prevention.

[See [Spanning-Tree Protocol Overview](#).]

### System Logging

- **Support for system log notifications over MQTT when configuring JET applications (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you can configure JET applications by using the CLI and then enable the applications to subscribe to system log events. The Junos OS Evolved eventd process sends system log notifications to the event broker. The broker then sends the notifications formatted as JSON messages to the JET application over the Message Queuing Telemetry Transport (MQTT) protocol.

To enable a JET application to subscribe to system log events, configure the **extension-service** statement at the **[edit system extensions]** hierarchy level. You can manage the JET application by using commands at the **[edit system extensions]** and **[edit system services extension-service]** hierarchy levels.

[See [JET Architecture](#), [request extension-service \(start | stop\)](#), [show extension-service status](#), [extensions](#), and [extension-service \(System Services gRPC\)](#).]

### System Management

- **Boundary clock and enterprise profile support for Precision Time Protocol (PTP) (QFX5220 devices)**—Starting in Junos OS Evolved Release 20.1R1, the enterprise profile, which is based on PTPv2, provides the ability for enterprise and financial markets to timestamp on different systems and to handle a range of latency and delays.

The enterprise profile supports the following options:

- IPv4 multicast transport
- Boundary clocks
- 512 downstream slave clocks

You can enable the enterprise profile at the **[edit protocols ptp profile-type]** CLI hierarchy.

[See [Understanding the Precision Time Protocol Enterprise Profile](#).]

- **Precision Time Protocol (PTP) transparent clock (QFX5220)**—Starting in Junos OS Release 20.1R1 EVO, you can use a transparent clock to update the PTP packets with the residence time as the packets pass through the switch. There is no master/slave designation. The switches support end-to-end transparent clocks, which include only the residence time. The transparent clock can update the residence time in a one-step process, which means it sends the timestamps in one packet.

To use a transparent clock, enable the **e2e-transparent** statement at the **[edit protocols ptp]**.

[See [Understanding Transparent Clocks in Precision Time Protocol](#).]

- **Packet Timestamping using PTP Reference Clock (QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, flow-based packets are timestamped with the recovered PTP clock. Packet timestamping is supported on family ethernet-switching, inet and inet6.

[See [Understanding Transparent Clocks in Precision Time Protocol](#).]

## What's Changed

### IN THIS SECTION

- [Changes in CLI Statements and Commands](#) | 71
- [What's Changed in Junos OS Evolved Release 20.1R2](#) | 71
- [What's Changed in Junos OS Evolved Release 20.1R1](#) | 71

Learn about what changed in Junos OS Evolved Release 20.1R2 for the QFX5220.

### Changes in CLI Statements and Commands

- To see the list of CLI statements and commands changed in Junos OS Evolved, see *How Junos OS Evolved Differs from Junos OS* in the [Introducing Junos OS Evolved Guide](#).

### What's Changed in Junos OS Evolved Release 20.1R2

#### General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (PTX10003 PTX10008, and QFX5220)**—Starting with Junos OS Evolved Release 20.1R2, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

[See [commit \(System\)](#).]

### What's Changed in Junos OS Evolved Release 20.1R1

#### Interfaces and Chassis

- **Modified Ethernet interface configuration (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you'll no longer see the **gether-options** statement at the **[edit interfaces**

**interface-name**] hierarchy level. To configure Ethernet interfaces and link aggregation groups (LAGs), use the **set interfaces interface-name ether-options** command instead.

The following options are added to the **ether-options** statement:

- **fec**
- **loopback-remote**

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

### *Junos Telemetry Interface (JTI)*

- **LLDP ON\_CHANGE statistics support with JTI (PTX Series, QFX Series)**—Enhanced telemetry ON\_CHANGE event support provides the following LLDP attributes: - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes. - ON\_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered

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

### *System Logging*

- **Removal of show system queues command (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.1R1, you'll no longer be able to use the **show system queues** command because the information in the command output is not relevant.

## Known Limitations

### IN THIS SECTION

- [Interfaces and Chassis | 72](#)

Learn about limitations in this release for the QFX5220. For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

### Interfaces and Chassis

- The **show interfaces** command to view traffic statistics does not show any traffic statistics for IPv6. [PR1421685](#)



- When PTP is enabled or disabled, the interface flaps momentarily while the PTP state is applied. [PR1487043](#)

## Open Issues

### IN THIS SECTION

- [Routing Policy and Firewall Filters | 73](#)
- [Routing Protocols | 73](#)
- [System Management | 73](#)
- [User Interface and Configuration | 74](#)

Learn about open issues in this release for the QFX5220. For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

### Routing Policy and Firewall Filters

- QFX5220 does not support the following matches which were supported on QFX5200: first-fragment, tcp-established, ls-fragment, and Hop-limit. [PR1499009](#)

### Routing Protocols

- Keychain authentication configuration is not supported on Junos OS Evolved platforms. [PR1490351](#)
- The **virtual-router** option is not supported under **routing-instance** in lean rpd image. [PR1494029](#)

### System Management

- Differences in XML tags for **show system virtual-memory**. [PR1438110](#)
- PTP-FPGA access error displayed as PTP-FPGA is not initialized. [PR1457094](#)
- **switch-options** configuration not available on lean RPD images. [PR1457825](#)
- When the SSH service is configured with certain key-exchange algorithms, the SSH server becomes unresponsive. For example: **set system services ssh key-exchange dh-group1-sha1** [PR1467467](#)

