

Release Notes: Junos[®] OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices

12 JaMay 2022

Contents	Introduction 3
	Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, and PTX10008 Devices 3
	What's New 4
	What's New in Junos OS Evolved Release 20.2R2 4
	What's New in Junos OS Evolved Release 20.2R1 4
	What's Changed 30
	Changes in CLI Statements and Commands 30
	What's Changed in Junos OS Evolved Release 20.2R2 30
	What's Changed in Junos OS Evolved Release 20.2R1 31
	Known Limitations 32
	General Routing 33
	Interfaces and Chassis 34
	MPLS 34
	User Interface and Configuration 34
	Open Issues 34
	General Routing 35
	Class of Service (CoS) 37
	Infrastructure 38
	Interfaces and Chassis 38
	MPLS 38
	Routing Protocols 38

User Interface and Configuration	38
Resolved Issues	38
Resolved Issues: 20.2R2	39
Resolved Issues: 20.2R1	42
Junos OS Evolved Release Notes for QFX5220 Devices	48
What's New	49
What's New in Junos OS Evolved Release 20.2R2	49
What's New in Junos OS Evolved Release 20.2R1	49
What's Changed	50
Changes in CLI Statements and Commands	51
What's Changed in Junos OS Evolved Release 20.2R2	51
What's Changed in Junos OS Evolved Release 20.2R1	52
Known Limitations	53
Interfaces and Chassis	54
Open Issues	54
Class of Service (CoS)	55
Routing Policy and Firewall Filters	55
Routing Protocols	55
System Management	55
User Interface and Configuration	56
Resolved Issues	56
Resolved Issues: 20.2R2	56
Resolved Issues: 20.2R1	57
Upgrade Your Junos OS Evolved Software	59
Finding More Information	61
Documentation Feedback	61
Requesting Technical Support	63
Self-Help Online Tools and Resources	63
Creating a Service Request with JTAC	64
Revision History	64

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. But 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 20.2R2.

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](#).

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, and PTX10008 Devices

IN THIS SECTION

- [What's New | 4](#)
- [What's Changed | 30](#)
- [Known Limitations | 32](#)
- [Open Issues | 34](#)
- [Resolved Issues | 38](#)

These release notes accompany Junos OS Evolved Release 20.2R2 for PTX10001-36MR, 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 20.2R2 | 4](#)
- [What's New in Junos OS Evolved Release 20.2R1 | 4](#)

Learn about new features introduced in Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, and PTX10008.

What's New in Junos OS Evolved Release 20.2R2

IN THIS SECTION

- [High Availability and Resiliency | 4](#)

Learn about new features introduced in Junos OS Evolved Release 20.2R2 for PTX10001-36MR, PTX10003, and PTX10008.

High Availability and Resiliency

- **Platform resiliency (PTX10008)**—Starting in Junos OS Evolved Release 20.2R2, we provide support for platform resiliency. The resiliency includes handling of failures and faults related to hardware components such as line cards, switch fabric, control boards, fan trays, fan tray controllers, and power supply units. Fault handling enables the router to detect and log errors, raise alarms, send SNMP traps, provide indication about the error through LEDs, perform self-healing, and take components out of service.

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

What's New in Junos OS Evolved Release 20.2R1

IN THIS SECTION

- [Hardware | 5](#)
- [Class of Service \(CoS\) | 12](#)

●	Flow-Based and Packet-Based Processing 12
●	High Availability (HA) and Resiliency 13
●	Interfaces and Chassis 14
●	IPv4 16
●	Junos Telemetry Interface 16
●	Layer 2 Features 18
●	MPLS 19
●	Multicast 22
●	Network Management and Monitoring 22
●	Port Security 23
●	Routing Policy and Firewall Filters 23
●	Routing Protocols 25
●	Services Applications 28
●	System Logging 28
●	System Management 28
●	VPNs 29

Learn about new features introduced in Junos OS Evolved Release 20.2R1 for PTX10001-36MR, PTX10003, and PTX10008.

Hardware

- **New PTX10001-36MR Packet Transport Router (PTX series)**—In Junos OS Evolved Release 20.2R1, we introduce the PTX10001-36MR. The PTX10001-36MR is a fixed-configuration router with 36 multirate ports that provide 100GbE and 400GbE ports in a low-profile 1-U form factor. The PTX10001-36MR network ports support:
 - Twelve 100GbE ports that support data rates of 10-Gbps, 25-Gbps, 40-Gbps, and 100-Gbps
 - Twenty-four 400GbE ports that support data rates of 10-Gbps, 25-Gbps, 40-Gbps, 100-Gbps, and 400-Gbps

On the PTX10001-36MR, you can configure the speed at port level only. [See [Port Speed](#).]

To install the PTX10001-36MR router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the [PTX10001-36MR Fixed Packet Transport Router Hardware Guide](#).

[Table 1 on page 6](#) summarizes the PTX10001-36MR features supported in Junos OS Evolved Release 20.2R1.

Table 1: Features Supported by the PTX10001-36MR

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> Standard CoS feature support, including classifiers (behavior aggregate [BA], fixed, and multifield [MF]), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, and drop-profiles. [See CoS Features and Limitations on PTX Series Routers.] Support for classification override configured under a forwarding policy. [See CoS Features and Limitations on PTX Series Routers and Overriding the Input Classification.]
DHCP	<ul style="list-style-type: none"> DHCPv4 and DHCPv6 stateless relay. DHCP relay agents are used to forward DHCP requests and replies between clients and servers when they are not on the same physical subnet. Stateless DHCPv4 and DHCPv6 relay includes support for option-82 (DHCPv4), option-18 and option-37 (DHCPv6), DHCPv4/v6 stateless relay over Layer 3 interfaces, and virtual router aware DHCPv4/v6 stateless relay. [See Extended DHCP Relay Agent Overview.]
Firewalls and policers	<ul style="list-style-type: none"> IPv4 and IPv6 support for firewall filters. [See Firewall Filter Match Conditions and Actions (PTX Series Routers).] GRE tunnel de-encapsulation using flexible tunnel interfaces (FTIs). [See Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling.]

Table 1: Features Supported by the PTX10001-36MR (*continued*)

Feature	Description
High availability (HA) and resiliency	<ul style="list-style-type: none"> Platform resiliency to handle faults and failures related to the switch fabric. The PTX10001-36MR provides logical instances of the switch fabric because it is a fixed-configuration router. With resiliency enabled, the router is able to detect and log the switch fabric errors, raise alarms, send SNMP traps, provide indication about the error through LEDs, self-heal, and take the component out of service. [See show system errors active.] Resiliency support on the Routing Engine. The PTX10001-36MR is a fixed configuration router and does not support a pluggable Routing Engine. The Routing Engine supports fault-handling actions such as logging errors, raising alarms, sending SNMP traps, and providing indication about the errors through LEDs. [See show system errors active.] Resiliency support on the Packet Forwarding Engine. The software detects, reports, and takes action on Packet Forwarding Engine faults. Actions are taken based on default configuration or user configuration available for the errors. [See show chassis alarms and show system firmware.] Fabric management support. The PTX10001-36MR does not have a pluggable line card or switch fabric because the PTX10001-36MR is a fixed configuration router. The router provides logical instances for the components that include one FPC, six Packet Forwarding Engines, and one switch fabric. Each packet forwarding engine has 32 fabric links to the switch fabric, each running at a speed of 56.25 Gbps. From a CLI perspective, each link is viewed as a plane (as shown in the show chassis fabric sibs output). The PTX10001-36MR router provides an overall bandwidth of 9.6 Tbps (1.6 Tbps per Packet Forwarding Engine) in non-oversubscribed mode and 10.8 Tbps in oversubscribed mode. The PTX10001-36MR router does not support fabric redundancy, per plane offline sequence, or the per Packet Forwarding Engine offline operation. If fabric link errors, such as the link training error occur, the link is disabled. The router does not support local repair or per link recovery. [See show chassis fabric sibs.]

Table 1: Features Supported by the PTX10001-36MR (*continued*)

Feature	Description
Interfaces and chassis	<ul style="list-style-type: none"> • Environmental monitoring, field-replaceable unit (FRU) management, and support for overriding the maximum power output value of a power supply module (PSM). [See show chassis hardware and max-power.] • Support for an inbuilt routing engine (model number RE-JNP10001-36MR). The PTX10001-36MR does not support a pluggable routing engine. The PTX10001-36MR does not support graceful Routing Engine switchover (GRES), because it does not have a redundant Routing Engine. The PTX10001-36MR router does not support these operational commands: request chassis routing-engine master acquire and request chassis routing-engine master release. [See show chassis hardware.]
Junos OS XML and scripting	<ul style="list-style-type: none"> • Support for XSLT, SLAX, and Python commit, event, op, and SNMP scripts. Python automation scripts enable you to take advantage of Python features and libraries as well as leverage Junos PyEZ APIs to perform operational and configuration tasks on the devices. [See Junos Automation Scripts Overview.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • Support for Packet Forwarding Engine statistics and sensor support for label-switched path (LSP statistics). [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).] • Support for Neighbor Discovery Protocol (NDP) statistics. You can stream telemetry data from the PTX10001-36MR router to an outside collector using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Sensors export statistics for NDP router-advertisement and NDP table state. [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]
Layer 2 features	<ul style="list-style-type: none"> • Support for LLDP including management interfaces. [See LLDP Overview.]
Layer 3 features	<ul style="list-style-type: none"> • Support for Layer 3 forwarding features: unicast routing or forwarding BFD, BFD-triggered local repair (distributed BFD), and 128-way ECMP. [See Layer 3 VPNs User Guide for Routing Devices.] • MPLS-based Layer 3 VPNs, IPv6 VPN Provider Edge routers (6VPE) and IPv6 Provider Edge routers (6PE). [See Layer 3 VPNs User Guide for Routing Devices and Example: Configuring an IPv6-over-IPv4 Tunnel.]

