

Release Notes

Published
2022-07-29

Junos[®] OS Evolved 20.3R1 Release Notes

SUPPORTED ON

- PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices

HARDWARE HIGHLIGHTS

- New PTX10004 Packet Transport Router (PTX Series)
- New PTX10K-LC1202-36MR line card (PTX10008)
- New QFX5130-32CD Ethernet Switch (QFX series)

SOFTWARE HIGHLIGHTS

- VLAN tag manipulation: pop, push, and swap (PTX10003)
- Juniper Agile Licensing (PTX10003-80C, and PTX10003-160C, QFX5220-32CD, and QFX5220-128C)
- Support for virtual routing and forwarding (VRF) and source address (PTX10008)

Release Notes: Junos[®] OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices

29 July 2022

Contents	Introduction 4
	Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, and PTX10008 Devices 4
	What's New 5
	What's New in Junos OS Evolved 20.3R1-S1 6
	What's New in Junos OS Evolved 20.3R1 6
	What's Changed 74
	What's Changed in Junos OS Evolved Release 20.3R1-S3 74
	What's Changed in Junos OS Evolved Release 20.3R1-S2 74
	What's Changed in Junos OS Evolved Release 20.3R1 75
	Known Limitations 78
	General Routing 78
	Interfaces and Chassis 80
	MPLS 81
	Software Licensing 81
	User Interface and Configuration 81
	Open Issues 81
	Class of Service (CoS) 82
	Fault Management 82

Forwarding and Sampling | 82

General Routing | 82

Infrastructure | 88

Interfaces and Chassis | 88

MPLS | 89

Network Management and Monitoring | 89

Network Management and Monitoring | 89

Routing Protocols | 89

User Interface and Configuration | 89

Resolved Issues | 90

Authentication and Access Control | 90

General Routing | 90

Infrastructure | 92

Interfaces and Chassis | 93

Routing Policy and Firewall Filters | 93

User Interface and Configuration | 93

Junos OS Evolved Release Notes for QFX5130-32CD and QFX5220 Devices | 93

What's New | 94

Hardware | 95

Junos OS XML API and Scripting | 102

Junos Telemetry Interface | 102

Routing Protocols | 103

Software Licensing | 105

What's Changed | 108

What's Changed in Junos OS Evolved 20.3R1-S2 | 109

What's Changed in Junos OS Evolved 20.3R1 | 109

Known Limitations | 111

Network Security | 112

Routing Protocols | 112

Software Licensing | 112

System Management | 112

User Interface and Configuration | 113

Open Issues | 113

Interfaces and Chassis | 114

Routing Protocols | 114

System Management | 114

Resolved Issues | 115

Infrastructure | 115

System Management | 115

Upgrade Your Junos OS Evolved Software | 117

Licensing | 119

Finding More Information | 119

Documentation Feedback | 120

Requesting Technical Support | 121

Self-Help Online Tools and Resources | 121

Creating a Service Request with JTAC | 122

Revision History | 122

Introduction

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

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 20.3R1.

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 | 74](#)
- [Known Limitations | 78](#)
- [Open Issues | 81](#)
- [Resolved Issues | 90](#)

These release notes accompany Junos OS Evolved Release 20.3R1 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.

What's New

IN THIS SECTION

- [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 20.3R1 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 1 on page 7](#) summarizes the PTX10004 features supported in Junos OS Evolved Release 20.3R1.

Table 1: Features Supported by the PTX10004 Routers

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> • 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. <p>[See CoS Features and Limitations on PTX Series Routers.]</p>
DHCP	<ul style="list-style-type: none"> • Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agents.]
Firewalls and policers	<ul style="list-style-type: none"> • 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"> • Forwarding IPv4 and IPv6 • Firewall filter • Load balancing • MPLS fast reroute • Host path • Egress peer engineering <p>[See Firewall Filter Match Conditions and Actions (PTX Series Routers).]</p>

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
High availability (HA) and resiliency	<ul style="list-style-type: none"> • Support for fabric management. You can perform these fabric management functions on the PTX10004: <ul style="list-style-type: none"> • Initialize, configure, and manage fabric ASIC and other fabric components on the switch fabric board • Configure and manage PFE ASIC fabric blocks and other fabric devices on the line card (FPC). • Fabric CCL link training and link error handling • Configure and manage fabric plane • System wide synchronization of fabric operations related to FPC online/offline, SIB offline/online, and removing failed link from data path. • Fabric ASIC initialization time error handling • Fabric ASIC runtime error monitoring and handling • Fabric ASIC initialization time error handling • Fabric statistics collection • Handling RE switchover related aspects for SIB [See show chassis fabric summary.] • 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. The router also supports firmware upgrade for FPGA and U-boot. 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 show system errors active.] • Resiliency support on the Packet Forwarding Engine. The software detects, reports, and takes action on Packet Forwarding Engine faults. Actions are taken based on default configuration or user configuration available for the errors. [See show system errors active.]

Table 1: Features Supported by the PTX10004 Routers *(continued)*

Feature	Description
Interfaces and chassis	

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • Support for the dedicated virtual routing and forwarding (VRF) instance <code>mgmt_junos</code> for the management Ethernet Interface. [See Management Interface in a Nondefault Instance and management-instance.] • Supported port speeds. You can configure port speed for the JNP10K-LC1201 line card by using the port profile configuration statements in the <code>[edit interfaces]</code> hierarchy. You can choose to configure the following port speeds on all 36 ports: <ul style="list-style-type: none"> • 4x10 Gbps, 4x25 Gbps, and 2x50 Gbps • 40 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps <p>To configure the port speed of a physical port in a PIC, use the speed statement at the <code>[edit interfaces interface-name]</code> hierarchy. To configure the port speed of a physical port with a specified number of channelized interfaces, use the speed statement and specify the number of channelized interfaces by using the number-of-sub-ports statement at the <code>[edit interfaces interface-name]</code> hierarchy. To control the number of interfaces (channelized or non-channelized) created on a physical port, use the unused statement at the <code>[edit interfaces interface-name]</code> hierarchy.</p> <p>NOTE: The following CLI statements are no longer supported:</p> <ul style="list-style-type: none"> • pic-mode statement at the <code>[edit chassis fpc fpc-slot-number pic pic-number]</code> hierarchy • speed statement at the <code>[edit chassis fpc fpc-slot-number pic pic-number port port-number]</code> hierarchy • Port profile configuration for the JNP10K-LC1201 line card at this hierarchy level. <p>[See Port Speed.]</p> • 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"> • Distributed LACP (AFT-based PPMAN) • Load balancing • LAG Links or Interfaces with limited LACP capability to be up

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
	<p>[See LACP.]</p> <ul style="list-style-type: none"> • Support for chassis environment policy and power budgeting. The functions supported are: <ul style="list-style-type: none"> • Active power management • On over-heating, only the FRU shuts down instead of the entire chassis • Power control management <p>[See show chassis environment.]</p> <ul style="list-style-type: none"> • Load balancing under enhanced-hash-key includes: <ul style="list-style-type: none"> • GRE key inclusion for transit IPv4 and IPv6 traffic • IP Layer 3 fields • IP Layer 4 fields • IPv6 flow label inclusion • MPLS labels • MPLS port data • MPLS pseudowire traffic • Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets • RSVP-TE load balancing in proportion to LSP bandwidth <p>[See enhanced-hash-key.]</p> <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 Resilient Hashing on LAGs and ECMP groups.]</p>

Table 1: Features Supported by the PTX10004 Routers *(continued)*

Feature	Description
Layer 3 features	

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • Support for Layer 3 forwarding features: <ul style="list-style-type: none"> • IPv4 and IPv6 • IPv6 accounting, includes interface statistics for explicit null, PTX10008 hop popping (UHP), 6PE, and BGP labeled unicast (BGP LU). • IPv6 subnet support on loopback interfaces • IPv4 and IPv6 forwarding – limited longest prefix match (LPM) entries for OSPFv3, ISIS, BGP, static routes • IPv4 and IPv6 traceroute and ping • IPv4 and IPv6 multicast and unicast traffic statistics • Junos Telemetry Interface (JTI), including routing process daemon (RPD) and RPD Task Memory • Interface ranges • Single and Flexible VLAN-tagged Layer 3 logical interfaces • Virtual router routing instances for unicast protocols • MPLS • LAG • ECMP • Path MTU discovery • MTU checks • Routing protocol next-hop acknowledgements and time intervals • ICMP • OSPF, including support for: <ul style="list-style-type: none"> • LDP synchronization • Multitopology • OSPF database protection for OSPF and OSPFv3 • OSPF export and import policies for network-summary LSAs • OSPFv2 and OSPFv3 • Policy based support for LFA in OSPF • Remote LFA support in OSPF (IGP and LDP) • Shortcuts • SNMP MIB support for OSPFv3 • Stub network overload and external route leaking • IS-IS, including support for: <ul style="list-style-type: none"> • IPv6 and IPV6 unicast topology

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • JTI adjacency, link state database (LSDB) streaming, state, interface, adjacency statistics, and LSDB export • LDP synchronization • LFA and LFA policy based support • Multitopology • Remote LFA support for LDP in IS-IS • Shortcuts • Stub network overload <p>[See IS-IS User Guide.]</p> <ul style="list-style-type: none"> • Route-filter walkup • ARP • Static ARP • NDP • BGP • BFD • LACP • LDP • RSVP • LLDP • VRF-lite • TTL expiry • IP options • IP Fragmentation • DDOS

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
MPLS	

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • 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 [edit protocols mpls] hierarchy level: <ul style="list-style-type: none"> • no-decrement-ttl - Decrements the IP time-to-live (TTL) by 1 before encapsulating the IP packet within an MPLS packet. • no-propagate-ttl - Disables normal TTL decrementing for all RSVP-signaled or LDP-signaled LSPs [See Fast Reroute Overview, no-decrement-ttl and no-propagate-ttl.] • 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"> • MPLS ping and traceroute for single IS-IS and OSPF node or prefix segment • BGP-LS segment routing extensions for IS-IS and OSPF • IS-IS: Adjacency SID, link and node protection for IPv4 addressing in TI-LFA • OSPF: Inter-area support, Static adjacency SID, link and node protection in TI-LFA • Advertising maximum link bandwidth and administrative color without RSVP-TE configuration for IS-IS and OSPF • 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 • IGP adjacency SID hold time • PCEP for segment routing LSPs • BGP IPv4 labeled-unicast resolution over: BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF, non-colored and static colored IPv4 SR-TE with segment routing using IS-IS and OSPF • BGP Layer 3 VPN over colored and non-colored SR-TE tunnels, and IPv4 protocol next hops • BGP-triggered dynamic SR-TE colored tunnels • Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs • First-hop label support for SID instead of an IP address • Path specification using router IP addresses

Table 1: Features Supported by the PTX10004 Routers (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • First-hop label support for SID instead of an IP address • SR-TE color mode: 00—Route resolution fallback to IGP path and 01—Route resolution fallback to color only null routes • Static LSPs with member-link next-hops for aggregated Ethernet bundles <p>[See Understanding Source Packet Routing in Networking (SPRING).]</p>
Network management and monitoring	<ul style="list-style-type: none"> • 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 Automation Scripting User Guide.] • Support for adding nonnative YANG modules to the Junos OS schema. [See Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS.] • 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 Puppet for Junos OS Administration Guide.] • Chef for Junos OS support. [See Chef for Junos OS Getting Started Guide.] • Supports third-party application Open-R.
Timing and synchronization	<ul style="list-style-type: none"> • Support for NTP over IPv4 and IPv6. [See Time Management Administration Guide.]