- On QFX5220 platforms with Junos OS Evolved images, when the lo0 firewall filter (**inet/inet6**) is used, Layer 3 forwarding traffic might be discarded by the lo0 filter. [PR1475620](#)
- **show firewall** cli shows error **error: communication failure with /re0/evo-pfemamd/** when unique filter is applied over all physical interfaces. [PR1499294](#)
- ACL is not getting applied to ingress when single filter is applied to both input and output direction of intf simultaneously. [PR1501015](#)
- When changing a physical interface from vlan-id-list to multiple single vlan ids or vlan-id-lists that fall within the original vlan range, we may see one or more of the new physical interfaces not notified. For example, in the current commit we have the following configuration **set interfaces ae100 unit 1 vlan-id-list 1-100**, in the next commit we do the following **set interfaces ae100 unit 1 vlan-id 1 set interfaces ae100 unit 2 vlan-id 2**. We may see that one of the two physical interfaces is not notified. [PR1504374](#)
- DHCPv6 clients may fail to bind after the dhcp-relay configuration is deleted and then quickly re-added a short time later in a given routing instance. [PR1505993](#)

## User Interface and Configuration

- On QFX5220-32CD, VLANs between 3968 to 4095 are reserved for Layer 3 interfaces by default. So, these VLANs cannot be used for Layer 2 interfaces. As of now there is no commit check added for this purpose. You need to take care of this while configuring VLANs for Layer 2. [PR1423468](#)

## Resolved Issues

### IN THIS SECTION

- [Class of Service \(CoS\) | 75](#)
- [Infrastructure | 75](#)
- [Interfaces and Chassis | 75](#)
- [Platform and Infrastructure | 75](#)
- [System Management | 75](#)
- [User Interface and Configuration | 75](#)

Learn which issues were resolved in Junos OS Evolved Release 20.1R2 for the QFX5220.

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

## Class of Service (CoS)

- cosd anomalies are observed after committing the EBGp baseline configuration. [PR1494573](#)

## Infrastructure

- fibd core files can be seen after commit when vlan-id of a vlan is changed. [PR1497790](#)

## Interfaces and Chassis

- ET interface linklocal self ping is not working. [PR1476971](#)

## Platform and Infrastructure

- A specific IPv4 packet can lead to FPC restart. (CVE-2020-1638) [PR1493176](#)

## System Management

- Sometimes dhclient generates a core file during ZTP retry when the dhclient process is exiting. [PR1460906](#)
- Failed to start MGD initialization of schema and database while upgrade from Junos OS Evolved Release 19.3R2.1 to Junos OS Evolved Release 19.4-202002240046.0 ] >>> mgd-init service fails due to clksyncd failure. [PR1491803](#)
- cmdd crashes when the invoked **request system shutdown reboot disk1** command is run. [PR1492955](#)
- Interface is not part of the member VLAN when configured. [PR1495207](#)
- QFX5220-32 and QFX5220-128C: evo-pfemamd core file is seen if a loopback filter is applied with **input-list** and after the switch is rebooted. [PR1497909](#)

## User Interface and Configuration

- Observed error **The command-handler subsystem is not responding to management requests** while upgrading from Junos OS Evolved 19.3R2.1 to Junos OS Evolved 19.4R2.3. [PR1499018](#)

# Upgrade your Junos OS Evolved Software

Products impacted: PTX10003, PTX10008, QFX5220-32CD, QFX5220-128C.

Follow these steps to upgrade your Junos OS Evolved software:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:  
<https://www.juniper.net/support/downloads/>
2. In the **Find a Product** box, enter the Junos OS platform for the software that you want to download.
3. Select **Junos 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.

**NOTE:** We don't recommend that you download the Services Profile 1 image to use the lean rpd profile. We will deprecate this image in Junos OS Evolved 21.4R1. For more information about the types of Junos OS Evolved installation package prefixes, see [Junos OS Evolved Installation Packages](#).

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

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

For more information about software installation and upgrade, see [Software Installation and Upgrade Overview \(Junos OS Evolved\)](#).

For more information about EOL releases and to review a list of EOL releases, see <https://support.juniper.net/support/eol/software/junosevo/>.

## Finding More Information

Learn about more information on Junos OS Evolved and other Juniper products.

- Feature Explorer—The Juniper Networks Feature Explorer is a Web-based app that helps you to explore and compare Junos OS and Junos OS Evolved feature information to find the right software release and hardware platform for your network. <https://pathfinder.juniper.net/feature-explorer/>
- PR Search Tool—Keep track of the latest and additional information about Junos OS Evolved open defects and issues resolved. [prsearch.juniper.net](https://prsearch.juniper.net)
- Hardware Compatibility Tool—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](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 for Juniper Networks products. [apps.juniper.net/compliance/](https://apps.juniper.net/compliance/)

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.

- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

# Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service 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 <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

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

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

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



## Creating a Service Request with JTAC

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

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

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

## Revision History

7 October 2021—Revision 10, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

16 July 2021—Revision 9, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

22 April 2021—Revision 8, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

13 January 2021—Revision 7, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

30 July 2020—Revision 6, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

16 July 2020—Revision 5, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

2 July 2020—Revision 4, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

26 June 2020—Revision 3, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

18 June 2020—Revision 2, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

5 June 2020—Revision 1, Junos OS Evolved Release 20.1R2 for the PTX10003, PTX10008, and QFX5220 Devices.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.