Table 1: Features Supported by the PTX10001-36MR (continued)

Feature	Description
MPLS	

Table 1: Features Supported by the PTX10001-36MR (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • CLI support for monitoring MPLS label usage • Inline MPLS and IPv6 lookup for explicit null • 128,000 transit LSPs • Explicit null support for MPLS LSPs • MPLS Label Block Configuration • MPLS over untagged Layer 3 interfaces • MPLS OAM - LSP ping • JTI: OCST: MPLS operational state streaming (v2.2.0) • 2000 ingress LSP support • 2000 egress LSP support • Entropy Label support • MPLS: JTI: Junos Telemetry Interface MPLS self-ping, TE++, and misc augmentation • LDP, including: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Bypass LSP static configuration • Ingress LSP statistics in a file • RSVP-TE Hitless-MBB with no artificial delays • 128,000 transit LSPs • Auto bandwidth • Class based forwarding (CBF) with 16 classes • CBF with next-hop resolution • Convergence and scalability • Graceful restart signalling

Table 1: Features Supported by the PTX10001-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • TE++: Dynamic ingress LSP splitting • Traffic engineering extensions (OSPF-TE and ISIS-TE) • Traffic engineering options: bgp, bgp-igp, bgp-igp-both-ribs, mpls-forwarding • no-decrement-ttl, and no-propagate-ttl statements <p>[See MPLS Applications User Guide.]</p>
Network management and monitoring	<ul style="list-style-type: none"> • sFlow technology support. [See Understanding How to Use sFlow Technology for Network Monitoring.] • Support for Open Routing (OpenR), a third-party application that runs on platforms running Junos OS Evolved. • Support for the dedicated virtual routing and forwarding (VRF) instance mgmt_junos for the management Ethernet Interface. [See Management Interface in a Nondefault Instance and management-instance.] • Support for adding custom YANG data models to the Junos OS schema. [See Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS.] • Support for Puppet as a third-party application. The Junos OS Evolved software image provides Puppet as part of the image and also supports running the Puppet agent as a Docker container. [See Puppet for Junos OS Administration Guide.] • Support for HealthBot an automated and programmable device-level diagnostics and network analytics tool. [See Healthbot.]

Table 1: Features Supported by the PTX10001-36MR (*continued*)

Feature	Description
Services applications	<ul style="list-style-type: none"> Support for active flow monitoring. Both IPFIX and V9 template formats are supported using an Output filter. [See Configuring Inline Active Flow Monitoring on PTX Series Routers.]
Software installation and upgrade	<ul style="list-style-type: none"> Support for zero-touch provisioning (ZTP) on the management interface. ZTP automates the provisioning of the device configuration and software upgrade over the management interface of the Routing Engine (RE0). [See Zero Touch Provisioning.] Support for ZTP on WAN interfaces. You can use either WAN interfaces or management interfaces to automatically download and install the appropriate software and the configuration file on your device during the bootstrap process. [See Zero Touch Provisioning Overview.]

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported across all platforms, see the [Hardware Compatibility Tool](#).

- **PTX10K-LC1201-36CD Line Card**—Starting with Junos OS Evolved Release 20.2R1, the PTX10K-LC1201-36CD line card now supports the Mellanox 10-Gbps pluggable adapter (model number: MAM1Q00A-QSA). You can use the QSA adapter to convert quad-lane-based ports to a single-lane-based SFP+ port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ module. You can plug a 10-Gbps SFP+ transceiver into the QSA adapter, which is then inserted into a QSFP56-DD port of the PTX10K-LC1201-36CD line card.

[See [The Hardware Compatibility Tool: Adapters.](#)]

Class of Service (CoS)

- **Layer 2 CoS support (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, we've extended support for Layer 2 (L2) class of service (CoS) to PTX10003 devices. You can create default and user-defined L2 classifiers and rewrite rules (IEEE 802.1 and IEEE 802.1ad) and apply them to **bridge**, **ethernet-switching**, and **vpls** logical interface families.

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

- **Classification override support (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, 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 and Packet-Based Processing

- **Load balancing under enhanced-hash-key (PTX10008)**—In Junos OS Evolved Release 20.2R1, we've introduced load balancing under the **[edit forwarding-options enhanced-hash-key]** hierarchy. The bit length of hash material is limited only 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 pseudowire traffic
- Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (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

[See [enhanced-hash-key](#).]

High Availability (HA) and Resiliency

- **Fabric resiliency (PTX10008)**—PTX10008 routers with Junos OS Evolved Release 20.2R1 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

- **Disable power budgeting (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can disable the power budgeting on the PTX10008 router, by using the **set chassis no-power-budget** command. If you disable the power budgeting, the system powers on all the field-replaceable units (FRUs) by default. The system does not take any FRU offline in case of insufficient power. However, in case of a power shortage, a power redundancy alarm is raised.

[See [no-power-budget](#).]

- **Monitoring, detecting threats, and taking action on degraded physical Ethernet links to minimize packet loss (PTX10003, PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can monitor physical link degradation (indicated by bit error rate (BER) threshold levels) on 10-Gigabit, 25-Gigabit, 40-Gigabit, 50-Gigabit, 100-Gigabit, and 400-Gigabit Ethernet interfaces. You can then take corrective actions if the BER threshold value drops below the configured threshold value.

We've introduced configurations at the **[edit interfaces interface-name]** hierarchy level to support the link monitoring and recovery feature on Junos OS Evolved. You can enable or disable link monitoring by using the **link-degrade-monitor-enable** or **no-link-degrade-monitor-enable** statement, respectively. You can view the link recovery status and the BER threshold values by using the **show interfaces interface-name** command.

[See [link-degrade-monitor](#).]

- **Fabric link autoheal (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, the PTX10008 router supports fabric link autoheal, a mechanism that attempts to recover faulty fabric links from a link error condition. Autoheal, attempted for both runtime and initialization time failures, involves bringing down the faulty fabric link and then training it. Junos OS Evolved attempts to recover a faulty link from a maximum of three link error instances (per link) within a span of 24 hours. The autoheal feature is enabled

by default. You can use the existing **show chassis fabric errors autoheal** command to view the details of the autoheal actions performed by the software.

NOTE: If you remove a line card or switch fabric card ungracefully and the FRU on the other side of the link reports a link fault, the software attempts to automatically heal the faulty link. Subsequently, it marks the autoheal status as unsuccessful.

[See [show chassis fabric errors autoheal](#).]

- **Support for DCU accounting and SCU accounting (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, the PTX10008 support destination class usage (DCU) and source class usage (SCU) accounting. SCU accounting provides breakdown of output interface traffic statistics for each prefix from which the traffic originated. DCU accounting provides breakdown of input interface traffic statistics for each prefix to which the traffic is destined.

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

- **Support for QSFP28 100GbE DWDM transceivers (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, 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 port profile configuration on JNP10K-LC1201 line cards using interfaces hierarchy (PTX10008 routers)**—Starting in Junos OS Evolved Release 20.2R1, 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 and 20.2R1, 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

- **Support for gRPC and gNMI services with JTI for SPRING sensors (PTX10003)**—Junos OS Evolved Release 20.2R1 introduces the following additional sensors for Source Packet Routing in Networking (SPRING), also known as segment routing:
 - `/junos/services/segment-routing/interface/ingress/usage/`
 - `/junos/services/segment-routing/interface/egress/usage/`
 - `/junos/services/segment-routing/sid/usage/`

SPRING statistics are exported to an outside collector using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services.

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

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

- **UDP (native) Packet Forwarding Engine sensor support for JTI (PTX10008)**—Junos OS Evolved Release 20.2R1 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 \(Junos Telemetry Interface\)](#).]

- **JTI support for physical and logical interface sensors with gRPC services, gNMI services, and UDP (PTX10008)**—Junos OS Evolved 20.2R1 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 \(Junos Telemetry Interface\)](#).]

- **gRPC dial-out support on JTI (PTX Series)**—Junos OS Evolved Release 20.2R1 provides remote procedure call (gRPC) dial-out support for telemetry. In this method, the target device (server) initiates a gRPC session with the collector (client) and, when the session is established, streams the telemetry data that is specified by the sensor-group subscription to the collector. This is in contrast to the gRPC network management interface (gNMI) dial-in method, in which the collector initiates a connection to the target device.

gRPC dial-out provides several benefits as compared to gRPC dial-in, including simplifying access to the target advice and reducing the exposure of target devices to threats outside of their topology.

To enable export of statistics, include the **export-profile** and **sensor** statements at the **[edit services analytics]** hierarchy level. The export profile must include the reporting rate, the transport service (for example, gRPC), and the format (for example, gbp-gnmi). The sensor configuration must include the name of the collector (the server's name), the name of the export profile, and the resource path. An example of a resource path is **/interfaces/interface[name='fxp0']**.

[See [Using gRPC Dial-Out for Secure Telemetry Collection](#).]

- **Support for BGP routing information base sensor on JTI (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, Junos telemetry interface (JTI) supports streaming BGP routing information base (RIB) statistics to an outside collector using remote procedure call (gRPC).

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

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (PTX10003)**—Junos OS Evolved Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

Layer 2 Features

- **Support for IRB (PTX10003)**—Integrated routing and bridging (IRB) enables routing of Layer 3 traffic between a bridge domain and another routed interface. To set up your IRB interface, use either enterprise-style configuration or service-provider-style configuration. (For a discussion of these two configuration styles, see [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port](#).) You can troubleshoot with the CLI operational mode command **show interfaces irb**.