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 2: Features Supported by the PTX10001-36MR Routers

Feature	Description

Table 2: Features Supported by the PTX10001-36MR Routers (*continued*)

Feature	Description
Firewalls and policers	<ul style="list-style-type: none"> • Support for independent micro BFD sessions are enabled on a per member link basis for a LAG bundle. [See Configuring Micro BFD Sessions for LAG.] • 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 Firewall Filter Match Conditions and Actions (PTX Series Routers).] • Support for matching ip-options in IPv4 packet headers. The ip-options any match condition is supported at the [edit firewall family inet filter name term name from] hierarchy level. [See Firewall Filter Match Conditions for IPv4 Traffic.]
High availability (HA) and resiliency	<ul style="list-style-type: none"> • 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 show chassis fabric errors autoheal.] • 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 show system errors active.] • 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 show system errors active.]

Table 2: Features Supported by the PTX10001-36MR Routers (*continued*)

Feature	Description
Interfaces and chassis	<ul style="list-style-type: none"> • 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, thus providing performance improvement over the existing process, which requires a recirculation path to support de-encapsulation. [See Configuring IP-IP Decapsulation by Tunnel Termination on FTI.] • 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 Unicast Reverse-Path-Forwarding.] • Support for port profile configuration. You can configure the port speed on the router by using the port profile configuration statements at the [edit interfaces] hierarchy level. [See Port Speed.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • 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 Guidelines for gRPC and gNMI Sensors.] • Support for Packet Forwarding Engine sensor. Supports UDP (native) streaming of Packet Forwarding Engine statistics using JTI. [See sensor (Junos Telemetry Interface).] • 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 Guidelines for gRPC and gNMI Sensors.]

Table 2: Features Supported by the PTX10001-36MR Routers (*continued*)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 2 VPN. You can configure the following Layer 2 VPN features: <ul style="list-style-type: none"> • Transport of Layer 2 frames over MPLS (LDP signaling) • Simple Ethernet and VLAN-based CCC. • Local and remote switching • Ethernet and VLAN-based CCC and single tagged CCC logical interfaces • Control word • Regular and aggregated Ethernet interfaces • 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 • Virtual circuit connectivity verification (VCCV) with type 2 and type 3 <p>[See Layer 2 VPNs and VPLS Feature Guide for Routing Devices.]</p>

Table 2: Features Supported by the PTX10001-36MR Routers (*continued*)

Feature	Description
MPLS	<ul style="list-style-type: none"> • Support for segment routing: <ul style="list-style-type: none"> • MPLS ping and traceroute for single IS-IS and OSPF node or prefix segment • BGP-LS segment routing extensions for IS-IS and OSPF • 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 • IS-IS: Adjacency SID, link and node protection for IPv4 addressing in TI-LFA • OSPF: Interarea support, Static adjacency SID, link and node protection in TI-LFA • Advertising maximum link bandwidth and administrative color without RSVP-TE configuration for IS-IS and OSPF • 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 • IGP adjacency SID hold time • PCEP for segment routing LSPs • BGP IPv4 labeled-unicast resolution over: BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF, non-colored and static colored IPv4 SR-TE with segment routing using IS-IS and OSPF • BGP Layer 3 VPN over colored and non-colored SR-TE tunnels, and IPv4 protocol next hops • BGP-triggered dynamic SR-TE colored tunnels • Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs • First-hop label support for SID instead of an IP address • Path specification using router IP addresses • SR-TE color mode: 00—route resolution fallback to IGP path 01—route resolution fallback to color only null routes • Static LSPs with member-link next hops for aggregated Ethernet bundles <p>[See Understanding Source Packet Routing in Networking (SPRING).]</p>

Table 2: Features Supported by the PTX10001-36MR Routers (*continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> • 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, 40 GbE, 100 GbE, and 400GbE interfaces. [See link-degrade-monitor.] • Support for DCU accounting and SCU accounting. SCU accounting provides breakdown of output interface traffic statistics for each prefix from which the traffic originated. DCU accounting provides breakdown of input interface traffic statistics for each prefix to which the traffic is destined. [See Understanding Source Class Usage and Destination Class Usage Options.] • Support for LACP features. LACP monitors the links to check whether both ends of the bundle are connected to the correct group. [See Configuring Aggregated Ethernet Interfaces.] • Support for load balance properties such as IPv6 flow label, GTP hashing and weighted ECMP. [See Load Balancing and Ethernet Link Aggregation.] • 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 Understanding Port Mirroring.] • 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 Passive Monitoring.] • 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 Tunneling LDP LSPs in RSVP LSPs Overview.] • Support for BGPLU FRR MPLS inter-AS link protection. Link protection ensures traffic restoration in case of an interface failure. [See BGP User Guide.] • 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 Link Protection for MPLS LSPs.] • 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 IPv6 Neighbor Discovery User Guide.]

Table 2: Features Supported by the PTX10001-36MR Routers (*continued*)

Feature	Description
Services applications	<ul style="list-style-type: none"> • 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 Configuring Inline Active Flow Monitoring.] • 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 [edit services monitoring rpm] hierarchy level. [See Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.]

- **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. The 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 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 3 on page 24](#) summarizes the PTX10K-LC1202-36MR features supported in Junos OS Evolved Release 20.3R1.

Table 3: Features Supported by the PTX10K-LC1202-36MR

Feature	Description
BGP	<ul style="list-style-type: none"> • 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 tunneled from CPE devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop. To configure a dynamic IPv4-over-IPv6 tunnel, include the dynamic-tunnels statement at the [edit routing-options] hierarchy level. [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]
Class of service (CoS)	<ul style="list-style-type: none"> • 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 CoS Features and Limitations on PTX Series Routers.] • Classification override support. Classification override is configured under a forwarding policy on PTX10K-LC1202-36MR line cards on PTX10008 devices. [See CoS Features and Limitations on PTX Series Routers and Overriding the Input Classification.] • Support for CoS-based forwarding. You can configure CoS-based forwarding (CBF) on PTX10K-LC1202-36MR line cards for PTX10008 devices. [See Forwarding Policy Options Overview.]
DHCP	<ul style="list-style-type: none"> • Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agents.]
Firewalls and policers	<ul style="list-style-type: none"> • 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"> • Firewall filter • DCU accounting and SCU accounting [See Firewall Filter Match Conditions and Actions (PTX Series Routers).]

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
High availability (HA) and resiliency	<ul style="list-style-type: none"> • 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 show system errors active.] • Support for micro BFD sessions and child links of aggregated Ethernet or LAG bundle with interface IP addresses. [See Configuring Micro BFD Sessions for LAG.] • Support for resilient hashing and consistent 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 Resilient Hashing on LAGS and ECMP and Load Balancing for a BGP Session.] • 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 show system errors active.]

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
Interfaces and chassis	

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • 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 [edit interfaces fti0 unit number tunnel encapsulation IPIP] hierarchy level. Tunnel termination is completed in a single pass of packet processing, thus providing performance improvement over existing process which requires a recirculation path to support decapsulation. [See Configuring GRE Tunnel Interfaces.] • GRE tunnel decapsulation using flexible tunnel interfaces (FTI). You can configure GRE tunnels on flexible tunnel interfaces (FTIs). When you enable the tunnel-termination statement at the [edit interfaces fti0 unit unit-number] CLI hierarchy, tunnels are terminated on the WAN interface before any other actions, such as sampling, port mirroring, or filtering, are applied. GRE tunnel encapsulation is not supported. [See Configuring GRE Tunnel Interfaces.] • You can monitor link degradation of the 10-Gigabit Ethernet interfaces and 100-Gigabit Ethernet interfaces. To enable your device to monitor the links, use the link-degrade-monitor statement at the [edit interfaces interface-name] hierarchy level. [See Link Degrade Monitoring Overview.] • Support for Media Access Control Security (MACsec), including AES-256 encryption, extended packet numbering, and fail-open mode. [See Configuring Media Access Control Security (MACsec) on Routers.] • 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. <ul style="list-style-type: none"> • 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. • 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. [See Port Speed on PTX10K-LC1202-36MR Overview.] • Support for the dedicated virtual routing and forwarding (VRF) instance mgmt_junos for the management Ethernet Interface. [See Management Interface in a Nondefault Instance and management-instance.]

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> 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"> Configuration synchronization for MC-LAG groups Distributed LACP (AFT-based PPMAN) Load balancing MC-LAG links or interfaces with limited LACP capability to be up Telemetry support (periodic streaming and on-change) [See LACP.] Load balancing under enhanced-hash-key includes: <ul style="list-style-type: none"> GRE key inclusion for transit IPv4 and IPv6 traffic IP Layer 3 fields IP Layer 4 fields IPv6 flow label inclusion MPLS labels MPLS port data MPLS pseudowire traffic Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets RSVP-TE load balancing in proportion to LSP bandwidth [See enhanced-hash-key.] Support for QSFP28 100GE DWDM optical transceiver module (shown in CLI as QSFP-100GE-DWDM2) for 100GbE applications. The transceiver is suited for data center interconnect which requires high-fiber-capacity for links up to 80 km. The features supported by QSFP28 100GE DWDM transceiver include the following: <ul style="list-style-type: none"> View the optics inventory information. [See show chassis hardware and show chassis pic fpc-slot slot pic-slot View the diagnostics data, warnings, and alarms for interfaces. [See show interfaces diagnostics optics.] Configure the optics loopback mode. [See optics-options.]

