

Release Notes

Published
2021-04-22

Junos OS Evolved Release 21.1R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices

SOFTWARE HIGHLIGHTS

- Support for BGP neighbor telemetry with sharding (PTX10003 and PTX10008)
- Support to change the default re-merge behavior on the P2MP LSP (PTX10008)
- Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP-labeled unicast (PTX10008)
- Support for flexible algorithms in IS-IS for segment segment routing-traffic engineering (SR-TE) (PTX10001-36MR, PTX10004, and PTX10008)
- Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC 6286 (PTX10001-36MR, PTX10003, PTX10008, and QFX5220)

Day One+

- Use this [new setup tool](#) to get your Junos OS Evolved up and running in three quick steps.

Table of Contents

Introduction | 1

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, and PTX10008 Devices

What's New | 2

What's New in 21.1R1 | 3

- Hardware | 3
- Juniper Extension Toolkit (JET) | 29
- Junos OS API and Scripting | 29
- Junos Telemetry Interface | 30
- Licensing | 31
- MPLS | 33
- Network Management and Monitoring | 33
- Routing Protocols | 34
- Segment Routing | 35

What's Changed | 36

- What's Changed in Release 21.1R1 | 37

Known Limitations | 39

Open Issues | 42

Resolved Issues | 49

- Resolved Issues: 21.1R1 | 49

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

What's New | 59

What's New in 21.1R1 | 59

- Hardware | 59
- MPLS | 62
- Multicast | 62
- Services Applications | 63

What's Changed | 63

| What's Changed in Release 21.1R1 | 63

Known Limitations | 66

Open Issues | 67

Resolved Issues | 69

| Resolved Issues: 21.1R1 | 69

Licensing | 71

Finding More Information | 72

Documentation Feedback | 73

Requesting Technical Support | 73

Upgrade Your Junos OS Evolved Software | 75

Revision History | 76

Introduction

With Junos OS Evolved, you can enable higher availability, accelerate your deployments, innovate rapidly, and operate your network more efficiently. We've aligned Junos OS Evolved with Junos OS so that you can seamlessly continue to manage and automate your network. There are, however, differences between Junos OS and Junos OS Evolved that you should know about before moving over to this powerful new OS. For details on these changes, see [How Junos OS Evolved Differs from Junos OS](#).

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 21.1R1.

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 | 2](#)
- [What's Changed | 36](#)
- [Known Limitations | 39](#)
- [Open Issues | 42](#)
- [Resolved Issues | 49](#)

These release notes accompany Junos OS Evolved Release 21.1R1 for PTX10001-36MR, PTX10003, PTX10004, and PTX10008 Packet Transport Routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

See [Table 1 on page 2](#) for a complete list of the Junos OS Evolved releases that we support on the PTX Series routing platforms.

Table 1: Junos OS Evolved Releases Supported on the PTX Series Routing Platforms

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

What's New

IN THIS SECTION

- [What's New in 21.1R1 | 3](#)

Learn about new features introduced in the Junos OS Evolved main and maintenance releases for the PTX10001-36MR, PTX10003, PTX10004, and PTX10008.

What's New in 21.1R1

IN THIS SECTION

- [Hardware | 3](#)
- [Juniper Extension Toolkit \(JET\) | 29](#)
- [Junos OS API and Scripting | 29](#)
- [Junos Telemetry Interface | 30](#)
- [Licensing | 31](#)
- [MPLS | 33](#)
- [Network Management and Monitoring | 33](#)
- [Routing Protocols | 34](#)
- [Segment Routing | 35](#)

Hardware

- We haven't introduced new hardware in Junos OS Evolved 21.1R1. Use the following tables to learn about some of the new features supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 routers in Junos OS Evolved Release 21.1R1.

We've added the following features to the PTX10001-36MR routers in Junos OS Evolved Release 21.1R1.

Table 2: Feature Support Added to the PTX10001-36MR Router

Feature	Description
EVPN	<ul style="list-style-type: none">• Support for EVPN-MPLS. The router supports MAC VRF routing instances on the EVPN-MPLS network. [See EVPN User Guide.]

Table 2: Feature Support Added to the PTX10001-36MR Router (*Continued*)

Feature	Description
Interfaces and chassis	<ul style="list-style-type: none"> Support for custom optics profile. You can now configure the optical module parameters of QSFP modules by using the CLI from the configuration prompt. In earlier releases, you could not configure the parameters from the configuration prompt. We do not support these module parameters for QSFP direct attach copper (DAC) cable modules or QSFP load modules. [See optics-options.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> Support for optics diagnostics, openconfig-platform-transceiver.yang, version 0.5.0, output power, input power, laser bias for JTI. We've enhanced support for physical and logical interface transceiver statistics. Use the new sensors under the resource paths /components/component/transceiver/ and /components/component/transceiver/physical-channels/ to export statistics to an outside collector using JTI and RPC or gNMI services. <p>To export statistics, include the following sensors in a subscription:</p> <ul style="list-style-type: none"> /components/component/transceiver/oc-instant:transceiver:output-power /components/component/transceiver/oc-instant:transceiver:input-power /components/component/transceiver/oc-instant:transceiver:laser-bias-current /components/component/transceiver/physical-channels/oc-instant:transceiver:output-power /components/component/transceiver/physical-channels/oc-instant:transceiver:input-power /components/component/transceiver/physical-channels/oc-instant:transceiver:laser-bias-current <p>[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]</p>

Table 2: Feature Support Added to the PTX10001-36MR Router (*Continued*)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> Support for the IRB interface, which enables routing of Layer 3 traffic between a bridge domain and another routed interface. [See Integrated Routing and Bridging.]
MPLS	<ul style="list-style-type: none"> Support for point-to-multipoint (P2MP) ping and P2MP label-switched path (LSP) traceroute. [See MPLS Applications User Guide.]
Network management and monitoring	<ul style="list-style-type: none"> Support for sFlow for IP over IP (IP-IP) traffic. You can use the sFlow technology to sample egress sFlow for IP-IP traffic at the tunnel originator, transit, and tunnel terminator nodes of a physical port. We support this feature for IP-IP tunnels (with an IPv4 outer header) that carry IPv4 or IPv6 traffic. You can randomly sample network packets from IP-IP tunnels and send the samples to a destination collector for monitoring. Devices that act as an IP-IP tunnel entry point, transit device, or tunnel endpoint support sFlow sampling. [See Overview of sFlow Technology and Configuring IP Tunnel Interfaces.] Support for SNMP for PFC and ECN. You can monitor the PFC and ECN status from the remote network using SNMP. [See Enterprise-Specific SNMP MIBs Supported by Junos OS.]

Table 2: Feature Support Added to the PTX10001-36MR Router (*Continued*)

Feature	Description
Routing protocols	<ul style="list-style-type: none"> • Support to label interfaces and match the interface-group ID on firewall filter. We support labeling of interfaces with specified group IDs from 1 to 255 and matching the interface-group ID on the firewall filter. The filter can identify which interface a packet comes from and perform actions specified only for a certain interface group. [See Understanding BGP Flow Routes for Traffic Filtering.] • Support for FIB compression. Forwarding information base (FIB) compression overcomes the hardware limitation of storing a limited number of uncompressed routes. We've enabled FIB compression by default. You can now store a larger number of compressed IPv4 and IPv6 routes. • Support for higher scale of flows in the BGP FlowSpec filter. We now support a higher scale of flows in BGPFlowSpec filter. You can achieve this by reducing the number of terms in the BGP FlowSpec filter using the filter optimization techniques without affecting the functionality. The per-term counter is not configured by default. To create room for term compression, the default per-term counters are disabled. However, to enable per-term counter a new CLI command, per-route-accounting has been introduced under the existing [edit routing-options flow] hierarchy level." [See "per-route-accounting".]

Table 2: Feature Support Added to the PTX10001-36MR Router (*Continued*)

Feature	Description
Routing policy and firewall filters	<ul style="list-style-type: none"> • Support for nested filters. This feature enables 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"> • <code>inet</code>—Input and output directions • <code>inet6</code>—Input and output directions • <code>mpls</code>—Input direction only <p>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.]</p> • Support for firewall filters. We've added support for Layer 2 firewall filter match conditions and actions. [See Understanding Firewall Filter Match Conditions.]
Services applications	<ul style="list-style-type: none"> • Support for configuring GRE tunnel encapsulation on flexible tunnel interfaces (FTIs). You can configure encapsulation by using the tunnel encapsulation udp source <i>address</i> destination <i>address</i> command at the [edit interfaces fti unit <i>unit</i>] hierarchy level. <ul style="list-style-type: none"> • If you add the tunnel-termination statement, it makes the tunnel a de-encapsulation-only tunnel and disables the encapsulation. • You must specify both the source and destination addresses if you do not configure tunnel-termination. • You cannot configure a variable prefix mask on the source address. <p>[See encapsulation.]</p>

Table 2: Feature Support Added to the PTX10001-36MR Router (*Continued*)

Feature	Description
	<ul style="list-style-type: none"> • Support for RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. TWAMP sends out probes to measure network performance and is often used to check compliance with service-level agreements. You can configure TWAMP at the [edit services monitoring twamp] hierarchy level. The support for this service is limited to the following: <ul style="list-style-type: none"> • 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 <p>[See Understanding Two-Way Active Measurement Protocol on Routers.]</p> • Support for multiple collectors in inline active flow monitoring. You can now configure inline active flow monitoring to export flow records to up to four different collectors. Previously, inline flow monitoring could export flow records only to a single collector. [See Configuring Inline Active Flow Monitoring on PTX Series Routers.]

Table 2: Feature Support Added to the PTX10001-36MR Router (*Continued*)

Feature	Description
	<ul style="list-style-type: none"> • Support for 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.] • Support for configuring UDP tunnel encapsulation on flexible tunnel interfaces (FTIs). You can configure encapsulation by using the tunnel encapsulation udp source <i>address</i> destination <i>address</i> command at the [edit interfaces fti unit <i>unit</i>] hierarchy level. <ul style="list-style-type: none"> • If you add the tunnel-termination statement, it makes the tunnel a de-encapsulation-only tunnel and disables the encapsulation. • You must specify both the source and destination addresses if you do not configure tunnel-termination. • You cannot configure a variable prefix mask on the source address. [See encapsulation.]
Software installation and upgrade	<ul style="list-style-type: none"> • Support for listing incompatible field-replaceable units (FRUs) in the output of the request system software add restart and request system software validate-restart commands, and giving users the opportunity to take the unsupported FRUs offline and continue with the restart or validation. [See request system software add and request system software validate-restart.]

We've added the following features to the PTX10003 routers in Junos OS Evolved Release 21.1R1.

Table 3: Feature Support Added to the PTX10003 Router

	Description
DHCP	<ul style="list-style-type: none"> Support for stateless DHCP relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See DHCP Relay Agent.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> Support for platform, interface, and alarm statistics using JTI and gRPC Network Management Interface (gNMI) services. You can use this feature to send ON_CHANGE statistics of a device to an outside collector. <p>This feature supports these OpenConfig models:</p> <ul style="list-style-type: none"> openconfig-platform.yang: oc-ext:openconfig-version 0.12.1 openconfig-interfaces.yang: oc-ext:openconfig-version 2.4.1 openconfig-alarms.yang: oc-ext:openconfig-version 0.3.1 <p>Use the following resource paths in a gNMI subscription:</p> <ul style="list-style-type: none"> /components/component (for each installed FRU) /interfaces/interface/state/ /interfaces/interface/subinterfaces/subinterface/state/ /alarms/alarm/ <p>[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]</p>
Layer 2 VPN	<ul style="list-style-type: none"> Support for virtual private wire service (VPWS) with single segment and multisegment pseudowires using FEC 129. [See "Example: Configuring FEC 129 BGP Autodiscovery for VPWS".]

Table 3: Feature Support Added to the PTX10003 Router (Continued)

	Description
Services applications	<ul style="list-style-type: none"> • Support for RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. TWAMP sends out probes to measure network performance and is often used to check compliance with service-level agreements. You can configure TWAMP at the [edit services monitoring twamp] hierarchy level. The support for this service is limited to the following: <ul style="list-style-type: none"> • 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 <p>[See Understanding Two-Way Active Measurement Protocol on Routers.]</p>

We've added the following features to the PTX10004 routers in Junos OS Evolved Release 21.1R1.