Supported items for IRB include:

- All Layer 2 protocols already supported on PTX10003
- Layer 3 protocols: BGP, IGMP, IS-IS, OSPF, PIM, and RIP
- Per-IRB logical interface MAC and statistics
- IRB Layer 3 multicast support with flooding only
- Address family support for IPv4 and IPv6, and support for IPv4 MTUs and IPv6 MTUs with different MTU values
- IRB interface in VRF routing instances
- Directed subnet broadcast support with IRB

The following features are *not* supported in this release:

- VPLS
- VXLAN or EVPN
- Tunnel interfaces as Layer 2 interfaces

- Pseudo-interfaces other than aggregated Ethernet interfaces as Layer 2 logical interfaces
- Management VLAN functionality
- IRB on private VLANs and IRB over MPLS-based cores

[See [Integrated Routing and Bridging](#).]

- **GRE tunnel decapsulation using flexible tunnel interfaces (FTI) (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can configure GRE tunnels on flexible tunnel interfaces (FTIs). When you enable the **tunnel-termination** statement at the **[edit interfaces fti0 unit unit-number]** 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.

- **Proxy ARP (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, 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, supported on et and ae interfaces, and supported on active routes.

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](#).]

MPLS

- **6PE, 6VPE, and virtual router routing instance support (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, 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 20.2R1, 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 segment identifier (SID) for segment routing–traffic engineering (SR-TE)
 - Binding SID for SR-TE [*draft-previdi-idr-segment-routing-te-policy*]
 - Programmable routing protocol process APIs for SR-TE policy provisioning
 - Static SR-TE policy with mandatory color specification
 - Static SR-TE policy without color specification
 - IS-IS:
 - Adjacency SID
 - Advertising maximum link bandwidth and administrative color without RSVP-TE configuration
 - Anycast and prefix SIDs
 - Configurable segment routing global block (SRGB)
 - Node and link SIDs
 - Segment Routing Mapping Server (SRMS) and client
 - Topology-independent loop-free alternate (TI-LFA):
 - 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 \(SPRING\)](#).]

- **Routing Engine-based S-BFD for segment-routing traffic engineering (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can run Routing Engine-based seamless BFD (S-BFD) over non-colored and colored label-switched paths (LSPs) with first-hop label resolution and use S-BFD as a fast mechanism to detect path failures.

First hop is the first destination of the packets that are being sent to.

[See [Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution](#).]

- **Support for MPLS ping and traceroute for segment routing (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, we extend the MPLS ping and traceroute support for all of types segment routing–traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

We also support the following features:

- FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
- ECMP traceroute support for all types of SR-TE paths.

We do not support the following:

- Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
- OAM for IPv6 prefix.
- BFD

[See [traceroute mpls segment-routing spring-te](#) and [ping mpls segment routing spring-te](#).]

Multicast

- **Support for multicast-related features (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, PTX10008 devices support IPv4 and IPv6 multicast including IGMP, MLD, PIM-SM, and PIM-SSM. Rendezvous point (RP), and first hop and last-hop router are not supported.

[See [Multicast Configuration Overview](#).]

Network Management and Monitoring

- **Local port mirroring support (PTX10008)**—Starting with Junos OS Evolved Release 20.2R1, 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 at both ingress and egress
- Sampling range of 1-16000000
- 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 and Analyzers](#).]

Port Security

- **MACsec support (PTX10008)**—The PTX10008 supports Media Access Control Security (MACsec) on all interfaces, including 400GbE interfaces. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. PTX10008 support for MACsec includes AES-256 encryption, extended packet numbering, and fail-open mode.

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

Routing Policy and Firewall Filters

- **Firewall filters support (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, the PTX10008 supports the following firewall features:
 - Forwarding IPv4 and IPv6
 - Firewall filter
 - Load balancing

- MPLS fast reroute
- Host path
- Egress peer engineering

Non-replication mode, 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 interface and class-usage accounting profiles (PTX10008)**—Starting in Junos OS Release 20.2R1, the PTX10008 supports both **interface-profile** and **class-usage** accounting profiles.

You can use accounting profiles to collect persistent (locally saved) traffic statistics as an alternative to SNMP. Data collected through **interface-profile** can include input and output byte and packet counts for various protocols, interface errors, and reverse-path-forwarding (RPF) checks. Data collected through source and destination **class-usage** profiles can include per-interface traffic counts for the prefixes specified, for example to track traffic received on a per-customer basis for billing.

Configure accounting profiles at the **[edit accounting-options]** hierarchy level.

[See [Understanding Accounting Profiles](#).]

- **Support for nested filters (PTX10003)**—Starting in Junos Evolved OS Release 20.2R1, you can reference a common firewall filter by attaching it to multiple firewall policies (a filter being one or more match conditions and corresponding actions). Nested filters can, in turn, be bound to the following interface types: **inet**, **inet6**, **mpls**, and **ccc**. They can also be bound to routing instances, and in the ingress direction, the egress direction, or both directions.

Configure the filter at the **[edit firewall family *family-name* filter *filter-name* term *term-name* filter *filter-name*]** hierarchy level.

[See [Guidelines for Nesting References to Multiple Firewall Filters](#) and [Example: Nesting References to Multiple Firewall Filters](#).]

- **Support for matching ip-options in IPv4 packet headers (PTX10003 and PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, both the PTX10003 and PTX10008 support 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](#).]

- **Support for additional route filter qualifiers in a policy statement (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, the PTX10003 supports the following list-level qualifiers: **exact**, **longer**, **orlonger**, **prefix-length-range**, and **upto**.

You can use route filter lists to group individual route filters created at the **[edit policy-options]** hierarchy level. Each item in a list consists of a complete route filter statement, including a destination prefix, a match type, and an optional action. Reuse the list in different policies, adding whatever qualifiers you need, instead of re-creating a different one for every use case.

[See [Understanding Route Filters for Use in Routing Policy Match Conditions](#).]

- **Support for the prefix-list match condition with apply-path option (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, you can use the **apply-path** statement to simplify the maintenance of group prefix lists used in policies and firewalls. When used, **apply-path** expands the scope of a given prefix list to include all prefixes based on the specified path. Enable the **apply-path** statement at the **[edit policy-options prefix-list name]** hierarchy level.

[See [apply-path](#).]

Routing Protocols

- **Support for RIP (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, PTX10008 routers support:
 - RIP version 1
 - RIP version 2
 - RIP next generation (RIPng)

[See [RIP and RIPng Overview](#).]

- **Support for BGP- LU FRR MPLS inter-AS link protection (PTX10008)**—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 20.2R1, 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 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 (JTI) 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](#).]

- **Support for BGP flowspec redirect to IP (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, we support BGP flow specification as described in BGP Flow-Spec Internet draft draft-ietf-idr-flowspec-redirect-ip-02.txt, *Redirect to IP Action* on PTX10003 devices. Redirect to IP action uses extended BGP community to provide traffic filtering options for DDoS mitigation in service provider networks. Legacy flow specification, as specified in the Internet draft draft-ietf-idr-flowspec-redirect-ip-00.txt, *BGP Flow-Spec Extended Community for Traffic Redirect to IP Next Hop*, redirect to IP uses the BGP next-hop attribute to support interoperability of devices. Junos OS advertises the redirect to IP flow specification action using the extended community by default. The redirect to IP action allows you to divert matching flow specification traffic to a globally reachable address. You need this feature to support service chaining in a virtual Service Control Gateway (vSCG).

To configure a static IPv4 or IPv6 flow specification route, include the **redirect *ipv4-address*** or the **redirect *ipv6-address*** statement at the **[edit routing-options flow route then]** hierarchy level.

To configure legacy flow specification, include the **legacy-redirect-ip-action** statement at the **[edit group *bgp-group* neighbor *bgp neighbor* family inet flow]** hierarchy level.

To configure BGP to use VRF.inet.0 table to resolve VRF flow specification routes, include the **secondary-independent-resolution** statement at the **[edit protocols bgp neighbor family flow]** hierarchy level.

[See [legacy-redirect-ip-action](#) and [Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic](#).]

- **Support for dynamic peer AS range for BGP groups (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, you can configure acceptable autonomous system (AS) ranges for external BGP (EBGP) groups that can be used for bringing up BGP peers while establishing a BGP session. BGP accepts a peer request based on the configured AS range and rejects a peer request if the AS does not fall into the specified range. This allows you to control BGP peering when the neighbor's exact IP address is not known.

To define the peer AS range for BGP groups through policy, you can include the **as-list** statement at the **[edit policy-options]** hierarchy level. To include the specified peer AS list, include the **peer-as-list** **peer-as-list** statement at the **[edit protocols bgp group *group-name*]** hierarchy level.

[See [peer-as-list](#) and [as-list](#).]

- **TI-LFA SRLG protection for IS-IS (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, you can configure Shared Risk Link Group (SRLG) protection for segment routing to choose a fast reroute path that does not include SRLG links in the topology-independent loop-free alternate (TI-LFA) backup paths. This is in addition to existing fast reroute options such as **link-protection**, **node protection**, and **fate-sharing protection** for segment routing. IS-IS computes the fast reroute path that is aligned with the post-convergence path and excludes the SRLG of the protected link. All local and remote links that are from the same SRLG as the protected link are excluded from the TI-LFA back up path. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface.

To enable TI-LFA SRLG protection with segment routing for IS-IS, include the **srlg-protection** statement at the **[edit protocols isis interface *name* level *number* post-convergence-lfa]** hierarchy level.

[See [Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS](#).]

- **Support for AIGP metric to MED translation (PTX10003)**—Starting in Release 20.2R1, Junos OS Evolved supports the translation of AIGP metric to MED. You can enable this feature when you want the end to end effective AIGP metric in order to choose the best path. Effective AIGP is the AIGP value advertised with the route plus the IGP cost to reach the next hop. This is especially useful in Inter-AS MPLS VPNs solution, where customer sites are connected via two different service providers, and customer edge routers want to take IGP metric based decision. You can configure a minimum-aigp to prevent unnecessary updation of route when effective-aigp changes past the previously known lowest value.