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
Junos telemetry interface (JTI)	

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • 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"> • Packet Forwarding Engine statistics (resource path <code>/junos/system/linecard/packet/usage</code>) • Network processing unit (NPU) memory statistics (resource path <code>/junos/system/linecard/npu/memory/</code>) • NPU utilization statistics (resource path <code>/junos/system/linecard/npu/utilization/</code>) <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p> • 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"> • Routing Engine logical interface statistics (resource path <code>/interfaces/interface/</code>). Applicable for management interfaces (re0:mgmt-0 and re1:mgmt-0). • Address Resolution Protocol (ARP) (resource path <code>/arp-information/</code>) • Network Discovery Protocol (NDP) (resource path <code>/nd6-information/</code>) • Operational state of AE with LACP (resource path <code>/lACP/</code>) • Operational state of LLDP (neighbors information) (resource path <code>/lldp/</code>) • Sensor for NDP router-advertisement statistics (resource path <code>/ipv6-ra/</code>) <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p> • 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"> • <code>/junos/system/linecard/interface/</code> • <code>/junos/system/linecard/interface/logical/usage/</code> <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface) and sensor (Junos Telemetry Interface).]</p>

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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>). • 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>). • NPU utilization statistics. This sensor exports Packet Engine utilization and packet load (sensor name is <code>/junos/system/linecard/npu/utilization/</code>). <p>[See sensor (Junos Telemetry Interface).]</p>
Layer 2 features	<ul style="list-style-type: none"> • Support for Link Layer Discovery Protocol (LLDP). LLDP support includes management interfaces and disabling of LLDP time, length, and value (TLV) messages. [See LLDP Overview.] • 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 Restricted and Unrestricted Proxy ARP Overview.] • Support for Layer 2 VPNs. You can configure the following Layer 2 VPNs: <ul style="list-style-type: none"> • Transport of Layer 2 frames over MPLS (LDP signaling) • Layer 2 VPNs over tunnels (BGP signaling) • Simple Ethernet and VLAN based cross connect (also known as connections) • Local and remote switching 5) Ethernet and VLAN-CCC • Single tagged CCC logical interfaces • Control word • Regular and aggregated Ethernet interfaces • Layer 2 protocol pass-through • Layer 2 circuit backup interface and backup neighbor • Layer 2 circuit statistics and CoS • VCCV with type 2 and type 3 • Layer 2 circuit statistics and CoS <p>[See Layer 2 VPNs and VPLS Feature Guide for Routing Devices.]</p>

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
Layer 3 features	

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • Support for Layer 3 forwarding features: <ul style="list-style-type: none"> • IPv4 and IPv6 • IPv6 accounting, includes interface statistics for explicit null, hop popping (UHP), 6PE, and BGP labeled unicast (BGP LU) • IPv6 subnet support on loopback interfaces • IPv4 and IPv6 forwarding – limited longest prefix match (LPM) entries for OSPFv3, ISIS, BGP, static routes • IPv4 and IPv6 traceroute and ping • IPv4 and IPv6 multicast and unicast traffic statistics • Junos Telemetry Interface (JTI), including routing process daemon (RPD) and RPD task memory • Interface ranges • Single and Flexible VLAN-tagged Layer 3 logical interfaces • Virtual router routing instances for unicast protocols • MPLS • LAG • ECMP • Path MTU discovery • MTU checks • Routing protocol next-hop acknowledgements and time intervals • ICMP • OSPF, including support for: <ul style="list-style-type: none"> • LDP synchronization • Multitopology • OSPF database protection for OSPF and OSPFv3 • OSPF export and import policies for network-summary LSAs • OSPFv2 and OSPFv3 • Policy based support for LFA in OSPF • Remote LFA support in OSPF (IGP and LDP) • Shortcuts • SNMP MIB support for OSPFv3 • Stub network overload and external route leaking • [See OSPF User Guide.] • IS-IS, including support for: <ul style="list-style-type: none"> • IPv6 and IPV6 unicast topology • JTI adjacency, link state database (LSDB) streaming, state,

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<p>interface, adjacency statistics, and LSDB export</p> <ul style="list-style-type: none"> • LDP synchronization • LFA and LFA policy based support • Multitopology • Remote LFA support for LDP in IS-IS • Shortcuts • Stub network overload <p>[See IS-IS User Guide.]</p> <ul style="list-style-type: none"> • Route-filter walkup • ARP • Static ARP • NDP • BGP • BFD • LACP • LDP • RSVP • LLDP • VRF-lite • TTL expiry • IP options • IP Fragmentation • DDOS

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
Layer 3 VPNs	<ul style="list-style-type: none"> • 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"> • MPLS over Layer 3 VLAN-tagged subinterfaces • Per-next-hop label allocation • Mapping of the label-switched interface (LSI) logical interface label to the VPN routing and forwarding (VRF) routing table using the vrf-table-label statement • ICMP tunneling and MPLS traceroute • Disabling time-to-live (TTL) decrementing using no-propagate-ttl • IPv6-over-MPLS using: <ul style="list-style-type: none"> • IPv6 VPN Provider Edge routers (6VPE) • IPv6 Provider Edge routers (6PE) <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 Layer 3 VPNs Feature Guide for Routing Devices and Tunneling IPv6 Traffic over MPLS IPv4 Networks.]</p>

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
MPLS	

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • 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"> • MPLS ping and traceroute for single IS-IS and OSPF node or prefix segment • BGP-LS segment routing extensions for IS-IS and OSPF • IS-IS: Adjacency SID, link and node protection for IPv4 addressing in TI-LFA • OSPF: Inter-area support, Static adjacency SID, link and node protection in TI-LFA • Advertising maximum link bandwidth and administrative color without RSVP-TE configuration for IS-IS and OSPF • 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 • IGP adjacency SID hold time • PCEP for segment routing LSPs • BGP IPv4 labeled-unicast resolution over: BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF, non-colored and static colored IPv4 SR-TE with segment routing using IS-IS and OSPF • BGP Layer 3 VPN over colored and non-colored SR-TE tunnels, and IPv4 protocol next hops • BGP-triggered dynamic SR-TE colored tunnels • Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs • First-hop label support for SID instead of an IP address • Path specification using router IP addresses • First-hop label support for SID instead of an IP address • SR-TE color mode: 00—Route resolution fallback to IGP path and 01—Route resolution fallback to color only null routes • Static LSPs with member-link next-hops for aggregated Ethernet bundles <p>[See Understanding Source Packet Routing in Networking (SPRING).]</p> <ul style="list-style-type: none"> • MPLS support. You can configure the following MPLS features: <ul style="list-style-type: none"> • CLI support for monitoring MPLS label usage • Inline MPLS and IPv6 lookup for explicit null • Transit LSP scale: 128,000 without FRR configuration and 32,000 with FRR configuration • Explicit null support for MPLS LSPs

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • MPLS Label Block Configuration • MPLS over untagged Layer 3 interfaces • MPLS OAM - LSP ping • JTI: OCST: MPLS operational state streaming (v2.2.0) • Ingress LSP scale: 16,000 without FRR configuration and 2,000 with FRR configuration • Egress LSP scale: 100,000 without FRR configuration • Entropy Label support • MPLS: JTI: Junos Telemetry Interface • MPLS self-ping, TE++, and misc augmentation • LDP including: <ul style="list-style-type: none"> • Configurable label withdraw delay • Egress policy • Explicit null • Graceful restart signalling • IGP synchronization • Ingress policy • IPv6 for LDP transport session • Strict targeted hellos • Track IGP metric • Tunneling (LDP over RSVP)

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> • RSVP++ • RSVP-TE including: <ul style="list-style-type: none"> • Bypass LSP static configuration • Ingress LSP statistics in a file • RSVP-TE Hitless-MBB with no artificial delays • 32,000 transit LSPs • Transit LSP scale: 2,000 without FRR configuration and 2,000 with FRR configuration • Auto bandwidth • Class based forwarding (CBF) with 16 classes • CBF with next-hop resolution • Convergence and scalability • Graceful restart signalling • JTI interface statistics and LSP event export • LSP next-hop policy • LSP self ping • MPLS fast reroute (FRR) • LSP self ping • Optimize adaptive teardown • Node/link protection • Refresh reduction • Soft preemption • Shared Risk Link Group (SRLG) • Static LSPs with IPv4 next-hop, IPv6 next-hop, and IPv6 next-hop with next-table support for bypass • Traffic engineering including: TE++: Dynamic ingress LSP splitting, Traffic engineering extensions (OSPF-TE and ISIS-TE), Traffic engineering options: bgp, bgp-igp, bgp-igp-both-ribs, and amps-forwarding <p>[See MPLS Applications User Guide.]</p>
Multicast	<ul style="list-style-type: none"> • 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 Multicast Overview.]