Table 4: Feature Support Added to the PTX10004 Router

	Description
EVPN	<ul style="list-style-type: none"> • Support for EVPN-MPLS. The router supports MAC VRF routing instances on the EVPN-MPLS network. [See EVPN User Guide.]

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
Interfaces and chassis	<ul style="list-style-type: none"> • Support for custom optics profile. You can now configure the optical module parameters of QSFP modules by using the CLI from the configuration prompt. In earlier releases, you could not configure the parameters from the configuration prompt. We do not support these module parameters for QSFP direct attach copper (DAC) cable modules or QSFP load modules. [See optics-options.] • Support for ALB on multiple Packet Forwarding Engines for aggregated Ethernet bundles. Adaptive load balancing (ALB) for aggregated Ethernet bundles evenly redistributes the traffic load across multiple ingress Packet Forwarding Engines on the same line card. [See Load Balancing on Aggregated Ethernet Interfaces.]
IP tunneling	<ul style="list-style-type: none"> • Support for IPv4 unicast tunneling. To support IP-over-IP on the PTX models, the routing protocol process (rpd) sends fully resolved next-hop-based tunnels to the Packet Forwarding Engine. [See Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • Support for optics diagnostics, openconfig-platform-transceiver.yang, version 0.5.0, output power, input power, laser bias for JTI. We've enhanced support for physical and logical interface transceiver statistics. Use the new sensors under the resource paths /components/component/transceiver/ and /components/component/transceiver/physical-channels/ to export statistics to an outside collector using JTI and RPC or gNMI services.

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
	<ul style="list-style-type: none"> • To export statistics, include the following sensors in a subscription: <ul style="list-style-type: none"> • /components/component/transceiver/oc-instant:transceiver:output-power • /components/component/transceiver/oc-instant:transceiver:input-power • /components/component/transceiver/oc-instant:transceiver:laser-bias-current • /components/component/transceiver/physical-channels/oc-instant:transceiver:output-power • /components/component/transceiver/physical-channels/oc-instant:transceiver:input-power • /components/component/transceiver/physical-channels/oc-instant:transceiver:laser-bias-current <p>[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]</p>

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for the following basic Layer 2 learning, bridging, and flooding features: <ul style="list-style-type: none"> • Enterprise-style bridging (trunk and access modes) • Service-provider-style bridging (also known as sub-interface mode) • BUM (broadcast, unknown unicast, and multicast) traffic handling, including split horizon • MAC learning and aging • Static MAC addresses • Trunk port, VLAN membership, and native-vlan-id • 802.1Q EtherType - 8100 • 802.1Q VLAN tagging: Single tagging with normalization • Clearing of all MAC address information • Global MAC limit • Global source MAC aging time • MAC moves • Disabling of MAC learning at global and interface levels. • Native VLAN ID for Layer 2 logical interfaces • Single VLAN-tagged Layer 2 logical interfaces • Interface statistics <p>NOTE: The show ethernet-switching statistics command and child logical interface statistics for aggregated Ethernet are not supported.</p>

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
	<ul style="list-style-type: none"> Flexible Ethernet services NOTE: Enterprise-style Layer 2 logical interfaces are not supported under flexible Ethernet services encapsulation. Virtual switch Persistent MAC learning (sticky MAC) Service provider bridging: <ul style="list-style-type: none"> Multiple logical interfaces (on a physical interface) which are part of same bridge domain Ethernet bridge encapsulation <p>[See Layer 2 Bridging, Address Learning, and Forwarding User Guide.]</p>
MPLS	<ul style="list-style-type: none"> Support for next-hop-based dynamic UDP tunnels, also known as MPLS-over-UDP tunnels. When you configure an MPLS-over-UDP tunnel, Junos OS dynamically creates a tunnel composite next hop, an indirect next hop, and a forwarding next hop to resolve the tunnel destination route. You can also use policy control to resolve MPLS-over-UDP tunnels over select IP prefixes. When you enable next hops by default, the MPLS-over-UDP feature provides a scaling advantage for the number of IP tunnels supported on the router. [See Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels.]

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
Network management and monitoring	<ul style="list-style-type: none"> Support for remote port mirroring with ToS or DSCP settings. This feature enables you to send sampled copies of incoming packets to remote network management software. You can send the packets through flexible tunnel interfaces (FTIs) using GRE encapsulation. You can set type of service (ToS) and DiffServ code point (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 in the [edit forwarding-options port-mirroring instance <instance-name> output] hierarchy. [See instance (Port Mirroring).]
Routing protocols	<ul style="list-style-type: none"> Support to label interfaces and match the interface-group ID on firewall filter. We support labeling of interfaces with specified group IDs from 1 to 255 and matching the interface-group ID on the firewall filter. The filter can identify which interface a packet comes from and perform actions specified only for a certain interface group. [See Understanding BGP Flow Routes for Traffic Filtering.] Support for higher scale of flows in the BGP FlowSpec filter. We now support a higher scale of flows in BGPFlowSpec filter. You can achieve this by reducing the number of terms in the BGP FlowSpec filter using the filter optimization techniques without affecting the functionality. The per-term counter is not configured by default. To create room for term compression, the default per-term counters are disabled. However, to enable per-term counter a new CLI command, per-route-accounting has been introduced under the existing [edit routing-protocols-flow] hierarchy level." [See "per-route-accounting".]

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
Routing policy and firewall filters	<ul style="list-style-type: none"> Support for nested filters. This feature enables 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—Input and output directions inet6—Input and output directions mpls—Input direction only <p>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.]</p> Support for interface and class-usage accounting profiles. Use accounting profiles to collect persistent (locally saved) traffic statistics as an alternative to SNMP. Data collected through interface profiles can include input and output byte and packet counts for various protocols, interface errors, and RPF checks. Data collected through source and destination class-usage profiles can include per-interface traffic counts for the prefixes specified, for example to track traffic received on a per-customer basis for billing. Configure accounting profiles at the [edit accounting-options] hierarchy level. [See Understanding Accounting Profiles.]

Table 4: Feature Support Added to the PTX10004 Router *(Continued)*

	Description
Services applications	<ul style="list-style-type: none"> • Support for RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. TWAMP sends out probes to measure network performance and is often used to check compliance with service-level agreements. You can configure TWAMP at the [edit services monitoring twamp] hierarchy level. The support for this service is limited to the following: <ul style="list-style-type: none"> • 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 <p>[See Understanding Two-Way Active Measurement Protocol on Routers.]</p> <ul style="list-style-type: none"> • Support for multiple collectors in inline active flow monitoring. You can now configure inline active flow monitoring to export flow records to up to four different collectors. Previously, inline flow monitoring could export flow records only to a single collector. [See Configuring Inline Active Flow Monitoring on PTX Series Routers.]

Table 4: Feature Support Added to the PTX10004 Router (*Continued*)

	Description
	<ul style="list-style-type: none"> • Support for 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.] • Support for configuring UDP tunnel encapsulation on flexible tunnel interfaces (FTIs). You can configure encapsulation by using the tunnel encapsulation udp source <i>address</i> destination <i>address</i> command at the [edit interfaces fti unit <i>unit</i>] hierarchy level. <ul style="list-style-type: none"> • If you add the tunnel-termination statement, it makes the tunnel a de-encapsulation-only tunnel and disables the encapsulation. • You must specify both the source and destination addresses if you do not configure tunnel-termination. • You cannot configure a variable prefix mask on the source address. [See encapsulation.] • Support for optics drivers upgrade. We've extended support to upgrade the optics drivers without a full Junos OS Evolved upgrade. You can upgrade the optics drivers by executing the request system software add variable>package_name</variable> command. [See Hardware Supported by Junos Continuity Software.] • Support for FTI-based encapsulation and de-encapsulation of IPv4 and IPv6 packets. You can configure IP-IP encapsulation and de-encapsulation on an FTI. [See ipip.]

We've added the following features to the PTX10008 routers in Junos OS Evolved Release 21.1R1.

Table 5: Feature Support Added to the PTX10008 Router

	Description
Class of service	<ul style="list-style-type: none"> • Support for priority-based flow control (PFC) of untagged traffic at Layer 3 and support for explicit congestion notification (ECN). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic and Understanding CoS Explicit Congestion Notification.] • Support for PFC watchdog. Use this feature to detect and mitigate PFC pause storms received for PFC-enabled queues. When a PFC-enabled port receives PFC pause frames for an extended period of time, the PFC watchdog mitigates the situation by disabling the affected queue for a configurable length of time called the recovery time. After the recovery time passes, the PFC watchdog re-enables PFC. [See PFC Watchdog.]
EVPN	<ul style="list-style-type: none"> • Support for EVPN-MPLS. The router supports MAC VRF routing instances on the EVPN-MPLS network. [See EVPN User Guide.]
Hardware	<ul style="list-style-type: none"> • Use the hardware compatibility matrix to know about the optical interfaces and transceivers, active optical cables (AOCs), and direct attach copper (DAC) cables that we support on the PTX10K-LC1202-36MR line card. [See the Hardware Compatibility Tool.]

Table 5: Feature Support Added to the PTX10008 Router *(Continued)*

	Description
Interfaces and chassis	<ul style="list-style-type: none"> • We support RSVP-based and LDP-based point-to-multipoint (P2MP) label-switched paths (LSPs). [See Port Speed on PTX10K-LC1201 Overview.] • Support for custom optics profile. You can now configure the optical module parameters of QSFP modules by using the CLI from the configuration prompt. In earlier releases, you could not configure the parameters from the configuration prompt. We do not support these module parameters for QSFP direct attach copper (DAC) cable modules or QSFP load modules. [See optics-options.] • Support for 10-Gbps speed on SFP+-10G-T-DWDM-ZR and SFP +-10G-CWDM optics. We support 10-Gbps speed on these optics modules through the QSA adapter (MAM1Q00A-QSA) plugged into the QSFP or QSFP+ ports. You can select the wavelength on the SFP +-10G-T-DWDM-ZR optics. To see the list of optics supported by the PTX10K-LC1202-36MR line card, visit the Hardware Compatibility Tool page. <p>NOTE: The ports 1, 3, 19, and 21 are disabled if the preceding ports (0, 2, 18, and 20) are not in 100-Gbps mode. This means, of the 36 ports on the line card, only 32 ports are available to be configured as 4x25-Gbps and 4x10-Gbps ports.</p> <p>[See PTX10K-LC1202-36MR Line Card.]</p> <ul style="list-style-type: none"> • Support for performance monitoring and TCA. We support transport performance monitoring and threshold-crossing alert (TCA) information for the QSFP28 100GE DWDM optical transceiver modules. You can view the diagnostic data, warnings, and alarms for transport performance monitoring interfaces. TCAs give the management system an early indication about the state of the associated entity when it crosses a certain threshold. [See show interfaces transport pm and tca.] • Support for ALB on multiple Packet Forwarding Engines for aggregated Ethernet bundles. Adaptive load balancing (ALB) for

Table 5: Feature Support Added to the PTX10008 Router *(Continued)*

	Description
	aggregated Ethernet bundles evenly redistributes the traffic load across multiple ingress Packet Forwarding Engines on the same line card. [See Load Balancing on Aggregated Ethernet Interfaces .]
IP tunneling	<ul style="list-style-type: none"> Support for IPv4 unicast tunneling. To support IP-over-IP on the PTX models, the routing protocol process (rpd) sends fully resolved next-hop-based tunnels to the Packet Forwarding Engine. [See Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation.]
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> Support for Jvision sensors cmerror configuration and cmerror counter for JTI. Junos telemetry interface (JTI) supports streaming chassis management error (cmerror) configuration and counters as well as Flexible PIC Concentrator (FPC) optics statistics to an outside collector using remote procedure calls (gRPC). <p>We support the following cmerror base resource paths:</p> <ul style="list-style-type: none"> /junos/chassis/cmerror/configuration /junos/chassis/cmerror/counters <p>[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]</p>

Table 5: Feature Support Added to the PTX10008 Router (*Continued*)