The following configuration statements are introduced at the **[edit protocols bgp group *<group-name>* metric-out]** hierarchy level:

- **effective-aigp** to track the effective AIGP metric
- **minimum-effective-aigp** to track the minimum effective AIGP metric.

[See [effective-aigp](#) and [minimum-effective-aigp](#).]

Services Applications

- **Support for active flow monitoring (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can use the export capabilities of inline active flow monitoring (inline J-Flow) to define a flow record template suitable for IPv4 or IPv6 traffic. Both IPFIX and V9 template formats are supported with output filters.

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

- **Support for FTI-based de-encapsulation of IPv4 and IPv6 packets (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can configure IP-IP de-encapsulation 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 during packet processing, thus providing performance improvement over the filter-based process, which requires a recirculation path to support de-encapsulation.

[See [Flexible Tunnel Interfaces Overview](#).]

System Logging

- **Support to track the maximum number of routing and forwarding (RIB/FIB) routes and VRFs (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, you can track and display the high-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system (RPD) using the **show route summary** CLI command. High-water mark refers to the maximum number of routing and forwarding (RIB/FIB) table routes and VRFs that was present in the RPD system. The high-water mark data can also be viewed in the syslog at the **LOG_NOTICE** level.

You can configure the interval of the high-water mark data using the **highwatermark-log-interval** CLI configuration statement at the `[edit routing-options]` hierarchy level. The minimum time gap at which the high-water mark data logged in the syslog is 30 seconds. You can configure the value for **highwatermark-log-interval** CLI configuration statement between 5 to 1200 seconds.

[See [routing-options](#) and [show route summary](#).]

System Management

- **Support for passive monitoring (PTX10008)**—Starting with Junos OS Evolved Release 20.2R1, 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](#).]

- **Support for the NTP restrict command (PTX10003)**—Starting in Junos OS Evolved Release 20.2R1, you can use the **nouery** option with the **set system ntp restrict** option to restrict incoming ntpq queries, which can be used for amplification attacks.

[See [Configuring NTP Access Restrictions for a Specific Address \(SRC CLI\)](#).]

VPNs

- **Layer 3 VPN on PTX10008**—Starting in Junos OS Evolved Release 20.2R1, you can configure MPLS-based Layer 3 VPNs on PTX10008 routers. The supported features include:

- MPLS over Layer 3 VLAN-tagged subinterfaces
- Per-next-hop label allocation
- Mapping of the label-switched interface (LSI) logical interface label to the VPN routing and forwarding (VRF) routing table using the **vrf-table-label** statement
- ICMP tunneling and MPLS traceroute
- Disabling time-to-live (TTL) decrementing using **no-propagate-ttl**

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

- **Layer 2 VPN support (PTX10008)**—Starting in Junos OS Evolved Release 20.2R1, you can configure the following Layer 2 VPN features 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 Changed

IN THIS SECTION

- [Changes in CLI Statements and Commands | 30](#)
- [What's Changed in Junos OS Evolved Release 20.2R2 | 30](#)
- [What's Changed in Junos OS Evolved Release 20.2R1 | 31](#)

Learn about what changed in Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, 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 20.2R2

EVPN

- **New output flag for the `show bridge mac-ip table` command**—The Layer 2 address learning daemon does not send updated MAC and IP Address advertisements to the Routing Protocol daemon when an IRB interface is disabled in an EVPN-VXLAN network. Junos has added the NAD flag in the output of the **`show bridge mac-ip-table`** command to identify the disabled IRB entries where the MAC and IP address advertisement will not be sent.

[See [show bridge mac-ip-table](#).]

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](#).]
- **Change in auto-bandwidth adjustment (PTX10001-36MR, PTX10003, and PTX10008)**—If auto-bandwidth adjustment fails because of bandwidth unavailable error, the router tries to bring up the LSP with the same bandwidth during the subsequent reoptimization. In earlier releases, when the auto-bandwidth adjustment fails, the current bandwidth is reset to the bandwidth that was already active.

[See [rsvp-error-hold-time](#).]

Routing Protocols

- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (PTX Series and QFX Series)**—We've made changes to export multiple loopback addresses into `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier releases, multiple secondary loopback addresses in TED were added into `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router-id.
- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the `arp-snoop` packet type option in the `edit system ddos-protection protocols arp` protocol group to `arp`. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

System Management

- **Support for exclude option under file archive (PTX Series and QFX Series)**—The `exclude` option is added under the command `file archive` that specifies the file pattern to exclude. This option helps to exclude files that delay compression or files that do not require compression.

[See [file archive](#).]

What's Changed in Junos OS Evolved Release 20.2R1

Class of Service (CoS)

- **PTX10001-36MR, 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 20.2R1, 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 \(System\)](#).]

- **Modified request system storage cleanup command (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.2R1, the `request system storage cleanup` command does not remove Junos OS Evolved images from the device. It removes all core files, log files from `/var/log/`, and all `/var/log/*` files. To remove old images from the device, use the `request system software delete` command.

[See [request system storage cleanup](#) and [request system software delete](#).]

Juniper Extension Toolkit (JET)

- **Updates to IDL for RIB service API bandwidth field (PTX10003, PTX10008, and QFX5220)**—The Interface Definition Language (IDL) file for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

[See [Juniper EngNet](#).]

- **Python 3 add-on modules (PTX Series)**— Junos OS Evolved includes additional Python 3 libraries and modules, which Python scripts can import and use.

[See [Overview of Python Modules on Devices Running Junos OS](#).]

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](#).]

System Logging

- **Removal of show system queues command (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.2R1, 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

- [General Routing | 33](#)
- [Interfaces and Chassis | 34](#)

Learn about limitations in this release for the PTX10001-36MR, PTX10003 and PTX10008.

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

General Routing

- On PTX Series devices running Junos OS Evolved 19.3R1, the command **show ddos-protection protocols eoam oam-cfm** returns blank output. This command is supposed to be disabled on PTX Series devices. There is no functional impact; it is just an extra visible command option that has no functionality. [PR1456043](#)
- PTX10008 graceful OIR leads to missing sensor data for power and temperature. Follow the optimized step given as a workaround to have OIR working. [PR1478951](#)
- Restarting the fabspoked-pfe application for the line card restarts the line card. [PR1486023](#)
- The local repair time for fast reroute is 50 ms. If the system has a scaling configuration or is heavily loaded for processing, the local repair time may be longer than 50 ms. In this case, it gets 65 ms local repair time. [PR1489139](#)
- On PTX10008 routers, the counter per aggregated Ethernet child member for egress statistics is not accurate for an aggregated Ethernet sub-interface (logical units) and it should be ignored. [PR1490130](#)
- For a faulty link, after autoheal is triggered by jresil and before the action is completed, if fabspoked-fchip restarts, the link will stay in fault state, after the app restarts. [PR1508915](#)
- If a ZF polled I/O fault is hit followed by a second ZF polled I/O fault on the same SIB before the SIB has completely recovered from the first one, it might result in non-deterministic state of fabric links to or from the impacted SIB. [PR1519855](#)
- For PTX10008 platforms, input or output bytes under ifd stats include the Layer 2 header. [PR1524650](#)
- Excess-rate configuration in non SPS mode scheduler map may not be completely honored in some scenarios where guaranteed rate is configured without any excess-rate value, but some other queues have excess rate explicitly programmed. In that case, the left-over bandwidth after honoring the guaranteed rate across all OQs might still go more to the guaranteed marked OQs with no excess rate

programmed, thus drawing more share from the excess region to this guaranteed-rate configured OQ than the intended limits. [PR1528124](#)

- When the MTU of an interface is set to greater than 9500B, and if all the traffic going out of the interface has only packets greater than 9500B, then the ARP resolution over that interface might fail and all the traffic might get dropped. Workaround for this problem is:
 1. Ensure MTU size of interface is less than 9500B
 2. If an MTU of greater than 9500B is required, then it must be ensured that the traffic going out of all the unresolved next hops on that interface has at least one packet that is less than 9500B. It can also be a simple ping packet. [PR1536747](#)

Interfaces and Chassis

- On PTX10003-80C and PTX10003-160C devices, the **show chassis fpc** command does not display the temperature, CPU utilization, and memory utilization information, as these are fixed-configuration routers that do not have physical FPCs with dedicated CPUs. [PR1367671](#)
- When highly accelerated life test is carried out on the PTX10003, FPC error message **egp_intr_pkt_trapcode** is seen on the console. There is no direct impact on the control plane protocol or to data traffic. [PR1425508](#)

MPLS

- The rpd might generate a core file if **srgb** is configured to take over the complete label space. [PR1480427](#)

User Interface and Configuration

- For large YANG files, augmentation might not work. [PR1416972](#)

Open Issues

IN THIS SECTION

- [General Routing | 35](#)
- [Class of Service \(CoS\) | 37](#)
- [Infrastructure | 38](#)
- [Interfaces and Chassis | 38](#)

- [MPLS | 38](#)
- [Routing Protocols | 38](#)
- [User Interface and Configuration | 38](#)

Learn about open issues in this release for the PTX10001-36MR, PTX10003, and PTX10008.