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
Network management and monitoring	

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
	<ul style="list-style-type: none"> 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. <p>NOTE: 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 Understanding Source Class Usage and Destination Class Usage Options.]</p> <ul style="list-style-type: none"> Load balancing under enhanced-hash-key includes: <ul style="list-style-type: none"> GRE key inclusion for transit IPv4 and IPv6 traffic IP Layer 3 fields IP Layer 4 fields IPv6 flow label inclusion MPLS labels MPLS port data MPLS pseudowire traffic Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets RSVP-TE load balancing in proportion to LSP bandwidth [See enhanced-hash-key.] Support for running third-party applications inside Linux containers. [See Overview of Third-Party Applications on Junos OS Evolved.] Support for third-party application Open-R. 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"> Adaptive sampling Adaptive sampling fallback which decreases the sampling load when the traffic load decreases after adaptive sampling has taken place Ability to configure forwarding class and DSCP values per collector [See Understanding How to Use sFlow Technology for Network Monitoring and SFlow.] 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

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
	<p>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:</p> <ul style="list-style-type: none"> • Interface filter on ingress • Interface filter on egress • Forwarding table filter (FTF) on ingress • Unicast traffic on ingress and egress • Multicast traffic on ingress and egress • Families inet and inet6 • Firewall filter action port-mirroring in both ingress and egress directions • Aggregated Ethernet interfaces at both ingress and egress • Sampling range of 1-16000000 • Packet clipping (maximum-packet-length) • run-length attribute to specify the number of samples following trigger events <p>Use the following CLI hierarchies to configure port mirroring:</p> <ul style="list-style-type: none"> • [edit interfaces] • [edit forwarding-options port-mirroring] • [edit firewall filter] <p>You can configure family inet and family inet6 in the [edit interfaces] and the [edit forwarding-options port-mirroring] hierarchies for this feature. The feature applies to global port mirroring only.</p> <p>[See Understanding Port Mirroring and Analyzers.]</p>

Table 3: Features Supported by the PTX10K-LC1202-36MR (continued)

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • 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 protection statement at the [edit protocols bgp group group-name family inet labeled-unicast] hierarchy level. Junos OS Evolved supports the following baseline BGP features: <ul style="list-style-type: none"> • Large communities • Link bandwidth community aggregation • Load balancing in proportion to available bandwidth community • local-as option • BGP Monitoring Protocol (BMP) version 3 • MTU discovery • BGP multipath • Multiple paths to IPv6 addresses • PIC core and inter-AS link protection for BGP-LU • Prefix prioritization • Route reflection • tcp-mss configuration • Route leaking between routing instances, including default (inet) routing instance through RIB-groups and instance import/export
Services applications	<ul style="list-style-type: none"> • 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 [edit services monitoring rpm] hierarchy level. [See Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.] • Support for inline active flow monitoring (J-Flow) for IPv4, IPv6, and MPLS traffic. Both IPFIX and Version 9 template formats are supported. [See Understanding Inline Active Flow Monitoring.]

Table 3: Features Supported by the PTX10K-LC1202-36MR (*continued*)

Feature	Description
System management	<ul style="list-style-type: none"> 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 [Passive Monitoring.]
Timing and synchronization	<ul style="list-style-type: none"> Support for NTP over IPv4 and IPv6.

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 4: Features Supported by the PTX10008

Feature	Description
Hardware	<ul style="list-style-type: none"> 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 PTX10008 Routing and Control Boards Components and Descriptions.] 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"> PTX10008-BASE3—One routing engines, 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. 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. <p>[See PTX10008 Configurations and Upgrade Options.]</p>

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
High availability (HA) and resiliency	<ul style="list-style-type: none"> 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 show system errors active.] Support for VRRP. The following features are not supported for VRRP on Junos OS Evolved: <ul style="list-style-type: none"> ISSU Proxy ARP MC-LAG Distribution support on aggregated Ethernet interface (ae) IRB Inline delegation <p>VRRP sessions might flap during GRES in centralized mode. [See Understanding VRRP.]</p>

Table 4: Features Supported by the PTX10008 (*continued*)

Feature	Description
Interfaces and chassis	

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
	<ul style="list-style-type: none"> 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 request system software add package_name command. [See Hardware Supported by Junos Continuity Software.] Unicast RPF support for both IPv4 and IPv6 traffic flows. [See Example: Configuring Unicast Reverse-Path-Forwarding Check.] Support for configuring GRE tunnel encapsulation on FTIs using the loopback interface. You can configure encapsulation by using the command tunnel encapsulation gre source address destination address at the [edit interfaces fti0 unit unit] hierarchy. Keep in mind the following when configuring this feature: <ul style="list-style-type: none"> Adding tunnel-termination makes the tunnel decap-only tunnel and encapsulation will be disabled. Both the source and destination address is mandatory when you don't configure the tunnel-termination command. Configuring a variable prefix mask on the source address isn't allowed. [See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.] 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 native-vlan-id.] Support for the discard interface. [See Discard Interfaces.] Support for the following transceivers: <ul style="list-style-type: none"> 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.

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
	<ul style="list-style-type: none"> • 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. <p>NOTE: These transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).</p> <p>[See the Hardware Compatibility Tool (HCT) for details.]</p> <ul style="list-style-type: none"> • Support for adaptive load balancing (ALB) on multiple Packet Forwarding Engines for aggregated Ethernet bundles. <p>NOTE: With ALB configured, the show interfaces ae0 extensive command displays adaptive statistics information under the physical interface, not the logical interface.</p> <p>[See Load Balancing on Aggregated Ethernet Interfaces.]</p> <ul style="list-style-type: none"> • 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 Configuration Guidelines for Aggregated Ethernet Interfaces.]

Table 4: Features Supported by the PTX10008 (*continued*)

Feature	Description
Junos telemetry interface (JTI)	

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
	<ul style="list-style-type: none"> 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. <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"> Transmitted packets and transmitted bytes Red drop packets and bytes Tail drop packets and bytes <p>This feature includes zero suppression support. It does not include support for summed up counters on aggregated Ethernet (ae) interfaces.</p> <p>[See sensor (Junos Telemetry Interface) and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p> <ul style="list-style-type: none"> 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. <p>The supported resource paths (sensors) are:</p> <ul style="list-style-type: none"> <code>/interfaces/interface/subinterfaces/subinterface/ipv4/neighbors/neighbor/state/</code> supporting leafs ip, link-layer-address, and origin <code>/interfaces/interface/subinterfaces/subinterface/ipv6/neighbors/neighbor/state/</code> supporting leafs ip, link-layer-address, origin, is-router, neighbor-state, and is-publish <code>/interfaces/interface/subinterfaces/subinterface/ipv4/addresses/address/state/</code> supporting leafs ip, prefix-length, and origin <code>/interfaces/interface/subinterfaces/subinterface/ipv4/addresses/address/state/</code> supporting leafs ip, prefix-length, and origin <code>/interfaces/interface/subinterfaces/subinterface/ipv4/unnumbered/state/</code> supporting leafs enabled, interface-ref/state/interface, and interface-ref/state/subinterface <code>/interfaces/interface/subinterfaces/subinterface/ipv4/state/</code> supporting leafs enabled and mtu <code>/interfaces/interface/subinterfaces/subinterface/ipv6/addresses/address/state/</code> supporting leafs ip, prefix-length, origin, and status <code>/interfaces/interface/subinterfaces/subinterface/ipv6/</code>

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
	<p>unnumbered/state/enabled</p> <ul style="list-style-type: none"> • /interfaces/interface/subinterfaces/subinterface/ipv4/unnumbered/state/ supporting leafs enabled, interface-ref/state/interface, and 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 <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p> <ul style="list-style-type: none"> • This release supports export of Source Packet Routing in Networking (SPRING) statistics to an outside collector using remote procedure call (gRPC) services. <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"> • For egress (per-interface egress) use the set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress at the [edit] hierarchy level. • For egress (per-SID egress) use the set protocols isis source-packet-routing sensor-based-stats per-sid egress at the [edit] hierarchy level. • For ingress (per-SID ingress) use the set protocols isis source-packet-routing sensor-based-stats per-sid ingress statement at the [edit] hierarchy level. <p>Use the following sensors to export statistics by means of gRPC services to an outside collector:</p> <ul style="list-style-type: none"> • /junos/services/segment-routing/interface/egress/usage/ for egress (per-interface egress) aggregate SPRING traffic. • /junos/services/segment-routing/sid/usage/ for egress (per-SID egress) and ingress (per-SID ingress) aggregate SPRING traffic. <p>[See source-packet-routing and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p>

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
Multicast	<ul style="list-style-type: none"> 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).
Network management and monitoring	<ul style="list-style-type: none"> 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 [edit forwarding-options port-mirroring instance <i>instance-name</i> output] hierarchy. [See instance (Port Mirroring).]
Port security	<ul style="list-style-type: none"> Support for fallback preshared key (PSK) for MACsec. [See Configuring MACsec with Fallback PSK.]
Routing policy and firewall filters	<ul style="list-style-type: none"> 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"> inet—Both input and output directions inet6—Both input and output directions mpls—Input direction only You can also bind them to routing instances, and in the input direction, the output direction, or both directions. [See Guidelines for Nesting References to Multiple Firewall Filters and Example: Nesting References to Multiple Firewall Filters.] Support for configuring priority for route prefixes through existing import policy in protocols. [See Configuring Priority for Route Prefixes in RPD Infrastructure.]

Table 4: Features Supported by the PTX10008 (continued)

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • Support for inline BFD. [See Understanding Bidirectional Forwarding Detection (BFD).] • BGP flow specification support. The following match conditions are not supported: <ul style="list-style-type: none"> • ICMP codes alone inet/inet6 • Source/destination prefix with offset for inet6 • Flow label for inet6 • Fragment for inet6 The following action is not supported: <ul style="list-style-type: none"> • Traffic marking [See Understanding BGP Flow Routes for Traffic Filtering.] • Support for BGP-LU over SR-TE for color-based mapping of VPN services. [See Understanding Static Segment Routing LSP in MPLS Networks.] • Support for BGP routes with n-multipath primary and 1-protection backup gateway. [See multipath (Protocols BGP), delay-route-advertisements, and egress-te.] • 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 BGP PIC Edge Using BGP Labeled Unicast Overview.]
Services applications	<ul style="list-style-type: none"> • 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 Configuring Inline Active Flow Monitoring on PTX Series Routers.] • 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 Understanding Inline Active Flow Monitoring and Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers.]

Table 4: Features Supported by the PTX10008 (*continued*)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> • 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 Junos OS Installation Package Names.] • 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 QFX10008 Configurations and Upgrades.] • 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 MX10008 Components and Configurations.] • Support for ZTP on WAN interfaces. [See Zero Touch Provisioning Overview.] • ZTP supports the DHCPv6 client. [See Zero Touch Provisioning Overview.]

