

# Release Notes: Junos<sup>®</sup> OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices

29 July 2022

<b>Contents</b>	<b>Introduction   4</b>
	<b>Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, and PTX10008 Devices   4</b>
	<b>What's New   5</b>
	What's New in Junos OS Evolved 20.3R2   5
	What's New in Junos OS Evolved 20.3R1-S1   6
	What's New in Junos OS Evolved 20.3R1   6
	<b>What's Changed   51</b>
	What's Changed in Junos OS Evolved Release 20.3R2   51
	What's Changed in Junos OS Evolved Release 20.3R1-S3   54
	What's Changed in Junos OS Evolved Release 20.3R1-S2   54
	What's Changed in Junos OS Evolved Release 20.3R1   55
	<b>Known Limitations   58</b>
	General Routing   58
	MPLS   60
	Platform and Infrastructure   60
	Routing Protocols   60

**Open Issues | 60**

- Authentication and Access Control | 61**
- General Routing | 61**
- Infrastructure | 64**
- Interfaces and Chassis | 65**
- Layer 2 Ethernet Services | 65**
- MPLS | 65**
- Network Management and Monitoring | 66**
- Routing Protocols | 66**
- User Interface and Configuration | 66**

**Resolved Issues | 67**

- Resolved Issues: 20.3R2 | 67**
- Resolved Issues: 20.3R1 | 71**

**Junos OS Evolved Release Notes for QFX5130-32CD and QFX5220 Devices | 74****What's New | 75**

- What's New in Junos OS Evolved 20.3R2 | 75**
- What's New in Junos OS Evolved 20.3R1 | 76**

**What's Changed | 84**

- What's Changed in Junos OS Evolved Release 20.3R2 | 85**
- What's Changed in Junos OS Evolved Release 20.3R1 | 86**

**Known Limitations | 87**

- Network Security | 88**
- Routing Protocols | 88**
- Software Licensing | 88**
- System Management | 88**
- User Interface and Configuration | 89**

**Open Issues | 89**

- General Routing | 90**
- Infrastructure | 90**
- Network Management and Monitoring | 91**
- Routing Protocols | 91**
- System Management | 91**
- User Interface and Configuration | 92**

Resolved Issues | 92

Resolved Issues: 20.3R2 | 92

User Interface and Configuration | 94

Resolved Issues: 20.3R1 | 94

Upgrade Your Junos OS Evolved Software | 95

Licensing | 96

Finding More Information | 96

Documentation Feedback | 97

Requesting Technical Support | 98

Self-Help Online Tools and Resources | 98

Creating a Service Request with JTAC | 99

Revision History | 99

# Introduction

Junos OS Evolved is the next generation Junos OS with a modernized infrastructure, enabling higher availability, accelerated deployment, greater innovation and improved operational efficiencies. Junos OS Evolved is aligned with Junos OS for management and automation, but as we make improvements to Junos OS Evolved there some changes to be aware of. For details on these changes, see [How Junos OS Evolved Differs from Junos OS](#).

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 20.3R2.

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, PTX10004, and PTX10008 Devices

### IN THIS SECTION

- [What's New | 5](#)
- [What's Changed | 51](#)
- [Known Limitations | 58](#)
- [Open Issues | 60](#)
- [Resolved Issues | 67](#)

These release notes accompany Junos OS Evolved Release 20.3R2 for PTX10001-36MR, PTX10003, PTX10004, 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.

Check out [Table 1 on page 5](#) to see a complete list of the Junos OS Evolved releases supported on routing platforms.

**Table 1: Junos OS Evolved releases supported on routing platforms**

Platform	Supported Junos OS Evolved Releases
PTX10001-36MR	20.2R1, 20.2R2, 20.3R1, 20.3R2, 20.4R1, and 21.1R1
PTX10003	19.1R1, 19.1R2, 19.2R1, 19.3R1, 19.3R2, 19.4R2, 20.1R1, 20.1R2, 20.2R1, 20.2R2, 20.3R1, 20.3R2, 20.4R1, and 21.1R1
PTX10004 with PTX10K-LC1201	20.3R1, 20.3R2, 20.4R1, and 21.1R1
PTX10004 with PTX10K-LC1202-36MR	20.4R1 and 21.1R1
PTX10008 with PTX10K-LC1201	20.1R1, 20.1R2, 20.2R1, 20.2R2, 20.3R1, 20.3R2, 20.4R1, and 21.1R1
PTX10008 with PTX10K-LC1202-36MR	20.3R1, 20.3R2, 20.4R1, and 21.1R1

## What's New

### IN THIS SECTION

- [What's New in Junos OS Evolved 20.3R2 | 5](#)
- [What's New in Junos OS Evolved 20.3R1-S1 | 6](#)
- [What's New in Junos OS Evolved 20.3R1 | 6](#)

Learn about new features introduced in Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008.

### What's New in Junos OS Evolved 20.3R2

There are no new features introduced in Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008.

## What's New in Junos OS Evolved 20.3R1-S1

### MPLS

- **RSVP and LDP based point-to-multipoint (P2MP) LSPs with graceful restart (PTX10008)**— Starting in Junos OS Evolved Release 20.3R1-S1, PTX10008 routers support RSVP and LDP based point-to-multipoint (P2MP) LSPs with graceful restart. In addition, the PTX10008 supports IP unicast traffic in a label-edge router (LER) role, and both IP unicast and multicast traffic in a label-switching router (LSR) role.

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

## What's New in Junos OS Evolved 20.3R1

### Hardware

- **New PTX10004 Packet Transport Router (PTX Series)**—In Junos OS Evolved Release 20.3R1, we introduce the PTX10004 router as the most compact, high-density, and power-efficient modular chassis in the PTX10000 line of modular packet-routing transport routers. At only 7U in height, the PTX10004 is designed for space-constrained facilities. Like the larger PTX10008 router, the PTX10004 supports Juniper's 400 GbE architecture with inline Media Access Control Security (MACsec) on all ports for point-to-point security on Ethernet links. Each PTX10K-1201-36CD line card has a throughput of 14.4 Tbps per line card, giving the chassis an effective switching capacity of 57.6 Tbps. That throughput means a fully equipped PTX10004 can support 576 10 GbE, 576 25 GbE, 144 40 GbE, 576 100 GbE, or 144 400 GbE interfaces in a single chassis. The line cards are enabled for optics and chassis MIBs support. The PTX10004 supports the same feature set as the PTX10008. To install the PTX10004 router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the [PTX10004 Packet Transport Router Hardware Guide](#).

[Table 2 on page 7](#) summarizes the PTX10004 features supported in Junos OS Evolved Release 20.3R1.

**Table 2: Feature Support Added to the PTX10004 Routers**

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> <li>Support for CoS features. These features include classifiers (behavior aggregate (BA), fixed, and multifield (MF)), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, and drop-profiles.</li> </ul> <p>[See <a href="#">CoS Features and Limitations on PTX Series Routers</a>.]</p>
DHCP	<ul style="list-style-type: none"> <li>Support for DHCPv4 and DHCPv6 stateless relay. [See <a href="#">DHCP Relay Agents</a>.]</li> </ul>
Firewalls and policers	<ul style="list-style-type: none"> <li>Support for IPv4 and IPv6 firewall filters. Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on the router from a source address to a destination address. Features include: <ul style="list-style-type: none"> <li>Forwarding IPv4 and IPv6</li> <li>Firewall filter</li> <li>Load balancing</li> <li>MPLS fast reroute</li> <li>Host path</li> <li>Egress peer engineering</li> </ul> </li> </ul> <p>[See <a href="#">Firewall Filter Match Conditions and Actions (PTX Series Routers)</a>.]</p>
High availability (HA) and resiliency	<ul style="list-style-type: none"> <li>Support for fabric management. You can perform these fabric management functions on the PTX10004: <ul style="list-style-type: none"> <li>Initialize, configure, and manage fabric ASIC and other fabric components on the switch fabric board</li> <li>Configure and manage PFE ASIC fabric blocks and other fabric devices on the line card (FPC).</li> <li>Fabric CCL link training and link error handling</li> <li>Configure and manage fabric plane</li> <li>System wide synchronization of fabric operations related to FPC online/offline, SIB offline/online, and removing failed link from data path.</li> <li>Fabric ASIC initialization time error handling</li> <li>Fabric ASIC runtime error monitoring and handling</li> <li>Fabric ASIC initialization time error handling</li> <li>Fabric statistics collection</li> <li>Handling RE switchover related aspects for SIB</li> </ul> </li> </ul> <p>[See <a href="#">show chassis fabric summary</a>.]</p>

Table 2: Feature Support Added to the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>Platform resiliency to handle failures and faults related to the hardware components such as line cards, switch fabric, control boards, fan trays, fan tray controllers, temperature sensors, and power supply units. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See <a href="#">show system errors active</a>.]</li> <li>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 <a href="#">show system errors active</a>.]</li> </ul>
Interfaces and chassis	<ul style="list-style-type: none"> <li>Support for the dedicated virtual routing and forwarding (VRF) instance mgmt_junos for the management Ethernet Interface. [See <a href="#">Management Interface in a Nondefault Instance</a> and <a href="#">management-instance</a>.]</li> <li>Supported port speeds. You can configure port speed for the JNP10K-LC1201 line card by using the port speed configuration statements in the <b>[edit interfaces]</b> hierarchy. You can choose to configure the following port speeds on all 36 ports: <ul style="list-style-type: none"> <li>4x10 Gbps, 4x25 Gbps, and 2x50 Gbps</li> <li>40 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps</li> </ul> <p>To configure the port speed of a physical port in a PIC, use the <b>speed</b> statement at the <b>[edit interfaces interface-name]</b> hierarchy. To configure the port speed of a physical port with a specified number of channelized interfaces, use the <b>speed</b> statement and specify the number of channelized interfaces by using the <b>number-of-sub-ports</b> statement at the <b>[edit interfaces interface-name]</b> hierarchy. To control the number of interfaces (channelized or non-channelized) created on a physical port, use the <b>unused</b> statement at the <b>[edit interfaces interface-name]</b> hierarchy.</p> <p>NOTE: The following CLI statements are no longer supported:</p> <ul style="list-style-type: none"> <li><b>pic-mode</b> statement at the <b>[edit chassis fpc fpc-slot-number pic pic-number]</b> hierarchy</li> <li><b>speed</b> statement at the <b>[edit chassis fpc fpc-slot-number pic pic-number port port-number]</b> hierarchy</li> <li>Port profile configuration for the JNP10K-LC1201 line card at this hierarchy level.</li> </ul> <p>[See <a href="#">Port Speed</a>.]</p> </li> <li>Support for Link Aggregation Control Protocol (LACP) features. You can configure LACP to bundle several physical interfaces to form one logical aggregated Ethernet interface. By default, LACP is disabled on aggregated Ethernet interfaces. You can configure LACP in active or passive mode. If LACP is configured, it is in passive mode, by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode. LACP monitors the links to check whether both ends of the bundle are connected to the correct group. You can also configure the following supported LACP features on PTX10004 routers: <ul style="list-style-type: none"> <li>Distributed LACP</li> <li>Load balancing</li> <li>LAG Links or Interfaces with limited LACP capability to be up</li> </ul> </li> </ul>



Table 2: Feature Support Added to the PTX10004 Routers (*continued*)

Feature	Description
	<p>[See <a href="#">LACP</a>.]</p> <ul style="list-style-type: none"> <li>• Support for chassis environment policy and power budgeting. The functions supported are: <ul style="list-style-type: none"> <li>• Active power management</li> <li>• On over-heating, only the FRU shuts down instead of the entire chassis</li> <li>• Power control management</li> </ul> </li> </ul> <p>[See <a href="#">show chassis environment</a>.]</p> <ul style="list-style-type: none"> <li>• Load balancing under enhanced-hash-key includes: <ul style="list-style-type: none"> <li>• GRE key inclusion for transit IPv4 and IPv6 traffic</li> <li>• IP Layer 3 fields</li> <li>• IP Layer 4 fields</li> <li>• IPv6 flow label inclusion</li> <li>• MPLS labels</li> <li>• MPLS port data</li> <li>• MPLS pseudowire traffic</li> <li>• Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets</li> </ul> </li> </ul> <p>[See <a href="#">enhanced-hash-key</a>.]</p> <p>Support for resilient hashing and consistent hashing. You can configure resilient hashing on PTX10004 routers to minimize flow remapping across link aggregation groups (LAGs) or equal cost paths. Resilient hashing works in conjunction with the default static hashing algorithm. When you configure resilient hashing on LAGs, the configuration is applicable to a specific aggregated Ethernet Interface. [See <a href="#">Resilient Hashing on LAGs and ECMP groups</a>.]</p>
Layer 3 features	<ul style="list-style-type: none"> <li>• Support for Layer 3 forwarding features: <ul style="list-style-type: none"> <li>• IPv4 and IPv6</li> <li>• IPv6 accounting, includes interface statistics for explicit null, hop popping (UHP), 6PE, and BGP labeled unicast (BGP LU)</li> <li>• IPv6 subnet support on loopback interfaces</li> <li>• IPv4 and IPv6 traceroute and ping</li> <li>• IPv4 and IPv6 multicast and unicast traffic statistics</li> <li>• Junos Telemetry Interface (JTI), including routing process daemon (RPD) and RPD task memory</li> <li>• Interface ranges</li> <li>• Single and Flexible VLAN-tagged Layer 3 logical interfaces</li> <li>• Virtual router routing instances for unicast protocols</li> <li>• MPLS</li> <li>• LAG</li> <li>• ECMP</li> <li>• Path MTU discovery</li> </ul> </li> </ul>

Table 2: Feature Support Added to the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• MTU checks</li> <li>• Routing protocol next-hop acknowledgements and time intervals</li> <li>• ICMP</li> <li>• OSPF, including support for: <ul style="list-style-type: none"> <li>• LDP synchronization</li> <li>• Multitopology</li> <li>• OSPF database protection for OSPF and OSPFv3</li> <li>• OSPF export and import policies for network-summary LSAs</li> <li>• OSPFv2 and OSPFv3</li> <li>• Policy based support for LFA in OSPF</li> <li>• Remote LFA support in OSPF (IGP and LDP)</li> <li>• Shortcuts</li> <li>• SNMP MIB support for OSPFv3</li> <li>• Stub network overload and external route leaking</li> </ul> </li> <li>[See <a href="#">OSPF User Guide</a>.]</li> <li>• IS-IS, including support for: <ul style="list-style-type: none"> <li>• IPv6 and IPV6 unicast topology</li> <li>• JTI adjacency, link state database (LSDB) streaming, state, interface, adjacency statistics, and LSDB export</li> <li>• LDP synchronization</li> <li>• LFA and LFA policy based support</li> <li>• Multitopology</li> <li>• Remote LFA support for LDP in IS-IS</li> <li>• Shortcuts</li> <li>• Stub network overload</li> </ul> </li> <li>[See <a href="#">IS-IS User Guide</a>.]</li> <li>• Route-filter walkup</li> <li>• ARP</li> <li>• Static ARP</li> <li>• NDP</li> <li>• BGP</li> <li>• BFD</li> <li>• LACP</li> <li>• LDP</li> <li>• RSVP</li> </ul>

Table 2: Feature Support Added to the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• LLDP</li> <li>• TTL expiry</li> <li>• IP options</li> <li>• IP Fragmentation</li> <li>• Control plane protection against DDoS attacks</li> </ul>
MPLS	<ul style="list-style-type: none"> <li>• MPLS fast reroute (FRR) support. MPLS FRR provides faster convergence time (less than 50 milliseconds) for RSVP tunnels. The Routing Engine creates backup paths and the Packet Forwarding Engine (PFE) installs the backup path labels and next hops. You can also configure the following configuration statements at the <b>[edit protocols mpls]</b> hierarchy level:             <ul style="list-style-type: none"> <li>• <b>no-decrement-ttl</b> - Decrements the IP time-to-live (TTL) by 1 before encapsulating the IP packet within an MPLS packet.</li> <li>• <b>no-propagate-ttl</b> - Disables normal TTL decrementing for all RSVP-signaled or LDP-signaled LSPs [See <a href="#">Fast Reroute Overview</a>, <a href="#">no-decrement-ttl</a> and <a href="#">no-propagate-ttl</a>.]</li> </ul> </li> <li>• Support for segment routing. You can configure the following Source Packet Routing in Networking (SPRING) or segment routing features on the PTX10004 router:             <ul style="list-style-type: none"> <li>• MPLS ping and traceroute for single IS-IS and OSPF node or prefix segment</li> <li>• BGP-LS segment routing extensions for IS-IS and OSPF</li> <li>• IS-IS: Adjacency SID, link and node protection for IPv4 addressing in TI-LFA</li> <li>• OSPF: Inter-area support, Static adjacency SID, link and node protection in TI-LFA</li> <li>• Advertising maximum link bandwidth and administrative color without RSVP-TE configuration for IS-IS and OSPF</li> <li>• Anycast and prefix SIDs, node and link SIDs, configurable SRGB, SRMS and client, TI-LFA protection for SRMS prefixes for both IS-IS and OSPF</li> <li>• IGP adjacency SID hold time</li> <li>• PCEP for segment routing LSPs</li> <li>• BGP IPv4 labeled-unicast resolution over: BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF, non-colored and static colored IPv4 SR-TE with segment routing using IS-IS and OSPF</li> <li>• BGP Layer 3 VPN over colored and non-colored SR-TE tunnels, and IPv4 protocol next hops</li> <li>• BGP-triggered dynamic SR-TE colored tunnels</li> <li>• Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs</li> <li>• First-hop label support for SID instead of an IP address</li> <li>• Path specification using router IP addresses</li> <li>• First-hop label support for SID instead of an IP address</li> <li>• SR-TE color mode: 00—Route resolution fallback to IGP path and 01—Route resolution fallback to color only null routes</li> </ul> </li> </ul>

Table 2: Feature Support Added to the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>Static LSPs with member-link next-hops for aggregated Ethernet bundles [See <a href="#">Understanding Source Packet Routing in Networking (SPRING)</a>.]</li> </ul>
Network management and monitoring	<ul style="list-style-type: none"> <li>Support for Python, SLAX, and XSLT scripting languages and for commit scripts and macros, event policy and event scripts, op scripts, and SNMP scripts. [See <a href="#">Automation Scripting User Guide</a>.]</li> <li>Support for adding nonnative YANG modules to the Junos OS schema. [See <a href="#">Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS</a>.]</li> <li>Puppet for Junos OS support. 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 <a href="#">Puppet for Junos OS Administration Guide</a>.]</li> <li>Chef for Junos OS support. [See <a href="#">Chef for Junos OS Getting Started Guide</a>.]</li> <li>Supports third-party application Open-R.</li> </ul>
Timing and synchronization	<ul style="list-style-type: none"> <li>Support for NTP over IPv4 and IPv6. [See <a href="#">Time Management Administration Guide</a>.]</li> </ul>

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

- We've added the following features to the PTX10001-36MR in Junos OS Release Evolved 20.3R1.

Table 3: Feature Support Added to the PTX10001-36MR Routers

Feature	Description
Firewalls and policers	<ul style="list-style-type: none"> <li>Support for independent micro BFD sessions is enabled on a per member link basis for a LAG bundle. [See <a href="#">Configuring Micro BFD Sessions for LAG</a>.]</li> <li>Support for firewall filters, which include forwarding IPv4 and IPv6, load balancing, MPLS fast reroute, host path, and egress peer engineering. Nonreplication mode, segmented filters, output filter binding on routing tables, and flexible offset filters are not supported in this release. [See <a href="#">Firewall Filter Match Conditions and Actions (PTX Series Routers)</a>.]</li> <li>Support for matching ip-options in IPv4 packet headers. The <b>ip-options any</b> match condition is supported at the <b>[edit firewall family inet filter name term name from]</b> hierarchy level. [See <a href="#">Firewall Filter Match Conditions for IPv4 Traffic</a>.]</li> </ul>
High availability (HA) and resiliency	<ul style="list-style-type: none"> <li>Support for fabric link autoheal. Fabric link autoheal involves bringing down the faulty fabric link and then training it in an attempt to recover faulty fabric links from a link error condition. [See <a href="#">show chassis fabric errors autoheal</a>.]</li> </ul>

Table 3: Feature Support Added to the PTX10001-36MR Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• Support for platform resiliency. With platform resiliency, the router can detect and log errors, raise alarms, send SNMP traps, provide error indication through LEDs, self-heal, and take components out of service. [See <a href="#">show system errors active</a>.]</li> <li>• Support for Routing Engine resiliency for the faults related to CPU and FPGA. Routing Engine resiliency enables the router to perform fault-handling actions such as logging the error, raising alarms, sending SNMP traps, and the error through LEDs. [See <a href="#">show system errors active</a>.]</li> </ul>
Interfaces and chassis	<ul style="list-style-type: none"> <li>• Support for flexible tunnel interface (FTI) based decapsulation of IPv4 and IPv6 packets. You can configure IP-IP de-encapsulation on an FTI by configuring tunnel termination. Tunnel termination is completed in a single pass of packet processing to support de-encapsulation. [See <a href="#">Configuring IP-IP Decapsulation by Tunnel Termination on FTI</a>.]</li> <li>• Support for unicast reverse-path forwarding (unicast RPF). Unicast RPF helps protect against DoS and DDoS attacks by verifying the unicast source address of packets arriving on a protected interface. Packets that are not from a valid path can be discarded. [See <a href="#">Unicast Reverse-Path-Forwarding</a>.]</li> <li>• Port speed support using the interfaces hierarchy. You can configure the port speed on the PTX10001-36MR router by using the statements in the <b>[edit interfaces]</b> hierarchy. To streamline the configuration, we've migrated the port speed configuration statements from the <b>[edit chassis]</b> hierarchy to the <b>[edit interfaces]</b> hierarchy for the PTX10001-36MR router. On the PTX10001-36MR router:             <ul style="list-style-type: none"> <li>• To configure the port speed of a PIC, use the speed statement at the <b>[edit interfaces interface-name]</b> hierarchy.</li> <li>• To configure the port speed, use the speed statement and specify the number of channelized interfaces using the number-of-sub-ports statement at the <b>[edit interfaces interface-name]</b> hierarchy.</li> <li>• To control the number of interfaces (channelized or non-channelized), use the unused statement at the <b>[edit interfaces interface-name]</b> hierarchy.</li> </ul>             [See <a href="#">Port Speed</a>.]           </li> </ul>
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> <li>• Support for ON_CHANGE export of Address Resolution Protocol (ARP) and Neighbor Discovery Protocol (NDP) table states to an outside collector using JTI and remote procedure call (gRPC) services. This feature adds interface address (IPv4, IPv6) telemetry data. [See <a href="#">Guidelines for gRPC</a> and <a href="#">gNMI Sensors</a>.]</li> <li>• Support for Packet Forwarding Engine sensor. Supports UDP (native) streaming of Packet Forwarding Engine statistics using JTI. [See <a href="#">sensor (Junos Telemetry Interface)</a>.]</li> <li>• Support for physical and logical interface sensors with gRPC services, gNMI services, and UDP. You can stream statistics using JTI and gRPC services, gRPC Network Management Interface (gNMI) services, or through UDP. [See <a href="#">Guidelines for gRPC</a> and <a href="#">gNMI Sensors</a>.]</li> </ul>
Layer 2 features	<ul style="list-style-type: none"> <li>• Support for Layer 2 VPN. You can configure the following Layer 2 VPN features:             <ul style="list-style-type: none"> <li>• Transport of Layer 2 frames over MPLS (LDP signaling)</li> </ul> </li> </ul>