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

General Routing

- Static link protections are not supported on the PTX10003. [PR1382695](#)
- Traffic loss of upto 400 ms is seen during MPLS FRR test with following scale: 600K IPv4 routes 40K IPv6 routes 19K Transit LSPs 1.6K ingress LSPs 10K BGP LU. [PR1472908](#)
- J-Flow sampling traffic is not rate-limited in host pipe. [PR1473844](#)
- On the PTX10001-36MR platform, fabric links might go to the fault state after an admin or automatic restart of the fabspoked-pfe application. [PR1473851](#)
- There is a minor data mismatch in RRO between PCE and PCC. [PR1481462](#)
- BGP started creating route pointing to list NH having indirect members. However, 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 the Packet Forwarding Engine as loop detection does not happen. [PR1484589](#)
- Counters for **show interfaces voq et-xxxx** can cause delay for more than 5 minutes and are not consistent for each CLI **show**. [PR1485459](#)
- In StrictPriority mode scheduler, low-priority traffic might get lesser than the expected share of the port bandwidth. For a good workaround, run strict-priority-scheduler with one of the following:
 - strict-high
 - high
 - medium-high
 - medium-low
 and avoid using priority-low. [PR1492985](#)
- When configured to use DHCP to configure IPv6 addresses on the management Ethernet the Junos OS Evolved, device ends up configuring the interface with a /64 subnet in all scenarios instead of a /128. This does not impact functionality. [PR1494823](#)

- The copying of files to the RCB over WAN ports is slow. This is observed across all platforms running Junos OS Evolved platforms. [PR1496895](#)
- On a PTX10008, six SIBs are required to carry line-rate traffic, with no fabric redundancy. Even when ingress traffic rate is such that five SIBs are sufficient to carry ingress traffic (for example, traffic is less than 1280 Gbps), ungraceful SIB failures result in transient loss of traffic, till system failure handling is triggered. In Junos OS Release 20.1R2, failure handling may result in about 4-6 seconds of traffic loss. We recommend that you take the fabric cards offline by using the **request chassis sib offline** command before removing the SIBs for maintenance. [PR1497212](#)
- Current design takes some time to program the BGP Flowspec filters on the device and that causes intermittent "error messages" for the **show firewall application routing** command. However, there is no data path functionality impact for this. And, it is a current design limitation and it requires optimization in the **dfwd-junos-relay** process and this is unlikely to get fixed in the Junos OS Evolved 20.2R1 branch. [PR1497926](#)
- When **fabric_down_condition_on_pfe** or **degraded_fabric_condition_on_pfe** errors are seen on a Packet Forwarding Engine, interfaces corresponding to the whole ASIC chip remain down. [PR1500969](#)
- If the number of flows is less, load balancing might not work properly on aggregated Ethernet members. If the traffic rate is over the line limit, tail drop is seen. [PR1501083](#)
- According to the current implementation, to display **switchover status** using the **show system switchover** command on backup, only the readiness of **configuration database**, **object database** and **applications: ready state** are considered. It does not check for the output of the **request chassis routing-engine master switch check** command. As a result, in some of the cases the output of both the commands could be different for 'switchover readiness' for a short duration of time. This is a day-1 behavior in Junos OS Evolved. [PR1503269](#)
- BGP-SRTE binding-sid with more than one label stack needs enhancement for PTX10003-80C and PTX10003-160C. [PR1512213](#)
- On a large scale routing on PTX10008, when an aggregated Ethernet interface flaps, during the next-hop convergence, you might see syslog filled with: **Possible out of order deleting of AftNode** error messages. [PR1513474](#)
- 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](#)
- In a scaled setup with 4k logical interfaces and 8k policers, when there is a reactivate and activate of firewall and interface configurations, aggregated Ethernet interfaces do not come up. [PR1518601](#)
- The queue Tx rate is incorrect for case with the **scheduler priority strict-high + tx-rate exact** configuration (l3_cos_schedulers.robot). [PR1519313](#)
- There is a delay enforced between back-to-back switchovers. This is 360 seconds unless overridden by platform-specific values. So, if the last switchover happened within the last 360 seconds, then you see the following error: **Not ready for mastership switch, try after x seconds**. The **show system switchover**

CLI output might show **Switchover Status: Ready**. However, this field indicates the ready status of all applications on the backup Routing Engine. This does not consider the back-to-back switchover delay. [PR1519364](#)

- In PTX10001-36MR, with multiple protocols and services configured with scale, the user may see available free memory reducing over time even without doing any changes in the configuration. [PR1520016](#)
- In PTX10001-36MR in rare cases with 10 s of repeated power-cycles, file system might encounter error or file system might not initialize, resulting in user login failure. [PR1523238](#)
- BGP switchover convergence rate is degraded by 42% on Junos OS Evolved 20.2R1. [PR1529365](#)
- Junos OS Evolved app "packetio" cannot be restarted like any other app available in the PTX10001-36MR. [PR1535979](#)

Class of Service (CoS)

- Layer 2 CoS is not supported but the classifier or rewrite configurations on Layer 2 interface is committed successfully. The commit constraint checks are not enforced. [PR1474086](#)
- CoS IEEE-802.1p default classifier is not working as expected, after deleting the customer IEEE-802.1p classifier or rewrite. [PR1478563](#)
- CoS default classifiers and rewrite rules are not working for both DSCP and DSCP-IPv6. [PR1479326](#)
- On all Junos OS Evolved platforms, CoS EXP classifier and rewrite with protocol option **mpls-inet-both-non-vpn** is not working as expected. [PR1479575](#)
- 32 DSCP classifier or rewrite functionality does not work after cosd restart (the PTX10008 platform supports only 16 DSCP classifier or rewrite). [PR1479676](#)
- On all Junos OS Evolved platforms, the cosd anomalies might be seen if CoS is configured on an interface using wildcard (like et-*, et-1/*/*, ae*, and fti*). These anomalies can cause traffic impact. The issue is due to some out-of-order internal events. [PR1494573](#)

Infrastructure

- Standard "issue" message is not printed for the initial login attempt, but if that login attempt fails for any reason and the user is re-prompted, then the message is printed. This seems to be an issue common to PTX10004. [PR1528996](#)

Interfaces and Chassis

- When there is a change in IPv6 address, ifmand crashes. There is no impact to system. [PR1496361](#)

MPLS

- When static MPLS LSPs are configured with different routing instances, and, interface configuration is deleted and added back, then these LSPs fail to resolve ARP due to incorrect routing-instance binding information. [PR1500521](#)

Routing Protocols

- Junos OS Evolved does not support configuring interface name as next hop. [PR1497012](#)

User Interface and Configuration

- CoS queue Tx rate change does not take effect even after successful commit. [PR1536615](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.2R2 | 39](#)
- [Resolved Issues: 20.2R1 | 42](#)

Learn which issues were resolved in Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, and PTX10008.

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

Resolved Issues: 20.2R2

Authentication and Access Control

- Random IPv6 address is displayed in NAS-IPv6-Address when IPv4 address is used as NAS-IP for IPv6 server. [PR1501146](#)

General Routing

- The aggregated Ethernet LACP distributed mode **show ppm adjacencies** does not show FPC slot number in the distribution address. [PR1458228](#)
- A configuration command to modify PCIe correctable and uncorrectable error thresholds on a PTX10008 FPC. [PR1462953](#)
- Block unified ISSU commands from the CLI for Junos OS Evolved Release 19.4R1. [PR1470943](#)
- Packet Forwarding Engine resiliency - interrupt handling on the ASIC corresponding to the PCIe link goes down. [PR1491128](#)
- Scheduler ingress Packet Forwarding Engine VOQ drop counters do not match egress queue drop counters, if the difference is greater than 100,000. [PR1494785](#)
- Sensor for physical Ethernet interface statistics is missing (for example, in-oversize-frame). [PR1496275](#)
- Colored packets are not policed as expected due to policer inaccuracy for color-blind single-rate-three-color policer at lower BW limits. [PR1497856](#)
- There is a discrepancy of 22 bytes for the same packets between the firewall filter in Junos OS Evolved and Junos OS. [PR1503145](#)
- PTX10008: Aggregated Ethernet (AE) interface flaps causing next hops to contain wrong encapsulation information, with the router having 800,000 IP routes, 2000 ingress LSPs, around 500 Layer 3 VPNs and LSPs have link protection enabled. [PR1503260](#)
- Junos OS Evolved versus Junos OS (serviceability) - Resilient hash seed configuration is not supported in Junos OS Evolved Release on PTX10003 or PTX10008. [PR1504461](#)
- Junos OS Evolved versus Junos OS (Serviceability) - CLI command is not available to display Packet Forwarding Engine hash-enabled fields in Junos OS Evolved. [PR1504544](#)
- For **show network-agent statistics**, Junos OS Evolved is reporting in milliseconds but Junos OS is reporting in seconds. [PR1504872](#)
- Packetio crashes during the initialization and this might result in a second reboot. [PR1505150](#)
- Shapers applied on interface output queues, either through the transmit-rate exact or rate-limit configuration, might not sometimes draw the expected output scheduler accuracy. [PR1506855](#)
- On a PTX10008 router, we observe small packet loss randomly during SIB offline. [PR1506866](#)