	Description
	<ul style="list-style-type: none"> Support for platform, interface, and alarm statistics using JTI and gRPC Network Management Interface (gNMI) services. You can use this feature to send ON_CHANGE statistics of a device to an outside collector. <p>This feature supports these OpenConfig models:</p> <ul style="list-style-type: none"> openconfig-platform.yang: oc-ext:openconfig-version 0.12.1 openconfig-interfaces.yang: oc-ext:openconfig-version 2.4.1 openconfig-alarms.yang: oc-ext:openconfig-version 0.3.1 <p>Use the following resource paths in a gNMI subscription:</p> <ul style="list-style-type: none"> /components/component (for each installed FRU) /interfaces/interface/state/ /interfaces/interface/subinterfaces/subinterface/state/ /alarms/alarm/ <p>[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]</p>

Table 5: Feature Support Added to the PTX10008 Router (*Continued*)

	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for the following basic Layer 2 learning, bridging, and flooding features: <ul style="list-style-type: none"> • Enterprise-style bridging (trunk and access modes) • Service-provider-style bridging (also known as sub-interface mode) • BUM (broadcast, unknown unicast, and multicast) traffic handling, including split horizon • MAC learning and aging • Static MAC addresses • Trunk port, VLAN membership, and native-vlan-id • 802.1Q EtherType - 8100 • 802.1Q VLAN tagging: Single tagging with normalization • Clearing of all MAC address information • Global MAC limit • Global source MAC aging time • MAC moves • Disabling of MAC learning at global and interface levels. • Native VLAN ID for Layer 2 logical interfaces • Single VLAN-tagged Layer 2 logical interfaces • Interface statistics <p>NOTE: The show ethernet-switching statistics command and child logical interface statistics for aggregated Ethernet are not supported.</p> <ul style="list-style-type: none"> • Flexible Ethernet services

Table 5: Feature Support Added to the PTX10008 Router (*Continued*)

	Description
	<p>NOTE: Enterprise-style Layer 2 logical interfaces are not supported under flexible Ethernet services encapsulation.</p> <ul style="list-style-type: none"> • Virtual switch • Persistent MAC learning (sticky MAC) • Service provider bridging: <ul style="list-style-type: none"> • Multiple logical interfaces (on a physical interface) which are part of same bridge domain • Ethernet bridge encapsulation <p>[See Layer 2 Bridging, Address Learning, and Forwarding User Guide.]</p>
MPLS	<ul style="list-style-type: none"> • Support for LDP tunneling over segment routing-traffic engineering (SR-TE). You can tunnel LDP LSPs over SR-TE in your network. Tunneling LDP over SR-TE provides consistency and coexistence of LDP LSPs and SR-TE LSPs. [See Tunneling LDP LSPs in RSVP LSPs Overview.] • Support for next-hop-based dynamic UDP tunnels, also known as MPLS-over-UDP tunnels. When you configure an MPLS-over-UDP tunnel, Junos OS dynamically creates a tunnel composite next hop, an indirect next hop, and a forwarding next hop to resolve the tunnel destination route. You can also use policy control to resolve MPLS-over-UDP tunnels over select IP prefixes. When you enable next hops by default, the MPLS-over-UDP feature provides a scaling advantage for the number of IP tunnels supported on the router. [See Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels.]
Multicast	<ul style="list-style-type: none"> • Support for PIM multicast features on the Layer 2 IRB interface. We support any-source multicast (ASM), PIM source-specific multicast (PIM SSM), and PIM dense mode (PIM DM). [See Multicast Routing Protocols.]

Table 5: Feature Support Added to the PTX10008 Router *(Continued)*

	Description
Network management and monitoring	<ul style="list-style-type: none"> Support for RFC 3635. Junos OS Evolved supports all objects and tables defined in RFC 3635, except dot3StatsRateControlAbility and dot3StatsRateControlStatus in the dot3StatsEntry table. [See Standard SNMP MIBs Supported by Junos OS.] On PTX10K-LC1202-36MR line cards, you can configure 50-Gbps speed on ports 0, 2, 18, and 20 through QSFP28 optics. When you configure the port speed of any of the even ports (0, 2, 18, 20) as 2 x 50 Gbps, you must configure the odd ports (1, 3, 19, 21) as unused. You cannot configure 4x10GbE, 2x50GbE, 1x40GbE, and 4x25GbE simultaneously on the following ports: <ul style="list-style-type: none"> 0 and 2 18 and 20 [See PTX10K-LC1202-36MR Port Speed.]
Routing protocols	<ul style="list-style-type: none"> Support to label interfaces and match the interface-group ID on firewall filter. We support labeling of interfaces with specified group IDs from 1 to 255 and matching the interface-group ID on the firewall filter. The filter can identify which interface a packet comes from and perform actions specified only for a certain interface group. [See Understanding BGP Flow Routes for Traffic Filtering.] Support for FIB compression. Forwarding information base (FIB) compression overcomes the hardware limitation of storing a limited number of uncompressed routes. We've enabled FIB compression by default. You can now store a larger number of compressed IPv4 and IPv6 routes.

Table 5: Feature Support Added to the PTX10008 Router *(Continued)*

	Description
	<ul style="list-style-type: none"> Support for higher scale of flows in the BGP FlowSpec filter. We now support a higher scale of flows in BGPFlowSpec filter. You can achieve this by reducing the number of terms in the BGP FlowSpec filter using the filter optimization techniques without affecting the functionality. The per-term counter is not configured by default. To create room for term compression, the default per-term counters are disabled. However, to enable per-term counter a new CLI command, per-route-accounting has been introduced under the existing [edit routing-protocols-flow] hierarchy level." [See "per-route-accounting".]
Segment routing	<ul style="list-style-type: none"> Support for color-based mapping of VPN services over SR-TE LSPs. You can use both color-attribute and IP protocol next hops to resolve transport tunnels over static colored and BGP segment routing-traffic engineering (SR-TE) label-switched paths (LSPs). [See Color-Based Mapping of VPN Services Overview.]

Table 5: Feature Support Added to the PTX10008 Router (*Continued*)

	Description
Services applications	<ul style="list-style-type: none"> • Support for RFC 5357 Two-Way Active Measurement Protocol (TWAMP) monitoring service. TWAMP sends out probes to measure network performance and is often used to check compliance with service-level agreements. You can configure TWAMP at the [edit services monitoring twamp] hierarchy level. The support for this service is limited to the following: <ul style="list-style-type: none"> • 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 <p>[See Understanding Two-Way Active Measurement Protocol on Routers.]</p> • Support for configuring UDP tunnel encapsulation on flexible tunnel interfaces (FTIs). You can configure encapsulation by using the tunnel encapsulation udp source <i>address</i> destination <i>address</i> command at the [edit interfaces fti unit <i>unit</i>] hierarchy level. <ul style="list-style-type: none"> • If you add the tunnel-termination statement, it makes the tunnel a de-encapsulation-only tunnel and disables the encapsulation. • You must specify both the source and destination addresses if you do not configure tunnel-termination. • You cannot configure a variable prefix mask on the source address. <p>[See encapsulation.]</p>

Table 5: Feature Support Added to the PTX10008 Router (*Continued*)

	Description
	<ul style="list-style-type: none"> Support for FTI-based encapsulation and de-encapsulation of IPv4 and IPv6 packets. You can configure IP-IP encapsulation and de-encapsulation on an FTI. [See ipip.]
Software installation and upgrade	<ul style="list-style-type: none"> Support for dynamically detecting the port speed of WAN interfaces during zero-touch provisioning (ZTP). [See Zero Touch Provisioning.] Support for either WAN interfaces or management interfaces to automatically download and install the appropriate software and the configuration file on your device during the ZTP bootstrap process. [See Zero Touch Provisioning.]

Juniper Extension Toolkit (JET)

- JET support for 64-bit infrastructure (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, Juniper Extension Toolkit (JET) applications use 64-bit infrastructure by default. Using 64-bit infrastructure instead of 32-bit infrastructure increases efficiency because 64-bit applications can handle larger file sizes and a greater number of files. Using 64-bit infrastructure also increases system security because 64-bit applications support additional security best practices.

[See [Develop On-Device JET Applications](#).]

- Python 3 support for JET applications (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, you can use Python 3 to write Juniper Extension Toolkit (JET) applications. To run a JET script that uses Python 3, you must configure the **language python3** statement at the **[edit system scripts]** hierarchy level.

[See [Develop On-Device JET Applications](#) and [Understanding Python Automation Scripts for Devices Running Junos OS](#).]

Junos OS API and Scripting

- Python 3 support for commit, event, op, and SNMP scripts (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, you can use Python 3 to execute commit, event, op, and SNMP scripts. To use Python 3, you must configure the **language python3** statement at the **[edit system scripts]** hierarchy level.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

- **Automation script support for 64-bit infrastructure (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, commit, event, op, and SNMP scripts use 64-bit infrastructure by default. Using 64-bit infrastructure instead of 32-bit infrastructure increases efficiency because 64-bit applications can handle larger file sizes and a greater number of files. Using 64-bit infrastructure also increases system security because 64-bit applications support additional security best practices.

[See [Automation Scripting User Guide.](#)]

Junos Telemetry Interface

- **BGP neighbor telemetry with sharding (PTX10003 and PTX10008)**—Starting in Junos OS Release 21.1R1, we support BGP neighbor telemetry with sharding (multithreading).

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

- **SPRING statistics sensor support for JTI (PTX10003)**—Starting in Junos OS Evolved Release 21.1R1, JTI enables support for the export of Source Packet Routing in Networking (SPRING) statistics to an outside collector using remote procedure call (gRPC) services and gRPC Network Management Interface (gNMI) services.

This feature provides SPRING statistics for ingress and egress per-interface and per-child member at the physical interface level.

To enable SPRING statistics, include the following statements on the client device:

- For egress (per-interface and per-child member at the physical interface level), use the **set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress** statement at the **[edit]** hierarchy level.
- For ingress (per-interface and per-child member at the physical interface level ingress), use the **set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress** statement at the **[edit]** hierarchy level.

Use the following sensors to export statistics using gRPC or gNMI services to an outside collector:

- **/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/in-octets/** for ingress (per-interface and per-child member at the physical interface level ingress) SPRING traffic.
- **/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/in-pkts/** for ingress (per-interface and per-child member at the physical interface level ingress) SPRING traffic.

- `/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/out-octets/` for egress (per-interface and per-child member at the physical interface level egress) SPRING traffic.
- `/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/out-pkts/` for egress (per interface and per-child member at the physical interface level egress) SPRING traffic.

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

Licensing

- **Juniper Agile Licensing (PTX10001-36MR, PTX10004, and PTX10008)—**

Starting in Junos OS Evolved Release 21.1R1, the listed PTX Series routers support Juniper Agile Licensing.

Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can use Juniper Agile Licensing to install and manage licenses for software features.

Juniper Agile Licensing supports soft enforcement of software feature licenses. With soft enforcement, if you configure a feature without a license, Junos OS displays a warning when you commit the configuration. However, the feature is operational. In addition, Junos OS generated periodic alarms indicating that you need the license to use the feature. You can see the list of alarms at [System Log Explorer](#).

If you are upgrading to Junos OS Evolved Release 21.1R1, 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 21.1R1.

"Table 6" on page 31 describes the licensing support for the PTX10001-36MR, PTX10004, and PTX10008 devices.