Table 3: Feature Support Added to the PTX10001-36MR Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• Simple Ethernet and VLAN-based CCC.</li> <li>• Local and remote switching</li> <li>• Ethernet and VLAN-based CCC and single tagged CCC logical interfaces</li> <li>• Control word</li> <li>• Regular and aggregated Ethernet interfaces</li> <li>• Layer 2 VPNs over tunnels (BGP signaling), Layer 2 protocol pass-through, Layer 2 circuit backup interface and backup neighbor, Layer 2 circuit statistics and CoS</li> <li>• Virtual circuit connectivity verification (VCCV) with type 2 and type 3</li> </ul> <p>[See <a href="#">Layer 2 VPNs and VPLS Feature Guide for Routing Devices</a>.]</p>
MPLS	<ul style="list-style-type: none"> <li>• Support for segment routing: <ul style="list-style-type: none"> <li>• MPLS ping and traceroute for single IS-IS and OSPF node or prefix segment</li> <li>• BGP-LS segment routing extensions for IS-IS and OSPF</li> </ul> </li> <li>• BGP: Binding SID for SR-TE, PRPD APIs for SR-TE policy provisioning, static SR-TE policy with mandatory color specification and without color specification</li> <li>• IS-IS: Adjacency SID, link and node protection for IPv4 addressing in TI-LFA</li> <li>• OSPF: Interarea support, Static adjacency SID, link and node protection in TI-LFA</li> <li>• Advertising maximum link bandwidth and administrative color without RSVP-TE configuration for IS-IS and OSPF</li> <li>• Anycast and prefix SIDs, node and link SIDs, configurable SRGB, SRMS and client, TI-LFA protection for SRMS prefixes for both IS-IS and OSPF</li> <li>• IGP adjacency SID hold time</li> <li>• PCEP for segment routing LSPs</li> <li>• BGP IPv4 labeled-unicast resolution over: BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF, non-colored and static colored IPv4 SR-TE with segment routing using IS-IS and OSPF</li> <li>• BGP Layer 3 VPN over colored and non-colored SR-TE tunnels, and IPv4 protocol next hops</li> <li>• BGP-triggered dynamic SR-TE colored tunnels</li> <li>• Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs</li> <li>• First-hop label support for SID instead of an IP address</li> <li>• Path specification using router IP addresses</li> <li>• SR-TE color mode: 00—route resolution fallback to IGP path 01—route resolution fallback to color only null routes</li> <li>• Static LSPs with member-link next hops for aggregated Ethernet bundles</li> </ul> <p>[See <a href="#">Understanding Source Packet Routing in Networking (SPRING)</a>.]</p>

Table 3: Feature Support Added to the PTX10001-36MR Routers (*continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>• Support for monitoring, detecting threats, and taking action on degraded physical Ethernet links to minimize packet loss. You can monitor link degradation of the 10GbE, 25GbE, 40GbE, 100GbE, and 400GbE interfaces. [See <a href="#">link-degrade-monitor</a>.]</li> <li>• Support for destination class usage (DCU) accounting and source class usage (SCU) accounting. SCU accounting provides breakdown of output interface traffic statistics that originates from specific prefixes. DCU accounting provides breakdown of input interface traffic statistics that is destined for specific prefixes. [See <a href="#">Understanding Source Class Usage and Destination Class Usage Options</a>.]</li> <li>• Support for LACP features. LACP monitors the links to check whether both ends of the bundle are connected to the correct group. [See <a href="#">Configuring Aggregated Ethernet Interfaces</a>.]</li> <li>• Support for load balance properties such as IPv6 flow label, GTP hashing and weighted ECMP. [See <a href="#">Load Balancing and Ethernet Link Aggregation</a>.]</li> <li>• Support for local port mirroring. Local port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. [See <a href="#">Understanding Port Mirroring</a>.]</li> <li>• Support for passive monitoring, which allows the device to accept and monitor traffic on the interface and forward those packets to monitoring tools such as IDS servers and packet analyzers, or other devices such as routers or end-node hosts. [See <a href="#">Passive Monitoring</a>.]</li> <li>• Support for scalability for LDP over RSVP and BGP-LU services. This feature enhances rpd to produce the chain nexthop for various services. The rpd is enhanced to build a translation layer between RIB and FIB to segment transport tunnels. [See <a href="#">Tunneling LDP LSPs in RSVP LSPs Overview</a>.]</li> <li>• Support for BGPLU FRR MPLS inter-AS link protection. Link protection ensures traffic restoration in case of an interface failure. [See <a href="#">BGP User Guide</a>.]</li> <li>• Support for higher scale of transit LSPs with link protection. Link protection enables the PTX10001-36MR to support a higher scale of transit RSVP-TE signaled MPLS LSPs. [See <a href="#">Link Protection for MPLS LSPs</a>.]</li> <li>• Support for IPv6. Junos OS Evolved supports the following IPv6 features: NDP, path MTU, traceroute, ping, neighbor discovery, router advertisements, stateless autoconfiguration, and link aggregation. [See <a href="#">IPv6 Neighbor Discovery User Guide</a>.]</li> </ul>
Services applications	<ul style="list-style-type: none"> <li>• Support for inline active flow monitoring. Inline active flow monitoring defines a flow record template for IPv4 or IPv6 traffic. Both IPFIX and Version 9 template formats are supported with output filters. [See <a href="#">Configuring Inline Active Flow Monitoring</a>.]</li> <li>• Support for real-time performance monitoring (RPM). RPM sends out probes to the network to give a measure of network performance. The scope of support is limited to probe generation and reception (client) as well as reflection (server) for the following RPM probe types: icmp-ping, icmp-timestamp, udp-ping, and udp-timestamp. Also supported are probe history management as well as reporting through syslog only. For Junos OS Evolved, RPM is configured at the <a href="#">edit services</a></li> </ul>

Table 3: Feature Support Added to the PTX10001-36MR Routers (*continued*)

Feature	Description
	<b>monitoring rpm</b> ] hierarchy level. [See <a href="#">Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.</a> ]

- **New PTX10K-LC1202-36MR line card (PTX10008)**—In Junos OS Evolved Release 20.3R1, we introduce the PTX10K-LC1202-36MR, a new line card which plugs into the PTX10008 router horizontally at the front of the chassis. The PTX10K-LC1202-36MR is a 36-port line card designed to support a pure 100-Gbps port speed configuration or a mixed port speed configuration of 100 Gbps and 400 Gbps. This line card, which provides a line rate throughput of 4.8 Tbps, has thirty-two QSFP28 ports, each capable of supporting a speed of 100 Gbps, and four QSFP56-DD ports, each capable of supporting a speed of up to 400 Gbps. [See [PTX10008 Line Card Components and Descriptions.](#)]

To install the PTX10008 router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see [PTX10008 Packet Transport Router Hardware Guide](#).

[Table 4 on page 16](#) summarizes the PTX10K-LC1202-36MR features supported in Junos OS Evolved Release 20.3R1.

Table 4: Feature Support Added to the PTX10K-LC1202-36MR

Feature	Description
BGP	<ul style="list-style-type: none"> <li>• Support of redistribution of IPv4 routes with IPv6 next hop into BGP. Devices running Junos OS Evolved, can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, IPv4 traffic is forwarded from CPE devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop. To configure a dynamic IPv4-over-IPv6 tunnel, include the dynamic-tunnels statement at the <b>[edit routing-options]</b> hierarchy level. [See <a href="#">Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.</a>]</li> </ul>
Class of service (CoS)	<ul style="list-style-type: none"> <li>• Support for CoS features. The COS features include classifiers (behavior aggregate (BA), fixed, and multifield (MF)), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, and drop-profiles. [See <a href="#">CoS Features and Limitations on PTX Series Routers.</a>]</li> <li>• Classification override support. Classification override is configured under a forwarding policy on PTX10K-LC1202-36MR line cards on PTX10008 devices. [See <a href="#">CoS Features and Limitations on PTX Series Routers</a> and <a href="#">Overriding the Input Classification.</a>]</li> <li>• Support for CoS-based forwarding. You can configure CoS-based forwarding (CBF) on PTX10K-LC1202-36MR line cards for PTX10008 devices. [See <a href="#">Forwarding Policy Options Overview.</a>]</li> </ul>
DHCP	<ul style="list-style-type: none"> <li>• Support for DHCPv4 and DHCPv6 stateless relay. [See <a href="#">DHCP Relay Agents.</a>]</li> </ul>



Table 4: Feature Support Added to the PTX10K-LC1202-36MR (*continued*)

Feature	Description
Firewalls and policers	<ul style="list-style-type: none"> <li>Support for IPv4 and IPv6 firewall filters. Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on the router from a source address to a destination address. Firewall filter features include: <ul style="list-style-type: none"> <li>Firewall filter</li> <li>DCU accounting and SCU accounting</li> </ul> </li> </ul> <p>[See <a href="#">Firewall Filter Match Conditions and Actions (PTX Series Routers)</a>.]</p>
High availability (HA) and resiliency	<ul style="list-style-type: none"> <li>Fabric resiliency support. Resiliency includes handling the faults related to fabric links (the links between the PTX10K-LC1202-36MR line card and the switch fabric JNP10008-SF3). Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See <a href="#">show system errors active</a>.]</li> <li>Support for micro BFD sessions and child links of aggregated Ethernet or LAG bundle with interface IP addresses. [See <a href="#">Configuring Micro BFD Sessions for LAG</a>.]</li> <li>Support for resilient hashing. You can configure resilient hashing to minimize flow remapping across link aggregation groups (LAGs) or equal cost paths. Resilient hashing works in conjunction with the default static hashing algorithm. When you configure resilient hashing on LAGs, the configuration is applicable to a specific aggregated Ethernet Interface. You can also configure consistent hashing across ECMP groups by using BGP multipath. You can configure a maximum of 64 member links for resilient hashing on LAG and consistent hashing on BGP. [See <a href="#">Resilient Hashing on LAGs and ECMP</a> and <a href="#">Load Balancing for a BGP Session</a>.]</li> <li>Platform resiliency support. Resiliency includes handling faults pertaining to the line card hardware and transceivers. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See <a href="#">show system errors active</a>.]</li> </ul>
Interfaces and chassis	<ul style="list-style-type: none"> <li>Support for FTI based decapsulation of IPv4 and IPv6 packets. You can configure IP-IP decapsulation on a flexible tunnel interface by configuring tunnel termination at the <b>[edit interfaces fti0 unit <i>number</i> tunnel encapsulation IPIP]</b> hierarchy level. Tunnel termination is completed in a single pass of packet processing to support decapsulation. [See <a href="#">Configuring GRE Tunnel Interfaces</a>.]</li> <li>GRE tunnel decapsulation using flexible tunnel interfaces (FTI). You can configure GRE tunnels on flexible tunnel interfaces (FTIs). When you enable the <b>tunnel-termination</b> statement at the <b>[edit interfaces fti0 unit <i>unit-number</i>]</b> CLI hierarchy, tunnels are terminated on the WAN interface before any other actions, such as sampling, port mirroring, or filtering, are applied. [See <a href="#">Configuring GRE Tunnel Interfaces</a>.]</li> <li>You can monitor link degradation of Gigabit Ethernet interfaces. To enable your device to monitor the links, use the <b>link-degrade-monitor</b> statement at the <b>[edit interfaces <i>interface-name</i>]</b> hierarchy level. [See <a href="#">Link Degrade Monitoring Overview</a>.]</li> <li>Support for Media Access Control Security (MACsec), including AES-256 encryption, extended packet numbering, and fail-open mode. [See <a href="#">Configuring Media Access Control Security (MACsec) on Routers</a>.]</li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> <li>Channelized interface support. The default port speed is 400 Gbps for ports 4, 10, 24 and 30 and 100 Gbps for the remaining ports. Each PTX10K-LC1202-36MR line card provides a maximum bandwidth of 4.8Tbps.</li> <li>To operate at 3.6T mode, you can configure all 36 ports to 100 Gbps speed. To operate at 4.8T mode, you can configure four 400GE ports (4, 10, 24, and 30) to operate at 400 Gbps speed and the remaining ports to operate at 100 Gbps speed.</li> <li>You can configure the 400GE capable ports to operate at 400 Gbps, 4x100 Gbps, 2x100 Gbps, 4x10 and 8x25Gbps speeds. You can configure the 100GE capable ports to operate at 100 Gbps, 4x25 Gbps, and 4x10 Gbps speeds.</li> </ul> <p>[See <a href="#">Port Speed on PTX10K-LC1202-36MR Overview</a>.]</p> <ul style="list-style-type: none"> <li>Support for the dedicated virtual routing and forwarding (VRF) instance mgmt_junos for the management Ethernet Interface. [See <a href="#">Management Interface in a Nondefault Instance and management-instance</a>.]</li> <li>Support for Link Aggregation Control Protocol (LACP) features. You can configure LACP to bundle several physical interfaces to form one logical aggregated Ethernet interface. By default, LACP is disabled on aggregated Ethernet interfaces. You can configure LACP in active or passive mode. If LACP is configured, it is in passive mode by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode. LACP monitors the links to check whether both ends of the bundle are connected to the correct group. You can also configure the following supported LACP features: <ul style="list-style-type: none"> <li>Distributed LACP</li> <li>Load balancing</li> <li>Telemetry support (periodic streaming and on-change)</li> </ul> </li> </ul> <p>[See <a href="#">LACP</a>.]</p> <ul style="list-style-type: none"> <li>Load balancing under enhanced-hash-key includes: <ul style="list-style-type: none"> <li>GRE key inclusion for transit IPv4 and IPv6 traffic</li> <li>IP Layer 3 fields</li> <li>IP Layer 4 fields</li> <li>IPv6 flow label inclusion</li> <li>MPLS labels</li> <li>MPLS port data</li> <li>MPLS pseudowire traffic</li> <li>Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets</li> </ul> </li> </ul> <p>[See <a href="#">enhanced-hash-key</a>.]</p> <ul style="list-style-type: none"> <li>Support for QSFP28 100GE DWDM optical transceiver module (part number: IN-Q2AY2-XX, where XX represents the channel ID) for 100GbE applications. This transceiver is shown in the show chassis hardware output as QSFP-100GE-DWDM2. The transceiver is suited for data center</li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (continued)

Feature	Description
	<p>interconnect which requires high-fiber-capacity for links up to 80 km. The features supported by QSFP28 100GE DWDM transceiver include the following:</p> <ul style="list-style-type: none"> <li>• View the optics inventory information. [See <a href="#">show chassis hardware</a> and <a href="#">show chassis pic fpc-slot slot pic-slot</a></li> <li>• View the diagnostics data, warnings, and alarms for interfaces. [See <a href="#">show interfaces diagnostics optics</a>.]</li> <li>• Configure the optics loopback mode. [See <a href="#">optics-options</a>.]</li> </ul>
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> <li>• JTI for Packet Forwarding Engine statistics. JTI supports export of Packet Forwarding Engine statistics using remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. Using gRPC or gNMI and JTI, data is streamed from devices to outside collectors at configurable intervals. Export the following statistics: <ul style="list-style-type: none"> <li>• Packet Forwarding Engine statistics (resource path <code>/junos/system/linecard/packet/usage</code>)</li> <li>• Network processing unit (NPU) memory statistics (resource path <code>/junos/system/linecard/npu/memory/</code>)</li> <li>• NPU utilization statistics (resource path <code>/junos/system/linecard/npu/utilization/</code>)</li> </ul> [See <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a>.] </li> <li>• JTI for Routing Engine sensors. JTI supports export to an outside collector of Routing Engine core statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Use the following resource paths (sensors): <ul style="list-style-type: none"> <li>• Routing Engine logical interface statistics (resource path <code>/interfaces/interface/</code>). Applicable for management interfaces (re0:mgmt-0 and re1:mgmt-0).</li> <li>• Address Resolution Protocol (ARP) (resource path <code>/arp-information/</code>)</li> <li>• Network Discovery Protocol (NDP) (resource path <code>/nd6-information/</code>)</li> <li>• Operational state of AE with LACP (resource path <code>/lACP/</code>)</li> <li>• Operational state of LLDP (neighbors information) (resource path <code>/lldp/</code>)</li> <li>• Sensor for NDP router-advertisement statistics (resource path <code>/ipv6-ra/</code>)</li> </ul> [See <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a>.] </li> <li>• JTI support for physical and logical interface sensors with gRPC services, gNMI services, and UDP. 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: <ul style="list-style-type: none"> <li>• <code>/junos/system/linecard/interface/</code></li> <li>• <code>/junos/system/linecard/interface/logical/usage/</code></li> </ul> [See <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a> and <a href="#">sensor (Junos Telemetry Interface)</a>.] </li> <li>• UDP (native) Packet Forwarding Engine sensor support for JTI. UDP streaming for these sensors is not supported over the management interface. The sensors supported are:</li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• Packet Forwarding Engine statistics. This sensor exports statistics for counters and provides visibility into Packet Forwarding Engine error and drop statistics (sensor name is <code>/junos/system/linecard/packet/usage/</code>).</li> <li>• 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 <code>/junos/system/linecard/npu/memory</code>).</li> <li>• NPU utilization statistics. This sensor exports Packet Engine utilization and packet load (sensor name is <code>/junos/system/linecard/npu/utilization/</code>).</li> </ul> <p>[See <a href="#">sensor (Junos Telemetry Interface)</a>.]</p>
Layer 2 features	<ul style="list-style-type: none"> <li>• Support for Link Layer Discovery Protocol (LLDP). LLDP support includes management interfaces and disabling of LLDP time, length, and value (TLV) messages. [See <a href="#">LLDP Overview</a>.]</li> <li>• Support for restricted and unrestricted ARP. By default, proxy ARP is unrestricted. Proxy ARP is supported on et and ae interfaces and interface routes only. [See <a href="#">Restricted and Unrestricted Proxy ARP Overview</a>.]</li> <li>• Support for Layer 2 VPNs. You can configure the following Layer 2 VPNs: <ul style="list-style-type: none"> <li>• Transport of Layer 2 frames over MPLS (LDP signaling)</li> <li>• Layer 2 VPNs over tunnels (BGP signaling)</li> <li>• Simple Ethernet and VLAN based cross connect (also known as connections)</li> <li>• Local and remote switching</li> <li>• Ethernet and VLAN-CCC</li> <li>• Single tagged CCC logical interfaces</li> <li>• Control word</li> <li>• Regular and aggregated Ethernet interfaces</li> <li>• Layer 2 protocol pass-through</li> <li>• Layer 2 circuit backup interface and backup neighbor</li> <li>• Layer 2 circuit statistics and CoS</li> <li>• VCCV with type 2 and type 3</li> <li>• Layer 2 circuit statistics and CoS</li> </ul> </li> </ul> <p>[See <a href="#">Layer 2 VPNs and VPLS Feature Guide for Routing Devices</a>.]</p>
Layer 3 features	<ul style="list-style-type: none"> <li>• Support for Layer 3 forwarding features: <ul style="list-style-type: none"> <li>• IPv4 and IPv6</li> <li>• IPv6 accounting, includes interface statistics for explicit null, hop popping (UHP), 6PE, and BGP labeled unicast (BGP LU)</li> <li>• IPv6 subnet support on loopback interfaces</li> <li>• IPv4 and IPv6 traceroute and ping</li> </ul> </li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• IPv4 and IPv6 multicast and unicast traffic statistics</li> <li>• Junos Telemetry Interface (JTI), including routing process daemon (RPD) and RPD task memory</li> <li>• Interface ranges</li> <li>• Single and Flexible VLAN-tagged Layer 3 logical interfaces</li> <li>• Virtual router routing instances for unicast protocols</li> <li>• MPLS</li> <li>• LAG</li> <li>• ECMP</li> <li>• Path MTU discovery</li> <li>• MTU checks</li> <li>• Routing protocol next-hop acknowledgements and time intervals</li> <li>• ICMP</li> <li>• OSPF, including support for: <ul style="list-style-type: none"> <li>• LDP synchronization</li> <li>• Multitopology</li> <li>• OSPF database protection for OSPF and OSPFv3</li> <li>• OSPF export and import policies for network-summary LSAs</li> <li>• OSPFv2 and OSPFv3</li> <li>• Policy based support for LFA in OSPF</li> <li>• Remote LFA support in OSPF (IGP and LDP)</li> <li>• Shortcuts</li> <li>• SNMP MIB support for OSPFv3</li> <li>• Stub network overload and external route leaking</li> </ul> <p>[See <a href="#">OSPF User Guide</a>.]</p> </li> <li>• IS-IS, including support for: <ul style="list-style-type: none"> <li>• IPv6 and IPV6 unicast topology</li> <li>• JTI adjacency, link state database (LSDB) streaming, state, interface, adjacency statistics, and LSDB export</li> <li>• LDP synchronization</li> <li>• LFA and LFA policy based support</li> <li>• Multitopology</li> <li>• Remote LFA support for LDP in IS-IS</li> <li>• Shortcuts</li> <li>• Stub network overload</li> </ul> <p>[See <a href="#">IS-IS User Guide</a>.]</p> </li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> <li>• Route-filter walkup</li> <li>• ARP</li> <li>• Static ARP</li> <li>• NDP</li> <li>• BGP</li> <li>• BFD</li> <li>• LACP</li> <li>• LDP</li> <li>• RSVP</li> <li>• LLDP</li> <li>• TTL expiry</li> <li>• IP options</li> <li>• IP Fragmentation</li> <li>• Control plane protection against DDoS attacks</li> </ul>
Layer 3 VPNs	<ul style="list-style-type: none"> <li>• Support for Layer 3 VPNs. You can configure MPLS-based Layer 3 VPNs. You can also configure Layer 3 IPv6 VPN Provider Edge router (6VPE) over MPLS. The features supported include: <ul style="list-style-type: none"> <li>• MPLS over Layer 3 VLAN-tagged subinterfaces</li> <li>• Per-next-hop label allocation</li> <li>• 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</li> <li>• ICMP tunneling and MPLS traceroute</li> <li>• Disabling time-to-live (TTL) decrementing using no-propagate-ttl</li> <li>• IPv6-over-MPLS using: <ul style="list-style-type: none"> <li>• IPv6 VPN Provider Edge routers (6VPE)</li> <li>• IPv6 Provider Edge routers (6PE)</li> </ul> </li> </ul> <p>6PE and 6VPE enable you to multiplex IPv6 as a service across an IPv4-only MPLS core using dual-stack PE routers. BGP is used to distribute IPv6 routing-information and IPv4 signaled MPLS LSPs are used to forward the IPv6 traffic across an IPv6-free-core.</p> <p>[See <a href="#">Layer 3 VPNs Feature Guide for Routing Devices</a> and <a href="#">Tunneling IPv6 Traffic over MPLS IPv4 Networks</a>.]</p> </li> </ul>
MPLS	<ul style="list-style-type: none"> <li>• Support for segment routing. You can configure the following Source Packet Routing in Networking (SPRING) or segment routing features: <ul style="list-style-type: none"> <li>• MPLS ping and traceroute for single IS-IS and OSPF node or prefix segment</li> <li>• BGP-LS segment routing extensions for IS-IS and OSPF</li> <li>• IS-IS: Adjacency SID, link and node protection for IPv4 addressing in TI-LFA</li> <li>• OSPF: Inter-area support, Static adjacency SID, link and node protection in TI-LFA</li> </ul> </li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> <li>• Advertising maximum link bandwidth and administrative color without RSVP-TE configuration for IS-IS and OSPF</li> <li>• Anycast and prefix SIDs, node and link SIDs, configurable SRGB, SRMS and client, TI-LFA protection for SRMS prefixes for both IS-IS and OSPF</li> <li>• IGP adjacency SID hold time</li> <li>• PCEP for segment routing LSPs</li> <li>• BGP IPv4 labeled-unicast resolution over: BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF, non-colored and static colored IPv4 SR-TE with segment routing using IS-IS and OSPF</li> <li>• BGP Layer 3 VPN over colored and non-colored SR-TE tunnels, and IPv4 protocol next hops</li> <li>• BGP-triggered dynamic SR-TE colored tunnels</li> <li>• Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs</li> <li>• First-hop label support for SID instead of an IP address</li> <li>• Path specification using router IP addresses</li> <li>• First-hop label support for SID instead of an IP address</li> <li>• SR-TE color mode: 00—Route resolution fallback to IGP path and 01—Route resolution fallback to color only null routes</li> <li>• Static LSPs with member-link next-hops for aggregated Ethernet bundles</li> </ul> <p>[See <a href="#">Understanding Source Packet Routing in Networking (SPRING)</a>.]</p> <ul style="list-style-type: none"> <li>• MPLS support. You can configure the following MPLS features: <ul style="list-style-type: none"> <li>• CLI support for monitoring MPLS label usage</li> <li>• Inline MPLS and IPv6 lookup for explicit null</li> <li>• Explicit null support for MPLS LSPs</li> <li>• MPLS Label Block Configuration</li> <li>• MPLS over untagged Layer 3 interfaces</li> <li>• MPLS OAM - LSP ping</li> <li>• JTI: OCST: MPLS operational state streaming (v2.2.0)</li> <li>• Entropy Label support</li> <li>• MPLS: JTI: Junos Telemetry Interface</li> <li>• MPLS self-ping, TE++, and misc augmentation</li> <li>• LDP including: <ul style="list-style-type: none"> <li>• Configurable label withdraw delay</li> <li>• Egress policy</li> <li>• Explicit null</li> <li>• Graceful restart signalling</li> <li>• IGP synchronization</li> </ul> </li> </ul> </li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• Ingress policy</li> <li>• IPv6 for LDP transport session</li> <li>• Strict targeted hellos</li> <li>• Track IGP metric</li> <li>• Tunneling (LDP over RSVP)</li> <li>• RSVP++</li> <li>• RSVP-TE including: <ul style="list-style-type: none"> <li>• Bypass LSP static configuration</li> <li>• Ingress LSP statistics in a file</li> <li>• RSVP-TE Hitless-MBB with no artificial delays</li> <li>• Auto bandwidth</li> <li>• Class based forwarding (CBF) with 16 classes</li> <li>• CBF with next-hop resolution</li> <li>• Convergence and scalability</li> <li>• Graceful restart signalling</li> <li>• JTI interface statistics and LSP event export</li> <li>• LSP next-hop policy</li> <li>• LSP self ping</li> <li>• MPLS fast reroute (FRR)</li> <li>• LSP self ping</li> <li>• Optimize adaptive teardown</li> <li>• Node/link protection</li> <li>• Refresh reduction</li> <li>• Soft preemption</li> <li>• Shared Risk Link Group (SRLG)</li> </ul> </li> <li>• Static LSPs with IPv4 next-hop, IPv6 next-hop, and IPv6 next-hop with next-table support for bypass</li> <li>• Traffic engineering including: TE++: Dynamic ingress LSP splitting, Traffic engineering extensions (OSPF-TE and ISIS-TE), Traffic engineering options: bgp, bgp-igp, bgp-igp-both-ribs, and mpls-forwarding</li> </ul> <p>[See <a href="#">MPLS Applications User Guide</a>.]</p>
Multicast	<ul style="list-style-type: none"> <li>• Support for IPv4 and IPv6 multicast, including Multicast Source Discovery Protocol (MSDP). This release also includes support for PIM-SM as the first hop router (FHR) or last hop router (LHR), and anycast, static, or local rendezvous point (RP). [See <a href="#">Multicast Overview</a>.]</li> </ul>