- The firewall process would take too a long to come online. [PR1507433](#)
- SIB Link error alarms get cleared after recovery of few fault planes by restarting the SIB even though there are some planes which are in fault state to different SIB. [PR1510766](#)
- PTX10008- policer is policing packets incorrectly for VLAN-tagged interfaces. [PR1511508](#)
- The evo-aftmand process might be stuck at 100% CPU usage in a scaled setup. [PR1511597](#)
- Packet Forwarding Engine generates MLPs with **global-no-mac-learning** enabled. [PR1511601](#)
- Firewall generates a core file after deleting and adding the filter back in a single commit. [PR1512065](#)
- Both degraded_fabric_reachability_to_peer_pfe and degraded_fabric_condition_on_pfe are seen on the same Packet Forwarding Engine. [PR1512270](#)
- 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 go offline and are brought online in quick succession, the device might crash and reboot after reporting a kernel msmi error or zookeeper session failed error. [PR1512633](#)
- Fabric cards might get stuck in the offline state sometimes if they go offline or are brought online in quick succession. [PR1512814](#)
- evo-cda-bt core file is seen on FPC0 after rebooting the device. [PR1513067](#)
- [Error] AftTelemetry : AftTelemetryHeaderGetComponentId: Unable to get FPC number error message. [PR1513306](#)
- The **show system errors active** command with or without FRU filters would give empty output when there have been no errors ever in the system. [PR1514105](#)
- On PTX10001-36MR transit traffic with DSCP EF marking seems to get rewritten with cs6 on the egress interface. [PR1514372](#)
- evo-aftmand-bt might crash if FPC offline is performed when the system comes up. [PR1514722](#)
- IPv6 neighbor entry is deleted in the **show ipv6 neighbors** table, thereby IPv6 traffic forwarding does not happen or gets dropped after interface flap while passing transit traffic with pktlen of 9000. [PR1515034](#)
- Continuous evo-cda-bt CDA syslog error messages are observed during the negative triggers (aggregated Ethernet configurations disable or enable and FPC restart). [PR1515806](#)
- SNMP traps are not seen for 'fpc_link_to_sib_fault' when CCL link is brought down from ZF to BT. [PR1516487](#)
- PTX10008: Do not configure confidentiality offset other than 0. [PR1517985](#)
- Traffic through channels 2 and 3 might be dropped when the interface et-0/2/0 is channelized to 4x10G or 4x25G speed on the PTX10001-36MR. [PR1517997](#)

- When all 5x400 g ports are used inside one portion of the chip, there can be impact on the traffic throughput performance. [PR1518368](#)
- **request system zeroize** does not reboot the device. [PR1518946](#)
- After **request system application app *application-name* node *node-name* restart**, the interface drop is observed and there is major dp_1_zfo_intr_dp1_fabcell_drop error. [PR1519402](#)
- A large number of publish-deleted hwdre anomalies are dumped after the second iteration of switchover from re1 (primary) to re0. [PR1519427](#)
- PTX10001-36MR: Junos OS Evolved 20.2R1 FDT image programs P1 coefficients for AC coupling instead of proto-2 parameters required for DC coupling. [PR1520962](#)
- The **show interfaces voq intf | ae non-zero** command is not working as expected and displays incorrect output. [PR1521281](#)
- With traffic running, if the FPC that is connected to the traffic generator reboots, it might lead to a stuck traffic scenario on certain ports post resumption. [PR1523066](#)
- rpd-agent daemon has a minor memory leak when it receives certain states. [PR1527399](#)
- Hwdre generates a core file when trying to take a faulty SIB to the offline state. [PR1527790](#)
- FPC does not boot if primary BIOS is corrupted (with FPC primary BIOS corruption golden BIOS failed to reprogram primary). [PR1528469](#)
- evo-aftmand-bt core file found at EalUnilistNh::getAlITargetElements (this=0x7f25a7ba2a40, child_nhs=0x7f2572065200, final_nhs=0x7f257242a000, ifl_indexs=0x7f257242a200, child_balances=0x7f257242a600. [PR1530696](#)
- The PTX10008 system might drop traffic in certain scenarios after a GRES event. [PR1532446](#)
- CLI issues are observed under the **show system** and **show chassis** hierarchies. [PR1535880](#)
- Physical interface does not come back to up state after configuration changes at peer end. [PR1536270](#)
- The **show chassis** environment reports some ZF internal temperatures as 0 degrees C. [PR1536497](#)
- LLDP might not be enabled after factory default. [PR1536531](#)
- Dependency-state details for the type net::juniper::hwd::fruHwStatus are seen after hwdfpc app restart with LSR core file profile configuration. [PR1537409](#)
- Channel 1 does not come up on channelizing 100G to 2x50G speed. [PR1539795](#)
- PTX10001-36MR :: Interface not coming up after enabling and disabling local loopback configuration. [PR1538393](#)
- On channelizing 400G DAC to 8x25G and 8x50G, ping is failing on the channels. [PR1539859](#)
- On channelizing 40G to 4x10G, ping is failing on the channels. [PR1539864](#)

- PTX10003: Multicast traffic drops after port is converted from trunk to access for a Layer 2 bridging case. Unicast and broadcast traffic are not affected. [PR1540495](#)
- PTX10008: SIB 0 ZFO internal temp sensor normal fan speed temp threshold is set to 0. [PR1540576](#)

Infrastructure

- Junos OS Evolved: **console log-out-on-disconnect** fails to terminate session on console cable disconnection (CVE-2020-1666). [PR1406238](#)
- Telnet login-related issue is observed with the template (TACACS+ and RADIUS). [PR1482363](#)

Interfaces and Chassis

- Commit time with 32,000 Layer 2 IFLs is 8 minutes. [PR1504294](#)
- Rarely can we see fan not getting detected post system reboot. [PR1517476](#)
- Minimum IFD MTU is recommended with IPv6 configuration. [PR1518692](#)
- Some of the SSD SMART attributes are reflecting as Unknown_Attribute. [PR1519264](#)
- Last flapped of virtual interfaces like aggregated Ethernet is not updated when aggregated Ethernet transitions from DOWN to UP state. [PR1521978](#)

Network Management and Monitoring

- SNMP response packets have [DF] Don't fragment flag set by default in routers running affected Junos OS Evolved release. [PR1514156](#)
- Multiple SNMP core files might be generated on Junos OS Evolved platforms. [PR1516161](#)
- SNMP MIB ifInErrors [OID 1.3.6.1.2.1.2.2.1.14] reports wrong values in routers running affected Junos OS Evolved release. [PR1534286](#)

Resolved Issues: 20.2R1

General Routing

- The aggregated Ethernet queue statistics are all 0 for more than 3 minutes (whereas rate, pps, or bps statistics are correct) after the **clear interfaces statistics all** command is issued and when one AE member link flaps (24x10G ae0). [PR1423134](#)
- 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](#)
- Operational and configuration CLIs changed. [PR1451455](#)
- The rpd core file might be generated and the process might not come up. [PR1453446](#)
- With 3000 LSPs, about 500 ms of traffic loss is observed for Routing Engine-driven local repair. [PR1459265](#)

- System software issue where PSM FW 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 command output. [PR1461691](#)
- UDP feature is not enabled for the Packet Forwarding Engine statistics on PTX10003-80C and PTX10003-160C. [PR1467333](#)
- SSH service is unresponsive after setting **system services ssh key-exchange dh-group1-sha1**. [PR1467467](#)
- Ports are in disabled state. Platform binding queue points to incomplete objects in evo-aftmand-zx with **Waiting usr inc : true** in one guide, and also it can be recovered only by reboot. [PR1467546](#)
- **xmlproxyd** binding queue complete-deleted anomalies seen for the type **net::juniper::config::services::agent**. [PR1470872](#)
- The master-only configuration in the management Ethernet configuration does not work. [PR1473388](#)
- On the PTX10001-36MR platform, fabric links may go to the fault state after an admin or automatic restart of the fabspoked-pfe application. [PR1473851](#)
- BGP sessions might flap in a large-scale scenario. [PR1474532](#)
- All images except current and rollback removed when /var is full on the backup Routing Engine. [PR1475655](#)
- Traffic is silently discarded (without notification) on PTX10008 after FPC physical OIR operations with aggregated Ethernet bundle. [PR1475694](#)
- The picd top-level periodic is hogging CPU. [PR1475753](#)
- The TCP connection is still established even after the interface is deleted or deactivated or disabled. The **show Agent sensors** still shows the subscription. [PR1476886](#)
- **IFL index availability is critical** alarm seen after multiple GRES operations with basic configuration. [PR1477215](#)
- FPC may continuously crash after deactivating or activating loopback filter or reboot the system after configuring the loopback filter. [PR1477740](#)
- PSM output power failure on specific Junos OS Evolved platforms. [PR1477975](#)
- Default election priority goes to backup state, when we configure disabled state. [PR1478717](#)
- Power off - power on of a PTX10008 Junos OS Evolved leads to SIB error (bootup_udev_create_error) in **show system error active** (no traffic). [PR1480354](#)
- USB upgrade fails from Junos OS Evolved Release 19.3R2.1 to Junos OS Evolved Release 19.4R1.8. [PR1480371](#)
- The configuration statement **name** not available after **show path-computation-client lsp** and **show path-computation-client lsp extensive**. [PR1481822](#)

- The hwdre might crash when removing and replacing FPCs after performing back-to-back Routing Engine mastership switchovers. [PR1482453](#)
- Deleting L2 logical interface does not clean up applied Ba idx configuration. [PR1484044](#)
- Regression counters for **show interfaces voq et-xxxx** are delayed by more than 5 minutes and are not consistent for each CLI show. [PR1485459](#)
- In strict-priority mode, with any form of rate-limiter/shaper configured, it might create some accuracy gaps in the scheduler across queues. [PR1485478](#)
- The 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](#)
- Scheduler in strict-priority mode, queue tx rate should not be affected by configured transmit-rate/remainder. [PR1485757](#)
- Junos OS Evolved apps show up in Junos OS Evolved 19.4R1-S1-202001272356.0 but not defined in FS. [PR1486018](#)
- The **Component(s)** field of **show agent sensors** is implementation-defined. [PR1487125](#)
- **set system processes app failover** configuration statement is not supported. [PR1487775](#)
- With 80% traffic, single SIB yank (ungraceful OIR) leads to silent dropping of traffic and major errors across other SIBs and some FPCs. [PR1489544](#)
- On PTX10008 running Junos OS Evolved Release 20.2R1, ping to any neighboring routers fails to work whenever the PTX10008 comes up with BGP sessions configured. [PR1489569](#)
- Rewrite at egress interface is not working as expected. After rewrite, BE queue is carrying the traffic instead of AF. [PR1490814](#)
- Packet Forwarding Engine resiliency - interrupt handling on the ASIC corresponding to the PCIe link down. [PR1491128](#)
- Traffic statistics are not updated for bypass LSP even though the traffic is flowing through the bypass LSP. [PR1491467](#)
- Unexpected incomplete object notified by application controller causes aftmand crash. [PR1491548](#)
- The evoaft-jvisiond core file is generated after GRES. [PR1492059](#)
- sFlow ingress sampling reports wrong next hop in case IPv6 traffic is routed through LSP at ingress provider edge (PE) router. [PR1492076](#)
- Behavior change in clients with multiple gRPC channels to same target. [PR1492088](#)
- l2ald core seen if VLAN member scale configuration is re-applied. [PR1492546](#)
- sFlow ingress sampling reports next hop as 0.0.0.0 in case IPv6 is label-swapped through segment routing LSP. [PR1492616](#)