Table 6: Licensed Features on PTX10001-36MR, PTX10004, and PTX10008 Devices

License Model	Supported Features	Scale
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), SNMP, and VLAN	64,000 forwarding information bases (FIBs) and 64,000 routing information bases (RIBs), 1000 firewall filters, and 4000 firewall terms

Table 6: Licensed Features on PTX10001-36MR, PTX10004, and PTX10008 Devices *(Continued)*

License Model	Supported Features	Scale
Advanced	A1	
	BGP, FBF, GRE, IS-IS, Junos telemetry interface (JTI), OSPF, VRF-lite, and VRRP	256,000 FIBs, 3 million RIBs, 1000 BGP peers, 1000 IP tunnels, 256 VRF instances (VRFs)/virtual routers, 16,000 PIMs, 4000 firewall filters, and 16,000 firewall terms
	A2	
	A1 features, connectivity fault management (CFM), EVPN-VXLAN, VRF (VXLAN), VXLAN, Q-in-Q, ESI-LAG, Precision Time Protocol (PTP), PTP-BC, and Synchronous Ethernet	1 million FIBs, 3 million RIBs, 1000 LSPs, 1000 BGP peers, 3000 IP tunnels, 1000 VRFs/virtual routers, 32,000 PIMs, 16,000 firewall filters, and 64,000 firewall terms
Premium	P1	
	A2 features, EVPN-MPLS, Layer 2 circuit, Layer 3 VPN, LDP, MPLS, RSVP, segment routing, Layer 3 Multicast, PCEP, PIM, and SR-TE	2 million FIBs, 6 million RIBs, 32,000 LSPs, 2000 BGP peers, 5000 IP tunnels, 1000 VRFs/virtual routers, 32,000 PIMs, 16,000 firewall filters, and 64,000 firewall terms
	P2	
	P1 features and H-QoS	FIBs and IP tunnels are up to the device capacity, 60 million RIBs, 200,000 LSPs, 4000 BGP peers, 1000+ VRFs/virtual routers, 100,000 PIMs, 64,000 firewall filters and terms

In addition, you can install additional port bandwidth usage license to increase the port bandwidth usage.

You cannot use network mode to install licenses on PTX10001-36MR, PTX10004, and PTX10008 devices because Juniper Agile License Manager is not supported on these devices.

[See [Flex Software License for PTX Series Routers](#), [Juniper Agile Licensing Guide](#), and [Configuring Licenses in Junos OS](#).]

MPLS

- **Support to change the default re-merge behavior on the P2MP LSP (PTX10008)**—You can change the default re-merge behavior on a RSVP point-to-multipoint (P2MP) label-switched path (LSP). The term re-merge refers to the case of an ingress or transit node (re-merge node) that creates a re-merge branch intersecting the P2MP LSP at another node in the network. This may occur due to events such as an error in path calculation, an error in manual configuration, or network topology changes during the establishment of the P2MP LSP.

[See [Re-merge Behavior on Point-to-Multipoint LSP Overview](#).]

- **Install prefixes for RSVP-TE LSPs using PCEP (PTX Series)**—Starting in Junos OS Evolved Release 21.1R1, you can configure different prefixes for Path Computation Element (PCE)-initiated and PCE-delegated RSVP-TE LSPs using the Path Computation Element Protocol (PCEP). Prior to this feature, for PCE-initiated LSPs, you could install prefixes as routes through templates and map the templates to the LSPs. For Path Computation Client (PCC)-configured LSPs, although you could install prefixes on the device, this information was not reported to the PCE.

With this feature, you can install prefixes for external RSVP-TE LSPs through PCEP communication, and enable the PCC to report installed prefixes for all local RSVP-TE LSPs to the PCE. This support provides you better traffic engineering capabilities and allows Junos OS to interoperate with other vendor's PCC or PCE.

[See [Support of the Path Computation Element Protocol for RSVP-TE Overview](#).]

Network Management and Monitoring

- **Python 3 support for YANG action and translation scripts (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, you can use Python 3 to execute YANG action and translation scripts. To use Python 3, you must configure the `language python3` statement at the `[edit system scripts]` hierarchy level.

[See [Understanding Python Automation Scripts for Devices Running Junos OS](#).]

- **YANG action and translation script support for 64-bit infrastructure (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, YANG action and translation scripts use 64-bit infrastructure by default. Using 64-bit infrastructure instead of 32-bit

infrastructure increases efficiency because 64-bit applications can handle larger file sizes and a greater number of files. Using 64-bit infrastructure also increases system security because 64-bit applications support additional security best practices.

[See [NETCONF XML Management Protocol Developer Guide](#).]

- **Enhanced sFlow support (PTX10003)**—Starting in Junos OS Evolved Release 21.1R1, you can perform sFlow sampling and detection of traffic profiles of GRE encapsulation and de-encapsulation scenarios.

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

- **Support for the pimNeighborLoss trap (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, we support the **pimNeighborLoss** trap as defined in RFC 2934. In Junos OS Evolved release, RFC 2934 is implemented based on a draft version of the Physical Interface Module (PIM) MIB as defined in the **mib-pimmib.txt** file in the Junos OS Standard MIBs package.

The **pimNeighborLoss** trap signifies the loss of an adjacency with a neighbor device. This trap is generated when the router has an IP address lower than that of the interface to which the neighbor is connected.

[See [SNMP MIB Explorer](#).]

- **Remote port mirroring with ToS or DSCP settings (PTX10004)**
—Starting in Junos OS Release 21.1R1, this feature enables you to send sampled copies of incoming packets to remotely connected network management software. You send the packets via GRE encapsulation, which is supported by flexible tunnel interfaces (FTI). 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 <instance-name> output]** hierarchy.

[See [instance \(Port Mirroring\)](#).]

Routing Protocols

- **Support for configuring multiple independent IGP instances of IS-IS (PTX10001-36MR, PTX10003, and PTX10008)**—You can configure and run multiple independent interior-gateway protocol (IGP) instances of IS-IS simultaneously on a router.

NOTE: You cannot configure the same logical interface in multiple IGP instances of IS-IS.

[See [How to Configure Multiple Independent IGP Instances of IS-IS](#).]

- **Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP-labeled unicast (PTX10008)**—You can configure BGP PIC Edge protection for Layer 2 circuit, Layer 2 VPN, and VPLS (BGP VPLS, LDP VPLS, and FEC 129 VPLS) services with BGP-labeled unicast as the transport protocol. BGP PIC Edge using the BGP-labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multidomain networks. Multidomain networks are typically used in metro-aggregation and mobile backhaul network designs.

[See [Load Balancing for a BGP Session](#).]

- **Support for flexible algorithms in IS-IS for segment routing–traffic engineering (SR-TE) (PTX10001-36MR, PTX10004, and PTX10008)**—You can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize the IGP metric and another flexible algorithm to compute a path based on the traffic engineering metric to divide the network into separate planes. This feature enables networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

To define a flexible algorithm, include the **flex-algorithm** statement at the **[edit routing-options]** hierarchy level. To configure a device to participate in a flexible algorithm include the **flex-algorithm** statement at the **[edit protocols isis segment routing]** hierarchy level.

[See [Understanding IS-IS Flexible Algorithms for Segment Routing](#).]

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC 6286 (PTX10001-36MR, PTX10003, PTX10008)**—You can establish a BGP connection using a BGP ID that is a 4-octet, unsigned, nonzero integer. The ID must be unique only within the autonomous system (AS) per RFC 6286.

[See [router-id](#)]

- **Support for implicit filter for default EBGp route propagation behavior without policies per RFC 8212 (PTX10001-36MR, PTX10004, and PTX10008)**—We've introduced a new configuration hierarchy **defaults ebgp no-policy** at the existing **[edit protocols bgp]** hierarchy level. The configuration option separates the default policy for receive and advertise into different clauses (accept, reject, or reject-always). As a result, the route propagation behavior of the EBGp speaker varies independently from its default behavior.

[See [Implicit filter for Default EBGp Route Propagation Behavior without Policies](#) and [defaults](#)]

Segment Routing

- **Support for flexible algorithm in OSPFv2 for segment routing traffic engineering (PTX10001-36MR, PTX10003, and PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, you can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path

based on traffic engineering metric to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

To define a flexible algorithm, include the **flex-algorithm** statement at the **[edit routing-options]** hierarchy level.

To configure a device to participate in a flexible algorithm, include the **flex-algorithm** statement at the **[edit protocols ospf source-packet-routing]** hierarchy level.

[See [How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering](#).]

- **Support for strict SPF and IGP shortcut (PTX10001-36MR, PTX10003, and PTX10008)**— Starting in Junos OS Evolved Release 21.1R1, you can configure segment routing algorithm 1 (strict SPF) and advertise its SIDs in IS-IS link-state PDU (LSPDU) and use these SIDs to create SR-TE tunnels to forward the traffic by using the shortest IGP path to reach the tunnel-end point while avoiding loops. You can also specify a set of prefixes in the import policy, based on which the tunnel can redirect the traffic to a certain destination. You can use SR-Algo 1 (strict SPF) along with SR-Algo 0 (default SPF) by default when SPRING is enabled.

[See [How to Enable Strict SPF SIDs and IGP Shortcut, prefix-segment, and source-packet-routing](#).]

- **Support for adjacency SID and prefix SID in pure IPv6 environment (PTX10008)**—Starting in Junos OS Evolved Release 21.1R1, you can configure adjacency segment identifier (SID) and prefix SID for MPLS segment routing using IS-IS in pure IPv6 networks.

[See [Static Adjacency Segment Identifier for ISIS](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 21.1R1](#) | 37

Learn about what changed in the Junos OS Evolved main and maintenance releases for the PTX10001-36MR, PTX10003, PTX10004, and PTX10008.

What's Changed in Release 21.1R1

IN THIS SECTION

- [Authentication, Authorization, and Accounting \(AAA\) | 37](#)
- [Junos XML API and Scripting | 37](#)
- [Layer 2 Features | 38](#)
- [Network Management and Monitoring | 38](#)
- [Operation, Administration and Maintenance | 39](#)
- [Routing Protocols | 39](#)
- [User Interface and Configuration | 39](#)

Authentication, Authorization, and Accounting (AAA)

- **SSH session connection attempt limits and connection limits (PTX10003 and PTX10008)**—We have introduced the **connection-limit** and **rate-limit** options at the **set system services ssh** hierarchy levels. The default connection limit value is 75 connections, and the default rate limit value is 3 connections per second. Junos OS measures the rate limit value per minute but Junos OS Evolved measures the rate limit value per second.

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, PTX Series, and QFX Series)**—The `jcs:invoke()` extension function supports the **no-login-logout** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, PTX Series, and QFX Series)**—The `jcs:invoke()` extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter,

the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

Layer 2 Features

- **Modification to `sync-reset` command (ACX Series, PTX Series, and QFX Series)**—Starting from this release, the `sync-reset` command is disabled by default on all the Junos OS Evolved platforms. The `Sync-reset` command enables the device to send the sync bit in the LACP packets on minimum-link failure. Previously the `sync-reset` command was enabled by default on QFX Series, while it was by default disabled on PTX Series and ACX series.

[See [sync-reset](#).]

Network Management and Monitoring

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, PTX Series, and QFX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the `[edit system services netconf hello-message yang-module-capabilities]`. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the `[edit system services netconf netconf-monitoring netconf-state-schemas]` hierarchy level.

[See [hello-message](#) and [netconf-monitoring](#).]

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, PTX Series, QFX Series)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `edit system services netconf ssh` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response, and the `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

[See [ssh \(NETCONF\)](#).]

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

[See [monitor traffic](#).]

Operation, Administration and Maintenance

- **Fabric OAM is disabled by default (PTX10003-80C and PTX10003-160C)**—We've disabled the fabric Operation, Administration, Maintenance (OAM) feature, which helps in detecting failures in fabric paths. This release does not support disabling this feature by using the **set chassis fabric oam detection-disable**. In Junos OS Evolved Release 20.4R1, the fabric OAM feature was enabled by default.

[See [Error Handling by Fabric OAM](#).]

Routing Protocols

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

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

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, PTX Series, and QFX Series)**—The Junos OS CLI exposes the **verbose** statement at the **[edit system export-format json]** hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from **verbose** to **ietf** in an earlier release. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **[edit system export-format json]** hierarchy level. Although the **verbose** statement is exposed in the Junos OS CLI as of the current release, you can configure this statement in earlier releases.

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

Known Limitations

IN THIS SECTION

● [EVPN](#) | 40

- General Routing | 40
- Interfaces and Chassis | 41
- MPLS | 42
- Routing Protocols | 42

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.