Table 4: Feature Support Added to the PTX10K-LC1202-36MR (*continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>Support for destination class usage (DCU) accounting and source class usage (SCU) accounting. SCU accounting provides breakdown of output interface traffic statistics that originates from specific prefixes. DCU accounting provides breakdown of input interface traffic statistics that is destined for specific prefixes.</li> </ul> <p><b>NOTE:</b> On PTX10K- LC1202-36MR line cards, when you assign SCU and DCU classes to system prefixes, they occupy more space in the forwarding information base (FIB) tables than regular routes. Limit the number of system prefixes that have a non-default class assigned to them. [See <a href="#">Understanding Source Class Usage and Destination Class Usage Options</a>.]</p> <ul style="list-style-type: none"> <li>Support for running third-party applications inside Docker containers. [See <a href="#">Overview of Third-Party Applications on Junos OS Evolved</a>.]</li> <li>Support for third-party application Open-R.</li> <li>sFlow support. The sFlow technology is a monitoring technology for high-speed switched or routed networks. You can configure the sFlow technology to sample, detect, and report MPLS traffic. The following sFlow features are supported: <ul style="list-style-type: none"> <li>Adaptive sampling</li> <li>Adaptive sampling fallback which decreases the sampling load when the traffic load decreases after adaptive sampling has taken place</li> <li>Ability to configure forwarding class and DSCP values per collector</li> <li>Interface statistics reporting</li> <li>sFlow ingress and egress sampling</li> <li>sFlow ingress and egress sampling on the Ultimate-hop-popping (UHP) node</li> </ul> [See <a href="#">Understanding How to Use sFlow Technology for Network Monitoring</a> and <a href="#">SFlow</a>.] </li> <li>Local port mirroring support. You can use port mirroring to copy packets entering or exiting a port or entering a VLAN and to send the copies to a local interface for local monitoring. 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: <ul style="list-style-type: none"> <li>Interface filter on ingress</li> <li>Interface filter on egress</li> <li>Forwarding table filter (FTF) on ingress</li> <li>Unicast traffic on ingress and egress</li> <li>Multicast traffic on ingress and egress</li> <li>Families <b>inet</b> and <b>inet6</b></li> <li>Firewall filter action <b>port-mirroring</b> in both ingress and egress directions</li> <li>Aggregated Ethernet interfaces at both ingress and egress</li> <li>Sampling range of 1-16000000</li> <li>Packet clipping (<b>maximum-packet-length</b>)</li> </ul> </li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> <li>• <b>run-length</b> attribute to specify the number of samples following trigger events</li> </ul> <p>Use the following CLI hierarchies to configure port mirroring:</p> <ul style="list-style-type: none"> <li>• <b>[edit interfaces]</b></li> <li>• <b>[edit forwarding-options port-mirroring]</b></li> <li>• <b>[edit firewall filter]</b></li> </ul> <p>You can configure <b>family inet</b> and <b>family inet6</b> in the <b>[edit interfaces]</b> and the <b>[edit forwarding-options port-mirroring]</b> hierarchies for this feature. The feature applies to global port mirroring only.</p> <p>[See <a href="#">Understanding Port Mirroring and Analyzers</a>.]</p>
Routing protocols	<ul style="list-style-type: none"> <li>• Support for BGP-LU FRR MPLS inter-AS link protection. Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. For an MPLS inter-AS environment, you can enable link protection when labeled-unicast is used to send traffic between autonomous systems (ASs). The ingress router then chooses an alternate link through another interface to send traffic to its destination. Starting in Junos Evolved Release 20.1R2, to configure link protection on an interface, you can include the <b>protection</b> statement at the <b>[edit protocols bgp group group-name family inet labeled-unicast]</b> hierarchy level. Junos OS Evolved supports the following baseline BGP features: <ul style="list-style-type: none"> <li>• Large communities</li> <li>• Link bandwidth community aggregation</li> <li>• Load balancing in proportion to available bandwidth community</li> <li>• <b>local-as</b> option</li> <li>• BGP Monitoring Protocol (BMP) version 3</li> <li>• MTU discovery</li> <li>• BGP multipath</li> <li>• Multiple paths to IPv6 addresses</li> <li>• PIC core and inter-AS link protection for BGP-LU</li> <li>• Prefix prioritization</li> <li>• Route reflection</li> <li>• <b>tcp-mss</b> configuration</li> <li>• Route leaking between routing instances, including default (inet) routing instance through RIB-groups and instance import/export</li> </ul> </li> </ul>
Services applications	<ul style="list-style-type: none"> <li>• Support for real-time performance monitoring (RPM). RPM sends out probes to the network to give a measure of network performance. The scope of support is limited to probe generation and reception (client) as well as reflection (server) for the following RPM probe types: icmp-ping, icmp-timestamp, udp-ping, and udp-timestamp. Also supported are probe history management as well as reporting through syslog only. For Junos OS Evolved, RPM is configured at the <b>[edit services</b></li> </ul>

Table 4: Feature Support Added to the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<p><b>monitoring rpm]</b> hierarchy level. [See <a href="#">Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.</a>]</p> <ul style="list-style-type: none"> <li>• Support for inline active flow monitoring (J-Flow) for IPv4, IPv6, and MPLS traffic. Both IPFIX and Version 9 template formats are supported. [See <a href="#">Understanding Inline Active Flow Monitoring.</a>]</li> </ul>
System management	<ul style="list-style-type: none"> <li>• Support for passive monitoring. When you enable passive monitoring on the interfaces, 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 <a href="#">[Passive Monitoring.]</a></li> </ul>
Timing and synchronization	<ul style="list-style-type: none"> <li>• Support for NTP over IPv4 and IPv6.</li> </ul>

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

- We've added the following features to the PTX10008 for JNP10008-SF3 fabric in Junos OS Evolved Release 20.3R1.

**Table 5: Feature Support Added to the PTX10008 Routers**

Feature	Description
Hardware	<ul style="list-style-type: none"> <li>• This release introduces a limited encryption version Routing Engine, the JNP10K-RE1-ELT. This Routing Engine supports 400-Tbps line cards and JNP10008-SF3 switch fabric. The Routing Engine runs a Junos OS Evolved limited image that does not support data plane encryption and is intended only for the countries of the Eurasian Customs Union (EACU). [See <a href="#">PTX10008 Routing and Control Boards Components and Descriptions</a>.]</li> <li>• Two new configuration models, and PTX10008-PREM2, are now available for sites that do not require a fully-populated chassis. These configurations allow the PTX10008 to operate with fewer Switch Interface Boards (SIBs). The configurations consist of:             <ul style="list-style-type: none"> <li>• PTX10008-BASE3—One routing engine, two fan trays, two fan tray controllers, six JNP10K-PWR-AC2 or JNP10K-PWR-DC2 power supplies, three SIBs, three SIB covers, and eight line-card slot covers.</li> <li>• PTX10008-PREM2—Two routing engines, two fan trays, two fan tray controllers, six JNP10K-PWR-AC2 or JNP10K-PWR-DC2 power supplies, four SIBs, two SIB covers, and eight line-card slot covers.</li> </ul> </li> </ul> <p>[See <a href="#">PTX10008 Configurations and Upgrade Options</a>.]</p>
High availability (HA) and resiliency	<ul style="list-style-type: none"> <li>• Platform resiliency enables the router to handle failures and faults related to the hardware components such as line cards, switch fabric, control boards, fan trays, fan tray controllers, and power supply units. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See <a href="#">show system errors active</a>.]</li> <li>• Support for VRRP. The following features are not supported for VRRP on Junos OS Evolved:             <ul style="list-style-type: none"> <li>• ISSU</li> <li>• Proxy ARP</li> <li>• MC-LAG</li> <li>• IRB</li> </ul> </li> </ul> <p>VRRP sessions might flap during GRES in centralized mode. [See <a href="#">Understanding VRRP</a>.]</p>
Interfaces and chassis	<ul style="list-style-type: none"> <li>• Support to upgrade the optic drivers on the PTX10K-LC1201 line card without a full Junos OS Evolved upgrade. You can upgrade the optics drivers by running the <b>request system software add package_name</b> command. [See <a href="#">Hardware Supported by Junos Continuity Software</a>.]</li> <li>• Unicast RPF support for both IPv4 and IPv6 traffic flows. [See <a href="#">Example: Configuring Unicast Reverse-Path-Forwarding Check</a>.]</li> <li>• Support for configuring GRE tunnel encapsulation on FTIs using the loopback interface. You can configure encapsulation by using the command <b>tunnel encapsulation gre source address destination address</b> at the <b>[edit interfaces fti0 unit unit ]</b> hierarchy. Keep in mind the following when configuring this feature:</li> </ul>

Table 5: Feature Support Added to the PTX10008 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• Adding <a href="#">tunnel-termination</a> makes the tunnel decap-only tunnel and encapsulation will be disabled.</li> <li>• Both the source and destination address is mandatory when you don't configure the tunnel-termination command.</li> <li>• Configuring a variable prefix mask on the source address isn't allowed.</li> </ul> <p>[See <a href="#">Tunnel and Encryption Services Interfaces User Guide for Routing Devices</a>.]</p> <ul style="list-style-type: none"> <li>• Native VLAN ID on Layer 3 interfaces enables the logical interface whose VLAN ID matches the native VLAN ID configured for that interface to accept untagged packets as well as tagged packets. The same logical interface with native VLAN ID enabled ensures that any packet going out of that interface does not have a tag attached. Packets can be outbound control packets or transit data packets. [See <a href="#">native-vlan-id</a>.]</li> <li>• Support for the discard interface. [See <a href="#">Discard Interfaces</a>.]</li> <li>• Support for the following transceivers: <ul style="list-style-type: none"> <li>• QSFP-100G-FR—These transceivers interoperate with the QDD-4X100G breakout optics. For example, the QDD-4X100G-FR interconnects with up to four QSFP-100G-FR transceivers. The QSFP-100G-FR transceivers interconnect in single links (QSFP-100G-FR to QSFP-100G-FR or to QSFP-100G-DR) and interoperate at the shortest link length.</li> <li>• QSFP-100G-DR—These transceivers interoperate with 400-Gbps breakout optics. For example, the QDD-400G-DR4 interconnects with up to four QSFP-100G-DR transceivers. The QSFP-100G-DR transceivers interconnect in single links (QSFP-100G-DR to QSFP-100G-DR or to QSFP-100G-FR) and interoperate at the shortest link length.</li> <li>• QSFP-100G-LR—These transceivers interoperate with the QDD-4X100G breakout optics. For example, the QDD-4X100G-LR interconnects with up to four QSFP-100G-LR transceivers. The QSFP-100G-LR transceivers interconnect in single links (QSFP-100G-LR to QSFP-100G-LR or to QSFP-100G-FR) and interoperate at the shortest link length.</li> </ul> <p><b>NOTE:</b> These transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).</p> <p>[See the <a href="#">Hardware Compatibility Tool (HCT)</a> for details.]</p> <li>• Support for adaptive load balancing (ALB) on multiple Packet Forwarding Engines for aggregated Ethernet bundles.</li> </li></ul> <p><b>NOTE:</b> With ALB configured, the <b>show interfaces ae0 extensive</b> command displays adaptive statistics information under the physical interface, not the logical interface.</p> <p>[See <a href="#">Load Balancing on Aggregated Ethernet Interfaces</a>.]</p> <ul style="list-style-type: none"> <li>• Aggregated Ethernet interfaces support mixed rates and mixed modes. The aggregated Ethernet supports member links of different modes (WAN and LAN) for 10GbE links and member links of different rates for WAN and LAN aggregated Ethernet bundles. [See <a href="#">Configuration Guidelines for Aggregated Ethernet Interfaces</a>.]</li> </ul>

Table 5: Feature Support Added to the PTX10008 Routers (*continued*)

Feature	Description
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> <li>Support for export of physical interface queue statistics to an outside collector using UDP (native) streaming, remote procedure call (gRPC) services, or gRPC network management interface (gNMI) services.</li> </ul> <p>To export statistics through UDP, gRPC, or gNMI, use the sensor <code>/junos/system/linecard/interface/queue/</code>.</p> <p>Each physical interface has 8 queues. The following counters are exported as part of this sensor for all configured physical interfaces:</p> <ul style="list-style-type: none"> <li>Transmitted packets and transmitted bytes</li> <li>Red drop packets and bytes</li> <li>Tail drop packets and bytes</li> </ul> <p>This feature includes zero suppression support. It does not include support for summed up counters on aggregated Ethernet (ae) interfaces.</p> <p>[See <a href="#">sensor (Junos Telemetry Interface)</a> and <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a>.]</p> <ul style="list-style-type: none"> <li>Supports ON_CHANGE export of ARP and NDP table states to an outside collector using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Junos OS Evolved already supports ARP and NDP periodic streaming and ON_CHANGE. This feature adds interface address (IPv4, IPv6) telemetry data.</li> </ul> <p>The supported resource paths (sensors) are:</p> <ul style="list-style-type: none"> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv4/neighbors/neighbor/state/</code> supporting leafs ip, link-layer-address, and origin</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv6/neighbors/neighbor/state/</code> supporting leafs ip, link-layer-address, origin, is-router, neighbor-state, and is-publish</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv4/addresses/address/state/</code> supporting leafs ip, prefix-length, and origin</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv4/addresses/address/state/</code> supporting leafs ip, prefix-length, and origin</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv4/unnumbered/state/</code> supporting leafs enabled, interface-ref/state/interface, and interface-ref/state/subinterface</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv4/state/</code> supporting leafs enabled and mtu</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv6/addresses/address/state/</code> supporting leafs ip, prefix-length, origin, and status</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv6/unnumbered/state/enabled</code></li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv4/unnumbered/state/</code> supporting leafs enabled, interface-ref/state/interface, and interface-ref/state/subinterface</li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv6/unnumbered/interface-ref/state/interface/subinterface/</code></li> <li><code>/interfaces/interface/subinterfaces/subinterface/ipv6/state/</code> supporting leafs enabled and mtu</li> </ul>

Table 5: Feature Support Added to the PTX10008 Routers (*continued*)

Feature	Description
	<p>[See <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a>.]</p> <ul style="list-style-type: none"> <li>• This release supports export of Source Packet Routing in Networking (SPRING) statistics to an outside collector using remote procedure call (gRPC) services.</li> </ul> <p>This feature provides the per-segment identifier (SID) level and interface level traffic counts for SPRING traffic. These statistics reflect the SPRING LSP utilization in the TED, which aids to correctly re-route the RSVP LSPs.</p> <p>To enable SPRING statistics include the following statements on the client device:</p> <ul style="list-style-type: none"> <li>• For egress (per-interface egress) use the <b>set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress</b> at the <b>[edit]</b> hierarchy level.</li> <li>• For egress (per-SID egress) use the <b>set protocols isis source-packet-routing sensor-based-stats per-sid egress</b> at the <b>[edit]</b> hierarchy level.</li> <li>• For ingress (per-SID ingress) use the <b>set protocols isis source-packet-routing sensor-based-stats per-sid ingress</b> statement at the <b>[edit]</b> hierarchy level.</li> </ul> <p>Use the following sensors to export statistics by means of gRPC services to an outside collector:</p> <ul style="list-style-type: none"> <li>• <b>/junos/services/segment-routing/interface/egress/usage/</b> for egress (per-interface egress) aggregate SPRING traffic.</li> <li>• <b>/junos/services/segment-routing/sid/usage/</b> for egress (per-SID egress) and ingress (per-SID ingress) aggregate SPRING traffic.</li> </ul> <p>[See <a href="#">source-packet-routing</a> and <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a>.]</p>
Multicast	<ul style="list-style-type: none"> <li>• IPv4 and IPv6 multicast support including MSDP. This release also includes support for PIM-SM as the first-hop router (FHR) or last-hop router (LHR), and anycast, static, or local rendezvous point (RP).</li> </ul>
Network management and monitoring	<ul style="list-style-type: none"> <li>• Remote port mirroring with ToS or DSCP settings enables you to send sampled copies of incoming packets to remotely connected network management software. You send the packets through GRE encapsulation, which is supported by FTIs. You can set ToS and DSCP values to provide necessary priorities in the network for these packets. You can also apply policing to sampled packets that are leaving the FTI. Configure the settings you need in the <b>[edit forwarding-options port-mirroring instance <i>instance-name</i> output]</b> hierarchy. [See <a href="#">instance (Port Mirroring)</a>.]</li> </ul>
Port security	<ul style="list-style-type: none"> <li>• Support for fallback preshared key (PSK) for MACsec. [See <a href="#">Configuring MACsec with Fallback PSK</a>.]</li> </ul>
Routing policy and firewall filters	<ul style="list-style-type: none"> <li>• Nested filters enable you to reference a common firewall filter by attaching it to multiple firewall policies (a filter being one or more match conditions and corresponding actions). You can bind nested filters to the following interface types: <ul style="list-style-type: none"> <li>• <b>inet</b>—Both input and output directions</li> <li>• <b>inet6</b>—Both input and output directions</li> <li>• <b>mpls</b>—Input direction only</li> </ul> </li> </ul>