- **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 (shown in the CLI as QSFP-100GE-DWDM2) for 100GbE applications. 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 5: Routing Protocol Features Supported by the PTX10003

Technology	Supported Features
BGP	

Table 5: Routing Protocol Features Supported by the PTX10003 (*continued*)

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

Table 5: Routing Protocol Features Supported by the PTX10003 (*continued*)

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

Table 5: Routing Protocol Features Supported by the PTX10003 (continued)

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

Table 5: Routing Protocol Features Supported by the PTX10003 (continued)

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

Table 5: Routing Protocol Features Supported by the PTX10003 (continued)

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

Table 5: Routing Protocol Features Supported by the PTX10003 (continued)

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

Table 5: Routing Protocol Features Supported by the PTX10003 (*continued*)

Technology	Supported Features
OpenConfig	<ul style="list-style-type: none"> • BGP configuration model (v2.1.1) <p>[See OpenConfig User Guide.]</p>
OSPF	<ul style="list-style-type: none"> • Policy-based inbound route filtering • Active backbone detection • Multitopology routing OSPF • Export external route to multiple area scoped type 7 LSAs • OSPFv3 SNMP MIB • Route install prioritization • Database protection • Per-prefix LFA • Option to overload stub networks through OSPF overload <p>[See OSPF User Guide.]</p>
OSPFv3	<ul style="list-style-type: none"> • Address family support • MIB support according to RFC5643
Path Computation Element Protocol (PCEP)	<ul style="list-style-type: none"> • MD5 authentication for PCC/PCEP • Point-to-multipoint • No-ERO or loose ERO support • Path computation for segment routing LSP • Support for PCEP MIB • Support for latest version of PCEP RFC 8231 • PCE support for RSVP-TE
Programmable Routing Protocol (PRPD)	<ul style="list-style-type: none"> • gRIBI RIB programming interface • BGP SR-TE policy AFI • BGP flowspec AF • Move programmed routes to hidden state upon next-hop interface failure <p>[See Juniper Extension Toolkit API Guide.]</p>

Table 5: Routing Protocol Features Supported by the PTX10003 (*continued*)

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

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 6 on page 72](#) describes the licensing support on the QFX5220-32CD and QFX5220-128C devices.

Table 6: 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	Advanced 1: BGP, FBF, GRE, IS-IS, JTI, MC-LAG, OSPF, sFlow, VRF, and VRRP
	Advanced 2: 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 7 on page 72](#) describes the licensing support for the PTX10003-80C and PTX10003-160C devices.

Table 7: Licensed Features on the PTX10003-80C and PTX10003-160C

Table 7: Licensed Features on the PTX10003-80C and PTX10003-160C (continued)

License Model	Detailed Features for Fabric Management	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	Advanced 1		
	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.	
	Advanced 2		
	Advanced 1 features, CFM, EVPN-VXLAN, Multicast, OAM, PTP, and Q-in-Q	256K FIB, 3M RIB, and 1K VRs/VRF (VXLAN)	
Premium	Premium 1		
	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	
	Premium 2		
	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 20.3R1-S3 | 74](#)
- [What's Changed in Junos OS Evolved Release 20.3R1-S2 | 74](#)
- [What's Changed in Junos OS Evolved Release 20.3R1 | 75](#)

Learn about what changed in Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, and PTX10008.

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><leaf-1>data</leaf-1><leaf-2>data</leaf-2><leaf-3>data</leaf-3><leaf1>data<leaf1>data<leaf2><leaf2><leaf3></leaf3></container>
will now appear correctly as:<container> <leaf-1> data </leaf-1>data <leaf2 data </leaf2 <leaf3 data
</leaf3 <container> <leaf1> data </leaf1> <leaf2> data </leaf2> <leaf3> data </leaf3> </container>
<container></leaf-1><leaf-2>data</leaf2><leaf3>data</leaf3><leaf1>data<leaf1>data<leaf2><leaf2><leaf3></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](#).]

General Routing

- **Inet6 is disabled in VT interface (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in this release, the `inet6` statement at the `[edit interfaces vt-interface-number unit unit-number family]` hierarchy level is disabled.

Interfaces and Chassis

- **Change in support for `interface-transmit-statistics` statement**—You cannot configure aggregated Ethernet interfaces to capture and report the actual transmitted load statistics by using the `interface-transmit-statistics` statement. Aggregated Ethernet interfaces do not support reporting of the transmitted load statistics. In Junos OS Evolved release 20.3R1, the `interface-transmit-statistics` statement is not supported in the aggregated Ethernet interfaces hierarchy. In earlier releases, the

interface-transmit-statistics statement was available in the aggregated Ethernet interfaces hierarchy but not supported.

- **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

- **Specify interpreter for on-device JET applications (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**—You can specify whether a device running Junos OS Evolved should run a daemonized on-device JET application using Bash, Python, or Python 3. Use the **interpreter (bash | python | python3)** statement at the **[edit system extensions extension-service application file filename]** hierarchy to specify the interpreter.

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

- **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 | 78](#)
- [Interfaces and Chassis | 80](#)
- [MPLS | 81](#)
- [Software Licensing | 81](#)
- [User Interface and Configuration | 81](#)

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

- Some XML tags for the **show system buffers** command are missing in Junos OS Evolved. [PR1429626](#)
- On PTX Series devices running Junos 19.3R1, the **show ddos-protection protocols eoam oam-cfm** command returns blank output. This command is supposed to be disabled on PTX Series devices. There is no functional impact; it is just an extra visible command option that has no functionality. [PR1456043](#)
- During fabric link bringup/fabspoked-fchip restart/sib offline or sib online, /re0/fabspoked-fchip might not respond to CLI commands (**show chassis sibs**). This is indicated by the following output on CLI error: **communication failure with /re0/fabspoked-fchip/**.
 1. The producer app is either down or unresponsive. Run the **show system processes node node | grep app** command to check if the app is running. Run the **show system application app** to check the state of the app.

2. If there is a change in mastership recently, wait for the switchover to complete for the app to be online. Run **request chassis routing-engine master switch check** to check the status of switchover. If the command is reissued after training or detrainig reaches a steady state, valid output shows up in the CLI. [PR1459430](#)

- On PTX10008 routers, the subsystem within the Packet Forwarding Engine continues to monitor the ASIC for new interrupts even for the ASIC for which all the interfaces are disabled. On an ASIC that has all the interfaces disabled due to a fatal error, all the new interrupts that are reported after the fatal event should be ignored. [PR1470391](#)
- PTX10008 graceful OIR leads to missing of sensor data for power and temperature. Follow the optimized step given as a workaround to have OIR working. [PR1478951](#)
- Restarting the fabspoked-pfe application for the line card restarts the line card. [PR1486023](#)
- The local repair time for fast reroute is 50 ms. If the system has a scaling configuration or is heavily loaded for processing, the local repair time may be longer than 50 ms. In this case, it gets 65 ms local repair time. [PR1489139](#)
- UDP feature does not work through Management IP. [PR1491755](#)
- Fabric drops counter shows 0 in the **show pfe statistics traffic** CLI command output. [PR1494226](#)
- On the PTX10003 router, it takes around 40 seconds for traffic to flow through all 4093 VLANs after the configuration is committed. [PR1496757](#)
- On PTX10003, with a MAC scale of 32,000, the MAC learning rate achieved is 15,000 MAC entries per second and with a MAC scale of 64,000, it is 6000 MAC entries per second. [PR1498568](#)
- SA and DA reject interface filters are not supported on PTX10008. [PR1500789](#)
- PCIe AER uncorrectable errors might be seen on the console when SIBs are powered off during system boot, when CLI is offline or when a new Routing Engine connects to the SIBs after mastership switchover. These errors are expected and are not indicative of underlying functional issues. [PR1501647](#)
- When an FPC is removed ungracefully, an alarm is generated indicating **Fpcx Node unreachable**. This alarm is cleared only when the FPC is plugged back into the same slot. [PR1503450](#)
- It was found that DSCP mark action for inet6 is not supported in Junos OS Evolved Packet Forwarding Engine. We might observe unexpected behavior if a filter is configured with IPv6 DSCP mark action. [PR1504463](#)
- MAC entries are stuck in I2ald or I2alm for some time after the **clear ethernet switching** command is executed. [PR1507812](#)
- Software works as designed. For a faulty link, after autoheal was triggered by jresil and before the action is completed, if fabspoked-fchip restarts, the link stays in fault state, after the app restarts. [PR1508915](#)
- Once the FPC restarts, the related FPC UDP statistics are not exported and the only way to recover from this state is to reboot the device under test and resubscribe to UDP sensor again. [PR1516432](#)
- MPLS ping does not work for RSVP LSP with UHP configured. [PR1517870](#)

- If a ZF pio fault is hit followed by a second ZF pio fault on the same sib before the sib has completely recovered from the first one, it may result in nondeterministic state of fabric links to/from the impacted sib. [PR1519855](#)
- For PTX10008 platforms, input or output bytes under ifd statistics include Layer 2 header. [PR1524650](#)
- In Junos OS Evolved 20.3, fabsopke-fchip generates a core file if fabsopke-fchip restarts and SIB offline happens one after other with in the same minute. Any previous alarm does not get cleared. [PR1525577](#)
- We see 34% when we set to have 50%. For example: the following configuration is meant to see 50% traffic mirrored; however, we see 34% mirrored:

set forwarding-options port-mirroring instance lns1 input run-length 10

set forwarding-options port-mirroring instance lns1 input rate 20 [PR1527129](#)

- Excess-rate configuration in non SPS mode scheduler map might not be completely honoured in some scenarios where guaranteed-rate is configured without any excess-rate value, but some other queues have excess-rate explicitly programmed. In that case the left over bandwidth after honouring the guaranteed rate across all OQs might still go more to the guaranteed marked OQs with no excess-rate programmed, thus drawing more share from the excess region to this guaranteed-rate configured OQ than the intended limits. [PR1528124](#)

Interfaces and Chassis

- When highly accelerated life test is carried out on PTX10003, FPC error messages **egp_intr_pkt_trapcode** are seen on the console. There is no direct impact on control plane protocol or to data traffic. [PR1425508](#)
- On PTX10003, when the Picd app restarts, sometimes the 10GbE link status remains down. The link can be recovered by disabling or enabling the remote end. [PR1488146](#)
- When the loopback configuration is applied to a LAG (aggregated Ethernet) interface, the interface goes down. [PR1497591](#)

MPLS

- LMP is not supported on Junos OS Evolved platforms [PR1524699](#)

Software Licensing

- When you upgrade to Junos OS Evolved 20.3R1 Release, the earlier license keys does 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.](#)]

User Interface and Configuration

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

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 82](#)
- [Fault Management | 82](#)
- [Forwarding and Sampling | 82](#)
- [General Routing | 82](#)
- [Infrastructure | 88](#)
- [Interfaces and Chassis | 88](#)
- [MPLS | 89](#)
- [Network Management and Monitoring | 89](#)
- [Network Management and Monitoring | 89](#)
- [Routing Protocols | 89](#)
- [User Interface and Configuration | 89](#)

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.