EVPN

- If a packet with an unknown inner **ether-type** is received at the device over an EVPN-MPLS tunnel, then such a packet is dropped. [PR1564431](#)

General Routing

- The license details are not fetched in broadcast mode when the device and the license server are in the same subnet. As a workaround, configure license server explicitly on the the device using the command **set system license server address *ip*** so as to not use broadcast mode. [PR1463584](#)
- Line rate traffic can be achieved on a PTX10008 router at 256B with 400K routes scale. [PR1468280](#)
- MAC entries are stuck in I2ald or I2alm for some time after the **clear ethernet switching** command is executed. [PR1507812](#)
- The excess-rate configuration in port schedulers might not be completely honored in certain scenarios, where there could be cases where the output queues configured with guaranteed-rate with no excess-rate might end up getting more share in the excess region of the bandwidth than the queues with explicit excess-rate configured. [PR1528124](#)
- When a scheduler-map binding is removed from an interface, then the default scheduler-map is bound to the interface. If default scheduler-map is an oversubscribed scheduler map for the interface, then that map is not applied to this interface and all the "interface queue" counters for this interface show statistics as 0. [PR1539052](#)

- PTX10008: IPv6 addressing is configured at the /64 subnet by default irrespective of the subnet configured on the DHCP server side. [PR1539839](#)
- Indirect next-hop acknowledgment configuration statement makes the rpd request forwarding next-hop acknowledgment from all the Packet Forwarding Engines before doing indirect next-hop change to point to the new forwarding next hop. However, this acknowledgment is not 100% guaranteed. The acknowledgment might not work if the same forwarding next-hop (which the indirect next hop is pointing to) installation is attempted through a route add or change path because in the route add or change flow, the next-hop acknowledgment is not enabled. Since the next-hop installation is already attempted through route flow, indirect next-hop change flow cannot request acknowledgment on a forwarding next hop which is already in queue to be sent to the Packet Forwarding Engine. One possible workaround might be to enable acknowledgment for all the next hops. But this may have scaling implications and performance penalty. [PR1547432](#)
- Egress sflow sampling with sampling ratios lower than 1:1 (e.g. 1:10, 1:100, etc.) could result in up to 17% reduction in sampling capacity. This issue is not applicable to ingress sflow sampling. Also, if sampling ratio of 1:1 is configured for egress sampling, the issue is not seen. [PR1567647](#)
- In case of fixed form factor chassis like PTX10003 and PTX10001-36MR, the samples on all the interfaces on all the virtual FPCs (if any) are considered to make the adaptive sampling decision. Due to this, the maximum adaptive sampling rate supported is 35000 samples/sec per device and not per-FPC in the case of fixed form factor chassis platforms. [PR1577229](#)
- After UDP tunnel header decapsulation, MPLS Payload with MPLS label 0 and Label 2 is dropped because this requires the decap of MPLS label and lookup of inner payload which is not supported without loopback in BT ASIC based products. We support only scenarios where we decapsulate tunnel header and forward the packet based on the exposed MPLS Label. [PR1580641](#)
- PTX10003-80C or PTX10003-160C interface queue and voq does not report drops when the low priority queue is slightly oversubscribed (seeing page timeouts). [PR1581490](#)

Interfaces and Chassis

- When there is over-subscription traffic across Packet Forwarding Engines in PTX10003-80C or PTX10003-160C, and one of the Fabric ASICs (ZF) goes bad, the software takes an automatic action to recover the system by issuing an automatic SIB offline or online. However, if the egress traffic is at close to line-rate, the traffic would take time to converge till which there would be fabric drops. [PR1580376](#)

MPLS

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

Routing Protocols

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

Open Issues

IN THIS SECTION

- [Authentication and Access Control | 43](#)
- [Class of Service \(CoS\) | 43](#)
- [EVPN | 43](#)
- [General Routing | 43](#)
- [Infrastructure | 47](#)
- [Interfaces and Chassis | 47](#)
- [Layer 2 Features | 48](#)
- [MPLS | 48](#)
- [User Interface and Configuration | 48](#)
- [VPNs | 49](#)

Learn about open issues in this release for the PTX10001-36MR, PTX10003, PTX10004, and PTX10008.

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

Authentication and Access Control

- Remote authentication through console is not supported. Local authentication through console and remote authentication through management interface is supported. This is the expected behaviour on Junos OS Evolved going forward. [PR1565251](#)

Class of Service (CoS)

- After multiple activate/deactivate of class-of-service configuration, and simultaneously running **show class-of-service** command can result in COSD core file. COSD recovers after the crash. [PR1580573](#)

EVPN

- Sometimes BUM traffic coming through evpn-mpls tunnel gets dropped or duplicated when going out of aggregated Ethernet interface after tunnel termination when aggregated Ethernet members are spanned across multiple Packet Forwarding Engines. This happens only when aggregated Ethernet has member on root Packet Forwarding Engine in BUM replication path and an additional member on another PPFE. [PR1578314](#)

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 Junos OS Evolved platforms running BGP with configuration statement **route-record** or **sampling**, the rpd process might crash. [PR1493963](#)
- On a PTX10008 device, 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, until system failure handling is triggered. We recommend that you take the fabric cards offline by using the **request chassis sib offline** command before removing the SIBs for maintenance. [PR1497212](#)
- Layer 2 circuit ping that uses CC type 1, also referred to as "PWE3 Control Word with 0001b as first nibble," is not supported. [PR1503724](#)

- On a large scale routing on PTX10008, when an aggregated Ethernet interface flaps, during the next-hop convergence, you might see syslog filled with **Possible out of order deleting of AftNode** error messages. [PR1513474](#)
- When upgrading from Junos OS Evolved 20.3R1 to a future release which has an FPC type which is incompatible for restart upgrade, the user is prompted in the CLI whether to offline the FPC or abort the upgrade. The CLI command waits for one hour for the user to respond. In case the user does not respond in an hour, the CLI command **request system software add var | tmp | ptx.iso restart** times out. Subsequent CLI software **add |delete |rollback** commands fail with the 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 *primary* re restart**. [PR1516384](#)
- Using the following sensor path: /mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate -sid-counter/mpls-label, **mpls-label** values do not get reaped out when configured for SR-sid ingress sensors. [PR1516811](#)
- The FPC generates a core file that gets stuck in /var/lib/ftp/in/. Sometimes when core files from multiple FPCs are generated, the core files from one FPC can get left behind in /var/lib/ftp/in and are not processed by the core-mgr. If the **show system core-dumps** shows some older cores in /var/lib/ftp/in/*fpc-hostname* and core files newer than those appear in /var/core/*node-name*, it indicates this issue was hit. [PR1522404](#)
- On any Junos OS Evolved system with SSD, in rare scenarios, with multiple abrupt and/or ungraceful power-cycles or reboots, the filesystem might encounter error or it might not get initialized resulting in user login failure. This issue does not occur on a graceful shutdown of the system. [PR1523238](#)
- On platform running Junos OS Evolved with dual Routing Engine support, when a new backup Routing Engine is inserted into the system, the **request system software sync** CLI command could not synchronize the software from the primary Routing Engine to the backup Routing Engine. [PR1529924](#)
- The SR-TE label with BSID traffic gets dropped in PTX10003. Please use the following configuration statement as a workaround:

```
protocols {
    source-packet-routing {
        no-chained-composite-next-hop;
    }
}
```

[PR1529933](#)

- Each filter contains one or more filter terms. Each filter term might contain a counter. If traffic hits the filter term then the counter (packets/bytes) increments. If a user tries to modify the unsupported scaled filter then counters (packets/bytes) in filter terms are cleaned up which means the counters do not retain the original values before the modification. [PR1530597](#)
- Any route or nexthops updates going during switch over process can appear out-order back on the primary. This can result into rpd core. [PR1531653](#)
- When multicast traffic is running over an aggregated Ethernet bundle, which spans multiple FPCs, if one of the FPCs is brought offline, the traffic on that FPC is impacted. It is expected to see some traffic drops. [PR1531859](#)
- In a scaled setup, it is possible that during GRES operation, the rpd on the primary Routing Engine m receive an ifstates update (DEL and CHANGE) and these are not expected during rpd re-syncing. In such a case the rpd can generate a core file and might recover automatically after sometime. [PR1537947](#)
- During PHP, exp rewrite does not happen on the exposed MPLS payload. [PR1538918](#)
- Software does not support GRES, when FPCs are triggered for restart. GRES with FPC triggered for offline or online is supported. [PR1539685](#)
- Available space in /data on RE0 (primary) and RE1 (backup) can become asymmetric due to the difference in size of the db files. This can cause the software upgrade to fail if there are more than 10 images on the system. Also if /data is only full in RE1 (backup) and not on RE0 (primary) then force option may not able to clean up the old software images to make space for new image to get installed. As a workaround, use the CLI command **request system software delete** to remove images from the system until there is adequate space on RE0 and RE1 to install the new software image. [PR1556992](#)
- On a PTX10008 running Junos OS Evolved, the **show interface** command combined with refresh configuration statement becomes slow after a **clear interface statistics** command. [PR1557843](#)
- Always use the **show system switchover** command on the backup Routing Engine to determine if the standby is ready to take over. [PR1557461](#)
- We have provided a configuration statement to control zero suppression behavior for streaming telemetry. The default behavior is still going to be suppressing zeroes so there will be no backward compatibility issues.

The new configuration statement is: **set services analytics zero-suppression no-zero-suppression**. Committing this configuration disconnects the existing collectors and users are required to establish streaming telemetry connections again from their collectors. [PR1559882](#)

- The asynchronous write APIs to CDA are fire-and-forget calls. If CDA clients make asynchronous write API calls to CDA server, the calls at rate higher than the CDA server can process them, this

leads to memory usage buildup in GRPC infrastructure, eventually leading to memory allocation failures in GRPC infrastructure code and process abort. [PR1565427](#)

- The command **sysmanctl app stop distributord node re0 trigger** is to be avoided in PTX10008 with NSR enabled as this may lead to flap of BGP sessions. [PR1570574](#)
- Per-BT-chip sampling rate limit reduced from 35000/seconds to 7000/seconds. [PR1571929](#)
- This issue is seen on PTX10008 with EVPN MPLS configuration, when there is an aggregated Ethernet interface associated with evpn-extended vlan. The aggregated Ethernet has multiple member links on different FPCs. It is observed that on some FPCs IFBD state for the aggregated Ethernet is out of sync with app controller for some reason. When the flood-nh update comes, it does not find the IFBD and registers for late binding. The IFBD object present in the controller, sends the flood-nh update again and this cycle keeps on repeating. This cycle keeps the Packet Forwarding Engine process busy and the Packet Forwarding Engine is unable to handle requests from CLI, and it causes the CLI timeout issue with error **communication failure with /fpc0/evo-aftmand-bt/**. Also as flood-nh programming is unable to finish due to missing aggregated Ethernet IFBD, this causes issues with L2 flooding within the bridge domain. The chance to see this issue is relatively low. [PR1571929](#)
- Changing Decap only tunnel destination address configuration after tunnel is up is not handled and previously configured tunnel destination address is used for decapsulation. Once system enters this state any further changes to tunnel configuration is not handled. [PR1575724](#)
- The Junos OS Evolved jkeys match the Junos OS jkeys util-asic-to-host-packet-rate -> util-wan-and-host-inject-pps-rate. [PR1578134](#)
- If a switchover is performed while an FPC is in the offline or online state, the FPC moves to the current state on the primary. [PR1580726](#)
- On the HwD app restart, certain propertyE objects might not be recreated. [PR1580735](#)
- In certain scenarios, shapers applied on a 10G interface might drop the traffic more than the configured max-rate. [PR1580795](#)
- Ping does not work when OSPF configuration is applied. [PR1581130](#)
- The config-sync service might appear offline on the backup Routing Engine in the case of a USB upgrade. [PR1582717](#)
- The sflow ingress sampling does not work for user IPv6 traffic with aggregated Ethernet ECMP case at last hop router with UHP LSP. [PR1582960](#)
- The OSPF-hello DDOS statistics pktCnt lists as 0. [PR1584458](#)

Infrastructure

- When using a source IP address as the management address of the box to ping a network address on a peer, the response for the ICMP ping from the peer can end up on the management interface of the box, which is dropped by the Linux kernel as the RPF check is set to strict by default on the Linux kernel used on Junos OS Evolved. The Linux kernel expects the path to the peer to be on the WAN side and so drops the packet when it is received on the management interface from the peer when the RPF check for the management interface is set to strict. [PR1498255](#)
- The standard "issue" message is not displayed for the initial login attempt, but if that login attempt fails for any reason and the user is re-prompted, then the message is displayed. This seems to be an issue common to PTX10004|PTX10008|PTX10001-36MR. [PR1528996](#)