Table 5: Feature Support Added to the PTX10008 Routers (*continued*)

Feature	Description
	<p>You can also bind them to routing instances, and in the input direction, the output direction, or both directions. [See <a href="#">Guidelines for Nesting References to Multiple Firewall Filters</a> and <a href="#">Example: Nesting References to Multiple Firewall Filters</a>.]</p> <ul style="list-style-type: none"> <li>• Support for configuring priority for route prefixes through existing import policy in protocols. [See <a href="#">Configuring Priority for Route Prefixes in RPD Infrastructure</a>.]</li> </ul>
Routing protocols	<ul style="list-style-type: none"> <li>• Support for inline BFD. [See <a href="#">Understanding Bidirectional Forwarding Detection (BFD)</a>.]</li> <li>• BGP flow specification support. The following match conditions are not supported: <ul style="list-style-type: none"> <li>• ICMP codes alone inet/inet6</li> <li>• Source/destination prefix for inet6</li> <li>• Flow label for inet6</li> <li>• Fragment for inet6</li> </ul> </li> </ul> <p>The following action is not supported:</p> <ul style="list-style-type: none"> <li>• Traffic marking</li> </ul> <p>[See <a href="#">Understanding BGP Flow Routes for Traffic Filtering</a>.]</p> <ul style="list-style-type: none"> <li>• Support for BGP-LU over SR-TE for color-based mapping of VPN services. [See <a href="#">Understanding Static Segment Routing LSP in MPLS Networks</a>.]</li> <li>• Support for BGP routes with n-multipath primary and 1-protection backup gateway. [See <a href="#">multipath (Protocols BGP)</a>, <a href="#">delay-route-advertisements</a>, and <a href="#">egress-te</a>.]</li> <li>• Support for BGP PIC edge with BGP Labeled Unicast (BGP-LU) as the transport protocol, which helps to protect traffic failures over border nodes (ABR and AS boundary router) in multidomain networks. [See <a href="#">BGP PIC Edge Using BGP Labeled Unicast Overview</a>.]</li> </ul>
Services applications	<ul style="list-style-type: none"> <li>• Support for multiple collectors in inline active flow monitoring. You can configure inline active flow monitoring to export flow records to up to four different collectors. Previously, inline flow monitoring could only export flow records to a single collector. [See <a href="#">Configuring Inline Active Flow Monitoring on PTX Series Routers</a>.]</li> <li>• Reporting of the true incoming interface for the sample packets for inline active flow monitoring. Inline active flow monitoring now reports the true incoming interface for the GRE-de-encapsulated packets entering the router for the configured inline active flow monitoring filter criteria. [See <a href="#">Understanding Inline Active Flow Monitoring</a> and <a href="#">Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers</a>.]</li> </ul>
Software installation and upgrade	<ul style="list-style-type: none"> <li>• Support for Junos OS Evolved limited image. The Junos OS Evolved limited image does not support data plane encryption and is intended only for the countries of the EACU. [See <a href="#">Junos OS Installation Package Names</a>.]</li> <li>• All models of the QFX10008 are now eligible for upgrade to PTX10008 Packet Transport Router models. Upgrade kits can be ordered to convert QFX10008 models to PTX10008-BASE3, PTX10008-PREM2, or PTX10008-PREM3. [See <a href="#">QFX10008 Configurations and Upgrades</a>.]</li> </ul>



Table 5: Feature Support Added to the PTX10008 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• All models of the MX10008 are now eligible for upgrade to PTX10008 Packet Transport Router models. Upgrade kits can be ordered to convert MX10008 models to PTX10008-BASE3, PTX10008-PREM2, or PTX10008-PREM3. [See <a href="#">MX10008 Components and Configurations</a>.]</li> <li>• Support for ZTP on WAN interfaces. [See <a href="#">Zero Touch Provisioning Overview</a>.]</li> <li>• ZTP supports the DHCPv6 client. [See <a href="#">Zero Touch Provisioning Overview</a>.]</li> </ul>

- **Support for QSFP-100G-FR, QSFP-100G-DR, and QSFP-100G-LR transceivers (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, we provide support for these transceivers:
  - QSFP-100G-FR—These transceivers interoperate with the QDD-4X100G breakout optics. For example, the QDD-4X100G-FR interconnects with up to four QSFP-100G-FR transceivers. The QSFP-100G-FR transceivers interconnect in single links (QSFP-100G-FR to QSFP-100G-FR or to QSFP-100G-DR) and interoperate at the shortest link length.
  - QSFP-100G-DR—These transceivers interoperate with 400-Gbps breakout optics. For example, the QDD-400G-DR4 interconnects with up to four QSFP-100G-DR transceivers. The QSFP-100G-DR transceivers interconnect in single links (QSFP-100G-DR to QSFP-100G-DR or to QSFP-100G-FR) and interoperate at the shortest link length.
  - QSFP-100G-LR—These transceivers interoperate with the QDD-4X100G breakout optics. For example, the QDD-4X100G-FR interconnects with up to four QSFP-100G-LR transceivers. The QSFP-100G-LR transceivers interconnect in single links (QSFP-100G-LR to QSFP-100G-LR or to QSFP-100G-FR) and interoperate at the shortest link length.

**NOTE:** These transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

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

- **Support for remote TACACS+ authorization for locally authenticated users (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, you can configure remote authorization on the TACACS+ server for locally authenticated users by using their locally configured parameters.

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

#### **Interfaces and Chassis**

- **Support for bringing line cards online parallelly (PTX10008)**—The PTX10008 with Junos OS Evolved Release 20.3R1 or later brings the line cards online parallelly (independent of each other), not sequentially. This change eliminates the dependency of a line card on a higher priority line card to come online.

However, the router allocates power to the line cards based on the priority configured in the system. The parallel bring-up behavior does not have a visible effect when there is sufficient power in the system. But when the system power capacity is degraded because of some reason, the system takes down the lowest priority line cards to power up the highest priority line cards. If the system is booting up with insufficient power, power is allocated only to the high-priority line cards.

[See [fru-poweron-sequence](#).]

- **VLAN tag manipulation: pop, push, and swap (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, you can configure your VLAN circuit cross-connect (CCC) logical interface on a Layer 2 circuit to handle single-tag packets. You can also use the **l2circuit-control-passthrough** statement at the **[edit forwarding-options]** hierarchy level to enable passthrough of certain Ethertype/DMAC-matched frames over the Layer 2 circuit after successful VLAN tag manipulation on the VLAN CCC logical interface. The VLAN CCC logical interface can be on a single Ethernet interface or on an aggregated Ethernet interface.

**NOTE:** You cannot configure **flexible-vlan-tagging** or **flexible-ethernet-services** on PTX10003.

[See [Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough](#).]

- **Support for QSFP28 100GE DWDM transceivers (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, the PTX10003 router supports the QSFP28 100GE DWDM optical transceiver module (part number: IN-Q2AY2-XX, where XX represents the channel ID) for 100GbE applications. This transceiver is shown in the show chassis hardware output as 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 100GE DWDM transceiver include the following:

- View the 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](#).]

### *Junos OS XML API and Scripting*

- **Routing instance support in Python 3 applications (PTX Series and QFX Series)**—Starting in Release 20.3R1, devices running Junos OS Evolved support specifying the routing instance used by a process or socket in Python 3 applications that are executed from the Linux shell. Python 3 applications can import the `libpyvrf` module and use the module's functions to set the Linux VRF corresponding to the Junos OS routing instance for a specific context.

[See [How to Specify the Routing Instance in Python 3 Applications on Devices Running Junos OS Evolved](#).]

### *Junos Telemetry Interface*

- **IS-IS sensor support for JTI (PTX10003 routers)**—Starting in Junos OS Evolved Release 20.3R1, Junos telemetry interface (JTI) supports OpenConfig Version v0.3.3 (from v0.2.1) for resource paths related to IS-IS link-state database (LSDB) streaming. The difference between the two versions results in changes, additions, deletions, or non-support for leaf devices related to the following IS-IS type length value (TLV) parameters and IS-IS areas:

- TLV 135: extended-ipv4-reachability
- TLV 236: ipv6-reachability
- TLV 22: extended-is-reachability
- TLV 242: router-capabilities
- IS-IS interface attributes
- IS-IS adjacency attributes

To stream data for the IS-IS routing protocol to an outside collector using remote procedure call (gRPC) services and Junos telemetry interface, include the following resource paths in a subscription:

- `/network-instances/network-instance[name_ 'instance-name']/protocols/protocol/isis/levels/level/`
- `/network-instances/network-instance[name_ 'instance-name']/protocols/protocol/isis/interfaces/interface/levels/level/`

To stream or export ON-CHANGE data for IS-IS adjacency and link-state database (LSDB) statistics to an outside collector using remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services and JTI, include the following resource paths in a subscription:

- `/network-instances/network-instance/protocols/protocol/isis/interfaces/interfaces/circuit-counters/state/` (stream)
- `/network-instances/network-instance/protocols/protocol/isis/interfaces/interface/levels/level/packet-counters/` (stream)
- `/network-instances/network-instance/protocols/protocol/isis/levels/level/system-level-counters/state/` (stream)

- /network-instances/network-instance/protocols/protocol/isis/interfaces/interfaces/levels/level/adjacencies/adjacency/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-ipv4-reachability/prefixes/prefix/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-ipv4-reachability/prefixes/prefix/subtlvs/subtlv/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-reachability/prefixes/prefix/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-reachability/prefixes/prefix/subtlvs/subtlv/ (stream)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-is-reachability/neighbors/neighbors/subTLVs/subTLVs/adjacency-sid/sid/state/ (ON-CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-is-reachability/neighbors/neighbors/subTLVs/subTLVs/lan-adjacency-sid/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-interfaces-addresses/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-srlg/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-te-router-id/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-interfaces-addresses/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/router-capabilities/router-capability/subtlvs/subtlv/segment-routing-capability/state/ (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/state (stream)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/area-address/state/address (stream)
- /network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/nlpid/state/nlpid (stream)

- `/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/lsp-buffer-size/state/size` (stream)
- `/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/hostname/state/hname` (stream)

[See [Configuring a Junos Telemetry Interface Sensor \(CLI Procedure\)](#) and [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **Support for BGP neighbor and MPLS sensors on JTI with gNMI (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, Junos telemetry interface (JTI) supports streaming BGP neighbors, label-switched path (LSP), and RSVP statistics to a remote collector. In prior releases, these statistics were supported on PTX10003 routers using remote procedure call (gRPC) services. This feature now adds support for streaming these statistics using gRPC network management interface (gNMI) services.

To stream data, include the following resource paths in a subscription:

- `/network-instances/network-instance[name='instance-name']/mpls/`
- `/network-instances/network-instance/protocols/protocol/bgp/`

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

- **Physical interface queue statistics sensor support for JTI (PTX10003)**—Junos OS Evolved Release 20.3R1 supports export of physical interface queue statistics to an outside collector using UDP (native) streaming, remote procedure call (gRPC) services, or gRPC network management interface (gNMI) services.

To export statistics through UDP, gRPC, or gNMI, use the sensor `/junos/system/linecard/interface/queue/`.

Each physical interface has 8 queues. The following counters are exported as part of this sensor for all configured physical interfaces:

- Transmitted packets and transmitted bytes
- Red drop packets and bytes
- Tail drop packets and bytes

This feature includes zero suppression support. It does not include support for summed up counters on aggregated ethernet (ae) interfaces.

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

## Layer 2 Features

- **Proxy ARP (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, PTX10003, 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](#).]

## Routing Policy and Firewall Filters

- **Filter-based GRE encapsulation and de-encapsulation and filter-based MPLS-in-UDP de-encapsulation (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, we've enabled the following encapsulation and de-encapsulation workflow:

1. An incoming packet matches a filter term with an encapsulate action. The packet is encapsulated in an IP+GRE header and is forwarded to the endpoint's destination.

```
set firewall tunnel-end-point tunnel-name ipv4|ipv6 source-address address
set firewall tunnel-end-point tunnel-name ipv4|ipv6 destination-address address
set firewall tunnel-end-point tunnel-name gre
set firewall family inet|inet6 filter name term name from source-address address
set firewall family inet|inet6 filter name term name then encapsulate tunnel-name
set firewall family inet|inet6 filter name term last then accept
set interfaces interface-name unit number family inet|inet6 filter input
set interfaces interface-name unit number family inet|inet6 address address # This source address differs
from the one for the tunnel endpoint.
```

2. At the destination, the packet matches a filter term with a de-encapsulate action. The GRE header or MPLS-in-UDP header is stripped from the packet. The inner packet is routed to its destination.

```
set firewall family inet|inet6 filter name term name from source-address address
```

```

set firewall family inet|inet6 filter name term name from protocol gre
set firewall family inet|inet6 filter name term name then decapsulate gre # Optionally de-encapsulate
mpls-in-udp.
set firewall family inet|inet6 filter name term last then accept
set interfaces interface-name unit number family inet|inet6 filter input filter-name
set interfaces interface-name unit number family inet|inet6 address address # This is the destination
address.

```

[See [Components of Filter-Based Tunneling Across IPv4 Networks](#) and [tunnel-end-point](#) .]

- **Support for unicast RPF strict mode and fail-filters (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, you can use unicast reverse path forwarding (RPF) strict mode, and fail filters, to prevent IP spoofing on IPv4 and IPv6 packet flows (unicast RPF loose mode is already supported).

With unicast RPF enabled, the PTX 10003 forwards packets from a valid path to the destination address, and either discards packets from an invalid path or sends them on to the **fail-filter** for further processing. This can be an effective way to mitigate denial-of service (DoS) attacks. In strict mode, the router interface only accepts packets if the source address matches a route (default or learned) that is reachable through the interface.

[See [Understanding How Unicast RPF Prevents Spoofed IP Packet Forwarding](#).]

### Routing Protocols

- **VRRP Support (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, PTX10003 routers support VRRP. The following features are not supported for VRRP on Junos OS Evolved: GRES, NSR, ISSU, ProxyArp, MC-LAG, IRB.

[See [Understanding VRRP](#)]

- **Support for multiple MD5 for RIPv2 (PTX10001-36MR, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 20.3R1, you can define multiple MD5 authentication keys for RIPv2. This feature supports adding of MD5 keys with their **start-time**. RIPv2 packets are transmitted with MD5 authentication using the first configured key. RIPv2 authentication switches to the next key based on its configured respective key **start-time**. This provides auto-key switching without user intervention to change the MD5 keys as in the case of having only one MD5 key.

To enable multiple MD5 support for RIPv2, include the **authentication-selective-md5** statement at the **[edit protocols rip]** hierarchy level.

- **Routing Protocol feature support (PTX10003)**—We've added the following routing protocols features to the PTX10003 routers in Junos OS Release Evolved 20.3R1.

**Table 6: Feature Support Added to the PTX10003**

Technology	Supported Features
BGP	<ul style="list-style-type: none"> <li>• BGP flowspec redirect to IP</li> <li>• BGP prefix-based outbound route filter (ORF)</li> <li>• BGP over IPv6</li> <li>• EBGp peering using link-local addresses (IPv6)</li> <li>• Regular expressions for BGP extended communities</li> <li>• BGP AS path lists</li> <li>• Option for link bandwidth in BGP multipath path selection</li> <li>• SNMP objects for BGP peer received prefix counters</li> <li>• BGP route target filtering</li> <li>• Label aggregation using BGP site of origin community attribute</li> <li>• Option to limit the number of active prefixes on BGP peering session</li> <li>• TCP MSS per BGP peer option</li> <li>• Timer-based solution to periodically advertise MED updates</li> <li>• Fast connectivity restoration using <b>add-path</b></li> <li>• Router reflector with dynamic policies</li> <li>• End-to-end restoration: BGP convergence in case of multihoming</li> <li>• Software-defined networking (SDN): BGP monitoring protocol v3 compliance</li> <li>• Static route target filtering</li> <li>• BGP link state distribution</li> <li>• End-to-end restoration: Tail end protection for BGP Labeled Unicast (BGP-LU)</li> <li>• End-to-end restoration: BGP-LU PE-CE link protection</li> <li>• End-to-end restoration: edge node-failure protection of BGP signaled pseudowires</li> <li>• BGP persistence</li> <li>• BGP graceful shutdown (RFC 6198)</li> <li>• BGP administrative shutdown communication (RFC 8203)</li> <li>• Advertise statically inactive route via BGP</li> <li>• End-to-end restoration: multihoming BGP</li> <li>• Entropy label support for BGP-LU</li> <li>• End-to-end restoration: BGP prefix-independent convergence in RSVP</li> <li>• Multiprotocol BGP over IPv6 (IPv4 over IPv6)</li> <li>• BGP prefix prioritization</li> <li>• BGP-LU FRR</li> <li>• BGP optimal route reflector with IS-IS</li> </ul>



Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
	<ul style="list-style-type: none"> <li>• BGP-LU support to include a stack of labels</li> <li>• BGP flowspec</li> <li>• BGP optimal route reflector with OSPF</li> <li>• BGP add path support for community</li> <li>• Paths to resolve a BGP prefix when using another BGP prefix for penultimate next-hop resolution</li> <li>• BGP 64-way add-path</li> <li>• BGP multipath</li> <li>• Enable BGP multipath configuration in global hierarchy</li> <li>• AS number count</li> <li>• BGP add-path support for VPN</li> <li>• BGP peer</li> <li>• BGP link bandwidth community aggregation</li> <li>• BGP LU (top label) statistics</li> <li>• BGP add path support for eBGP</li> <li>• Import IGP topology into BGP-LS</li> <li>• Performance enhancement for BGP reconfiguration</li> <li>• BGP remote next hop support for single-hop EBGP peers</li> <li>• Enterprise-specific BGP trap support for BGP clients with IPv6 address</li> <li>• Advertising multiple paths in BGP (upto 20 BGP add-paths for prefixes)</li> <li>• EBGP route server functionality</li> <li>• BGP route API support for EBGP</li> <li>• Add-path or multipath optimization to improve RIB learning rate</li> </ul> <p>[See <a href="#">BGP User Guide</a>.]</p>

Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
Bidirectional Forwarding Detection (BFD)	<ul style="list-style-type: none"> <li>• Static routes and MPLS PE to CE links</li> <li>• MPLS LSPs</li> <li>• PE-PE with ECMP awareness</li> <li>• PPMD and single-hop BFD to Packet Forwarding Engine</li> <li>• MIB</li> <li>• IPv6 static route</li> <li>• OSPFv3</li> <li>• Distributed BFD over aggregated interface</li> <li>• BFD-triggered fast reroute</li> <li>• Distributed BFD for BGP multihop</li> <li>• BFD over child links of aggregated Ethernet interface (ae) or LAG bundle</li> <li>• Control plane scaling</li> <li>• Support for IS-IS IPv6</li> <li>• Dampening for OSPF</li> </ul> <p>[See <a href="#">High Availability Feature Guide</a>.]</p>
Infrastructure	<ul style="list-style-type: none"> <li>• 64-bit RPD support</li> <li>• End-to-end restoration: host fast reroute (HFRR)</li> </ul>
Interior Gateway Protocol (IGP)	<ul style="list-style-type: none"> <li>• Prefix limit of imported external routes</li> <li>• Shared Risk Link Group (SRLG) support</li> <li>• IS-IS prioritized route installation in FIB</li> <li>• Policy-based support for loop-free alternate (LFA) in IS-IS and OSPF</li> <li>• IS-IS flooding group</li> <li>• Remote LFA (rLFA) support in OSPF</li> <li>• OSPFv3 for IPv6</li> </ul>
IPv6	<ul style="list-style-type: none"> <li>• BGP flowspec for IPv6</li> </ul> <p>[See <a href="#">Multiprotocol BGP</a>.]</p>

Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
IS-IS	<ul style="list-style-type: none"> <li>• Option to turn off IS-IS hello and/or SNP authentication</li> <li>• Route tagging</li> <li>• IPv6 multitopology extensions</li> <li>• LFAs</li> <li>• MIB according to RFC4444</li> <li>• Per-prefix LFA</li> <li>• FRR route convergence</li> <li>• Link down microloop avoidance</li> <li>• Option to overload stub networks through IS-IS overload</li> <li>• IGP shortcut selection by protocol for weighted equal-cost multipath (WECMP)</li> </ul> <p>[See <a href="#">IS-IS User Guide</a>.]</p>
Layer 2 Circuit	<ul style="list-style-type: none"> <li>• LSP ping for Layer 2 VPN and Layer 2 circuits</li> <li>• Zero control word processing for Layer 2 frames over MPLS</li> <li>• Traffic engineering of Layer 2 circuits over multiple LSPs</li> </ul> <p>[See <a href="#">Layer 2 VPNs and VPLS User Guide for Routing Devices</a>.]</p>
Layer 2 VPN	<ul style="list-style-type: none"> <li>• End-to-end restoration: Layer 2 VPN service mirroring</li> <li>• Inet and inet6 family support on PS (cross-functional)</li> <li>• FEC 129 multisegment pseudowire</li> <li>• FAT pseudowire support for BGP L2VPN and VPLS</li> <li>• Chained composite next hop for L2CKT and L2VPN</li> <li>• MPLS converged services: stitching of pseudowire segments (multisegment pseudowires with BGP-L2VPN)</li> </ul> <p>[See <a href="#">Layer 2 VPNs and VPLS User Guide for Routing Devices</a>.]</p>

Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
Layer 3 VPN	<ul style="list-style-type: none"> <li>• OSPF sham links for MPLS VPNs</li> <li>• MIBs for RFC 2547 VPNs</li> <li>• Loopback interfaces for each VRF</li> <li>• Simplified policy configuration for BGP community per-VRF</li> <li>• LDP-IGP as PE-CE device routing for carrier-of-carriers</li> <li>• IPv6 VPNs</li> <li>• Automatic configuration of route distinguishers on VPN PE routers</li> <li>• vrf-table-label for Layer 3 VPNs and Frame Relay uplinks</li> <li>• BGP and VRF option for <b>vpn-apply-export</b> statement</li> <li>• vrf-table-label for (non-VLAN) Ethernet P-PE uplinks</li> <li>• Path MTU discovery on IPv4 and IPv6 VPNs</li> <li>• IPv6 and OSPFv3 on VRF instances and running over tunnels</li> <li>• Peering with family <b>inet unicast</b> and <b>inet label-unicast</b> on the same session</li> <li>• VPN load-balancing between PE devices that have different route distinguishers</li> <li>• vrf-table-label for Layer 3 VPNs and ATM uplinks</li> <li>• Support for no-propagate-ttl on a per-VRF basis</li> <li>• CLI usability features for Layer 3 VPN</li> <li>• GRE tunnel through VRF</li> <li>• BGP policy to control VPN label allocation mode</li> <li>• MIB for route entries in VRF (RFC 4382)</li> <li>• PWT into VRF: logical tunnel redundancy</li> <li>• Increasing the number of Layer 3 VPNs with vrf-table-label configured</li> <li>• IRB in MPLS VRF</li> <li>• Multicast VRF route leaking</li> <li>• LSP ping for VPN LSPs</li> <li>• LSP ping for labeled BGP</li> <li>• CLI command <b>show route bgp route-distiguisher</b> for prefix</li> </ul> <p>[See <a href="#">Layer 3 VPNs User Guide for Routing Devices</a>.]</p>
Label Distribution Protocol (LDP)	<ul style="list-style-type: none"> <li>• Client or server LDP mapping</li> <li>• LDP native IPv6 support</li> </ul> <p>[See <a href="#">MPLS Applications User Guide</a>.]</p>

Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
MPLS	<ul style="list-style-type: none"> <li>• LSP ping for CCC and CCC secondary standby LSP</li> <li>• LSP ingress traceroute</li> <li>• Advertising MPLS labels in IS-IS</li> <li>• Flexible MPLS label stack depth for segment routing</li> <li>• Point-to-point traceroute (support RSVP FEC at ingress and transit)</li> <li>• Leaking MPLS.0 routes between virtual routers and VRFs</li> <li>• Cross-connect logical interface to two unsigaled virtual circuits terminating on different egresses</li> <li>• RSVP automatic mesh: full mesh based on the need to resolve a BGP next hop</li> <li>• Edge node-failure protection of LDP signaled pseudowires</li> <li>• Label operations: push and swap push</li> <li>• Static Ethernet pseudowires double-label operation</li> <li>• PE devices</li> <li>• CLI support for monitoring MPLS label usage</li> </ul> <p>[See <a href="#">MPLS Applications User Guide</a>.]</p>
MPLS-RSVP Point-to-Multipoint	<ul style="list-style-type: none"> <li>• Traffic-engineered LSPs with static explicit route object (ERO)</li> <li>• Traffic-engineered LSPs with link protection</li> <li>• Ultimate-hop popping for LSPs</li> <li>• Ingress PE redundancy for LSPs</li> <li>• Traffic engineering MIB</li> <li>• Interarea point-to-multipoint LSP</li> <li>• Load balancing over aggregated links</li> <li>• Multicast make-before-break</li> <li>• Admin-group for bypass LSPs</li> </ul> <p>[See <a href="#">MPLS Applications User Guide</a>.]</p>
Multipoint LDP	<ul style="list-style-type: none"> <li>• Inband signalling</li> <li>• MIB</li> </ul> <p>[See <a href="#">MPLS Applications User Guide</a>.]</p>
OCST: OpenConfig	<ul style="list-style-type: none"> <li>• Mapping of OpenConfig BGP commands for network-instance compliant (v4.0.1)</li> </ul> <p>[See <a href="#">OpenConfig User Guide</a>.]</p>
OpenConfig	<ul style="list-style-type: none"> <li>• BGP configuration model (v2.1.1)</li> </ul> <p>[See <a href="#">OpenConfig User Guide</a>.]</p>

Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
OSPF	<ul style="list-style-type: none"> <li>• Policy-based inbound route filtering</li> <li>• Active backbone detection</li> <li>• Multitopology routing OSPF</li> <li>• Export external route to multiple area scoped type 7 LSAs</li> <li>• OSPFv3 SNMP MIB</li> <li>• Route install prioritization</li> <li>• Database protection</li> <li>• Per-prefix LFA</li> <li>• Option to overload stub networks through OSPF overload</li> </ul> <p>[See <a href="#">OSPF User Guide</a>.]</p>
OSPFv3	<ul style="list-style-type: none"> <li>• Address family support</li> <li>• MIB support according to RFC5643</li> </ul>
Path Computation Element Protocol (PCEP)	<ul style="list-style-type: none"> <li>• MD5 authentication for PCC/PCEP</li> <li>• Point-to-multipoint</li> <li>• No-ERO or loose ERO support</li> <li>• Path computation for segment routing LSP</li> <li>• Support for PCEP MIB</li> <li>• Support for latest version of PCEP RFC 8231</li> <li>• PCE support for RSVP-TE</li> </ul>
Programmable Routing Protocol (PRPD)	<ul style="list-style-type: none"> <li>• gRIBI RIB programming interface</li> <li>• BGP SR-TE policy AFI</li> <li>• BGP flow-specification subsequent address family identifier (SAFI)</li> <li>• Move programmed routes to hidden state upon next-hop interface failure</li> </ul> <p>[See <a href="#">Juniper Extension Toolkit API Guide</a>.]</p>

Table 6: Feature Support Added to the PTX10003 (*continued*)

Technology	Supported Features
Segment Routing	<ul style="list-style-type: none"> <li>• Advertising MPLS labels in OSPF</li> <li>• BGP-LS with SPRING extensions</li> <li>• LDP mapping client</li> <li>• Advertising MPLS labels in IS-IS</li> <li>• RSVP interoperability</li> <li>• Static LSP with label stack</li> <li>• Adjacency SID support for different use cases (IS-IS)</li> <li>• Anycast and prefix segments and interarea (OSPF)</li> <li>• TI-LFA procedures for link and node protection (IS-IS and OSPF)</li> <li>• Static adjacency SID support for different use cases (OSPF)</li> <li>• Enabling first hop as segment ID instead of IP address</li> <li>• LDP mapping server</li> <li>• BGP triggered dynamic creation of colored SR-TE tunnels</li> <li>• Policy-based multipath routes</li> <li>• TI-LFA using SRMS routes (IS-IS)</li> <li>• Advertise traffic engineering attributes for segment routing irrespective of RSVP-TE</li> <li>• Flexible algorithm (IS-IS only)</li> <li>• MPLS-SR for IPv6 prefix and adjacency SID (IS-IS)</li> <li>• Segment routing global block (SRGB) for OSPF</li> <li>• MPLS ping and traceroute (IS-IS and OSPF for IPv4 only)</li> <li>• BGP binding SID (draft-previdi-idr-segment-routing-te-policy)</li> <li>• Segment list path ERO support using IP address as next hop and loose mode</li> </ul> <p>[See <a href="#">Understanding Source Packet Routing in Networking (SPRING)</a>.]</p>
VPN	<ul style="list-style-type: none"> <li>• Graceful restart for CCC</li> <li>• Generalized VPN MIB</li> </ul>

### Services Applications

- **Support for multiple collectors in inline active flow monitoring (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, you can configure inline active flow monitoring to export flow records to up to four different collectors. Previously, inline active flow monitoring could only export flow records to a single collector.  
[See [Configuring Inline Active Flow Monitoring on PTX Series Routers](#).]
- **Reporting of the true incoming interface for the sample packets for inline active flow monitoring (PTX10003)**—Starting in Junos OS Evolved Release 20.3R1, inline active flow monitoring reports the

true incoming interface for the GRE-de-encapsulated packets entering the router for the configured inline active flow monitoring filter criteria.

[See [Understanding Inline Active Flow Monitoring](#) and [Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers](#).]

- **Support for RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service (PTX10003)**—Starting in Junos OS Evolved 20.3R1, you can configure the TWAMP monitoring service on PTX10003 routers. This service sends out probes to measure network performance. TWAMP is often used to check compliance with service-level agreements. For Junos OS Evolved, TWAMP is configured at the **[edit services monitoring twamp]** hierarchy level. The support for this service is limited to the following:
  - IPv4 traffic only for control sessions and test sessions
  - Probe statistics and history
  - Control and test session status
  - Test session probe generation and reception, as well as reflection
  - Timestamps set by the Routing Engine or the Packet Forwarding Engine
  - Error reporting through system log messages only
  - Unauthenticated mode only

[See [Understanding Two-Way Active Measurement Protocol on Routers](#).]



## Software Licensing

- **Juniper Agile Licensing (QFX5220-32CD, QFX5220-128C, PTX10003-80C, and PTX10003-160C)**

—Starting in Junos OS Evolved Release 20.3R1, we're moving toward license-based software features. We now use Juniper Agile Licensing to support soft enforcement for software features on the listed devices. Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can install and manage licenses for hardware and software features using Juniper Agile Licensing. From this release onwards, you can now opt to use the Juniper Agile License Manager to significantly improve the ease of license management for an entire network of supported devices. If you are upgrading to this release, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS releases earlier than Junos OS Evolved Release 20.3R1.

[Table 7 on page 49](#) describes the licensing support on the QFX5220-32CD and QFX5220-128C devices.

**Table 7: Licensed Features on the QFX5220-32CD and QFX5220-128C**

QFX Switch License Model	Detailed Features
Standard license for integrated SKUs (standard hardware and software platform)	Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), QoS (Layer 2 and Layer 3), and SNMP
Advanced license for integrated and advanced SKUs	<b>Advanced 1:</b> BGP, FBF, GRE, IS-IS, JTI, MC-LAG, OSPF, sFlow, VRF, and VRRP
	<b>Advanced 2:</b> Includes Advanced 1 features + CFM, Layer 2 and Layer 3 multicast, OAM, Packet Timestamping, PTP, and Q-in-Q
Premium license for integrated and premium SKUs	Includes Advanced 2 features + EVPN-MPLS, MPLS, Layer 2 circuit, Layer 3 VPN, LDP, RSVP, segment routing, and SR-TE

Table 8 on page 50 describes the licensing support for the PTX10003-80C and PTX10003-160C devices.

**Table 8: Licensed Features on the PTX10003-80C and PTX10003-160C**

License Model	Supported Features	Scale	Security License
Standard	Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), quality of service or QoS (Layer 2 and Layer 3), and SNMP	64K FIB  K stands for 1000.	The PTX10003 router supports the MACsec feature, but you must purchase a license separately to use the feature.
Advanced	<b>Advanced 1</b>		
	BGP, FBF, GRE, IS-IS, Junos telemetry interface (JTI), OSPF, sFlow, VRF, and VRRP	256K FIB, 3M RIB, and 1K VR  M stands for million.	
	<b>Advanced 2</b>		
	Advanced 1 features, CFM, EVPN-VXLAN, Multicast, OAM, PTP, and Q-in-Q	256K FIB, 3M RIB, and 1K VRs/VRF (VXLAN)	
Premium	<b>Premium 1</b>		
	Advanced 2 features, EVPN-MPLS, IPFIX, Layer 2 circuit, Layer 3 VPN, LDP, RSVP, SR, and SR-TE	2M FIB, 6M RIB, 1K VRs/VRF (VXLAN), 32 VRF (MPLS Layer 3 VPN), and 32K LSP	
	<b>Premium 2</b>		
	Premium 1 features and fine-grained QoS	2M to 4M FIB, 60M to 80M RIB, 1K+ VRs/VRF (VXLAN), 32+ VRF (MPLS Layer 3 VPN), and 32K+ LSP	

\* Scaling is based on the device capacity.

[See [Supported Features on QFX5220-32CD and QFX5220-128C device](#), [Supported Features on PTX10003-80C and PTX10003-160C device](#), [Juniper Agile Licensing Guide](#), and [Configuring Licenses in Junos OS](#).]

### *User Interface and Configuration*

- **Support for virtual routing and forwarding (VRF) and source address (PTX10008)**—Starting in Junos OS Evolved release 20.3R1, VRF and source address is supported for NTP.

[See [Virtual Routing Instances](#) and [source-address](#).]

## What's Changed

### IN THIS SECTION

- [What's Changed in Junos OS Evolved Release 20.3R2 | 51](#)
- [What's Changed in Junos OS Evolved Release 20.3R1-S3 | 54](#)
- [What's Changed in Junos OS Evolved Release 20.3R1-S2 | 54](#)
- [What's Changed in Junos OS Evolved Release 20.3R1 | 55](#)

Learn about what changed in Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, and PTX10008.

## What's Changed in Junos OS Evolved Release 20.3R2

### *General Routing*

- **The write-file option at the monitor traffic interface hierarchy level takes precedence**—The **write-file** option at the **monitor traffic interface** hierarchy level takes precedence over the **extensive** option when you configure them simultaneously. If you try to configure these options at the same time, Junos OS Evolved gives you a warning message that the options are not compatible, and it only runs the **monitor traffic interface write-file** command.

[See [monitor traffic](#).]

### *Interfaces and Chassis*

- **The show chassis hardware output displays 1 x 10-Gbps support (PTX10008 routers with PTX10K-LC1201-36CD line card)**—The **show chassis hardware** output displays 1 x 10 Gbps support for the PTX10K-LC1201-36CD line card as follows:

FPC 4	REV 28	750-093524	BCBD3174	JNP10K-LC1201
CPU	REV 06	711-072984	BCBD0436	JNP10K-LC1201 PMB Board
PIC 0		BUILTIN	BUILTIN	JNP10K-36QDD-LC-PIC
Xcvr 0	REV 01	740-084673	SS1TY2A413002	QSFP28-DD-2X100GBASE-SR4
Xcvr 1	REV 01	740-084673	SS1TY2A411003	QSFP28-DD-2X100GBASE-SR4
Xcvr 3	REV 01	740-031980	AR22ANR	SFP+-10G-SR
Xcvr 4	REV 01	740-061405	1ACQ11370LC	QSFP-100GBASE-SR4-T2
Xcvr 5	REV 01	740-031980	AK80LXU	SFP+-10G-SR

The 10-Gbps support is provided by converting a 40GbE port to a 10GbE port with the help of a QSFP-to-SFP adapter (QSA).

[See [PTX10K-LC1201-36CD for PTX10008 Routers.](#)]

- **Adaptive load balancing support (PTX10008)**—You can configure adaptive load balancing on PTX10008 routers. When you configure adaptive load balancing, the **show interfaces ae0 extensive** command displays adaptive statistics under the physical interface.
- **PTX10003 routers do not support set chassis fpc fpc-slot power on**—The PTX10003-80C and PTX10003-160C routers do not support the **set chassis fpc fpc-slot power on** command. Executing this command on an FPC which is offline could cause unintended reboots of the router.
- **Warning message when taking an FPC offline**—PTX10003-80C and PTX10003-160C devices do not support the **request chassis fpc slot slot-number online** command. The only way to bring up an FPC (MPC) that is offline is by rebooting the chassis. So, when you take an FPC offline by using the **request chassis fpc slot slot-number offline** command, the screen displays the following message: **'Warning : FPC slot cannot be made online using a CLI command. You need to perform router reboot using "request system reboot" to online the FPC slot. Do you wish to continue ? [yes,no] (no).'**

[See [request chassis fpc.](#)]

### *Layer 2 Features*

- **New commit check for MC-LAG (PTX Series)**—We've introduced a new commit check to check the values assigned to the redundancy group identification number on the MC-AE interface ( **redundancy-group-id** ) and ICCP peer ( **redundancy-group-id-list** ) when you configure multichassis aggregation groups (MC-LAGs). If the values are different, the system reports a commit check error. In previous releases, if the configured values were different, the l2ald process would crash.

[See [iccp](#) and [mc-ae](#).]

### *MPLS*

- **The show mpls lsp extensive and show mpls lsp detail commands display next-hop gateway LSPid** — When you use the **show mpls lsp extensive** and **show mpls lsp detail** commands, you'll see next-hop gateway LSPid in the output.

### *Routing Protocols*

- Loading of the default configurations in a RIFT package causes the following changes:
  - Output of the **show rift node status** command displays the node ID in hexadecimal number even though the node ID is configured in decimal, hexadecimal, or octal number.

- Some of the DDoS default configurations change because of the DDoS protection interferes with the RIFT BFD operation.

### *Software Licensing*

- **License key format (PTX10003-160C, PTX10003-80C, QFX5220-32CD, QFX5220-128C)**—When you are upgrading from Junos OS Evolved Release 20.3R1 to Junos OS Evolved Release 20.3R2 or later releases, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS Evolved Release 20.3R2 or later.

### *User Interface and Configuration*

- **Verbose format option to export JSON configuration data (PTX Series and QFX Series)**—The Junos OS CLI exposes the **verbose** statement at the **edit system export-format json** hierarchy level. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **edit system export-format json** hierarchy level.

[See [export-format](#).]

## **What's Changed in Junos OS Evolved Release 20.3R1-S3**

### *General Routing*

- **Fault alarm generated for feed failure on a DC power supply (PTX10008)**—A fault alarm is generated when only one of the feeds on a DC power supply (A0 and B0 or A1 and B1) is faulty.

## **What's Changed in Junos OS Evolved Release 20.3R1-S2**

### *General Routing*

- **The write-file option at the monitor traffic interface hierarchy level takes precedence**—The **write-file** option at the **monitor traffic interface** hierarchy level takes precedence over the **extensive** option when you configure them simultaneously. If you try to configure these options at the same time, Junos OS Evolved gives you a warning message that the options are not compatible, and it only runs the **rmonitor traffic interface write-file** command.

[See [monitor traffic](#).]

- **Deprecated command show system buffers**—This command is not applicable in Junos OS Evolved because the command displays the status of kernel mbufs, which are not used in Linux-based systems like Junos OS Evolved. In releases before it was deprecated, the **show system buffers** command returns NULL.

[See [show system buffers](#).]

## What's Changed in Junos OS Evolved Release 20.3R1

### *Class of Service (CoS)*

- We've corrected the output of the **show class-of-service interface | display xml** command. Output of the following sort:

```
<container>
  <leaf1> data </leaf1>
  <leaf2> data </leaf2>
  <leaf3> data </leaf3>
  <leaf1> data </leaf1>
  <leaf2> data </leaf2>
  <leaf3> data </leaf3>
</container>
```

now appears correctly as:

```
<container>
  <leaf1> data </leaf1>
  <leaf2> data </leaf2>
  <leaf3> data </leaf3>
</container>
<container>
  <leaf1> data </leaf1>
  <leaf2> data </leaf2>
  <leaf3> data </leaf3>
</container>
```

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

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

## 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](#).]

## Interfaces and Chassis

- **Adaptive load balancing support (PTX10008)**—You can configure adaptive load balancing on PTX10008 routers. When you configure adaptive load balancing, the **show interfaces ae0 extensive** command displays adaptive statistics under the physical interface.
- **The show chassis power command displays the power supply state (PTX10008 and PTX10004)**—The **show chassis power** command displays the information regarding the state of the power supply (for instance, Online or Empty). This enhancement makes the **show chassis power** command output in Junos OS Evolved software consistent with that in Junos OS software.

See [show chassis power](#).

## Juniper Extension Toolkit (JET)

- **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](#).]

## Junos OS XML API and Scripting

- **Changes to Junos XML RPC request tag names (PTX Series, QFX Series)**—The Junos XML request tag name for some operational RPCs has been updated to ensure consistency across the Junos XML API. Devices running Junos OS will still accept the old request tag names, but we recommend using the new names going forward. The following changes have been made:
  - Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
  - Any uppercase characters are converted to lowercase.



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

### Layer 3 Features

- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (PTX Series, 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.

### MPLS

- **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.](#)]

- **Disable back-off behavior on PSB2 (PTX Series and QFX Series)**— We've introduced the `cspf-backoff-time` statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

[See [cspf-backoff-time..](#)]

### Routing Protocols

- **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 \(DDoS\) \(PTX Series and QFX Series\)..](#)]

- **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 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](#).]

## Known Limitations

### IN THIS SECTION

- [General Routing | 58](#)
- [MPLS | 60](#)
- [Platform and Infrastructure | 60](#)
- [Routing Protocols | 60](#)

Learn about limitations in this release for the PTX10001-36MR, PTX10003, PTX10004, 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

- License details are not fetched in broadcast mode when the device and license server are in the same subnet. As a workaround, configure license server explicitly on the device using the command **set system license server address ip** so as to not use broadcast mode. [PR1463584](#)
- Excess-rate configuration in port schedulers might not be completely honoured in certain scenarios, where there could be cases where output queues configured with guaranteed-rate with no excess-rate might end up getting more share in the excess region of the bandwidth than the queues with explicit excess-rate configured. [PR1528124](#)
- Double-fault scenarios are not handled by link auto-heal feature and fabric links remain down if Routing Engine switchover is attempted while auto-heal recovery is in progress. [PR1529599](#)
- 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 that the MTU size of interface is less than 9500B, and 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](#)

- When a scheduler-map binding is removed from an interface then default scheduler-map is bound to interface. If default scheduler-map is an oversubscribed scheduler map for the interface, then that map is not applied to this interface and all "interface queue" counters for this interface show statistics as 0. [PR1539052](#)
- When a scheduler-map binding is removed from an interface then default scheduler-map is bound to interface. If default scheduler-map is an oversubscribed scheduler map for the interface, then that map is also not applied to this interface. In this case scheduler might work based on previous scheduler configuration (prior to over-subscribed configuration). [PR1539060](#)
- PTX10008: IPv6 addressing is configured at the /64 subnet by default irrespective of the subnet configured on the DHCP server side. [PR1539839](#)
- In scenario of queue congestion, interface enable/disable or in PFC mode this interrupt can come on this release. [PR1553943](#)

## MPLS

- The rpd process might crash after network service configuration changed (like changing the range of MPLS labels) without rebooting all the Routing Engines (which is a system mandatory step). [PR1461468](#)

## Platform and Infrastructure

- In case of Junos OS RPM, the RPM client never sets the DF bit. Therefore, we don't see the issue between the Junos OS RPM client and the Junos OS RPM server. Whereas in case of Junos OS Evolved, the Junos OS Evolved RPM client sets the DF bit while sending the RPM probes to the RPM server. In case of Junos OS TVP based platforms, the RPM server is not able to decode the DF bit properly. The issue is not applicable for non TVP based Junos OS platforms acting as RPM server. [PR1508127](#)

## Routing Protocols

- BGP does not advertise route with Entropy Label Capability (ELC) and load-balance-label as stitching point on specific platforms. [PR1529387](#)

## Open Issues

### IN THIS SECTION

- [Authentication and Access Control | 61](#)
- [General Routing | 61](#)
- [Infrastructure | 64](#)
- [Interfaces and Chassis | 65](#)
- [Layer 2 Ethernet Services | 65](#)
- [MPLS | 65](#)
- [Network Management and Monitoring | 66](#)
- [Routing Protocols | 66](#)
- [User Interface and Configuration | 66](#)

Learn about open issues in this release for the PTX10001-36MR, PTX10003, PTX10004, 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.