- After flapping the aggregated Ethernet interface through child link flap 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](#)
- cmd is crashing when the **request system shutdown reboot disk1** command is executed. [PR1492955](#)
- For PTX10K-LC1101 Line Card, Routing Engine queue statistics show low value for tail-dropped bits per second. [PR1493787](#)
- JNP10K-LC1201 interfaces do not show up after the line card was taken offline, brought online, went through OIR, or restarted and then brought online.. [PR1494793](#)
- Basic IPv6 and IPv4 CBF (CoS-based forwarding) is broken in the latest Junos OS Evolved 20.1 dailies: 20.1-202003090136.0-EVO and 20.1-202003101244.0-EVO. After CBF action, the packets are not transmitted using correct queues. [PR1495119](#)
- Block software upgrade when image that is not compatible with the platform is used. [PR1495655](#)
- The evo-aftmand-bt generates a core file at JzphrHandleMplsLkupEntry::getFdbMplsEntKey. [PR1496057](#)
- Consistent hash information is missing : **error: timeout communicating with Evo-Aft BT daemon.** [PR1496077](#)
- Longevity: aggregated Ethernet logical interface disappear after switchover. [PR1496887](#)
- The picd generates a core file in net::juniper::picd_cmd::initialize_optic_port_information_DataIterator::receive_request]. [PR1497285](#)
- Tracking PR to optimize the publishing of flowspec scale filter (Between Bulk Start and Bulk End Marker). [PR1497926](#)
- Traceroute in an MPLS network might not show the complete path. [PR1498949](#)
- Packet drops following **rsvp load-balance** configuration on PTX10003. [PR1500711](#)
- Any new unicast NH token does not get updated into the indirect NH token. [PR1501284](#)
- Routes may be installed in Packet Forwarding Engine and used for forwarding even if the interface is down or disabled. [PR1501321](#)
- For the scheduler with oversubscribed traffic, CLI **show interface queue et-x/x/x** still shows no traffic drop, and **show interface voq et-x/x/x** shows drop (l3_lag_cos_voq_stats.robot l3_cos_voq_stats.robot). [PR1501802](#)
- Each FPC is exporting data for all interfaces. [PR1505521](#)
- The **set system processes app failover other-routing-engine** configuration is not supported. [PR1506480](#)
- PTX10003: Traffic drop is seen with 96,000 ND scale at dlu.ucode.discard and egp.ups_err. [PR1506590](#)
- PTX10008: Observe small packet loss randomly when the SIB is offline. [PR1506866](#)

- IP prefix-based traffic load balancing is not happening on PTX10003-80C and PTX10003-160C. [PR1507410](#)
- 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 get cleared. [PR1510763](#)
- SIB LINK Error alarms are getting cleared after recovery of few fault planes by restarting SIB even though there are some planes which are in fault state to different SIB. [PR1510766](#)
- The aggregated Ethernet interfaces are down after running scale tests with policer configuration. [PR1511597](#)
- FirewallD generates a core file after deleting and adding the filter back in single commit. [PR1512065](#)
- Both **degraded_fabric_reachability_to_peer_pfe** and **degraded_fabric_condition_on_pfe** are seen on a same Packet Forwarding Engine. [PR1512270](#)
- Fabric cards might get stuck in the offline state some times if it goes offline and brought online in quick succession. [PR1512814](#)
- evo-cda-bt core is seen on FPC0 after rebooting the DUT. [PR1513067](#)
- AftTelemetryHeaderGetComponentId: Unable to see FPC number error message. [PR1513306](#)
- SIB fault errors and alarms are not cleared after SIB recovers following fabspoked-fchip abrupt kill and restart. [PR1514668](#)
- IPv6 neighbor entry is deleted in **show ipv6 neighbors** table thereby IPv6 traffic forwarding not happening or dropped after interface flap while passing transit traffic with pktlen of 9000. [PR1515034](#)

Class of Service (CoS)

- Recommended FC configuration to aid CBF and classifier override. [PR1454501](#)
- cosd core file is seen at 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](#)
- CoS EXP classifier and rewrite with protocol option **mpls-inet-both-non-vpn** is not working as expected. [PR1479575](#)
- The 32 DSCP classifier or rewrite functionality does not work after the cosd restart. [PR1479676](#)
- cosd anomalies are observed after committing the EBGp baseline configuration. [PR1494573](#)
- CoS interface bind - interface name. [PR1500722](#)

Flow-Based and Packet-Based Processing

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

Infrastructure

- Ping is not working when we set **record-route** [PR1474649](#)
- fibd daemon crashes and generates a core file. [PR1487437](#)

Interfaces and Chassis

- The 400GbE link might take a longer time (about 50 seconds) to come up. [PR1450606](#)
- IPv4 and IPv6 unicast RPF test cases failed on aggregated Ethernet interface after restarting ifmand. [PR1453461](#)
- The CLI **show interfaces terse** operation is very slow in physical interface scaling case. [PR1454017](#)
- Deprecated gigether-options in Junos OS Evolved. [PR1473492](#)
- After application restart - false optics alarms continue to persist although underlying optics does not have the failures. [PR1493230](#)
- PTX10008 with LC1201 - the ifmand process restarts at IFAManager::findIfaoInSameSubnet. [PR1496361](#)
- Direct attach copper (DAC) cables on PTX10003 might experience delay in coming up or might fail to come up. [PR1504622](#)

Juniper Extension Toolkit (JET)

- The support of JET IDL packages for the EVO platform. [PR1471173](#)

Network Management and Monitoring

- The 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](#)
- Timeticks is all zeros for mplsL3VpnVrfCreationTime. [PR1499757](#)

Platform and Infrastructure

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

Routing Policy and Firewall Filters

- The system crashes after configuring filter with ICMP match conditions. [PR1496740](#)

Routing Protocols

- The rpd might be crashed after BGP peer flapping. [PR1482551](#)

User Interface and Configuration

- CMDD application restart issue: command-handler subsystem does not respond. [PR1434409](#)
- Initiate an empty commit from the master Routing Engine when a different Routing Engine is inserted in backup node. [PR1465291](#)
- On modifying filter terms after confined restart, all counters are not displayed under **show firewall**. [PR1484340](#)
- The aggregated Ethernet bundle is not coming up after configuring the **vlan-tagging** and then rolling back the configuration. [PR1491955](#)

Junos OS Evolved Release Notes for QFX5220 Devices

IN THIS SECTION

- [What's New | 49](#)
- [What's Changed | 50](#)
- [Known Limitations | 53](#)
- [Open Issues | 54](#)
- [Resolved Issues | 56](#)

These release notes accompany Junos OS Evolved Release 20.2R2 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 20.2R2 | 49](#)
- [What's New in Junos OS Evolved Release 20.2R1 | 49](#)

Learn about new features introduced in Junos OS Evolved Release 20.2R2 for the QFX5220.

What's New in Junos OS Evolved Release 20.2R2

There are no new features or enhancements to existing features for the QFX5220 switches in Junos OS Evolved Release 20.2R2.

What's New in Junos OS Evolved Release 20.2R1

Class of Service (CoS)

- **Priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic (QFX5220)**—Starting in Junos OS Evolved Release 20.2R1, to support lossless traffic across Layer 3 connections to Layer 2 subnetworks on QFX5220 switches, you can configure PFC to operate using 6-bit DSCP values from Layer 3 headers of untagged VLAN traffic, rather than IEEE 802.1p priority values in Layer 2 VLAN-tagged packet headers. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).

[See [Understanding PFC Using DSCP at Layer 3 for Untagged Traffic](#).]

Multicast

- **IGMP snooping support (QFX5220)**—Starting with Junos OS Evolved Release 20.2R1, you can configure IGMP snooping.

Support includes:

- IGMP snooping with IGMPv1, IGMPv2, and IGMPv3
- Any-source multicast (ASM) and source-specific multicast (SSM) modes
- IGMP querier at Layer 2

Limitations include:

- IGMP proxy mode is not supported.
- Devices receiving IGMP group-specific queries on a multicast router interface forward those queries to all other interfaces in the VLAN.
- Make before break (MBB) is not supported for multicast receivers on existing Layer 2 aggregated ethernet (**ae**x) or link aggregation group (LAG) interfaces, including for member additions, deletions, and link up or down events.
- The devices forward all unregistered IPv4 and IPv6 multicast packets to the multicast router interfaces in the VLAN even if you configured an interface as a multicast router interface only for IGMP snooping.
- IGMP snooping using IRB interfaces is not supported.
- IGMP snooping is not supported in EVPN-VXLAN or EVPN-MPLS networks.

[See [IGMP Snooping Overview](#).]

What's Changed

IN THIS SECTION

- [Changes in CLI Statements and Commands | 51](#)
- [What's Changed in Junos OS Evolved Release 20.2R2 | 51](#)
- [What's Changed in Junos OS Evolved Release 20.2R1 | 52](#)

Learn about what changed in Junos OS Evolved Release 20.2R2 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 20.2R2

Class of Service (CoS)

- On QFX5130 and QFX5220, you must map all PFC-enabled IEEE 802.1P code-points to a lossless (no-loss) forwarding class. If a CNP has code-points that are mapped to a lossy forwarding class, the entire CNP will not be programmed in hardware..