Interfaces and Chassis

- When routing instance primary IFF does not have an address, ping packets do not have source filled. This causes the ping packet reply to be not received. As a workaround, configure the lo0 address. [PR1547622](#)
- LACP force up option does not work. [PR1553997](#)
- There is no functional impact on resiliencyd. This core file is generated from the test command using cli-pfe. Once cored resiliencyd recovers, it behaves normally as expected. [PR1578822](#)
- The format specifier of alarms is changed from string to boolean for following leafs.
lane_tx_loss_of_signal_alarm lane_rx_loss_of_signal_alarm lane_tx_laser_disabled_alarm. [PR1580113](#)
- The output and receiver power should be displayed with "- Inf" string if it 0mW as the following:
Laser receiver power : 0.000 mW / - Inf dBm. This is not working in 21.1R1. [PR1582541](#)
- The ifmand core file seen at IFAManager::handleIFACModify.
 - Dual Routing Engine system configured with 2 management interfaces having addresses with primary-only configuration.
 - Address primary-only flag is removed in later commit.
 - Ifmand crashes and recovers.
 - Issue is only applicable in 21.1R1 release and 20.3.
 - [PR1583681](#)

Layer 2 Features

- It was observed rarely that issuing a **request system zeroize** does not trigger Zero Touch Provisioning (ZTP). A simple workaround is to re-initiate ZTP. [PR1529246](#)

MPLS

- The in-progress change of weighted ECMP to ECMP and vice versa does not take full effect for all routes using the LSPs until all the LSPs currently signaled with non-zero bandwidth have gone through at least one of the MBB/clear/disable events. The traffic distribution is not as expected until all the LSPs finish the transition. This is because without any of those triggers, the LSP does not automatically update its next hop based on weighted ECMP configuration. The next-hop content of each participation LSP affects the traffic distribution among ECMP paths. The weighted ECMP distribution ratio is not as expected until the adjustment of all LSPs is finished. [PR1501860](#)
- On PTX10008 with scaled LSPs (10,000 ingress, 40,000 transit), the **show mpls lsp statistics** command might take up to 55 minutes. [PR1551765](#)

User Interface and Configuration

- On-box validation of the large YANG files on NETCONF through request-yang-validate RPC might not work as the validation process might reach out of memory state while loading such YANG files. [PR1416972](#)
- When a user tries to deactivate the MPLS-related configuration, the commit fails on the backup Routing Engine. [PR1519367](#)
- In the dual Routing Engine scenario, the backup Routing Engine does not sync the configuration change while deleting a deactivating interface or routing-instances from the primary. So, after the operation, the deactivated interface or routing-instances still exist on the backup Routing Engine. [PR1552467](#)
- If an inet filter is configured with or without the **family** statement, then changing its configuration to without or with the **family** statement must not be done. [PR1556424](#)
- If an inet filter is configured with or without a **family** statement, then changing its configuration to remove or add a **family** statement might cause the filter process (firewalld) to restart unexpectedly. [PR1556426](#)

VPNs

- When Layer 2 VPN is configured with warm-standby, then upon Routing Engine switchover the new primary generates a core file. [PR1584642](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 21.1R1 | 49](#)

Learn which issues were resolved in the Junos OS Evolved main and maintenance releases for the PTX10001-36MR, PTX10003, PTX10004, and PTX10008.

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

Resolved Issues: 21.1R1

IN THIS SECTION

- [Authentication and Access Control | 50](#)
- [Class of Service \(CoS\) | 50](#)
- [General Routing | 50](#)
- [Infrastructure | 55](#)
- [Interfaces and Chassis | 56](#)
- [Routing Policy and Firewall Filters | 57](#)
- [Routing Protocols | 57](#)
- [User Interface and Configuration | 57](#)

Authentication and Access Control

- Junos OS Evolved versus Junos OS: The **strict-authorization** configuration option is missing under **set system tacplus-options**. [PR1557052](#)

Class of Service (CoS)

- PTX10008: scheduler with all queues oversubscribed, maximum latency is different on different queues 21ms ~ 29ms. [PR1478811](#)
- While configuring WRED profile to a scheduler, you can use either of the following: any or any not-any or not-any combination of protocol and loss priority. [PR1524259](#)
- The cosd might not come up after FPC restart. [PR1544531](#)
- When configuring Class of Service on a Junos OS Evolved system, the cosd process might restart unexpectedly. [PR1548014](#)
- cosd crash might be observed with classifier and rewrite bindings. [PR1566161](#)

General Routing

- LSP statistics CLI is slow in a scaled scenario during installation time. [PR1416363](#)
- Differences in XML tags for **show system virtual-memory**. [PR1438110](#)
- Differences in XML tags for **show system processes**. [PR1438129](#)
- 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](#)
- Rate-limiting might not work for J-Flow-sampled traffic on PTX10008. [PR1473844](#)
- PTX10008: No cmerror is raised for certain PIO errors when accessing the Packet Forwarding Engine ASIC. [PR1491130](#)
- The interface might not be added to BD after VLAN change. [PR1504374](#)
- Junos OS Evolved versus Junos OS (serviceability): Resilient hash seed configuration is not supported in Junos OS Evolved Release on PTX10003 or PTX10008. [PR1504544](#)
- Transit v4 traffic forwarding over BGP SR-TE might not work. [PR1505592](#)

- On a Junos OS Evolved platform, the output of aggregated Ethernet (AE) interface statistic does not include its member links' statistics. [PR1505596](#)
- OSPF might flap when the line rate traffic is sent at smaller packet sizes. [PR1511563](#)
- Incorrect warning message for the **show chassis fabric errors** command. [PR1511915](#)
- Junos OS Evolved (PTX10008): observing reboot delay in software add reboot, rollback reboot commands. [PR1525286](#)
- PTX10001-36MR :: Junos OS Evolved :JDI-RCT:clockd.default anomalies are seen at producer:re0::clockd::50331749 for the type net::juniper::resild::errorItem. [PR1527309](#)
- Host path corruption might be observed after output filters are configured on some Junos OS Evolved devices. [PR1528368](#)
- PTX10008 - FPC reboots if DDoS violations are seen during FPC restart. [PR1529847](#)
- [PTX10001-36MR] When the fabric hub process is restarted with the blackhole detection disable and fabric degrade detection enable configuration, it might result in traffic loss. [PR1530484](#)
- **show pfe statistics error** should print counters for meaningful errors. [PR1530710](#)
- After FPC restart traffic drop is observed for multicast streams sent to the PTX10008 device from other nodes. [PR1531429](#)
- PTX10008: Need support for **show chassis fabric summary** output.
[PR1532163](#)
- Traffic loss might happen after performing GRES in the Junos OS Evolved enabled chassis-based systems.[PR1532446](#)
- **show chassis aer-statistics** includes uncorrectable error threshold as well. [PR1533197](#)
- Interfaces might take longer to come up after loading baseline and rollback configurations.
[PR1534996](#)
- SNMP does not send a response when requested with partial FRU information on certain PTX Series platforms running Junos OS Evolved. [PR1535204](#)
- PTX10008: Multicast - After FPC restart, both traffic flood and loss are observed toward downstream aggregated Ethernet receiver interfaces. [PR1535545](#)
- CLI issues are observed under the **show system** and **show chassis** hierarchies. [PR1535880](#)
- The physical interface does not come up after configuration changes at the peer end. [PR1536270](#)
- The cosd might crash when multiple configuration changes are made in a single commit. [PR1536320](#)

- The **show chassis environment** command reports ZF-based switch fabric internal temperatures as 0 degrees celsius. [PR1536497](#)
- cosd core file is generated while configuring duplicate code point bits/alias. [PR1537289](#)
- On a PTX10008 running Junos OS Evolved, running PIM Dense Mode (DM) might cause an FPC to restart unexpectedly. [PR1537700](#)
- Next-hop addCollision errors are seen with scaled multicast routes. [PR1538849](#)
- On RE0 CLI node reboot, the rpd process is unable to connect to snmpd and hence we see scheduler slips in rpd and the protocol sessions flap. [PR1539705](#)
- On PTX10001-36MR on ports et-0/1/2 and et-0/2/2, channel 1 does not come up on channelizing 100-Gbps to 2x50-Gbps speed. [PR1539795](#)
- On PTX10001-36MR ping fails on all channels when some ports with 400GbE DAC are channelized to 8x25GbE and 8x50GbE. [PR1539859](#)
- On PTX10001-36MR traffic might not flow through channels 2 and 3 of some ports with 40G optics that are channelizing to 4x10GbE [PR1539864](#)
- Multicast traffic silent packet drop on downstream adjacent node seen for approximately 5-6 seconds as dlu.unicode.inv_start_pc trap after aggregated Ethernet member link deleted or added on PTX10008 node. [PR1539912](#)
- PTX10003: Multicast traffic drops after port is converted from trunk to access for a Layer 2 bridging case. The unicast and broadcast traffic are not affected. [PR1540495](#)
- In PTX10008, for some temperature sensors the threshold temperature to return cooling fans to normal speed is 0 degrees celsius. [PR1540576](#)
- PTX10008: Core file generated when a combination of hop limit and packet length is configured in the filter. [PR1540625](#)
- Installed licenses are getting deleted after an image upgrade. [PR1540881](#)
- PTX10008: Component Upgrade: Software rollback timed out when moving from image based on Junos OS Evolved Release 20.3R1 (mod in some apps) to any other Junos OS Evolved Release. [PR1541509](#)
- DHCP discover packet might be dropped if DHCP inform packet is received first. [PR1542400](#)
- [cos] [scheduler] : PTX10008: Junos OS Evolved CoS: The **show interfaces voq ae2 non-zero** command does not display the non-zero queue statistics, whereas the other aggregated Ethernet bundle ae1 displays it correctly (both ae1 and ae2 have 4x400G member links each from the same BT chip). [PR1543034](#)

- PTX10008: Junos OS Evolved picd generates a core file in reboot scenarios. [PR1543045](#)
- Any modification made in the middle of an existing firewall filter might lead to all the host-bound traffic getting discarded. [PR1544502](#)
- PTX10001-36MR - IP-IP: Routing Engine initiated tracroute packets are not using IP-IP encapsulation. [PR1545049](#)
- PTX10008:PTX10K-LC1202-36MR: PTX10001-36MR:L2:classifier classifier on aggregated Ethernet interface does not work correctly (cos_l2_cls_combine_mts_004.robot). [PR1546150](#)
- [PTX10008]: For 2x100G QSFP-DD optics, CTLE settings are not getting applied sometimes. [PR1546236](#)
- Backup RE vmcore may be seen due to absence of NH ACK Infra. [PR1547164](#)
- PTX10001-36MR:: **show ddos-protection protocols bgp statistics brief** throws error **communication failure with /re0/evo-aftmand-bt/**. [PR1547491](#)
- The drop profile (WRED) feature under class of service might not work on PTX Series Junos OS Evolved platforms. [PR1549007](#)
- With a firewall policer configured, after GRES, the firewalld application might crash. firewalld restores itself after the crash. [PR1549856](#)
- On Junos OS Evolved PTX10001-36MR, PTX10004, or PTX10008 devices, some host-received packets are handled in a different host path queue than the one being assigned due to the host-path queue default setting. [PR1551154](#)
- Traffic drops might be seen after egress FPC restarts and comes back online. [PR1551363](#)
- SNMP MIB jnxOperatingRestartTime output is in string format instead of OCTET string. [PR1553533](#)
- Junos OS Evolved:PTX10008:INDB:Compatible: Offline FPCs with user option no ends after the FPC is brought online after the upgrade with sysman as one of the apps. [PR1553667](#)
- The filter behavior is unchanged after deactivating the filter on the management port. [PR1553791](#)
- jnxFruLastPowerOn value is incorrect for FPCs. [PR1553924](#)
- PTX10004 :: The error *cmerror dp_1_viqm_intr_req_enq_err* is seen and it is not getting cleared even after loading the latest version Junos OS Evolved 20.4. [PR1553943](#)
- CoS WRED Curve: Create Expr Curve: No curve data points. Errors are seen when interpolate is configured under drop profile. [PR1554220](#)