## Authentication and Access Control

- With this change, when configuration statement **strict-authorization** is configured under tacplus-options and authorization fails, user is not allowed to login. [PR1557052](#)

## General Routing

- The **irp\_intr\_smp\_trap** error messages are seen on PTX10003 console while bouncing the IGP protocols and restarting routing the daemon. [PR1422881](#)
- The following XML tags are missing in Junos OS Evolved for the **show system virtual-memory** CLI:
   
`<system-virtual-memory-information>`, `<vmstat-memstat-malloc>`, `<vmstat-memstat-info>`,
   
`<memstat-name>`, `<inuse>`, `<memuse>`, `<high-use>`, `<memstat-req>`, `<memstat-size>`,
   
`<vmstat-memstat-zone>`, `<vmstat-zone-info>`, `<zone-name>`, `<zone-size>`, `<count-limit>`, `<used>`, `<free>`,
   
`<zone-req>`, `<vmstat-sumstat>`, `<cpu-context-switch>`, `<dev-intr>`, `<soft-intr>`, `<traps>`, `<sys-calls>`,
   
`<kernel-thrds>`, `<fork-calls>`, `<vfork-calls>`, `<rfork-calls>`, `<swap-pageins>`, `<swap-pagedin>`,
   
`<swap-pageouts>`, `<swap-pagedout>`, `<vnode-pageins>`, `<vnode-pagedin>`, `<vnode-pageouts>`,
   
`<vnode-pagedout>`, `<page-daemon-wakeup>`, `<page-daemon-examined-pages>`, `<pages-reactivated>`,
   
`<copy-on-write-faults>`, `<copy-on-write-optimized-faults>`, `<zero-fill-pages-zeroed>`,
   
`<zero-fill-pages-prezeroed>`, `<transit-blocking-page-faults>`, `<total-vm-faults>`, `<page-faults-requiring-io>`,
   
`<pages-affected-by-kernel-thrd-creat>`, `<pages-affected-by-fork>`, `<pages-affected-by-vfork>`,
   
`<pages-affected-by-rfork>`, `<pages-cached>`, `<pages-freed>`, `<pages-freed-by-daemon>`,
   
`<pages-freed-by-exiting-proc>`, `<pages-active>`, `<pages-inactive>`, `<pages-in-vm-cache>`,
   
`<pages-wired-down>`, `<pages-free>`, `<bytes-per-page>`, `<total-name-lookups>`, `<positive-cache-hits>`,
   
`<negative-cache-hits>`, `<pass2>`, `<cache-deletions>`, `<cache-falsehits>`, `<toolong>`, `<vmstat-intr>`,
   
`<vmstat-intr-info>`, `<intr-name>`, `<intr-cnt>`, `<intr-rate>`, `<vm-kernel-state>`, `<vm-kmem-map-free>`.
   
[PR1438110](#)
- Because ISSU is not supported in Junos OS Release Evolved 20.2R1 for PTX10008, the following commands are removed or blocked from the CLI. 1. **request system software validate-restart** 2. **request system software add restart** 3. **show system software add-restart**. [PR1470943](#)
- Traffic loss of up to 400 ms is seen during an MPLS FRR test with the following scale: 600K IPv4 routes 40K IPv6 routes 19K Transit LSPs 1.6K ingress LSPs 10K BGP-LU. [PR1472908](#)
- On PTX10008 routers, the sub-system within the Packet Forwarding Engine (PFE) reports a series of new interrupts for the ASIC for which the PCIe address bus or PCIe link is down. On an ASIC which has issues with the PCIe address bus or PCIe link, all the new interrupts that are reported are false alarms and must be ignored. [PR1491128](#)
- On PTX10008 routers, there might be incorrect memory handling causing memory leak and FPC reboot when continuous auto-bandwidth adjustment happens in the background. [PR1492667](#)

- L2 circuit ping that uses CC type 1, also referred to as "PWE3 Control Word with 0001b as first nibble," is not supported. [PR1503724](#)
- When changing a physical interface from a VLAN ID list to multiple single VLAN IDs or VLAN ID lists that fall within the original VLAN range, we might see one or more of the new physical interfaces not notified. For example, in the current commit we have the following configuration: **set interfaces ae100 unit 1 vlan-id-list 1-100**. In the next commit, if we run the **set interfaces ae100 unit 1 vlan-id 1 set interfaces ae100 unit 2 vlan-id 2** command, we might see that one of the two physical interfaces is not notified. [PR1504374](#)
- On Junos OS Evolved platforms, the **show chassis power** command does not display the state of the PSM. This could be confusing for customers as they do not know if a certain PSM that is plugged in, is online or faulty. [PR1505219](#)
- 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, graceful SIB offline might result in small transient loss of traffic, till system re-routes traffic around fabric paths go offline. [PR1506866](#)
- The command suggested to replace the Junos OS CLI command **show chassis fabric errors** is incorrectly stated on PTX10008 when the **show chassis fabric errors** command is issued. [PR1511915](#)
- There is a small window during BGP neighbor or BGP RIB periodic streaming when the background telemetry job gets suspended, there is a possibility that the node that is being rendered will be deleted or modified. If we are in the middle of rendering value for that particular node, the rpd might crash when the telemetry job resumes because the node is no more valid. [PR1512773](#)
- When upgrading from 20.3R1 to a future release which has an FPC type which is incompatible for restart upgrade, the user is prompted in the CLI whether to offline the FPC or abort the upgrade. The CLI command waits for one hour for the user to respond. In case the user does not respond in an hour, the CLI command **request system software add var | tmp | ptx.iso restart** times out. Subsequent CLI software add/delete/rollback commands fail with software upgrade in progress. The workaround for this is to restart orchestrator, the application responsible for managing CLI upgrade operations. Restarting orchestrator has no impact on functioning of the system. The command is **request system application app orchestrator node master restart**. [PR1516384](#)
- While doing MBB for the encapsulation due to MAC change or MAC refresh, transient traffic loss can be experienced on PTX10008 platforms. [PR1517274](#)
- On PTX10008 platforms, whenever MAC change or MAC refresh is in progress and the Packet Forwarding Engine uses make-before-break approach for encapsulating or creating new encapsulation for a transient period of time, traffic loss can be experienced.
- The rpd does not recover after restart in scaled configuration scenarios when Junos Telemetry Interface monitoring is enabled using asymmetric configurations. The issue is not seen with: -
  - 1) Higher polling frequency (> 15 seconds).
  - 2) When subscribed at top level xpath /interface/sub-interface level.

3) When the number of IPv4 and IPv6 entries are almost same that is, mostly symmetric configurations.

Example 1: Symmetric scale of 4000 IPv4 and 4000 IPv6 address => Streaming works fine.

Example 2: Asymmetric scale of 4000 IPv4 and 4 IPv6 address => when the number of IPv6 addresses on an interface is sparse, the time taken to stream data is affected and has impact on the system. - et-0/0/1- 4000 IPv4, 4 IPv6 - et-0/0/3, et-0/0/5- 1 IPv6. [PR1520977](#)

- On any Junos OS Evolved system with SSD, in rare scenarios, with multiple "abrupt and/or ungraceful" power-cycles or reboots, filesystem might encounter error or it may not get initialized resulting in user login failure. This issue does not occur on a graceful shutdown of the system. [PR1523238](#)
- Junos OS Evolved supports **set system services ssh connection-limit** and **set system services ssh rate-limit**. [PR1525691](#)
- Each filter contains one or more filters terms. Each filter term might contain counter. If traffic hits the filter term then counter (packets/bytes) increments. If the user tries to modify the unsupported scaled filter then counters (packets/bytes) in filter terms are cleaned up which means the counters do not retain the original values before the modification. [PR1530597](#)
- When initial sync is enabled, the sensors export all possible leaves to give the ubuntu collector the schema for the sensor. If there is no data for a given field, the value zero is exported to the collector. During Evo-aftman's initial sync for the /interface/ sensor, Evo-aftman exports the carrier transition field for backwards capability purposes. Please disregard the carrier transition field for evo-aftmand. Mib2d is only responsible for the carrier transition field for both aggregated Ethernet and normal interface (this field is not valid for evo-aftman or picd). [PR1530864](#)
- The rpd corefile is seen after switching over to the backup Routing Engine in some cases. [PR1531653](#)
- When multicast traffic is running over and aggregated Ethernet bundle, which spans multiple FPCs, if one of the FPCs is brought offline, the traffic on that FPC is impacted. It is expected to see some traffic drops. [PR1531859](#)
- In scaled setup with a large number of routes and next-hops, and reloading the baseline configuration; interfaces and specifically aggregated Ethernet interfaces take longer to come up. After loading baseline and rollback traffic loss is expected. The router comes back into service after 5 minutes on Junos OS Evolved whereas it takes 1.5 minutes on Junos OS. This a router reconfiguration or upgrade scenario. [PR1534996](#)
- In rare cases, when interface statistics of an individual interface is cleared, then the queue counter for that interface might show large number of incorrect 20-digit counter values. [PR1535523](#)
- **show version** and **show version detail** missing few tags in Junos OS Evolved compared to Junos OS. [PR1536027](#)
- Port mirroring stops working for FTI interface when gre source is changed. [PR1536223](#)
- It takes time for traffic to converge after restart routing. [PR1537376](#)
- In a scaled setup it is possible that during GRES operation, the rpd on the new primary Routing Engine can receive an ifstates update (DEL and CHANGE) and these are not expected during RPD re-syncing.

In such cases the RPD can generate a core file and might recover automatically after sometime.

[PR1537947](#)

- After configuration **global system name-server**, configuration commit must fail but commit succeeds. [PR1538514](#)
- Software does not support GRES, when FPCs are triggered for restart. GRES with FPC triggered for offline or online is supported. [PR1539685](#)
- PTX10008 JNP10K-LC1201: Not getting expected output bps after applying a filter-specific policer egress filter. [PR1541978](#)
- When one SIB is offline and traffic into one Packet Forwarding Engine data path exceeds fabric capacity, and traffic to a "priority high" queue is a major share of port speed, some high-priority traffic might be dropped. [PR1545061](#)
- On Junos OS Evolved PTX platforms, the throughput might fall below the advertised line-rate if multiple traffic flows use the same egress interface. Hence traffic loss may be observed. [PR1551363](#)
- On a PTX Series device running Junos OS Evolved software, the output of **show interface queue <>** always shows **Forwarding classes: 16 supported, 4 in use** with customized configuration. [PR1554370](#)
- This is applicable to Junos OS Evolved platforms. The PHP router installs "clone routes" - explicit null at egress router. [PR1556626](#)
- Available space in /data on RE0 (primary) and RE1 (backup) can become asymmetric due to difference in size of db files. This can cause software upgrade to fail if there are more than ten images on the system. Also if /data is only full in RE1 (backup) and not on RE0 (primary), then force option might not be able to cleanup old software images to make space for new image to get installed. As a workaround, use the CLI command **request system software delete** to remove images from the system until there is adequate space on RE0 and RE1 to install the new software image. [PR1556992](#)
- Issue is seen when bulk configuration is used with cosd trace level set to debug and blocking. Impact of the issue: Cosd gets stuck due to high CPU utilization. [PR1557136](#)
- The throughput might fall below the advertised line-rate if multiple traffic flows use the same egress interface. To recover, the traffic flows must be spread across multiple egress interfaces. [PR1557444](#)
- On a PTX10008 running Junos OS Evolved, the **show interface** command combined with refresh configuration statement becomes slow after a **clear interface statistics** command. [PR1557843](#)

## Infrastructure

- Rebooting the PTX10003 during a broadcast storm on the management port may cause a fault on the PTX10003. [PR1423216](#)
- When using a source IP address as the management address of the box to ping a network address on a peer, the response for the ICMP ping from the peer, can end up on the management interface of the box, which is dropped by the Linux kernel as the RPF check is set to strict by default on the Linux kernel



used on Junos OS Evolved. The Linux kernel expects the path to the peer to be on the WAN side and so drops the packet when it is received on the management interface from the peer when the RPF check for the management interface is set to strict. [PR1498255](#)

- Standard **issue** message is not displayed for the initial login attempt, but if that login attempt fails for any reason and the user is re-prompted, then the message is displayed. This seems to be an issue common to PTX10004|PTX10008|PTX10001-36MR. [PR1528996](#)
- The GRES triggered through **request chassis routing-engine master switch** starts showing connector driver overlay messages: {master} root@ptx10004-05-re0> [ 1185.081257] gpio-jnx-i2cs gpio-jnx-i2cs.50: Asserting power\_status irq 59 [ 1185.125182] OF: overlay: overlay\_is\_topmost: #9 clashes #10 @/ftc0/i2c-bus/i2cs@54/fan\_hwmon [ 1185.125183] OF: overlay: overlay #9 is not topmost [PR1539232](#)
- PTX10008: Unexpected high packet loss duration - ungraceful FPC offline. [PR1547288](#)

## Interfaces and Chassis

- Fabric related logs (fabspoke\_fchip.log, fabspoke\_pfe.log) get collected in the /var/tmp directory unlike other logs in the /var/log directory. [PR1373609](#)
- The lacpd process might generate a core file while a router is shutting down in the following scenarios:
  - a. router reboot
  - b. router upgrade since the router is shutting down, this does not have an impact on forwarding traffic. [PR1553196](#)
- LACP force up option not working. [PR1553997](#)

## Layer 2 Ethernet Services

- The copying of files to the RCB over WAN ports is slow. This is observed across all platforms running Junos OS Evolved platforms. [PR1496895](#)

## MPLS

- The in-progress change of weighted ECMP to ECMP and vice versa does not take full effect for all routes using the LSPs until all the LSPs currently signaled with non-zero bandwidth have gone through at least one of the MBB/clear/disable event. The traffic distribution is not as expected until all the LSPs finish the transition. This is because without any of those triggers, the LSP does not automatically update the its next hop based on weighted ECMP configuration, next-hop content of each participation LSP affects the traffic distribution among ECMP paths. The weighted ECMP distribution ratio is not as expected until the adjustment of all LSPs are finished. [PR1501860](#)
- For a high scale of LSPs on Junos OS Evolved platforms, with high cpu load **clear mpls lsp statistics** command takes more time than expected. The time taken is more if the number of LSPs are more. [PR1504184](#)

- After graceful switchover, ingress RSVP sessions in new primary device go down and take up to 5 minutes to come up. [PR1518968](#)
- For a high scale of LSPs on Junos OS Evolved platforms, with high CPU load statistics fetching takes more time than expected. The time taken is more if the number of LSPs are more. Due to this container LSP sampled aggregate bandwidth does not get updated and splitting does not happen. [PR1540287](#)

## Network Management and Monitoring

- If some application spawns too many threads, trace-relay daemon might generate a core file if it reaches the upper cap of memory usage. The trace infra recovers and there is no impact to service. No action needed from user. [PR1556040](#)

## Routing Protocols

- Junos OS Evolved does not support configuring interface name as next hop. [PR1497012](#)
- PTX10008: traffic loss seen after Graceful switchover with LSR core MPLS\_DSCP profile configuration. [PR1540400](#)

## User Interface and Configuration

- The default shell used in Junos OS Evolved is /bin/sh. This is not preloaded with libsi.so which is required for ping functionality to work properly. Hence it is recommended to use utilities like ping from CLI instead of directly invoking them from the shell. [PR1430324](#)
- On Juniper device running Junos OS Evolved, NETCONF Service over SSH with dedicated TCP port (It is configured with system services netconf ssh and the default port is 830) might not work if in-band management is used (that is, connection is established through network interface or loopback interface etc.). [PR1517160](#)
- On Junos OS Evolved PTX Series devices, the firewall filter for both IPv4 and IPv6 does not work when it is applied through apply-groups on the subinterface. [PR1534858](#)
- If an inet filter is configured with or without **family** statement then changing its configuration to without/with **family** statement must not be done. [PR1556424](#)
- If an inet filter is configured with or without a **family** statement then changing its configuration to remove or add a **family** statement might cause the filter process (firewalld) to restart unexpectedly. [PR1556426](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.3R2 | 67](#)
- [Resolved Issues: 20.3R1 | 71](#)

Learn which issues were resolved in Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, 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: 20.3R2

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

- [cos] [scheduler] PTX10008:scheduler with all queues oversubscribed, maximum latency is different on different queues - 21ms ~ 29ms. [PR1478811](#)
- While configuring WRED profile to a scheduler, you can use either of the following: any/any not-any/not-any combination of protocol and loss priority. [PR1524259](#)
- The COSd might not come up after the FPC restart with COS configuration. [PR1544531](#)

#### *General Routing*

- Rate-limiting might not work for J-Flow-sampled traffic on PTX10008. [PR1473844](#)
- PTX10008: No cmerror is raised for certain PIO errors when accessing Packet Forwarding Engine ASIC. [PR1491130](#)
- Colored packets are not policed as expected due to policer inaccuracy for color-blind single-rate three-color policer at lower BW limits. [PR1497856](#)
- Junos OS Evolved- Aggregate Ethernet interfaces do not display member links' statistic. [PR1505596](#)
- The OSPF might flap when the line rate traffic is sent at smaller packet sizes. [PR1511563](#)
- When all the 4x400GbE and 4x100GbE of PFE0 ports are used inside one portion of the chip, there can be impact on the traffic throughput performance. [PR1518368](#)
- 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](#)
- The BFD session flaps might be seen on Routing Engine switchover. [PR1523505](#)

- Junos OS Evolved:PTX10008: observing reboot delay in software add reboot, rollback reboot commands. [PR1525286](#)
- [resiliency] [generic\_evo] PTX10001-36MR ::clockd.default anomalies are seen at producer:re0::clockd::50331749 for the type net::juniper::resild::errorItem. [PR1527309](#)
- The **show interfaces diagnostics optics** command does not work when the IFD name is not given or when a wild card is used. [PR1527821](#)
- Host path corruption may be observed after output filters are configured on some Junos OS Evolved devices. [PR1528368](#)
- FPC does not boot if the primary BIOS is corrupted (with FPC primary BIOS corruption golden BIOS failed to reprogram primary). [PR1528469](#)
- PTX10008 Doon RCB does not raise alarm Mixed Primary and Backup Routing Engine types when re0 and re1 have different FRU types. [PR1529042](#)
- BGP switchover convergence rate is degraded by 42% on Junos OS Evolved 20.3R1. [PR1529365](#)
- PTX10008 - FPC reboots if DDoS violations are seen during FPC restart. [PR1529847](#)
- [PTX10001-36MR] When the fabric hub process is restarted with the blackhole detection disable and fabric degrade detection enable configuration, it might result in traffic loss. [PR1530484](#)
- **show pfe statistics error** should print counters for meaningful errors. [PR1530710](#)
- After FPC restart, traffic drop is observed for multicast streams sent to PTX10008 device from other nodes. [PR1531429](#)
- Traffic loss might happen after performing GRES in Junos OS Evolved enabled chassis-based systems. [PR1532446](#)
- Running SNMP MIB walk and executing the **show interfaces** command might cause the picd process to crash. [PR1533766](#)
- PTX10008 : Multicast - After FPC restarts, we observe both traffic flood and loss toward downstream aggregated Ethernet receiver interfaces. [PR1535545](#)
- Physical interface does not come back to up state after configuring changes at the peer end. [PR1536270](#)
- The cosd might crash when multiple configuration changes are made in a single commit. [PR1536320](#)
- The **show chassis environment** command reports some ZF internal temperatures as 0 degrees C. [PR1536497](#)
- On PTX10001-36MR, PTX10000 with LC1201/LC1202 card and with Junos OS Evolved running, multicast does not work correctly on some of the ports, all multicast traffics are dropped on these ports. [PR1536963](#)
- Dependency-state details for the type net::juniper::hwd::fruHwStatus are seen after hwdfpc app restart with LSR core file profile configuration. [PR1537409](#)

- On a PTX10008 running Junos OS Evolved, running PIM Dense Mode (DM) might cause an FPC to restart unexpectedly. [PR1537700](#)
- [interface] PTX10001-36MR: Interface does not come up after enabling and disabling local loopback configuration. [PR1538393](#)
- Next-hop addCollision errors are seen with scaled multicast routes. [PR1538849](#)
- Auto-RP is not learned when a PTX10008 node is in discovery mode. [PR1539152](#)
- On re0 cli node reboot, rpd is unable to connect to snmpd and hence we see scheduler slips in rpd and protocol sessions flap. [PR1539705](#)
- On PTX10001-36MR on ports et-0/1/2 and et-0/2/2, Channel 1 does not come up on channelizing 100-Gbps to 2x50-Gbps speed. [PR1539795](#)
- On PTX10001-36MR ping fails on all channels when some ports with 400GbE DAC are channelized to 8x25GbE and 8x50GbE. [PR1539859](#)
- On PTX10001-36MR traffic might not flow through channels 2 and 3 of some ports with 40G optics that are channelizing to 4x10GbE. [PR1539864](#)
- PTX10008 : Multicast traffic silent packet drop on downstream adjacent node seen for approximately 5-6 seconds as dlu.ucode.inv\_start\_pc trap after aggregated Ethernet member link is deleted from or added to PTX10008 node. [PR1539912](#)
- PTX10003: Multicast traffic drops after port is converted from trunk to access for a Layer 2 bridging case. Unicast and broadcast traffic is not affected. [PR1540495](#)
- In PTX10008, for some temperature sensors the threshold temperature to return cooling fans to normal speed is 0 degrees C. [PR1540576](#)
- PTX10008 : Core file is generated when the combination of hop limit and packet length is configured in the filter. [PR1540625](#)
- PTX10008 JNP10K-LC1202 : savecore.sh fails to transfer core and journal files from FPCs due to directory create errors. [PR1540807](#)
- Installed licenses are deleted after image upgrade. [PR1540881](#)
- PTX10008: Component Upgrade: Software rollback timed out when moving from the image based on Junos OS Evolved Release 20.3R1 to any other Junos OS Evolved Release. [PR1541509](#)
- [cos] [scheduler] : PTX10008 : Junos OS Evolved CoS: The **show interfaces voq ae2** non-zero does not display the non-zero queue statistics, whereas the other aggregated Ethernet bundle ae1 displays it correctly (both ae1 and ae2 has 4 x 400G member links each from the same BT chip). [PR1543034](#)
- Picd core file is seen when the FPC or system is rebooted. This could happen for either of the CLI commands:

**request system shutdown reboot**

**request chassis fpc slot 0 restart**

Picd core files are generated during cleanup as part of reboot. There is no effect on the new incarnation of picd after reboot. [PR1543045](#)