Class of Service (CoS)

- Support import - Classifier/Rewrite. [PR1483505](#)
- If you have to take an interface out of aggregated Ethernet bundle and configure it to operate in stand alone mode, then doing this in single commit may render the operation ineffective and could lead to connectivity issues. There is a known issue around this and this is seen due to a race condition between Routing Engine processes (CoSD, DCD/Chassisd), Packet Forwarding Engine and kernel. The below document link speaks of this issue when there is explicit COS configuration to be made on the interface - [Limitations on CoS for Aggregated Interfaces](#). However, the problem can be seen without explicit CoS too as there is default CoS that is always present. In some cases, it is possible that a single shot commit will send out multiple operational messages down to kernel and might confuse the kernel to do unintended optimisation that could lead to a message being consumed at kernel and not being sent to Packet Forwarding Engine. The result is the same even in this case. [PR1504287](#)
- While configuring WRED profile to a scheduler, use either of
 - a) any/any
 - b) not-any/not-any combination of protocol and loss priority. [PR1524259](#)
- PTX10008 JNP10K-LC1202 : CoS rewrite not working after router reboot. [PR1542070](#)

Fault Management

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

Forwarding and Sampling

- When GRES is triggered by SSD hardware failure, the syslog error of **rpdd[2191]: krt_flow_dfwd_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out** might be seen. [PR1397171](#)

General Routing

- The convergence time is of the order of 3508 ms for OSPF or IS-IS (50,000 routes scale, which is equivalent to 14,253 routes per second). In case of BGP, it is of the order of around 17,000 routes per second. [PR1379961](#)
- On disabling both the primary and secondary commit and then rollback and getting them up, traffic does not resume on primary. [PR1382695](#)

- No application detected during unified ISSU in case the same application is in offline state in the base image. [PR1438686](#)
- Traffic loss of up to 400 ms can be seen in MPLS FRR scenario. [PR1472908](#)
- The JFlow sampled traffic might not get rate limited in host-pipe and cause drop in throughput of interfaces on PTX10008 routers. This issue is seen due to the missing code for rate-limiters specifically on PTX10008. [PR1473844](#)
- Use traffic statistics instead of Packet Forwarding Engine statistics to understand the flow of traffic. [PR1478244](#)
- In PTX10001-36MR when all CoS queues are over-subscribed, then maximum latency may vary from 21ms to 29ms for different queues. [PR1478811](#)
- UDP sensor streaming does not currently support FPC restart and Junos OS Evolved apps restart. [PR1492096](#)
- On PTX10003-80C and PTX10003-160C systems, with sampling enabled and a FIB scale of 2 million, the rpd agent takes 8 minutes to complete restart after restart. [PR1493882](#)
- The copying of files to the RCB over WAN ports is slow. This is observed across all platforms running Junos OS Evolved platforms.. [PR1496895](#)
- On a PTX10008, six SIBs are required to carry line-rate traffic, with no fabric redundancy. Even when ingress traffic rate is such that five SIBs are sufficient to carry ingress traffic (for example, traffic is less than 1280 Gbps), ungraceful SIB failures result in transient loss of traffic, till system failure handling is triggered. In Junos OS Release 20.1R2, failure handling may result in about 4-6 seconds of traffic loss. We recommend that you take the fabric cards offline by using the request chassis sib offline command before removing the SIBs for maintenance.
[PR1497212](#)
- The **hwdr: HWD_FIRMWARE_VERSION_READ_ERROR_NOTICE: Failed to get firmware version for fru Sib** message is seen for each SIB after it is gracefully taken offline and brought online. This is expected behavior as the SIB version is read correctly during an online transition. Check that the SIB versions are being displayed correctly using **show system firmware** after the SIBs are online. [PR1504156](#)
- Resilient hash seed configuration is not supported in PTX10003 and PTX10008. [PR1504461](#)
- CLI command not available to display Packet Forwarding Engine hashing enabled fields in Junos OS Evolved. [PR1504544](#)
- When you add or delete IPv6 hash-key, the IPv4 and IPv6 source+destination address pair hashing is also disabled. Hence traffic load balancing does not work correctly. [PR1509694](#)
- BGP-SRTE binding-sid with more than one label stack needs enhancement for PTX10003-80C PTX10003-160C. [PR1512213](#)
- 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 which is being rendered is deleted or

modified. If we are in middle of rendering value for that particular node, RPD might crash when telemetry job resumes since the node is no more valid. [PR1512773](#)

- PTX10004 :: STS LED goes off for ~1 min while linecard is still booting up
- PTX10004:: STS LED goes off for ~1 min while linecard is still booting up. [PR1514884](#)
- IPv6 neighbor entry is deleted in the **show ipv6 neighbors** table; therefore IPv6 traffic forwarding does not happen or traffic gets dropped after interface flap while passing transit traffic with packet length of 9000. [PR1515034](#)
- 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 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 or delete or 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 the functioning of the system. The command is **request system application app orchestrator node master re restart**. [PR1516384](#)
- **set routing-options forwarding-table chained-composite-next-hop** is hidden on PTX10003 and PTX10008 and visible on PTX10001-36MR. [PR1516384](#)
- Once the FPC restarts the related FPC UDP stats will not be exported and only way to recover from this state is to reboot the DUT and resubscribe to UDP sensor again. [PR1516432](#)
- Traffic loss is seen for L3VPN streams in segment routing configuration. [PR1516604](#)
- **set routing-options forwarding-table chained-composite-next-hop** is hidden on PTX10008 or PTX10003. [PR1518237](#)

When all 5X400 g ports are used inside 1 portion of Chip, there can be impact on the traffic throughput performance. [PR1518368](#)

- On all Junos OS Evolved systems with the scaled setup of 4k logical interfaces on the aggregated Ethernet interface and 8k policers on the firewall filter, the aggregated Ethernet interface might stay down after deactivating and then activating the firewall filter and the aggregated Ethernet interface. Traffic loss might be seen due to the aggregated Ethernet interface is down. [PR1518601](#)
- [cos] [scheduler] PTX10001-36MR:When **exact-rate** is used along with **strict- high** scheduler priority, then interface queue rate might not be as expected. [PR1519313](#)
- With all queues oversubscribed, latency time varies from 5ms ~ 27 ms for queue#0~#7 (l3_cos_jitter.robot). [PR1519345](#)
- There is a delay enforced between back to back switchover. This is 360 seconds unless overridden by platform specific values. So, if the last switchover happened within the last 360 seconds, then it will give this error: **Not ready for mastership switch, try after x seconds?. The show system switchover CLI output might show switchover Status: Ready**. However, this field indicates the ready status of all applications on backup Routing Engine. This does not consider the back-to-back switchover delay. [PR1519364](#)

- Sometimes we observe system reboot delay after **request system software add reboot request system software rollback reboot** commands. [PR1525286](#)
- Command **request system application restart app packetio-bt** is no longer available. The way to restart packet I/O is by restarting entire line card. [PR1527140](#)
- The MLD control packet is not forwarded out to next router, the packet is trapped to RE. The packet is terminated on the local router. [PR1527674](#)
- **request system software sync** command will fail if junos-evo-install-ptx-x86-64-19.4R2-S2.5-EVO is one of the installed version on master Routing Engine. [PR1528163](#)
- RPD KPI improvement: RPD agent took 24 seconds after switchover to become active (All done from show platform ha-controller app rpdagent history). [PR1528264](#)
- PTX10008 Doon RCB does not raise alarm "Mixed Master and Backup RE 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](#)
- If ddos violations are seen during fpc restart, crash is seen and fpc restarts. [PR1529847](#)
- PTX10008 does not deploy on 19.4. With later images this issue is not seen. [PR1529876](#)
- Software rollback from junos-evo-install-ptx-x86-64-19.4R2-S2.5-EVO is not supported to latest image. New release should be scratch installed even if image is available as rollback image. [PR1529884](#)
- SR-TE label with BSID traffic gets dropped in PTX10003. Use the following configuration protocols **source-packet-routing no-chained-composite-next-hop**. [PR1529933](#)
- For initial synchronization, carrier transition field is only valid for Mib2D. [PR1530864](#)
- HTTP file transfers through wan ports could be slow resulting in longer file transfer times. [PR1531192](#)
- FPC vmcore files can be stored at /var/lib/ftp/in/fpc_slot/ on RE0/RE1. [PR1531214](#)
- PTX10008: [Error] Jexpr: NhList:- JexprNhlist: Update failed for token: 11536. [PR1531287](#)
- When multicast traffic is running over AE bundle, which spans multiple FPCs, if one of FPCs is brought offline, the traffic on that FPC is impacted. It is expected to see some traffic drops. [PR1531859](#)
- In rare cases, when interface statistics of an individual interface is cleared, then queue counter for that interface may show garbage value (a 20 digit counter value). [PR1535523](#)
- When multicast is over aggregated Ethernet bundle which spans multiple FPCs, if one FPC is restarted, the multicast traffic can be duplicated or loss on the receivers. [PR1535545](#)
- After rollback to baseline configuration, it brings down all interfaces including aggregated Ethernet logical interfaces. The ifcomplist associated with this aggregated Ethernet is removed too. However, RPD is in middle of adding an comp-nh but it sends a wrong ifcomplist which is no more valid. Comp->Uncast->IFL(ae), till unicast is freed, logical interface is struck in down state. [PR1535680](#)
- PTX10003 : Ingress Traffic loss is seen when PIC is offline or brought online on another FPC. [PR1535913](#)
- **show chassis clocks** support required on PTX10008. [PR1536024](#)