- PTX10008 PFC: This is a display issue. Global Ethernet flow-control must be disabled when PFC CNP is enabled on an interface. [PR1554345](#)
- The output of **show interface queue** always shows **Forwarding classes: 16 supported, 4 in use** with customized configuration. [PR1554370](#)
- MACsec session remains down after the CA name is changed to a newer name. [PR1555736](#)
- Cleaning up the unsupported configuration statement **set chassis redundancy keepalive-time and failover** from Junos OS Evolved software as it is not yet supported in Junos OS Evolved. [PR1556600](#)
- Upgrade or downgrade: Serviceability: In case of upgrade failure, the log files capturing upgrade fail reason do not give meaningful information to user. [PR1556807](#)
- The device might not boot up after performing the cleanup operation. [PR1557020](#)
- The configuration command **set routing-instances inst1 l3-interface** is not supported on PTX10008. [PR1558285](#)
- PTX1008-Error messages are seen @Error] Jexpr: CoS Scheduler Express Handle: Destructor: Interface Physical Handle is NULL. streamIndex:1147 schedNodeToken:4508 while executing 2000 static SRTE scale in PTX10008. [PR1558328](#)
- [firewall] [filter_installation] : Junos OS Evolved-PTX10003 :: firewalld:Anomalies are seen in firewalld app for Publish publish-deleted.[PR1559046](#)
- The FPC might reboot in a high-scale configuration scenario on Junos OS Evolved PTX10008 platform. [PR1560757](#)
- PTX10008: After sync all followed by rollback and then reboot, RE1 booted on snapshot. [PR1562189](#)
- Complete ingress multicast traffic loss might be seen on interfaces that are flapped using Packet Forwarding Engine offline/online command. [PR1562452](#)
- Interface loopback might not work if there are no optics connected to the port on PTX10008. [PR1562471](#)
- For topologies involving high ingress and transit LSP scale, error messages can be seen in journalctl when tearing down the ingress and transit LSPs. This also leads to slow memory leak. [PR1562503](#)
- FPC is not powered on using **request node power-on fpc**. [PR1562981](#)
- PTX10001-36MR - In scaled logical interface scenario when an interface connecting the customer with many logical interfaces in the same port is flapped, ARP might not resolve and the traffic might be dropped. [PR1563684](#)

- PTX10008: RE0 went into reboot loop continuously during validate restart (INDB unsupported). [PR1563742](#)
- The issue happens when there are multiple FPCs. Any of the FPCs might have scale configuration. While restarting an FPC (fpc[x]), it syncs configuration from other FPCs in the system. Evo-aftmand on fpc[x] takes a lot of time (depending on the scale) in reconciling the huge number of objects. It can only create the IFD, after it has reconciled the existing objects. The issue is in the infrastructure code, which takes a lot of time in reconciliation of the objects. [PR1564156](#)
- MACsec-Encrypted packets counter displays 0 under **Secure Association transmitted** of MACsec statistics output when AN rollovers with **sak-rekey-interval** is configured. [PR1566665](#)
- User folders are not created when snapshot is taken. [PR1567880](#)
- Minor memory leak in ndpd (NDP daemon) when **show ipv6 neighbors** or **clear ipv6 neighbors** is executed or on subscribing for '/nd6information' from collector using telemetry. [PR1568370](#)
- The firewalld crash might be seen if GRES is executed as soon as the firewall is activated (for example, commit is done). [PR1569427](#)
- Junos Telemetry Interface: PTX10008 :: NPU Memory KHT DLU IDB value. [PR1572704](#)
- The hash-key label-1-exp CLI configuration statement does not take effect. [PR1573109](#)
- Silent switchover is triggered on restarting routing with the CLI configuration: **set system processes routing failover other-routing-engine** and **set system switchover-on-routing-crash**. [PR1570993](#)

Infrastructure

- PTX10008 - with graceful-restart enabled (IS-IS or LDP or BGP), packet loss is observed during RE switchover. [PR1518609](#)
- Continuous log messages: ttp_update_with_mark_tlv:390] loss priority not supported yet. [PR1536732](#)
- When the CLI is launched, the system generates a syslog message at the notice level stating **cli[<pid>]: 32-bit application cli is not compatible with SI**. This message does not indicate a functional problem and can be ignored.. [PR1553897](#)
- [mgd_infra] [mgdtag] : Junos OS Evolved-PTX10003 :: Default multicast ff00::/8 route not available while verifying IPv6 multicast routes.
- The local port specified in an address-less bind() call is not applied on a TCP socket that issues a connect(). [PR1569696](#)

Interfaces and Chassis

- [Junos OS Evolved Changes]: XML tags mismatch for the command **show interfaces diagnostics optics**. [PR1529316](#)
- [Junos OS Evolved] [PTX10008] ALB scan-interval configuration not taking effect. [PR1538854](#)
- FTI tunnel commit fails for MTU change on its family (inet, inet6). [PR1540431](#)
- CLI command **show interface descriptions** does not show Link and admin status of the logical interfaces in routers running affected Junos OS Evolved release. [PR1545787](#)
- The **set chassis fpc power on** command in PTX10003 is removed as it is not supported. [PR1550862](#)
- Logs not being written in /var/log/messages on certain PTX Series platforms running Junos OS Evolved. [PR1551374](#)
- PTX10001-36MR: Control logical interface might not be present for ports et-0/0/11 and et-0/2/11. [PR1566752](#)
- Junos OS Evolved-PTX10003 :: Junos Telemetry Interface (JTI) optics sensor's alarm data type changed from **bool_val** to **str_val**. [PR1580113](#)
- Dynamic and persistent aggregated Ethernet MAC address allocation changes. [PR1575131](#)

MPLS

- On all Junos OS Evolved platforms, the rpd crash might be observed in a rare scenario when the configuration statement **no-propagate-ttl** is configured under the apply-groups along with link protection for static LSP. This issue is seen when activating or deactivating the MPLS configuration and might lead to traffic loss. [PR1528460](#)

Network Management and Monitoring

- A memory leak in the mib2d and snmpd processes might result in SNMP being unresponsive to SNMP queries on platforms running the affected Junos OS Evolved image. [PR1543508](#)
- The syslog messages might not be sent with the correct port. [PR1545829](#)
- When SNMP polling from NMS is in progress and the policy prefix configuration associated with SNMP community client-list-name is modified to add or delete a prefix, then snmpd may generate a core file. snmpd comes up fine after the core file is generated and no action is required from the user. [PR1548595](#)
- FTI: IP-in-IP tunneling: SNMP: ifMtu is wrong in snmp mib walk against fti0 interface. [PR1549220](#)

- Core file for trace-relay process is seen. [PR1556040](#)

Routing Policy and Firewall Filters

- [pfe_firewall] [policer] PTX10008 :: A 100 pps traffic is seen if two color policers are applied with a pps limit and burst limit of 1 pps. [PR1541194](#)
- BGP-LU fallback-reject event is not generating error messages. [PR1550023](#)
- Firewall: Use of filter configured without address-family as nested filter causes the configuration commit to fail. [PR1552641](#)
- Junos OS Evolved - changing the interface-specific field of the referenced filter in a single commit is not supported. [PR1556982](#)
- Syslog as an action of filter by default dump logs in syslog in Junos OS Evolved which is different from Junos OS. [PR1564088](#)
- Error while applying the filter as an output-list with DSCP action. [PR1569691](#)

Routing Protocols

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

User Interface and Configuration

- The firewall filter for both IPv4 and IPv6 might not work when it is applied through apply-groups. [PR1534858](#)
- The Junos OS Evolved operational state would be incorrect on the system and CoS schedulers configuration change might not take effect. [PR1536615](#)
- File archive source /var/log/* destination /var/log/.tgz is not working in Junos OS Evolved. [PR1549340](#)
- CLI command **show system uptime** throws error: **invalid xml tag (date: invalid date '@#')** from **command-handler daemon** on certain PTX Series platforms running Junos OS Evolved. [PR1566166](#)

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

IN THIS SECTION

- [What's New | 59](#)
- [What's Changed | 63](#)
- [Known Limitations | 66](#)
- [Open Issues | 67](#)
- [Resolved Issues | 69](#)

These release notes accompany Junos OS Evolved Release 21.1R1 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.

See [Table 7 on page 58](#) table for a complete list of the Junos OS Evolved releases supported on the QFX Series switching platforms.

Table 7: Junos OS Evolved Releases Supported on the QFX Switching Platforms

Platform	Supported Junos OS Evolved Releases
QFX5130-32CD	20.3R1, 20.3R2, 20.4R1, and 21.1R1
QFX5220-32CD	19.1R2, 19.2R1, 19.3R1, 19.3R2, 19.4R2, 20.1R1, 20.1R2, 20.2R1, 20.2R2, 20.3R1, 20.3R2, 20.4R1, and 21.1R1
QFX5220-128C	19.1R2, 19.2R1, 19.3R1, 19.3R2, 19.4R2, 20.1R1, 20.1R2, 20.2R1, 20.2R2, 20.3R1, 20.3R2, 20.4R1, and 21.1R1

What's New

IN THIS SECTION

- [What's New in 21.1R1 | 59](#)

Learn about new features introduced in the Junos OS Evolved main and maintenance releases for the QFX5130-32CD and QFX5220.

What's New in 21.1R1

IN THIS SECTION

- [Hardware | 59](#)
- [MPLS | 62](#)
- [Multicast | 62](#)
- [Services Applications | 63](#)

Hardware

- We haven't introduced new hardware in Junos OS Evolved 21.1R1. Use the following table to learn about some of the new features supported on the QFX5130-32CD switch.

Table 8: Feature Support Added to the QFX5130-32CD Switches

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none">• Support for priority-based flow control (PFC) of untagged traffic at Layer 3 using DiffServ Services Code Points (DSCPs). This enables lossless traffic across Layer 3 routers that connect Layer 2 subnets. DSCP-based PFC is required to support Remote Direct Memory Access over Converged Ethernet version 2 (RoCEv2). <p>[See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.]</p>
Forwarding options	<ul style="list-style-type: none">• Support for storm control, which enables the switch to:<ul style="list-style-type: none">• Monitor traffic levels.• Drop broadcast, multicast, and unknown unicast (BUM) packets before they can cause a traffic storm. <p>[See Understanding Storm Control.]</p>

Table 8: Feature Support Added to the QFX5130-32CD Switches *(Continued)*

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for Q-in-Q tunneling with a service-provider-style configuration in addition to the enterprise style that is already supported. You can enable the configuration for a physical interface using the set encapsulation flexible-ethernet-services command. Q-in-Q tunneling with service-provider-style configuration has these restrictions: <ul style="list-style-type: none"> • Untagged traffic is supported only if you use the native-vlan-id command for the physical interface. Untagged traffic will egress the network-to-network interface (NNI) or service VLAN interface as a single-tagged S-VLAN. • Insertion of native-vlan-id as the inner tag is not supported. The input-native-vlan-push configuration command is not supported. • Swap-push/pop-swap operations on single-tagged interfaces and swap-swap operations on double-tagged interfaces are not supported. • Multiple user-to-network interfaces (UNIs, or C-VLAN interfaces) with different VLAN IDs or with different Q-in-Q operations on the same bridge-domain are not recommended. C-VLAN traffic on one logical interface might leak into another. • Swap operation with the inner-vlan-id configuration command on the UNI is not supported. • Custom Tag Protocol Identifier (TPID) configuration and translation is not supported. • IRB logical interfaces and C-VLAN (UNI) logical interfaces cannot coexist on the same VLAN.

Table 8: Feature Support Added to the QFX5130-32CD Switches (Continued)

Feature	Description
	<ul style="list-style-type: none"> The vlan-id-list configuration is supported only on interfaces with valid I/O maps. <p>[See Flexible Ethernet Services Encapsulation and Configuring Q-in-Q Tunneling.]</p>
Timing and synchronization	<ul style="list-style-type: none"> Support for Precision Time Protocol (PTP) transparent clock. <p>[See Understanding Transparent Clocks in Precision Time Protocol.]</p>

MPLS

- Install prefixes for RSVP-TE LSPs using PCEP (QFX Series)**—Starting in Junos OS Evolved Release 21.1R1, you can configure different prefixes for Path Computation Element (PCE)-initiated and PCE-delegated RSVP-TE LSPs using the Path Computation Element Protocol (PCEP). Prior to this feature, for PCE-initiated LSPs, you could install prefixes as routes through templates and map the templates to the LSPs. For Path Computation Client (PCC)-configured LSPs, although you could install prefixes on the device, this information was not reported to the PCE.