- [cos] [classifier] PTX10008:PTX10K-LC1202-36MR: PTX10001-36MR:L2:classifier classifier on aggregated Ethernet interface does not work correctly (cos\_l2\_cls\_combine\_mts\_004.robot). [PR1546150](#)
- [PTX10008]: For 2x100G Qsfp-DD optics, CTLE settings are not applied sometimes. [PR1546236](#)
- The drop profile (WRED) feature under class-of-service might not work on PTX Junos OS Evolved platforms. [PR1549007](#)
- PTX10008:Add support for 1x10GE in **show chassis** output. [PR1549436](#)
- The first four forwarding classes might disappear if other forwarding classes are added. [PR1551154](#)
- The SNMP MIB jnxOperatingRestartTime output is in string format instead of OCTET string. [PR1553533](#)
- The jnxFruLastPowerOn value is incorrect for FPCs. [PR1553924](#)
- When a filter is applied on Routing Engine management interface and is deactivated, the filter behavior is unchanged after deactivation. [PR1553791](#)
- The software rollback might not work correctly on Junos OS Evolved platforms. [PR1556641](#)
- The upgrade/downgrade: serviceability: In case of upgrade failure, log files capturing upgrade fail reason do not give meaningful information to user. [PR1556807](#)
- The device might not boot up after performing the cleanup operation. [PR1557020](#)
- PTX10008:Junos OS Evolved plugged-in FPCs are automatically turning into online state, post performing primary-role switchover. [PR1557419](#)
- PTX10008 PDT: could not open user interface connection. Management process does not respond. Unable to go to edit or configuration mode. [PR1559070](#)
- Interface loopback might not work if there is no optics connected to the port on PTX10008. [PR1562471](#)
- PTX10008: After sync all followed by rollback and then reboot, RE1 booted on snapshot. [PR1562189](#)
- FPC is not powered on using **request node power-on fpc**. [PR1562981](#)
- PTX10001-36MR - In scaled logical interface scenario when interface connecting customer with many logical interfaces in the same port is flapped, ARP might not resolve and traffic might be dropped. [PR1563684](#)
- PTX10008: RE0 goes into reboot loop continuously during validate restart (INDB unsupported). [PR1563742](#)

### **Infrastructure**

- Launching the CLI produces a syslog notice. [PR1553897](#)

### **Interfaces and Chassis**

- [Junos OS Evolved] [PTX10008] ALB scan-interval configuration does not take effect. [PR1538854](#)

- FTI tunnel commit fails for MTU change on its family (inet, inet6). [PR1540431](#)

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

- 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](#)
- A memory leak in the mib2d and snmpd process might result in SNMP being unresponsive to SNMP queries on platforms running affected Junos OS Evolved image. [PR1543508](#)

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

- Firewall: Use of filter configured without address-family as nested filter causes configuration commit to fail. [PR1552641](#)
- Junos OS Evolved- changing the interface-specific field of the referenced filter in a single commit is not supported. [PR1556982](#)

#### ***Routing Protocols***

- The rpd might crash after reboot when MSDP is configured. [PR1536593](#)

#### ***User Interface and Configuration***

- Changing domain name does not effect. [PR1526660](#)
- File archive source /var/log/\* destination /var/log/.tgz does not work in Junos OS Evolved. [PR1549340](#)
- CLI command **show system uptime** throws error: **invalid xml tag (date: invalid date '@#')** from command-handler daemon on certain PTX Series platforms running Junos OS Evolved. [PR1566166](#)

### **Resolved Issues: 20.3R1**

#### ***Authentication and Access Control***

- Random IPv6 address displayed in NAS-IPV6-Address when IPv4 address is used as nas-ip for IPv6 server. [PR1501146](#)

#### ***General Routing***

- LSP statistics CLI is slow in a scaled scenario during installation time. [PR1416363](#)
- **show license** command not working after enabling network-mode licensing. [PR1459687](#)
- A configuration command to modify PCIe correctable and uncorrectable error thresholds on PTX10008 FPC. [PR1462953](#)
- The telemetry cannot export the data of MPLS LSP. [PR1489605](#)
- 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](#)
- Traffic loss might be observed when CBF is configured on Junos OS Evolved PTX Series platforms. [PR1492707](#)
- cmd crashes when **request system shutdown reboot disk1** command is executed. [PR1492955](#)
- [cos] [scheduler] PTX10008:scheduler ingress Packet Forwarding Engine VOQ drop counters does not match egress queue drop counters, if difference is greater than 100,000. [PR1494785](#)
- The CBF functionality does not work correctly on PTX10008 platform with Junos OS Evolved 20.1R1. [PR1495119](#)
- Block software upgrade when image that is not compatible with the platform is used. [PR1495655](#)
- The aftmand might crash when MPLS is enabled. [PR1496057](#)
- Consistent hash information is missing : error: **timeout communicating with Evo-Aft BT daemon**. [PR1496077](#)
- Sensor for physical Ethernet interface statistics is missing, for example: in-oversize-frame. [PR1496275](#)
- The logical interface might disappear after switchover. [PR1496887](#)
- Longevity: aggregated Ethernet logical interface disappears after switchover. [PR1497285](#)
- PTX10004 DPT fails for all interface extensive details are matching for 400G loopbacks in longevity test, FEC uncorrected errors found and high T\_max is observed. [PR1500587](#)
- There is a discrepancy of 22 bytes for the same packets between the firewall filter in Junos OS Evolved and in Junos OS. [PR1503145](#)
- PTX10008: Aggregated Ethernet (AE) interface flaps causing next hops to contain wrong encapsulation information, with router having 800,000 IP routes, 2000 ingress LSPs, around 500 L3VPNs and LSPs have link protection enabled. [PR1503260](#)
- The packetio crashes during the initialization and this might result in a second reboot. [PR1505150](#)
- **set system processes app failover other-routing-engine** configuration not supported. [PR1506480](#)
- 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 time to come online. [PR1507433](#)
- PTX10008 Junos OS Evolved packet loss around 1% was seen from Spirent and regular discard increase (**show pfe statistics traffic**) after GRES check pass and then GRES. [PR1511539](#)
- The evo-aftmand process might be stuck at 100% CPU usage in a scaled setup. [PR1511597](#)
- Packet Forwarding Engine generates MLP's with **global-no-mac-learning** enabled. [PR1511601](#)



- FirewallD generates a core file after deleting and adding the filter back in a single commit. [PR1512065](#)
- 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](#)
- The evo-aftmand-bt might crash if FPC offline is performed when the system comes up. [PR1514722](#)
- Continuous evo-cda-bt CDA syslog error messages are observed during the negative triggers (AE configurations disable/enable and FPC restart). [PR1515806](#)
- **show system License detail** does not reflect correct license usages after reboot. [PR1515896](#)
- SNMP traps are not seen for 'fpc\_link\_to\_sib\_fault' when CCL link is brought down from ZF->BT. [PR1516487](#)
- PTX10008: Do not configure confidentiality offset other than 0. [PR1517985](#)
- **request system zeroize** command does not reboot the device. [PR1518946](#)
- After **request system application node re0 app fabricHub restart** , interface drop and major **dp\_1\_zfo\_intr\_dp1\_fabcell\_drop** error. [PR1519402](#)
- A large number of publish-deleted hwdre anomalies are dumped after 2nd iteration of switchover from re1 (primary) to re0. [PR1519427](#)
- The **show interfaces voq intf | ae non-zero** command is not working as expected and displays incorrect output. [PR1521281](#)
- Junos OS Evolved: Grid:lfstatsd core file observed on longevity testbed with multiple triggers at net::juniper::evlruntime::evoobject\_publish\_validate. [PR1522703](#)
- With traffic running, if the FPC that is connected to the traffic generator reboots, it might lead to stuck traffic scenario on certain ports post resumption. [PR1523066](#)
- Hwdre generates a core file when trying to take faulty SIB to offline state. [PR1527790](#)
- Multiple l2cpd core files are seen during commit. [PR1528024](#)
- FPC does not boot if primary BIOS is corrupted (With FPC Primary BIOS corruption Golden BIOS failed to reprogram Primary). [PR1528469](#)
- On PTX Junos OS Evolved platforms with large scale ECMP (equal cost multi path), the evo-aftmand-bt process might crash and result in complete system crash when a lot of ECMP next-hop add/delete/change events take place. This issue is caused due to an uninitialized pointer. [PR1530696](#)
- The PTX10008 system might drop traffic in certain scenarios after a GRES event. [PR1532446](#)
- CoSd might crash when multiple configuration changes are made in a single commit. [PR1536320](#)

### **Infrastructure**

- ping does not work when we set record-route. [PR1474649](#)
- Telnet login related issue with template (TACACS and Radius). [PR1482363](#)

### **Interfaces and Chassis**

- ssd information gets removed from **show chassis hardware detail** after hwdre app restart. [PR1488706](#)
- After application restart - false optics alarms continue to persist although underlying optics does not have failures. [PR1493230](#)
- PTX10008 with LC1201 - the ifmand process restarts at IFAManager::findIfaInSameSubnet. [PR1496361](#)
- Rarely can see fan not getting detected post system reboot. [PR1517476](#)
- Minimum IFD MTU recommended with IPv6 configuration. [PR1518692](#)
- Last flapped of virtual interfaces like aggregated Ethernet is not updated when aggregated Ethernet transitions from DOWN to UP state [PR1521978](#)

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

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

### **User Interface and Configuration**

- To have a no-op **system configuration-database extend-size** available in Junos OS Evolved. [PR1533439](#)
- Remove support for CLI ping command vpn-interface attribute for Junos OS Evolved. [PR1533805](#)

## **Junos OS Evolved Release Notes for QFX5130-32CD and QFX5220 Devices**

### **IN THIS SECTION**

- [What's New | 75](#)
- [What's Changed | 84](#)
- [Known Limitations | 87](#)
- [Open Issues | 89](#)
- [Resolved Issues | 92](#)

These release notes accompany Junos OS Evolved Release 20.3R2 for QFX5130-32CD, QFX5220-32CD and QFX5220-128C Switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Check out [Table 9 on page 75](#) to see a complete list of the Junos OS Evolved releases supported on switching platforms.

**Table 9: Junos OS Evolved releases supported on switching platforms**

Platform	Supported Junos OS Evolved Releases
QFX5130-32CD	20.3R1, 20.3R2, and 20.4R1
QFX5220-32CD	19.1R2, 19.2R1, 19.3R1, 19.3R2, 19.4R2, 20.1R1, 20.1R2, 20.2R1, 20.2R2, 20.3R1, 20.3R2, and 20.4R1
QFX5220-128C	19.1R2, 19.2R1, 19.3R1, 19.3R2, 19.4R2, 20.1R1, 20.1R2, 20.2R1, 20.2R2, 20.3R1, 20.3R2, and 20.4R1

## What's New

### IN THIS SECTION

- [What's New in Junos OS Evolved 20.3R2 | 75](#)
- [What's New in Junos OS Evolved 20.3R1 | 76](#)

Learn about new features introduced in Junos OS Evolved Release 20.3R2 for the QFX5130-32CD and QFX5220.

### What's New in Junos OS Evolved 20.3R2

There are no new features introduced in Junos OS Evolved Release 20.3R2 for the QFX5130-32CD and QFX5220.

## What's New in Junos OS Evolved 20.3R1

### Hardware

- **New QFX5130-32CD Ethernet Switch (QFX series)**—In Junos OS Evolved Release 20.3R1, we introduce the QFX5130-32CD switch. The QFX5130-32CD is a fixed configuration switch that provides a maximum bandwidth of 12.8 Tbps. The QFX5130-32CD provides 32 QSFP56-DD user ports (32 ports of 400GbE or 64 ports of 200Gbe or 128 ports of 100Gbe) and 2 SFP+ (2 ports of 10Gbe) user ports. Four models are available: two featuring AC power supplies and front-to-back or back-to-front airflow, and two featuring DC power supplies and front-to-back or back-to-front airflow. To install the QFX5130-32CD switch hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the [QFX5130 Switch Hardware Guide](#).

[Table 10 on page 76](#) summarizes the QFX5130-32CD features supported in Junos OS Evolved Release 20.3R1.

**Table 10: Feature Support Added to the QFX5130-32CD**

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> <li>• Support for CoS features. The QFX5130-32CD supports most class of service (CoS) features. 802.3X Ethernet PAUSE is not supported. Also, CoS flexible hierarchical scheduling (ETS) is not supported. MPLS EXP ingress packet classification and egress rewrite rules are also not supported. Classifiers and rewrite rules are applied to logical interfaces instead of physical interfaces. [See <a href="#">CoS Support on QFX Series Switches, EX4600 Line of Switches, and QFabric Systems</a>.]</li> </ul>
DHCP	<ul style="list-style-type: none"> <li>• Support for DHCPv4 and DHCPv6 stateless relay. The DHCP relay agent is supported over L3 interfaces only. DHCP relay is not supported over IRB interfaces. [See <a href="#">DHCP Relay Agent</a>.]</li> </ul>
Firewalls and policers	<ul style="list-style-type: none"> <li>• Firewall filter support on Layer 3 interfaces. [See <a href="#">Firewall Filter Match Conditions and Actions (QFX5220)</a>.]</li> <li>• Support for Firewall Service and Interfaces Service JET APIs. [See <a href="#">Juniper Engineering Network website</a>.]</li> </ul>
Interfaces and chassis	<ul style="list-style-type: none"> <li>• Support for channelizing interfaces. The QFX5130-32CD contains a total of 34 ports, of which 32 are QSFP/QSFP28 ports and two are 10G SFP+ ports. The switch supports native 400-, or 100-, or 40-Gigabit Ethernet configuration. You can channelize the 400-Gbps QSFP/QSFP28 ports into: <ul style="list-style-type: none"> <li>• Four 100Gbe ports</li> <li>• Four 25Gbe ports</li> <li>• Four 10Gbe ports</li> </ul> </li> </ul> <p>[See <a href="#">Channelizing Interfaces on QFX5130-32CD Switches</a>.]</p>

Table 10: Feature Support Added to the QFX5130-32CD (*continued*)

Feature	Description
	<ul style="list-style-type: none"> <li>• Support for GRE tunneling. Tunneling provides a private, secure path for transporting packets through an otherwise public network by encapsulating packets inside a transport protocol known as an IP encapsulation protocol. GRE is an IP encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel. GRE tunneling is accomplished through routable tunnel endpoints that operate on top of existing physical and other logical endpoints. GRE tunnels connect one endpoint to another and provide a clear data path between the endpoints. Filter-based GRE encapsulation and decapsulation is not supported on QFX5130-32CD switches. [See <a href="#">Generic Routing Encapsulation</a>.]</li> <li>• Support for the dedicated virtual routing and forwarding (VRF) instance mgmt_junos for the management Ethernet interface. Virtual routing Instance of type VRF and Virtual-router. [See <a href="#">Management Interface in a Nondefault Instance</a> and <a href="#">management-instance</a>.]</li> <li>• Support for link aggregation groups (LAGs).</li> <li>• Layer 3 logical interfaces. [See <a href="#">Layer 3 Logical Interfaces</a>.]</li> </ul>
JET APIs	<ul style="list-style-type: none"> <li>• New IDL files that adhere to standardized guidelines. API files are to be made consistent with each other. Support of this feature is restricted to FE, GE, XE, and AE interfaces. Only the following functionality is included: Add, Delete, Modify, and GET for supported interface objects. There is also a Unit Permissions Get API, which is used to get the following information on an already created logical interface: was the logical interface was configured using APIs or CLI and, if it was configured using APIs, which client created the logical interface. [See <a href="#">Juniper Engineering Network website</a>.]</li> </ul>
Junos OS XML API and scripting	<ul style="list-style-type: none"> <li>• XSLT, SLAX, and Python scripts (which can trigger CLI operations).</li> </ul>
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> <li>• Support for Junos Telemetry Interface (JTI). OpenConfig support through gRPC and JTI is available for the following telemetry sensors: <ul style="list-style-type: none"> <li>• Sensor for Routing Engine internal interfaces (physical interface state and statistics)</li> <li>• Sensor for Routing Engine internal Logical interfaces (logical interface state and statistics)</li> <li>• Sensor for power supply module (PSM) FRU model number and fan trays</li> <li>• Sensor for hardware and chassis</li> <li>• Sensor for FEBs, FPCs, and PICs</li> </ul> </li> </ul> <p>To provision the sensor to export data through gRPCstreaming, use the telemetry Subscribe RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the Open Config and Network Agent packages. [See <a href="#">Configuring a JunosTelemetry Interface Sensor (CLI Procedure)</a>.]</p>
Layer 2 features	<ul style="list-style-type: none"> <li>• Support for Layer 2 bridging and forwarding. The QFX5130-32CD also supports Layer 2 control protocols (xSTP, LACP, and LLDP), including support for IRB.</li> </ul>

Table 10: Feature Support Added to the QFX5130-32CD (*continued*)

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> <li>• Support for L3 unicast forwarding features. The following Layer 3 unicast features on the QFX5130-32CD: <ul style="list-style-type: none"> <li>• IPv4 BGP</li> <li>• IPv4 MBGP</li> <li>• BGP 4 byte ASN support</li> <li>• BGP ADD Path (supporting 128 paths)</li> <li>• BGP Add Path community based selection</li> <li>• IS-IS</li> <li>• BFD (OSPF, ISIS, BGP)</li> <li>• Virtual router (VRF-lite) - ISIS, OSPF, BGP</li> <li>• IPv4 over GRE</li> <li>• IPv6 neighbor discovery protocol</li> <li>• Path MTU discovery</li> <li>• IPv6 ping</li> <li>• IPv4 and IPv6 static routing</li> <li>• IPv6 stateless auto configuration</li> <li>• IPv6 OSPFv3</li> <li>• IPv6 IS-IS</li> <li>• IPv6 multi-protocol BGP</li> <li>• 128 Way ECMP</li> <li>• Graceful Restart (BGP, ISIS, OSPF)</li> <li>• BGP-LS</li> </ul> </li> </ul>
Network management and monitoring	<ul style="list-style-type: none"> <li>• Analyzers and port mirroring enable you to mirror a copy of a packet to a configured destination and are supported on Layer 2 and Layer 3 interfaces. Features include: <ul style="list-style-type: none"> <li>• Native analyzer</li> <li>• Local port mirroring</li> <li>• Remote port mirroring</li> <li>• Remote port mirroring to an IP address (with GRE encapsulation)</li> <li>• Port mirroring with firewall filter action port-mirror</li> </ul> [See <a href="#">Understanding Port Mirroring and Analyzers</a>.] </li> <li>• sFlow support. sFlow is a monitoring technology for high-speed switched or routed networks. An sFlow monitoring system consists of an sFlow agent embedded in the router and a monitoring station called a collector. sFlow randomly samples network packets and sends the samples as UDP datagrams to the IP address and UDP destination port of the collector. [See <a href="#">Overview of sFlow Technology</a>.]</li> </ul>

Table 10: Feature Support Added to the QFX5130-32CD (continued)

Feature	Description
	<ul style="list-style-type: none"> <li>• Support for NETCONF XML management protocol. NETCONF is an XML-based protocol that client applications use to request and change configuration information on routing, switching, and security devices. [See <a href="#">NETCONF XML Management Protocol and Junos XML API Overview</a>.]</li> <li>• Management features supported:             <ul style="list-style-type: none"> <li>• Chef and Puppet provided as third-party applications, as part of the Junos OS Evolved image. [See <a href="#">Chef for Junos OS Getting Started Guide</a>.]</li> <li>• JET APIs (only management and rpd API sand gRPC client request-response—not Thrift—are supported)</li> <li>• YANG modeling</li> </ul> </li> <li>• Support for Standard MIBs. For information about Standard MIB objects, see the SNMP MIB Explorer. [See <a href="#">SNMP MIB Explorer</a>.]</li> </ul>
Routing protocols	<ul style="list-style-type: none"> <li>• Bidirectional forwarding detection (BFD) support for BGP, IS-IS, and OSPF/OSPFv3. [See <a href="#">Example: Configuring BFD for BGP</a> and <a href="#">Example: Configuring BFD for IS-IS</a>.]</li> <li>• Support for BGP Monitoring Protocol (BMP) Version 3 and IPv6 BGP standards. You can configure BMP, which sends BGP route information from the switch to a monitoring application, or station, on a separate device. [See <a href="#">Understanding the BGP Monitoring Protocol</a> and <a href="#">Supported IPv6 Standards</a>.]</li> <li>• Neighbor discovery protocol (NDP) proxy support. NDP functionality enables packet forwarding among the hosts that are in the same subnet and are restricted from communicating directly with each other. [See <a href="#">IPv6 Neighbor Discovery User Guide</a>.]</li> <li>• Unified Forwarding Table (UFT). [See <a href="#">Understanding the Unified Forwarding Table</a>.]</li> <li>• Support for VRF and virtual-router instances. [See <a href="#">Understanding Multicast Route Leaking for VRF and Virtual-Router Instances</a>.]</li> <li>• Static routing. [See <a href="#">Protocol-Independent Routing Properties User Guide</a>.]</li> </ul>
Software installation and upgrade	<ul style="list-style-type: none"> <li>• Secure boot support. The secure boot implementation is based on the UEFI2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected and thus safeguarded from tampering or modification. Secure boot is enabled by default on supported platforms. [See <a href="#">Software Installation and Upgrade Guide</a>.]</li> <li>• Support for ZTP using WAN interfaces. [See <a href="#">Zero Touch Provisioning</a>.]</li> </ul>

**NOTE:** To view the hardware compatibility matrix for optical interfaces and transceivers supported on the QFX5130-32CD, see the [Hardware Compatibility Tool](#).

### *Junos OS XML API and Scripting*

- **Routing instance support in Python 3 applications (PTX Series and QFX Series)**—Starting in Release 20.3R1, devices running Junos OS Evolved support specifying the routing instance used by a process or socket in Python 3 applications that are executed from the Linux shell. Python 3 applications can import the `libpyvrf` module and use the module's functions to set the Linux VRF corresponding to the Junos OS routing instance for a specific context.

[See [How to Specify the Routing Instance in Python 3 Applications on Devices Running Junos OS Evolved](#).]

### *Junos Telemetry Interface*

- **ARP and NDP table state, ON\_CHANGE support for JTI (QFX5220 switches)**—Junos OS Evolved Release 20.3R1 supports ON\_CHANGE export of Address Resolution Protocol (ARP) and Neighbor Discovery Protocol (NDP) table states to an outside collector using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. ARP and NDP periodic streaming and ON\_CHANGE support is already available in Junos OS Evolved. This feature adds interface address (IPv4, IPv6) telemetry data.