- **show version** and **show version detail** missing few tags on Junos OS Evolved compared to Junos OS. [PR1536027](#)
- Sometimes, an individual line card restart might render some of the interfaces on that FPC getting blocked. [PR1536037](#)
- **show chassis environment** reports some ZF internal temperatures as 0 degrees C. [PR1536497](#)
- When 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 may fail and all the traffic may get dropped. Work around for this problem is: 1. Ensure MTU size of interface is less than 9500B 2. If 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](#)
- Triton multicast software is incomplete up to Junos OS Evolved Release 20.3R1 and does not work correctly for the BT.dp1 half of ports per BTchip. [PR1536963](#)
- It takes time for traffic to converge after restart routing. [PR1537376](#)
- Dependency-state details for the type net::juniper::hwd::fruHwStatus seen after hwdfpc app restart with lsr core profile configuration. [PR1537409](#)
- In the map of node to handle, if the keys collide, an incorrect handle might be fetched and causes crash. [PR1537700](#)
- In a scaled setup its possible that during GRES operation, the rpd on new master Routing Engine can receive an ifstates update (DEL and CHANGE) and these are not expected during rpd. Re-syncing in such case can core the rpd can core and it might recover automatically after sometime. [PR1537947](#)
- In PTX10001-36MR, while removing loopback configurations on port 5-to-port 7 (of each PIC) with optics connected, the interface link might go down. [PR1538393](#)
- When port mirror is configured for an interface and its MTU is changed, it results in a evo-aftmand core that brings down the node. [PR1538626](#)
- There might be fake collision when installing multicast egress next hop. It prints out error messages, also slow down multicast next hop installation. [PR1538849](#)
- In some network topology, auto-RP might not be learnt when PTX10008 is in discovery mode. [PR1539152](#)
- FT Stress: Packet Forwarding Engines on FPC stuck in "READY" state on new master Routing Engine with GRES switch while FPC is restarting. [PR1539685](#)
- 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](#)
- IPv6 addressing is configured at /64 subnet by default irrespective of subnet configured on the DHCP server side. [PR1539839](#)

- In PTX10001-36MR when ports are used with 40G optics in 4x10G channelized mode, then ping fails on channel2 and cahnnel3 for following ports - et-0/1/9, et-0/1/10, et-0/1/11, et-0/2/1, et-0/2/2, et-0/2/3 , [PR1539864](#)
- When new member is added to the AE bundle which has other members on the same fpc, traffic loss can happen if multicast traffic is running over the aggregated Ethernet bundle. [PR1539912](#)
- sysmand core seen with lsr core profile configuration @OfpDomainZeroClient::processNodeStatus. [PR1539998](#)
- PTX10008 : traffic loss seen after graceful switchover with LSR Core MPLS_DSCP profile configuration. [PR150400](#)
- When the IRB is configured in a bridge-domain, the multicast traffic drop might be seen after modifying the port configuration. The unicast and broadcast traffic is not affected. [PR1540495](#)
- PTX10008 : 'SIB 0 ZFO internal temp' sensor normal fan speed temp threshold is set to 0. [PR1540576](#)
- PTX10008 : core generated while testing aggregated Ethernet filters. [PR1540625](#)
- PTX10008 JNP10K-LC1202 : savecore.sh fails to transfer core and journal files due to directory create errors. [PR1540807](#)
- PTX10008: Component Upgrade: Software rollback timed out when moving from image basedon 20.3R1 (mod in some apps) to Junos OS Evolved 20.3R1-202009111855.0 same image. [PR1541509](#)
- PTX10001-36MR : :: PTX10001-36MR:PDT:rpdc generates a core file, upgrade to Junos OS Evolved 20.3-202009152333.0.[PR1541718](#)
- Commit error is seen even when gre- and underlying interfaces are configured in same routing instances. [PR1541833](#)
- PTX10008 : [Resiliency PSM] PSM fault alarm not shown in show system alarms after GRES swtich over. [PR1541951](#)
- PTX10008 JNP10K-LC1201: Not getting expected output bps after applying a filter-specific policer egrees filter. [PR1541978](#)
- PTX10008 : commit fails on RE1 and succeeds on RE0. [PR1542064](#)
- PTX10008 JNP10K-LC1201 : rpd core seen @ krt_reset_krtq_op_mcnh with lsr core profile configuration. [PR1542511](#)
- PTX10008 JNP10K-LC1201 : One of the ingress side 400G interfaces keeps flapping, when specific CoS configuration is loaded with the Junos OS Evolved 20.3R1.14 image (The config uses all the 8x400G interfaces from the same BT chip. i.e 8x400G ingress traffic going to 8x400G Egress aggregated Ethernet bundles). [PR1542681](#)

Infrastructure

- When using source IP as 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 gets dropped by the Linux kernel as the rpf check is set to strict by default on the linux kernel used on Junos OS Evolved. Linux kernel expects the path to the peer to be on WAN side and so the packet gets dropped when it is received on the management interface from the peer when the rpf check for the management interface is set to strict. [PR1498255](#)
- Junos OS Evolved LDP graceful-restart in large topologies might result in traffic loss. [PR1518609](#)
- 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 with PTX10004. [PR1528996](#)
- PTX10K-LC1202-36MR MS: Continuously log messages: **ttp_update_with_mark_tlv:390] Loss priority not supported yet.** [PR1536732](#)
- PTX10K-LC1202-36MR MS: Continuously log messages: **jtd_add_ifl:112] ifl count exceeded the max(10000),increase /proc/sys/net/jarvis/jtd_max_ifl** [PR1536740](#)
- In PTX10001-36MR when ports are used with 40G optics in 4x10G channelized mode, then ping fails on channel2 and cahnnel3 for following ports - et-0/1/9, et-0/1/10, et-0/1/11, et-0/2/1, et-0/2/2, et-0/2/3. [PR1539864](#)

Interfaces and Chassis

- [chassisd] [generic_evo] Junos OS Evolved PTX10003-80C and PTX10003-160C - fabspoked generates a core file on configuration and deletion fpc power off or on. [PR1395788](#)
- XML tags mismatch for the command **show interfaces diagnostics optics** [PR1529316](#).
- scan-interval configuration statement under **interfaces ae-x aggregated-ether-options load-balance adaptive** does not take effect. [PR1538854](#)
- **rx-signal-avg-optical-power rx-signal-avg-optical-power-dbm** ->Regarding these two ,in Junos OS it gets displayed only for 10G SFP .However for 100g QSFP in Junos OS, laser_rx_optical_power is displayed which is same across Junos OS Evolved and Junos OS. For 10G SFP in Junos OS Evolved laser_rx_optical_power and laser_rx_optical_power_dbm is displayed which is same for QSFP and SFP .And internally calculation of rx-signal-avg-optical-power in Junos OS and laser_rx_optical_power is same and only naming difference is there. Because in Junos OS Evolved irrespective of SFP/QSFP same naming convention is followed. [PR159316](#)
- scan-interval configuration statement under interfaces ae-x aggregated-ether-options load-balance adaptive does not take effect. [PR159316](#)
- On Junos OS Evolved platforms, FTI (Flexible Tunnel Interfaces) tunnel MTU change fails with family inet or inet6 on PTX10001-36MR and PTX10008. [PR1540431](#)

MPLS

- PTX10008 - seeing mpls.0 table in IP-BGP configuration. [PR1526266](#)
- When no-propagate-ttl configuration is present on the router and we are doing an activate and deactivate of the MPLS configuration, RPD can core rarely at times generate a core file if the interface change messages arrive before TTL configuration change could take effect on the ingress Static LSP. This is due to the mismatch between the no-propagate-ttl values in the gateway in the route table and the gateway in the Next-Hop attached to the Static LSP . [PR1528460](#)

Network Management and Monitoring

- PTX10001-36MR supports puppet agent version 3.6.1. Puppet does not work if master version is not 3.x since this is not backward compatible. [PR1491329](#)

Network Management and Monitoring

- PTX10001-36MR supports puppet agent version 3.6.1. Puppet does not work if Master version is not 3.x since this is not backward compatible. [PR1491329](#)

Routing Protocols

- Junos OS Evolved does not support for configuring interface name as next hop. [PR1497012](#)
- BGP did not advertise route with ELC and load-balance-label is not as stitching point. [PR1529387](#)
- rpd crash @ msdp_listen_start during upgrade with MSDP configuration loaded on the device. [PR1536593](#)
- After performing GRES, drop for a few multicast streams is seen till the rpd comes up on the new master (after which traffic loss starts to recover) which might be due to rpd or rpd-agent sending unexpected RT_DELETE notification to the Packet Forwarding Engine. [PR1538176](#)

User Interface and Configuration

- In Junos OS Evolved 19.4R1, this is a product limitation. Workaround is to use **commit** without any configuration changes from master Routing Engine in this use case. [PR1465291](#)
- When changing the configuration between scale configuration and configuration containing no filter or apply-path, sometimes there is a delay in deleting the apply-path addresses. [PR1492765](#)
- Compressed /var/log/ creation through file archive compress might fail on certain PTX Series platforms running Junos Evolved. [PR1522339](#)

Resolved Issues

IN THIS SECTION

- [Authentication and Access Control | 90](#)
- [General Routing | 90](#)
- [Infrastructure | 92](#)
- [Interfaces and Chassis | 93](#)
- [Routing Policy and Firewall Filters | 93](#)
- [User Interface and Configuration | 93](#)

Learn which issues were resolved in Junos OS Evolved Release 20.3R1 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

- 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](#)
- cmddd 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 firewalld 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 (master) 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 longetivity 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](#)
- FCP does not boot if primary BIOS is corrupted (With FPC Primary BIOS corruption Golden BIOS failed to reprogram Primary). [PR1528469](#)
- The crash is triggered by uninitialized pointer. In getAllTargetElements() function, initialize the targetsNh to nullptr to prevent use targetsNh with uninitialized random value. Adding check if targetsNh is nullptr before use it. [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 | 94](#)
- [What's Changed | 108](#)
- [Known Limitations | 111](#)
- [Open Issues | 113](#)
- [Resolved Issues | 115](#)

These release notes accompany Junos OS Evolved Release 20.3R1 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.