With this feature, you can install prefixes for external RSVP-TE LSPs through PCEP communication, and enable the PCC to report installed prefixes for all local RSVP-TE LSPs to the PCE. This support provides you better traffic engineering capabilities and allows Junos OS to interoperate with other vendor's PCC or PCE.

[See [Support of the Path Computation Element Protocol for RSVP-TE Overview](#).]

Multicast

IN THIS SECTION

- [Routing Protocols](#) | 63

- Support for IGMP and IRB elaboration snooping. You can configure IGMP snooping with IGMPv1, IGMPv2, and IGMPv3 and IRB. [See [IGMP Snooping Overview](#) and [Integrated Routing and Bridging](#).]

Routing Protocols

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC 6286 (QFX5220)**—You can establish a BGP connection using a BGP ID that is a 4-octet, unsigned, nonzero integer. The ID must be unique only within the autonomous system (AS) per RFC 6286.

[See [router-id](#)]

Services Applications

- **Support for Packet Forwarding Engine DDoS protection (QFX5220)**—Starting in Junos OS Evolved Release 21.1R1, you can configure and install policers at the Packet Forwarding Engine (PFE) level for defense from DDoS attacks. By default, DDoS protection is enabled for many protocols on QFX5220-32CD and QFX5220-128C switches. You can disable DDoS protection or change default policer parameters for a protocol group or supported packet types in a protocol group. You can also configure the bandwidth or burst limit at the protocol level.

[See [Security Services Administration Guide](#), [show ddos-protection statistics](#), and [show ddos-protection version](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 21.1R1](#) | 63

Learn about what changed in the Junos OS Evolved main and maintenance releases for the QFX5130-32CD and QFX5220.

What's Changed in Release 21.1R1

IN THIS SECTION

- [EVPN](#) | 64
- [Junos XML API and Scripting](#) | 64

- Layer 2 Features | 64
- Network Management and Monitoring | 65
- Routing Protocols | 65
- User Interface and Configuration | 66

EVPN

- **Unresolved hosts identified in MAC-IP address entries (QFX5130-32CD, QFX5220)**—When you use the `show ethernet-switching mac-ip-table` command to display the MAC-IP entries in the Ethernet switching table, unresolved hosts are identified with a `Ur` flag.

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, PTX Series, and QFX Series)**—The `jcs:invoke()` extension function supports the **`no-login-logout`** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, PTX Series, and QFX Series)**—The `jcs:invoke()` extension function supports the **`no-login-logout`** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

Layer 2 Features

- **Modification to `sync-reset` command (ACX Series, PTX Series, and QFX Series)**—Starting from this release, the **`sync-reset`** command is disabled by default on all the Junos OS Evolved platforms. The **`Sync-reset`** command enables the device to send the sync bit in the LACP packets on minimum-link

failure. Previously the **sync-reset** command was enabled by default on QFX Series, while it was by default disabled on PTX Series and ACX series.

[See [sync-reset](#).]

Network Management and Monitoring

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, PTX Series, and QFX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the **[edit system services netconf hello-message yang-module-capabilities]**. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the **[edit system services netconf netconf-monitoring netconf-state-schemas]** hierarchy level.

[See [hello-message](#) and [netconf-monitoring](#).]

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, PTX Series, QFX Series)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the **client-alive-interval** and **client-alive-count-max** statements at the **edit system services netconf ssh** hierarchy level. The **client-alive-interval** statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response, and the **client-alive-count-max** statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

[See [ssh \(NETCONF\)](#).]

Routing Protocols

- **Recommendation to include the local-address statement when configuring IBGP and multihop EBG**
—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 EBG 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](#).]

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, PTX Series, and QFX Series)**—
The Junos OS CLI exposes the **verbose** statement at the **[edit system export-format json]** hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from **verbose** to **ietf** in an earlier release. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **[edit system export-format json]** hierarchy level. Although the **verbose** statement is exposed in the Junos OS CLI as of the current release, you can configure this statement in earlier releases.
[See [export-format](#).]

Known Limitations

IN THIS SECTION

- [Infrastructure](#) | 66
- [Routing Protocols](#) | 66
- [System Management](#) | 67

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.

Infrastructure

- When the USB is plugged in, the image is not picked from the USB in the boot-order by default.
[PR1399311](#)

Routing Protocols

- When the multicast stream has IRB as receivers with IGMP/MLD snooping enabled in the particular bridge or VLAN, there can be traffic loss for 3 minutes by default up to 5 minutes (based on the GR

configuration for MCSNOOPD) when the snooping configuration is removed. The multicast traffic streams recover and flooding for IRB receiver in the VLAN is seen as expected. [PR1550523](#)

System Management

- On the QFX5130-32CD platform, the port status LED remains off for a channelized interface, when one of the channels in the channelized port is down or is disabled. The port LED is lit only when all the channels that are part of the port are up. [PR1526532](#)
- The all statistics in **show chassis fpc** is missing as in following `root@host> show chassis fpc`

Temp CPU Utilization (%) CPU Utilization (%) Memory Utilization (%) Slot State (C) Total Interrupt
1min 5min 15min DRAM (MB) Heap Buffer 0 Online. [PR1563506](#)

Open Issues

IN THIS SECTION

- [Infrastructure | 68](#)
- [Interfaces and Chassis | 68](#)
- [Layer 2 Features | 68](#)
- [System Management | 68](#)
- [User Interface and Configuration | 69](#)

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.

Infrastructure

- The standard "issue" message is not displayed for the initial login attempt, but if that login attempt fails for any reason and the user is re-prompted, then the message is displayed. This seems to be an issue common to QFX5130-32CD. [PR1528996](#)

Interfaces and Chassis

- This issue seen from 20.4R1-EVO in QFX5130-32CD platform. OSPF routes are not learnt when neighborship is formed over IRB interfaces. [PR1570498](#)

Layer 2 Features

- It was observed rarely that issuing a **request system zeroize** does not trigger Zero Touch Provisioning (ZTP). A simple workaround is to re-initiate ZTP. [PR1529246](#)

System Management

- On the QFX5130-32CD platform running Junos OS Evolved, you cannot clear or reset the disk option specified in the scheduled **request node reboot** command. The node reboots with the disk option last specified. [PR1517596](#)
- When the evo-pfemand process is restarted due to any error conditions, the channelized interfaces (if configured) flap two times: first during the initialization and the second time after the configuration is programmed. [PR1526984](#)
- QFX5130-32CD Junos OS Evolved-pfemand binding queue complete-deleted objects are seen with hwd::pcsE. [PR1531820](#)
- Intermittently 100G interfaces will not come up between PTX10003-80C and QFX5220-128C after a configuration change and commit. [PR1531820](#)

User Interface and Configuration

- On box validation of the large YANG files on NETCONF through request-yang-validate RPC might not work as validation process might reach out of memory state while loading such YANG files. [PR1416972](#)

Resolved Issues

IN THIS SECTION

- Resolved Issues: 21.1R1 | [69](#)

Learn which issues were resolved in the Junos OS Evolved main and maintenance releases 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.

Resolved Issues: 21.1R1

IN THIS SECTION

- Interfaces and Chassis | [69](#)
- Network Management and Monitoring | [70](#)
- Routing Protocols | [70](#)
- System Management | [70](#)
- User Interface and Configuration | [71](#)

Interfaces and Chassis

- Provision control logical interface when none is configured for auto neighbor discovery. [PR1507347](#)

Network Management and Monitoring

- Trace file is not created under SNMP. [PR1546784](#)

Routing Protocols

- The virtual-router option is not supported under a routing-instance in a lean rpd image. [PR1494029](#)
- The IPv6 traffic might get null-route filtered when falling back from IP-in-IP tunnel to inet.0/inet6.0. [PR1508631](#)
- The rpd might crash after reboot when MSDP is configured. [PR1536593](#)

System Management

- Differences in XML tags for **show system virtual-memory**. [PR1438110](#)
- Differences in XML tags for **show system processes**. [PR1438129](#)
- QFX5220 does not support the following matches, which were supported on QFX5200: **first-fragment**, **tcp-established**, **ls-fragment**, and **hop-limit**. [PR1499009](#)
- The interface might not be added to BD after VLAN change. [PR1504374](#)
- 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](#)
- The **aggregate member links** field in the **show interfaces extensive** command output for aggregated Ethernet interfaces does not reflect the number of member-links. [PR1517841](#)
- On the QFX5130-32CD platform, when the management port speed is configured to 100 Mbps, the port's status LED does not light up. However, the port forwards traffic normally. [PR1521510](#)
- On QFX5130-32CD platforms with a large amount of ARP resolutions happening on IRBs in a very short time, the ARPD process usage can shoot to 100%. This issue does not happen with Layer 3 interfaces. [PR1523300](#)
- Ingress policer scale is limited to 128 due to a known issue in the Junos OS Evolved Release 20.3R1. [PR1525525](#)
- On the QFX5130-32CD platform, the Encapsulated Remote Switch Port ANalyzer (ERSPAN) status continues to show status as "Up" even when the destination IP address is unreachable. [PR1527505](#)

- 160,000 MAC flushing with traffic running takes close to 35 minutes after **clear ethernet-switching table**, without traffic flushed in 3-4 minutes. [PR1528511](#)
- On the QFX5130-32CD platform, the CLI option to configure a range of VLANs as the input for the analyzer is not supported in this release. [PR1529419](#)
- Traffic loss might happen after performing GRES in the Junos OS Evolved enabled chassis-based systems. [PR1532446](#)
- The cosd might crash when multiple configuration changes are made in a single commit. [PR1536320](#)
- The evo-pfemamd process might keep crashing on the QFX5220-32CD or QFX5220-128C platforms. [PR1536588](#)
- The interfaces might not come up in some instances after a power cycle or a soft reboot. [PR1538284](#)
- With port mirror, evo-pfemamd might restart unexpectedly on QFX Series platforms running Junos OS Evolved. [PR1538626](#)
- QFX5130-32CD: DHCPv6 packets are going to unknown multicast queue instead of DHCP queue, and DHCPV6 relay is not working. [PR1545754](#)
- When the system has a scaled number of ARP entries learned over IRB and the underlying logical interface of IRB is flapped, traffic drop is expected to happen for sometime for the ARP destinations and arp is expected to exhibit 100% CPU for some time. [PR1554151](#)
- The device might not boot up after performing the cleanup operation. [PR1557020](#)
- EVO:: JDI_FT_REGRESSIONS:: QFX5130:: Tunnel:: GRE:: tunneld core hit at unneld core was hit at IflibCommonData. [PR1563399](#)
- "ICMP destination unreachable" message is not sent from the QFX5130 when a firewall filter action rejects the packet. [PR1563404](#)

User Interface and Configuration

- The EVO operational state would be incorrect on the system and CoS scheduler configuration change might not take effect. [PR1536615](#)

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the

multiple product-driven licensing and packaging approaches that Juniper Networks has developed over the past several years.

The major components of the framework are:

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

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

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore 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 open defects and issues resolved.

<https://prsearch.juniper.net/InfoCenter/index?page=prsearch>

- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.

<https://apps.juniper.net/hct/home>

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


- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#).

<https://pathfinder.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

IN THIS SECTION

- [Self-Help Online Tools and Resources | 74](#)
- [Creating a Service Request with JTAC | 74](#)

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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/>.

Upgrade Your Junos OS Evolved Software

Products impacted: PTX10001-36MR, PTX10003, PTX10004, 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 OS Evolved from the OS drop-down list.
4. Select the relevant release number from the Version drop-down list.
5. In the **Install Package** section, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.

NOTE: Download the Services Profile 1 image to use the lean rpd profile. For more information about the types of Junos OS installation package prefixes, see [Installation Packages Prefixes](#).

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

Revision History

22 April 2021—Revision 4, Junos OS Release 21.1R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

16 April 2021—Revision 3, Junos OS Release 21.1R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

12 April 2021—Revision 2, Junos OS Release 21.1R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

26 March 2021—Revision 1, Junos OS Release 21.1R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the 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. Copyright © 2021 Juniper Networks, Inc. All rights reserved.