The supported resource paths (sensors) are:

- `/interfaces/interface/subinterfaces/subinterface/ipv4/neighbors/neighbor/state/` supporting leafs ip, link-layer-address. origin,
- `/interfaces/interface/subinterfaces/subinterface/ipv6/neighbors/neighbor/state/` supporting leafs ip, link-layer-address. origin, is-router, neighbor-state, and is-publish.
- `/interfaces/interface/subinterfaces/subinterface/ipv4/addresses/address/state/` supporting leafs ip, prefix-length, and origin.
- `/interfaces/interface/subinterfaces/subinterface/ipv4/addresses/address/state/` supporting leafs ip, prefix-length, and origin.
- `/interfaces/interface/subinterfaces/subinterface/ipv4/unnumbered/state/` supporting leafs enabled, interface-ref/state/interface, interface-ref/state/subinterface,
- `/interfaces/interface/subinterfaces/subinterface/ipv4/state/` supporting leafs enabled and mtu.
- `/interfaces/interface/subinterfaces/subinterface/ipv6/addresses/address/state/` supporting leafs ip, prefix-length, origin, and status.
- `/interfaces/interface/subinterfaces/subinterface/ipv6/unnumbered/state/enabled.`
- `/interfaces/interface/subinterfaces/subinterface/ipv4/unnumbered/state/` supporting leafs enabled, interface-ref/state/interface, interface-ref/state/subinterface,



- `/interfaces/interface/subinterfaces/subinterface/ipv6/unnumbered/interface-ref/state/interface/subinterface/`.
- `/interfaces/interface/subinterfaces/subinterface/ipv6/state/` supporting leafs enabled and mtu.

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

### Routing Protocols

- **Support for Unified Forwarding table and Layer 3 sub-interface (QFX5130-32CD)**—Starting with Junos OS Evolved Release 20.3R1, the following Unified Forwarding table and Layer 3 sub-interface features are supported:
  - **Unified Forwarding Table**— Unified Forwarding Table stores both Layer 2 and Layer 3 entries that allow you to set default sizes for different entries. The following applications share Unified Forwarding Table search banks:
    - Layer 2 MAC addresses
    - IPv4 and IPv6 host routes
    - Layer 2 and Layer 3 multicast routes
    - IPv4 and IPv6 LPM or prefix entries that are supported in Unified Forwarding Table through Algorithmic Longest Prefix Match (ALPM)

The following Unified Forwarding Table Profiles are supported:

Profile	Layer 2 MAC	Layer 3 Host Unicast	Layer 3 Host Multicast	IPv4 and IPv6 LPM
lpm-profile	32,000	32,000	16,000	IPv4 720,000 and IPv6 550,000
host-profile	160,000	160,000	32,000	IPv4 65,000 and IPv6 50,000
host-acl-profile	160,000	160,000	32,000	IPv4 65,000 and IPv6 22,000
default-profile (same as lpm-profile)	32,000	32,000	16,000	IPv4 720,000 and IPv6 550,000

**NOTE:** The **evo-pfemamd** Packet Forwarding Engine application restarts when one of the Unified Forwarding Table profile is configured. The default profile is active when no configuration is added. The configuration guidelines to delete the routes or MACs before changing the Unified Forwarding Table configuration, might not install all the entries from the previous Unified Forwarding Table profile due to the table size.

- **Layer 3 sub-interfaces**— You can configure Layer 3 sub-interfaces to route traffic between the VLANs. The following interfaces scales are supported:
  - Layer 2 VLANs or IRB interfaces: 1 to 4095
  - Layer 3 sub-interfaces VLAN-token: 1 to 4095

**NOTE:** The maximum VLAN scale is 4095 due to which the interfaces and VLANs share the same token space.

### Software Licensing

- **Juniper Agile Licensing (QFX5220-32CD, QFX5220-128C, PTX10003-80C, and PTX10003-160C)**  
—Starting in Junos OS Evolved Release 20.3R1, we're moving toward license-based software features. We now use Juniper Agile Licensing to support soft enforcement for software features on the listed devices. Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can install and manage licenses for hardware and software features using Juniper Agile Licensing. From this release onwards, you can now opt to use the Juniper Agile License Manager to significantly improve the ease of license management for an entire network of supported devices. If you are upgrading to this release, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS releases earlier than Junos OS Evolved Release 20.3R1.

[Table 11 on page 82](#) describes the licensing support on the QFX5220-32CD and QFX5220-128C devices.

**Table 11: Licensed Features on the QFX5220-32CD and QFX5220-128C**

QFX Switch License Model	Detailed Features
Standard license for integrated SKUs (standard hardware and software platform)	Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), QoS (Layer 2 and Layer 3), and SNMP
Advanced license for integrated and advanced SKUs	<b>Advanced 1:</b> BGP, FBF, GRE, IS-IS, JTI, MC-LAG, OSPF, sFlow, VRF, and VRRP
	<b>Advanced 2:</b> Includes Advanced 1 features + CFM, Layer 2 and Layer 3 multicast, OAM, Packet Timestamping, PTP, and Q-in-Q
Premium license for integrated and premium SKUs	Includes Advanced 2 features + EVPN-MPLS, MPLS, Layer 2 circuit, Layer 3 VPN, LDP, RSVP, segment routing, and SR-TE

Table 12 on page 83 describes the licensing support for the PTX10003-80C and PTX10003-160C devices.

**Table 12: Licensed Features on the PTX10003-80C and PTX10003-160C**

License Model	Supported Features	Scale	Security License
Standard	Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), quality of service or QoS (Layer 2 and Layer 3), and SNMP	64K FIB  K stands for 1000.	The PTX10003 router supports the MACsec feature, but you must purchase a license separately to use the feature.
Advanced	<b>Advanced 1</b>		
	BGP, FBF, GRE, IS-IS, Junos telemetry interface (JTI), OSPF, sFlow, VRF, and VRRP	256K FIB, 3M RIB, and 1K VR  M stands for million.	
	<b>Advanced 2</b>		
	Advanced 1 features, CFM, EVPN-VXLAN, Multicast, OAM, PTP, and Q-in-Q	256K FIB, 3M RIB, and 1K VRs/VRF (VXLAN)	
Premium	<b>Premium 1</b>		
	Advanced 2 features, EVPN-MPLS, IPFIX, Layer 2 circuit, Layer 3 VPN, LDP, RSVP, SR, and SR-TE	2M FIB, 6M RIB, 1K VRs/VRF (VXLAN), 32 VRF (MPLS Layer 3 VPN), and 32K LSP	
	<b>Premium 2</b>		
	Premium 1 features and fine-grained QoS	2M to 4M FIB, 60M to 80M RIB, 1K+ VRs/VRF (VXLAN), 32+ VRF (MPLS Layer 3 VPN), and 32K+ LSP	

\* Scaling is based on the device capacity.

[See [Supported Features on QFX5220-32CD and QFX5220-128C device](#), [Supported Features on PTX10003-80C and PTX10003-160C device](#), [Juniper Agile Licensing Guide](#), and [Configuring Licenses in Junos OS](#).]

## What's Changed

### IN THIS SECTION

- [What's Changed in Junos OS Evolved Release 20.3R2 | 85](#)
- [What's Changed in Junos OS Evolved Release 20.3R1 | 86](#)

Learn about what changed in Junos OS Evolved Release 20.3R2 for the QFX5130-32CD and QFX5220.

## What's Changed in Junos OS Evolved Release 20.3R2

### General Routing

- **The write-file option at the monitor traffic interface hierarchy level takes precedence**—The **write-file** option at the **monitor traffic interface** hierarchy level takes precedence over the **extensive** option when you configure them simultaneously. If you try to configure these options at the same time, Junos OS Evolved gives you a warning message that the options are not compatible, and it only runs the **monitor traffic interface write-file** command.

[See [monitor traffic](#).]

### Layer 2 Features

- **New commit check for MC-LAG (QFX Series)**—We've introduced a new commit check to check the values assigned to the redundancy group identification number on the MC-AE interface ( **redundancy-group-id** ) and ICCP peer ( **redundancy-group-id-list** ) when you configure multichassis aggregation groups (MC-LAGs). If the values are different, the system reports a commit check error. In previous releases, if the configured values were different, the l2ald process would crash.

[See [iccp](#) and [mc-ae](#).]

### MPLS

- **The show mpls lsp extensive and show mpls lsp detail commands display next-hop gateway LSPid** — When you use the **show mpls lsp extensive** and **show mpls lsp detail** commands, you'll see next-hop gateway LSPid in the output.

### Routing Protocols

- Loading of the default configurations in a RIFT package causes the following changes:
  - Output of the **show rift node status** command displays the node ID in hexadecimal number even though the node ID is configured in decimal, hexadecimal, or octal number.

- Some of the DDoS default configurations change because of the DDoS protection interferes with the RIFT BFD operation.

### *Software Licensing*

- **License key format (PTX10003-160C, PTX10003-80C, QFX5220-32CD, QFX5220-128C)**—When you are upgrading from Junos OS Evolved Release 20.3R1 to Junos OS Evolved Release 20.3R2 or later releases, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS Evolved Release 20.3R2 or later.

### *User Interface and Configuration*

- **Verbose format option to export JSON configuration data (PTX Series and QFX Series)**—The Junos OS CLI exposes the **verbose** statement at the **edit system export-format json** hierarchy level. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **edit system export-format json** hierarchy level.

[See [export-format](#).]

## **What's Changed in Junos OS Evolved Release 20.3R1**

### *Junos OS XML API and Scripting*

- **Changes to Junos XML RPC request tag names (PTX Series, QFX Series)**—We've updated the Junos XML request tag name for some operational RPCs to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include::
  - Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
  - Uppercase characters are converted to lowercase.

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

### Layer 2 Features

- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (PTX Series, 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.

### 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](#).]

## Known Limitations

### IN THIS SECTION

- [Network Security | 88](#)
- [Routing Protocols | 88](#)
- [Software Licensing | 88](#)
- [System Management | 88](#)
- [User Interface and Configuration | 89](#)

Learn about limitations in this release for the QFX5130-32CD and QFX5220. For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

## Network Security

- Configuring an egress ACL with 1000 terms and attaching the ACL to more than one interface results in exhaustion of the TCAM space. The user can attach an egress RACL to multiple interfaces as long as the numbers of terms and the number of interfaces attached remains within the maximum TCAM scale supported (1000). [PR1536550](#)

## Routing Protocols

- Few OSPF sessions are not established with 500,000 routes across 10 neighbors. [PR1536115](#)

## Software Licensing

- When you upgrade to Junos OS Evolved 20.3R1 Release, the earlier license keys do not persist on QFX5220-32CD, QFX5220-128C, PTX10003-80C, and PTX10003-160C devices. As a workaround, you must reinstall the license keys.

[See [Configuring Licenses in Junos OS using Network Licensing Mode.](#)]

## System Management

- Fragmentation does not happen for IPv4 packets on transit Junos OS Evolved based devices when the outgoing MTU is less than the packet size. The packet is dropped. [PR1429634](#)
- Due to a hardware limitation, ICMP type and ICMP code matches are not supported for egress ACLs. The complete filter does not work even if a single term has an ICMP type or ICMP code match. [PR1503168](#)
- In QFX5130-32CD, packet-length match configured in any firewall terms do not work. [PR1514578](#)
- In QFX5130-32CD, the dscp action configured in firewall filter does not work and does not rewrite the dscp bits. [PR1514580](#)
- In QFX5130-32CD, when a policer is configured as an action for a loopback filter, it does not take effect. [PR1514601](#)
- On the QFX5130 platform, the source MAC address of the IPv6 neighbour advertisement packet is not learnt and populated in the MAC-address-table. This issue is seen only when the Layer 3 gateway node is configured with IRB and the host is connected through an intermediate Layer 2 switch. This does not impact or prevent packet forwarding for the affected hosts. The MAC address of the host is learnt by the gateway node when the actual traffic from the host reaches the gateway node. [PR1524966](#)



- Ingress policer scale is limited to 128 due to known issue in Junos OS Evolved Release 20.3R1. [PR1525525](#)
- On the QFX5130-32CD platform, the port status LED remains off for a channelized interface, when one of the channels in the channelized port is down or is disabled. The port LED lights up only when all the channels part of the port are up. [PR1526532](#)

## User Interface and Configuration

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

## Open Issues

### IN THIS SECTION

- [General Routing | 90](#)
- [Infrastructure | 90](#)
- [Network Management and Monitoring | 91](#)
- [Routing Protocols | 91](#)
- [System Management | 91](#)
- [User Interface and Configuration | 92](#)

Learn about open issues in this release for the QFX5130-32CD and QFX5220. For the most complete and latest information about known Junos OS Evolved defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

## General Routing

- The following XML tags are missing in Junos OS Evolved for the **show system virtual-memory** CLI:  
`<system-virtual-memory-information>`, `<vmstat-memstat-malloc>`, `<vmstat-memstat-info>`,  
`<memstat-name>`, `<inuse>`, `<memuse>`, `<high-use>`, `<memstat-req>`, `<memstat-size>`,  
`<vmstat-memstat-zone>`, `<vmstat-zone-info>`, `<zone-name>`, `<zone-size>`, `<count-limit>`, `<used>`, `<free>`,  
`<zone-req>`, `<vmstat-sumstat>`, `<cpu-context-switch>`, `<dev-intr>`, `<soft-intr>`, `<traps>`, `<sys-calls>`,  
`<kernel-thrds>`, `<fork-calls>`, `<vfork-calls>`, `<rfork-calls>`, `<swap-pageins>`, `<swap-pagedin>`,  
`<swap-pageouts>`, `<swap-pagedout>`, `<vnode-pageins>`, `<vnode-pagedin>`, `<vnode-pageouts>`,  
`<vnode-pagedout>`, `<page-daemon-wakeup>`, `<page-daemon-examined-pages>`, `<pages-reactivated>`,  
`<copy-on-write-faults>`, `<copy-on-write-optimized-faults>`, `<zero-fill-pages-zeroed>`,  
`<zero-fill-pages-prezeroed>`, `<transit-blocking-page-faults>`, `<total-vm-faults>`, `<page-faults-requiring-io>`,  
`<pages-affected-by-kernel-thrd-creat>`, `<pages-affected-by-fork>`, `<pages-affected-by-vfork>`,  
`<pages-affected-by-rfork>`, `<pages-cached>`, `<pages-freed>`, `<pages-freed-by-daemon>`,  
`<pages-freed-by-exiting-proc>`, `<pages-active>`, `<pages-inactive>`, `<pages-in-vm-cache>`,  
`<pages-wired-down>`, `<pages-free>`, `<bytes-per-page>`, `<total-name-lookups>`, `<positive-cache-hits>`,  
`<negative-cache-hits>`, `<pass2>`, `<cache-deletions>`, `<cache-falsehits>`, `<toolong>`, `<vmstat-intr>`,  
`<vmstat-intr-info>`, `<intr-name>`, `<intr-cnt>`, `<intr-rate>`, `<vm-kernel-state>`, `<vm-kmem-map-free>`.  
[PR1438110](#)
- On Junos OS Evolved platforms based on TD4 chipset (for example, QFX5130-32CD), MAC learning rate is less than 1kpps. [PR1529754](#)
- QFX5130-32CD evo-pfemamd binding queue complete-deleted objects are seen with `hwd::pcsE`.  
[PR1531820](#)
- With transit traffic sent from scaled hosts under IRB more than 4000 to QFX5130-32CD, NDP is shooting to 100 percent. [PR1551644](#)

## 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 QFX5130-32CD. [PR1528996](#)
- Few VR BGP sessions out of 1000 are taking 6 to 8 minutes time to clear the "InboundConvergencePending" flag after interface flap. [PR1547012](#)

## Network Management and Monitoring

- Trace option is required to capture SNMP traps and also helps in debugging. [PR1546784](#)

## Routing Protocols

- The virtual router option is not supported under routing-instance in lean rpd image. [PR1494029](#)
- The instance-type option is missing under **set routing-instances foo** in lean RPD image. [PR1548297](#)

## System Management

- When changing a physical interface from a VLAN ID list to multiple single VLAN IDs or VLAN ID lists that fall within the original VLAN range, we might see one or more of the new physical interfaces not notified. For example, in the current commit we have the following configuration: **set interfaces ae100 unit 1 vlan-id-list 1-100**. In the next commit, if we run **the set interfaces ae100 unit 1 vlan-id 1 set interfaces ae100 unit 2 vlan-id 2** command, we might see that one of the two physical interfaces is not notified. [PR1504374](#)
- On the QFX5130-32CD platform running Junos OS Evolved software, you cannot clear or reset the disk option specified in the scheduled request node reboot command. The node will reboot with the disk option last specified. [PR1517596](#)
- On the QFX5130-32CD platform running Junos OS Evolved, when you try to change the port speed of the member links that are part of an already configured aggregated Ethernet bundle, the member link might remain unattached to the aggregated Ethernet after changing the port speed. [PR1526221](#)
- When the evo-pfemamd process is restarted due to any error conditions, the channelized interfaces (if configured) flaps two times (once during the initialisation) and the second time after the configuration is programmed. [PR1526984](#)

## 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: 20.3R2 | 92](#)
- [User Interface and Configuration | 94](#)
- [Resolved Issues: 20.3R1 | 94](#)

Learn which issues were resolved in Junos OS Evolved Release 20.3R2 for the QFX5130-32CD and 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: 20.3R2

#### *General Routing*

- On Junos OS Evolved platforms with policer configuration, the colored packets might not get policed due to policer inaccuracy for color-blind single-rate-three-color policer at lower bandwidth limits. [PR1497856](#)
- QFX5220 does not support the matches that were supported on QFX5200:
  - first-fragment
  - tcp-established
  - ls-fragment
  - Hop-limit[PR1499009](#)
- On the QFX5130-32CD platform, when the management port speed is configured to 100 Mbps, the port's status LED does not light up. However, the port forwards traffic normally. [PR1521510](#)

- On QFX5130-32CD platforms with a large amount of ARP resolutions happening on IRBs in a very short time, the arpd process usage can shoot to 100 percent. This issue might not occur with Layer 3 interfaces. [PR1523300](#)
- On the QFX5130-32CD platform, the ERSPAN status continues to show status as "Up" even when the destination IP address is unreachable. [PR1527505](#)
- Problem can be detected when the **show interfaces diagnostics optics** command is executed with a wild card (example et-4/0/\*) or when a physical interface name is not given at all. Execute the "show interfaces diagnostics optics" command with a specific physical interface name for all physical interfaces above PIC-port 10. [PR1527821](#)
- 160,000 MAC flushing with traffic running, takes close to 35 minutes after **clear ethernet-switching table** without traffic flushed in 3-4 minutes. [PR1528511](#)
- On the QFX5130-32CD platform, the CLI option to configure a range of VLANs as the input for the analyzer is not supported in this release. [PR1529419](#)
- Traffic loss might happen after performing GRES in Junos OS Evolved enabled chassis-based systems. [PR1532446](#)
- On all Junos OS Evolved platforms with large number of physical interfaces and logical interfaces, the picd might crash after running SNMP MIB walk and executing the **show interfaces** command multiple times (for example, 1000 times), further flap the physical interface. [PR1533766](#)
- On the platforms running Junos OS Evolved, committing classifiers configuration followed by restarting cosd might cause the cosd process to crash. [PR1536320](#)
- QFX5130-32CD: Interfaces not coming up in some instances during power cycle test and soft reboot test. [PR1538284](#)
- With a port mirror function on a QFX Series platform running Junos OS Evolved, changing the interface MTU might cause "evo-pfemand" process to restart unexpectedly. [PR1538626](#)
- On all Junos OS Evolved platforms, the software rollback might not work correctly when "request system software sync all" is run. [PR1556641](#)
- The device might not boot up after performing the cleanup operation. [PR1557020](#)

### ***Interfaces and Chassis***

- Control logical interface will be auto provisioned when there is none configured. [PR1507347](#)

### ***Routing Protocols***

- The rpd might crash after reboot when MSDP is configured. [PR1536593](#)

### **User Interface and Configuration**

- Changing domain name does not take effect. [PR1526660](#)

### **Resolved Issues: 20.3R1**

#### ***Infrastructure***

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

#### ***System Management***

- cmd crashes when the **request system shutdown reboot disk1** command is executed. [PR1492955](#)
- The interface is not part of a member VLAN when configured. [PR1495207](#)
- QFX5220-32CD and QFX5220-128C: evo-pfmand core file is seen if the loopback filter is applied with input-list and after switch is rebooted. [PR1497909](#)
- 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](#)
- QFX5130-32CD: In LPM and default UFT profile, MACs are learnt more than 32,000 in Routing Engine. [PR1528508](#)
- CoSd might crash when multiple configuration changes are made in a single commit. [PR1536320](#)

# Upgrade Your Junos OS Evolved Software

Products impacted: PTX10003, PTX10008, QFX5130-32CD, QFX5220-32CD, QFX5220-128C.

Follow these steps to upgrade your Junos OS Evolved software:

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

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

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

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

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

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

## Licensing

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

The major components of the framework are:

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

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

## Finding More Information

Learn about more information on 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](https://prsearch.juniper.net)
- Hardware Compatibility Tool—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](https://apps.juniper.net/hct/home)




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

- Juniper Networks Compliance Advisor—Review regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6 for Juniper Networks products.  
[apps.juniper.net/compliance/](https://apps.juniper.net/compliance/)

## Documentation Feedback

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

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



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

# Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

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

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

## Creating a Service Request with JTAC

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

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

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

## Revision History

29 July 2022—Revision 11, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

12 May 2022—Revision 10, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

5 May 2022—Revision 9, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

28 January 2022—Revision 8, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

7 October 2021—Revision 7, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

16 September 2021—Revision 6, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

16 July 2021—Revision 5, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

22 April 2021—Revision 4, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

12 April 2021—Revision 3, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

12 March 2021—Revision 2, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

24 February 2021—Revision 1, Junos OS Evolved Release 20.3R2 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

18 December 2020—Revision 9, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

10 December 2020—Revision 8, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

2 December 2020—Revision 7, Junos OS Evolved Release 20.3R1 for the PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

12 November 2020—Revision 6, Junos OS Evolved Release 20.3R1-S1 for the PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

29 October 2020—Revision 5, Junos OS Evolved Release 20.3R1 for the PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

15 October 2020—Revision 4, Junos OS Evolved Release 20.3R1 for the PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

8 October 2020—Revision 3, Junos OS Evolved Release 20.3R1 for the PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

30 September 2020—Revision 2, Junos OS Evolved Release 20.3R1 for the PTX10003, PTX10008, QFX5130-32CD, and QFX5220 Devices.

23 September 2020—Revision 1, Junos OS Evolved Release 20.3R1 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.

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