What's New

IN THIS SECTION

- [Hardware | 95](#)
- [Junos OS XML API and Scripting | 102](#)
- [Junos Telemetry Interface | 102](#)
- [Routing Protocols | 103](#)
- [Software Licensing | 105](#)

Learn about new features introduced in Junos OS Evolved Release 20.3R1 for the QFX5130-32CD and QFX5220.

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 8 on page 96](#) summarizes the QFX5130-32CD features supported in Junos OS Evolved Release 20.3R1.

Table 8: Features Supported by the QFX5130-32CD

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> • 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 CoS Support on QFX Series Switches, EX4600 Line of Switches, and QFabric Systems.]
DHCP	<ul style="list-style-type: none"> • 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 DHCP Relay Agent.]
Firewalls and policers	<ul style="list-style-type: none"> • Firewall filter support on Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions (QFX5220).] • Support for Firewall Service and Interfaces Service JET APIs. [See Juniper Engineering Network website.]

Table 8: Features Supported by the QFX5130-32CD (continued)

Feature	Description
Interfaces and chassis	<ul style="list-style-type: none"> 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"> Four 100Gbe ports Four 25Gbe ports Four 10Gbe ports [See Channelizing Interfaces on QFX5130-32CD Switches.] 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 Generic Routing Encapsulation.] 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 Management Interface in a Nondefault Instance and management-instance.] Support for link aggregation groups (LAGs). Layer 3 logical interfaces. [See Layer 3 Logical Interfaces.]
JET APIs	<ul style="list-style-type: none"> 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 Juniper Engineering Network website.]
Junos OS XML API and scripting	<ul style="list-style-type: none"> XSLT, SLAX, and Python scripts (which can trigger CLI operations).

Table 8: Features Supported by the QFX5130-32CD (*continued*)

Feature	Description
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • 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"> • Sensor for Routing Engine internal interfaces (physical interface state and statistics) • Sensor for Routing Engine internal Logical interfaces (logical interface state and statistics) • Sensor for power supply module (PSM) FRU model number and fan trays • Sensor for hardware and chassis • Sensor for FEBs, FPCs, and PICs <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 Configuring a Junos Telemetry Interface Sensor (CLI Procedure).]</p>
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 2 bridging and forwarding. The QFX5130-32CD also supports Layer 2 control protocols (xSTP, LACP, and LLDP), including support for IRB.

Table 8: Features Supported by the QFX5130-32CD (*continued*)

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • Support for L3 unicast forwarding features. The following Layer 3 unicast features on the QFX5130-32CD: <ul style="list-style-type: none"> • IPv4 BGP • IPv4 MBGP • BGP 4 byte ASN support • BGP ADD Path (supporting 128 paths) • BGP Add Path community based selection • IS-IS • BFD (OSPF, ISIS, BGP) • Virtual router (VRF-lite) - ISIS, OSPF, BGP • IPv4 over GRE • IPv6 neighbor discovery protocol • Path MTU discovery • IPv6 ping • IPv4 and IPv6 static routing • IPv6 stateless auto configuration • IPv6 OSPFv3 • IPv6 IS-IS • IPv6 multi-protocol BGP • 128 Way ECMP • Graceful Restart (BGP, ISIS, OSPF) • BGP-LS

Table 8: Features Supported by the QFX5130-32CD (continued)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> • 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"> • Native analyzer • Local port mirroring • Remote port mirroring • Remote port mirroring to an IP address (with GRE encapsulation) • Port mirroring with firewall filter action port-mirror [See Understanding Port Mirroring and Analyzers.] • 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 Overview of sFlow Technology.] • 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 NETCONF XML Management Protocol and Junos XML API Overview.] • Management features supported: <ul style="list-style-type: none"> • Chef and Puppet provided as third-party applications, as part of the Junos OS Evolved image. [See Chef for Junos OS Getting Started Guide.] • JET APIs (only management and rpd API sand gRPC client request-response—not Thrift—are supported) • YANG modeling • Support for Standard MIBs. For information about Standard MIB objects, see the SNMP MIB Explorer. [See SNMP MIB Explorer.]

Table 8: Features Supported by the QFX5130-32CD (continued)

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • Bidirectional forwarding detection (BFD) support for BGP, IS-IS, and OSPF/OSPFv3. [See Example: Configuring BFD for BGP and Example: Configuring BFD for IS-IS.] • 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 Understanding the BGP Monitoring Protocol and Supported IPv6 Standards.] • 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 IPv6 Neighbor Discovery User Guide.] • Unified Forwarding Table (UFT). [See Understanding the Unified Forwarding Table.] • Support for VRF and virtual-router instances. [See Understanding Multicast Route Leaking for VRF and Virtual-Router Instances.] • Static routing. [See Protocol-Independent Routing Properties User Guide.]
Software installation and upgrade	<ul style="list-style-type: none"> • 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 Software Installation and Upgrade Guide.] • Support for ZTP using WAN interfaces. [See Zero Touch Provisioning.]

Table 8: Features Supported by the QFX5130-32CD (continued)

Feature	Description
---------	-------------

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-pfemand** 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 9 on page 106](#) describes the licensing support on the QFX5220-32CD and QFX5220-128C devices.

Table 9: 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	Advanced 1: BGP, FBF, GRE, IS-IS, JTI, MC-LAG, OSPF, sFlow, VRF, and VRRP
	Advanced 2: 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 7 on page 72](#) describes the licensing support for the PTX10003-80C and PTX10003-160C devices.

Table 10: Licensed Features on the PTX10003-80C and PTX10003-160C

Table 10: Licensed Features on the PTX10003-80C and PTX10003-160C (continued)

License Model	Detailed Features for Fabric Management	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	Advanced 1		
	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.	
	Advanced 2		
	Advanced 1 features, CFM, EVPN-VXLAN, Multicast, OAM, PTP, and Q-in-Q	256K FIB, 3M RIB, and 1K VRs/VRF (VXLAN)	
Premium	Premium 1		
	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	
	Premium 2		
	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 20.3R1-S2 | 109](#)
- [What's Changed in Junos OS Evolved 20.3R1 | 109](#)

Learn about what changed in Junos OS Evolved Release 20.3R1 for the QFX5130-32CD and QFX5220.

What's Changed in Junos OS Evolved 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](#).]

What's Changed in Junos OS Evolved 20.3R1

Class of Service (CoS)

- **Priority-based flow control (PFC) support (QFX5120-32C)**—Starting with Junos OS 20.3R1, QFX5120-32C switches support priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic.

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](#).]

General Routing

- **Inet6 is disabled in VT interface (QFX5130-32CD and QFX5220)**—Starting in this release, the **inet6** statement at the **[edit interfaces vt-interface-number unit unit-number family]** hierarchy level is disabled.

Juniper Extension Toolkit

- **Specify interpreter for on-device JET applications (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**—You can specify whether a device running Junos OS Evolved should run a daemonized on-device JET application using Bash, Python, or Python 3. Use the **interpreter (bash | python | python3)** statement at the **[edit system extensions extension-service application file filename]** hierarchy to specify the interpreter.

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

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 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, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**—If auto-bandwidth adjustment fails because of bandwidth unavailable error, the router tries to bring up the LSP with the same bandwidth during the subsequent reoptimization. In earlier releases, when the auto-bandwidth adjustment fails, the current bandwidth is reset to the bandwidth that was already active.

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

- **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

- [Network Security | 112](#)
- [Routing Protocols | 112](#)
- [Software Licensing | 112](#)
- [System Management | 112](#)
- [User Interface and Configuration | 113](#)

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 interface attached remains with the max 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 the 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 the Junos OS Evolved Release 20.3R1. [PR1525525](#)
- On the QFX 5130-32CD platform, the port status LED remains off for a channelized interface, when one of the channel in the channelized port is down or when disabled. The port LED lits 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

- [Interfaces and Chassis | 114](#)
- [Routing Protocols | 114](#)
- [System Management | 114](#)

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.

Interfaces and Chassis

- The aggregate member links field in the **show interfaces extensive** command output for an aggregated Ethernet interfaces does not reflect number of member-link in the aggregated Ethernet bundle when the aggregated Ethernet interface is configured with multiple logical interfaces. [PR1517841](#)

Routing Protocols

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

System Management

- When a link goes up or down, Tx/Rx loss of signal alarm from diagnostics optics might take a long time (around 10 minutes) to show up. [PR1499741](#)
- 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, the user shall not be able to 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, when the management port speed is configured to 100Mbps, the port's status LED does not light up. However, the port forwards traffic normally. [PR1521510](#)
- On the QFX5130-32CD platform running Junos OS Evolved, when user tries to change change the port speed of the member links which are part of an already configured aggregated Ethernet bundle, the member link may remain unattached to the aggregated Ethernet after changing the port speed. [PR1526221](#)
- On the QFX5130-32CD platform, the ERSPAN status continues to show status as "Up" even when the destination IP address is unreachable. [PR1527505](#)
- Relearning of MACs post flush with running traffic takes 30-35 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](#)
- fsck gets automatically run but is not guaranteed to recover disk corruption after an abrupt power cycle. Secondary disk or partition and snapshot functionality is available to help recover from such issues. [PR1533232](#)

- The maximum power consumption for FPC is not exported in the JTI sensor for the FPC due to a known issue on the QFX5130 platform. The same can be viewed using the CLI command **show chassis fpc detail**. [PR1533630](#)
- When port mirror is configured for an interface and its MTU is changed, it results in an evo-aftmand core that brings down the node. [PR1538626](#)
- Commit error is seen even when gre and underlying interfaces are configured in same routing instances. [PR1541833](#)

Resolved Issues

IN THIS SECTION

- [Infrastructure | 115](#)
- [System Management | 115](#)

Learn which issues were resolved in Junos OS Evolved Release 20.3R1 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.

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-pfemamd 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 have been 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
- Hardware Compatibility Tool—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home


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

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

Documentation Feedback

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

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



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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

Revision History

29 July 2022—Revision 17, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

12 May 2022—Revision 16, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

5 May 2022—Revision 15, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

28 January 2022—Revision 14, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

7 October 2021—Revision 13, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

16 July 2021—Revision 12, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

22 April 2021—Revision 11, Junos OS Evolved Release 20.3R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

13 January 2021—Revision 10, Junos OS Evolved Release 20.3R1 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 © 2020 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.