EVPN

- **New output flag for the show bridge mac-ip table command**—The Layer 2 address learning daemon does not send updated MAC and IP Address advertisements to the Routing Protocol daemon when an IRB interface is disabled in an EVPN-VXLAN network. Junos has added the NAD flag in the output of the **show bridge mac-ip-table** command to identify the disabled IRB entries where the MAC and IP address advertisement will not be sent.

[See [show bridge mac-ip-table](#).]

MPLS

- **Change in auto-bandwidth adjustment (QFX5130-32CD and QFX5220)**—If auto-bandwidth adjustment fails because of bandwidth unavailable error, the router tries to bring up the LSP with the same bandwidth during the subsequent reoptimization. In earlier releases, when the auto-bandwidth adjustment fails, the current bandwidth is reset to the bandwidth that was already active.

[See [rsvp-error-hold-time](#).]

Routing Protocols

- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (PTX Series and QFX Series)**—We've made changes to export multiple loopback addresses into Isdist.0 and Isdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier releases, multiple secondary loopback addresses in TED were added into Isdist.0 and Isdist.1 routing tables as part of node characteristics and advertised them as the router-id.
- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

System Management

- **Support for exclude option under file archive (PTX Series and QFX Series)**—The **exclude** option is added under the command **file archive** that specifies the file pattern to exclude. This option helps to exclude files that delay compression or files that do not require compression.

[See [file archive](#).]

What's Changed in Junos OS Evolved Release 20.2R1

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 20.2R1, 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 \(System\)](#).]

- **Modified request system storage cleanup command (PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 20.2R1, the **request system storage cleanup** command does not remove Junos OS Evolved images from the device. It removes all core files, log files from **/var/log/**, and all **/var/log/*** files. To remove old images from the device, use the **request system software delete** command.

[See [request system storage cleanup](#) and [request system software delete](#).]

High Availability (HA) and Resiliency

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

Interfaces and Chassis

- **Support for 40-Gbps speed on QSFP28 ports (QFX5220-128C)**—The QFX5220-128C consists of 128 high-density 100-Gigabit Ethernet QSFP28 ports. Starting in Junos OS Evolved Release 20.2R1, you can configure all the QSFP28 ports to operate at 40-Gbps speed.

In earlier Junos OS Evolved releases, you can configure 40-Gbps speed only on even-numbered QSFP28 ports.

[See [QFX5220-128C Port Panel](#).]

Juniper Extension Toolkit (JET)

- **Updates to IDL for RIB service API bandwidth field (PTX10003, PTX10008, and QFX5220)**—The Interface Definition Language (IDL) file for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

[See [Juniper EngNet](#).]

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](#).]

Known Limitations

IN THIS SECTION

- [Interfaces and Chassis | 54](#)

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

- QFX5220 does not support mixed-speed aggregated Ethernet. [PR1515041](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 55](#)
- [Routing Policy and Firewall Filters | 55](#)
- [Routing Protocols | 55](#)
- [System Management | 55](#)
- [User Interface and Configuration | 56](#)

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.

Class of Service (CoS)

- On all Junos OS Evolved platforms, the cosd anomalies might be seen if CoS is configured on an interface using wildcard (like et-*, et-1/*/*, ae*, and fti*). These anomalies can cause traffic impact. The issue is due to some out-of-order internal events. [PR1494573](#)

Routing Policy and Firewall Filters

- 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](#)

Routing Protocols

- The **virtual-router** option is not supported under routing-instance in lean rpd image. [PR1494029](#)

System Management

- CLI changed in Junos OS Evolved 20.1R1 release to align with Junos OS:

All the boot device options and events that are available at **request system shutdown reboot** are moved under the **request node** options.

Boot device options: usb, network, disk1 and disk2.

Event options: reboot, halt, power-off, power-on, offline, and online. [PR1456037](#)

- PTP-FPGA access error displayed as **PTP-FPGA is not initialized**. [PR1457094](#)
- When link goes up or down, Tx/Rx loss of signal alarm from diagnostics optics might take a long time (around 10 minutes) to show up. [PR1499741](#)

User Interface and Configuration

- On QFX5220-32CD, VLANs between 3968 and 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

- [Resolved Issues: 20.2R2](#) | [56](#)
- [Resolved Issues: 20.2R1](#) | [57](#)

Learn which issues were resolved in Junos OS Evolved Release 20.2R2 for the QFX5220.

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

Resolved Issues: 20.2R2

Routing Policy and Firewall Filters

- The JET APIs related to firewall filter might fail on QFX5220-32CD and QFX5220-128C with Junos OS Evolved 20.1R1 and later releases. [PR1515850](#)

Routing Protocols

- The protocol LLDP might not be enabled after factory default. [PR1536531](#)

System Management

- QFX5220 does not support the following matches, which were supported on QFX5200: first-fragment, tcp-established, ls-fragment, and hop-limit. [PR1499009](#)
- SFP+-40G-CU1M's media type appears as fiber. [PR1499408](#)
- The evo-aftmand process might be stuck at 100% CPU usage in a scaled setup. [PR1511597](#)
- The temperature-threshold settings for yellow alarms are not reset back to original values after changing temperature thresholds on sensor components. [PR1513097](#)

- QFX5220-32CD: CRC errors are seen with 4x100G FR4 and DR4 optics in the environmental stress test. [PR1520133](#)
- Scheduling variance will be seen when queues 3 and 4 are congested on QFX5220. [PR1521531](#)

Resolved Issues: 20.2R1

Class of Service (CoS)

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

Infrastructure

- fibd generates core files that can be seen after commit when the VLAN ID of a VLAN is changed. [PR1497790](#)

Interfaces and Chassis

- Physical interface statistics retain old values after statistics are cleared and interfaces deactivated and then activated. [PR1403606](#)
- The 400 GbE link might take a longer time (about 50 seconds) to come up. [PR1450606](#)
- The interface is not a part of member VLAN when configured. [PR1495207](#)

Platform and Infrastructure

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

Routing Protocols

- SFTP does not connect properly. The following error is displayed **Received message too long**. [PR1475255](#)

System Management

- Output MTU Error statistics do not increase when the **show pfe statistics traffic** command is issued. [PR1434860](#)
- Event options and device options are moved from **request system shutdown reboot** to request node. [PR1456037](#)
- The **switch-options** configuration not available on lean rpd images. [PR1457825](#)
- On platforms running Junos OS Evolved, multiple syslog servers are not being configured on the device when multiple syslog servers are configured on the DHCP server. [PR1458931](#)
- Sometimes dhclient generates a core file during ZTP retry when the dhclient process is exiting. [PR1460906](#)
- Output 1PPS and 10MHz measurement ports on QFX5220-32CD are not fully available. [PR1464452](#)
- The timingd process generates a core file rarely after a switch reboot. [PR1464493](#)
- SSH service is unresponsive after setting **system services ssh key-exchange dh-group1-sha1**. [PR1467467](#)

- High CPU consumption is seen when a scaled VLAN configuration having 3965 VLANs is committed for 30 minutes. [PR1474374](#)
- The lo0 firewall filter might affect Layer 3 forwarding traffic on QFX5220 platforms. [PR1475620](#)
- The QFX5220-128C, INPHI firmware upgrade requires power cycling. [PR1477124](#)
- ZTP: Image upgrade using HTTPS does not work in Junos OS Evolved 20.2R1. [PR1479125](#)
- Failed to start MGD initialization of schema and database while upgrade from Junos OS Evolved 19.3R2.1 to 19.4-202002240046.0-EVO]: mgd-init service fails due to clksyncd failure. [PR1491803](#)
- cmdd crashes when **request system shutdown reboot disk1** command is executed. [PR1492955](#)
- QFX5220-32CD and QFX5220-128C: evo-pfemamd core file is seen if loopback filter is applied with input-list and after switch is rebooted. [PR1497909](#)
- On a QFX5220 :: **show firewall cli** shows error **error: communication failure with /re0/evo-pfemamd/** when unique filter is applied over all physical interfaces. [PR1499294](#)
- ACL does not get applied to ingress when single filter is applied to both input and output direction of intf. [PR1501015](#)

User Interface and Configuration

- CMDD application restart issue: command-handler subsystem does not respond. [PR1434409](#)
- The following error is observed **The command-handler subsystem is not responding to management requests** while upgrading from Junos OS Evolved 19.3R2.1-EVO to Junos OS Evolved 19.4R2.3. [PR1499018](#)

Upgrade Your Junos OS Evolved Software

Products impacted: PTX10001-36MR, PTX10003, PTX10008, QFX5220-32CD, and 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 more about Junos OS Evolved and other Juniper products.

- Feature Explorer—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network.
<https://apps.juniper.net/feature-explorer/>
- PR Search Tool—Keep track of the latest and additional information about Junos OS Evolved open defects and issues resolved. prsearch.juniper.net
- Hardware Compatibility Tool—Determine optical interfaces and transceivers supported across all platforms. 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/

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. 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

12 May 2022—Revision 8, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

5 May 2022—Revision 7, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

28 January 2022—Revision 6, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

7 October 2021—Revision 5, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

16 July 2021—Revision 4, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

22 April 2021—Revision 3, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

13 January 2021—Revision 3, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

10 December 2020—Revision 2, Junos OS Evolved Release 20.2R2 for the PTX10001-36MR, PTX10003, PTX10008, and QFX5220 Devices.

2 November 2020—Revision 1, Junos OS Evolved Release 20.2R2 for the PTX10003, PTX10008, and QFX5220 Devices.

Copyright © 2022 Juniper Networks, Inc. All rights reserved.

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.

