Release Notes

Published
2023-08-10

Junos OS Evolved Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices

SOFTWARE HIGHLIGHTS

- Support for Layer 2 VPN (ACX7100-32C)
- Support for Layer 3 VPN (ACX7100-32C)
- Support for Layer 2 VPN (ACX7100-48L)
- Support for Layer 3 VPN (ACX7100-48L)
- Support for performance monitoring and TCA (PTX10001-36MR)
- Support for Data Center Bridging Capability Exchange protocol (DCBX) (PTX10004, PTX10008, PTX10001-36MR)
- Support for per PFE-Restart (PTX10008)
- Support for Non-Stop Routing (NSR) (PTX10004, PTX10008, PTX10001-36MR)
- Support for Routing Engine Switchover (PTX10004, PTX10008)
- Support for proxy ARP on IRB (PTX10004, PTX10008, PTX10001-36MR)
- Support for IEEE 802.3ah Link Fault Management (LFM) (PTX10004, PTX10008, PTX10001-36MR)
- Support for multicast-only fast reroute (MoFRR) (PTX10004, PTX10008, PTX10001-36MR)
- Support for Public Key Infrastructure (PKI) (PTX10004, PTX10008, PTX10001-36MR)
- Support for BFD (PTX10004, PTX10008, PTX10001-36MR)
- Support for class-based firewall filters (PTX Series)
- Support for RPKI Sharding (PTX10003, PTX10004, PTX10008, PTX10001-36MR)

- Support for Distributed Constrained Shortest Path First (CSPF) for Segment Routing Traffic Engineering (SR-TE) paths (PTX10004, PTX10008)
- Support for CoS features on EVPN VXLAN (QFX5130-32CD)
- Support for EVPN-VXLAN support with MAC VRF routing instances (QFX5130-32CD)
- Support for EVPN Type 5 routes (QFX5130-32CD)
- Support for port mirroring in EVPN-VXLAN fabrics (QFX5130-32CD)
- Support for firewall filtering and policing (QFX5130-32CD)
- Support for sFlow technology with VXLAN (QFX5130-32CD)

HARDWARE HIGHLIGHTS

- New ACX7100-32C router
- New QFX5700 switch

Table of Contents

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 21.2R1.

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 ACX7100-32C and ACX7100-48L Devices

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 21.2R1 for ACX7100-32C and ACX7100-48L routers. They describe new features, limitations, and known problems in the hardware and software.

What's New

IN THIS SECTION

Learn about new features introduced in this release for the ACX Series Routers.

What's New in 21.2R1

IN THIS SECTION

To view features supported on the ACX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 21.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- ACX7100-32C
- ACX7100-48L

The following sections highlight the key features in this release.

Hardware

- **New ACX7100-32C Router (ACX Series)—**In Junos OS Evolved Release 21.2R1, we introduce the ACX7100-32C, a fixed-configuration and high-performance router with deep buffer capabilities. With high port density and MACsec-ready ports, the ACX7100-32C routers are ideal for secure and high-end service aggregation, large enterprises, and top-of-rack, spine-and-leaf data center applications.

  The ACX7100-32C provides several capabilities that include a comprehensive set of features, advanced programmability, and seamless integration with emerging distributed edge compute architectures and spine-and-leaf data center applications.

  With a 1-U, energy-efficient, and fixed-configuration design, the ACX7100-32C has a 4.8-Tbps-forwarding-capacity ASIC and the following port configurations:

- Thirty-two QSFP28 ports (ports 0 through 31) that operate at 100-Gbps or 40-Gbps speeds.
- Four double-density QSFP56-DD ports (ports 32 through 35) that operate at a default speed of 400-Gbps.

We ship the ACX7100-32C routers with redundant power supply modules (PSMs). You can order router models with either AC or DC PSMs. The router supports only front-to-back airflow.

To install the ACX7100-32C router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the ACX7100-32C Router Hardware Guide.

Table 1 on page 3 summarizes the ACX7100-32C features supported in Junos OS Evolved Release 21.2R1.

Table 1: Feature Support on the ACX7100-32C

Feature	Description
Chassis	• Support for a built-in Routing Engine. The ACX7100-32C doesn't support pluggable or redundant Routing Engines. As a result, you cannot use GRES on this router. [See show chassis hardware.]
	• Support for environmental monitoring and field-replaceable unit (FRU) management. [See show chassis hardware.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

	• Interface support. The ACX7100-32C has 32 built-in QSFP28 ports that can operate at a default speed of 100 Gbps and 4 built-in QSFP56-DD ports that can operate at a default speed of 400 Gbps. You can channelize the QSFP28 ports into:
	  • Four 25GbE interfaces
	  • Four 10GbE interfaces
	  • Two 50GbE interfaces
	You can channelize the QSFP56-DD ports into:
	  • Four 100GbE interfaces
	  • Two 100GbE interfaces
	  • Eight 50GbE interfaces
	  • Two 50GbE interfaces
	  • Eight 25GbE interfaces
	  • Four 10GbE interfaces
	See [Port Speed on ACX7100-32C Router Overview.]

Class of service (CoS)	• Support for classification and rewrite rules of all types (Inet-Prec, DSCP, DSCP-v6, IEEE-802.1p, and IEEE-802.1ad) at the logical interface level. [See Classifiers and Rewrite Rules at the Global, Physical, and Logical Interface Levels Overview.]
	• Support for port shaping and scheduling with eight virtual output queues (VoQs) per port and two scheduling priority levels (strict-high and low). This release also supports for multiple strict-high priority queues (round-robin scheduling), multiple low-priority queues (weighted fair queuing, or WFQ, scheduling), and default deep buffers. [See Schedulers Overview for ACX Series Routers and Shared and Dedicated Buffer Memory Pools on ACX Series Routers.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description
DHCP	• Support for DHCP server and DHCP relay configuration for IPv4 and IPv6 services. [See DHCP Overview.]

Ethernet OAM	• Support for Operations, Administration, and Maintenance (OAM). You can configure connectivity fault management (CFM), BFD, the ITU-T Y.1731 standard for Ethernet service OAM, and the following link-fault management (LFM) features:
	  • Discovery
	  • Link monitoring
	  • Remote fault detection
	  [See ITU-T Y.1731 Ethernet Service OAM Overview.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description
Firewall filters	• Support for firewall filters and policers. You can configure firewall filters with packet match conditions for the `bridge domain`, `IPv4`, and `IPv6` families. This release introduces support for the following match conditions in the egress direction for IPv4 and IPV6 traffic: For IPv4: • ttl • tcp-flags For IPv6: • destination-ip • next-header • traffic-class • hop-limit • L4 ports • icmp-type/code • tcp-flags The ACX71000-32C also supports firewall filter actions that include `count`, `discard`, `log`, `syslog`, and `policer`. [See Overview of Firewall Filter Match Conditions and Actions on ACX Series Routers.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

High availability	• Support for the following BFD features:
	  • BFD for IPv4 and IPv6 routes
	  • Single-hop BFD in inline mode with an interval range of 10 milliseconds to 1 second
	  • Single-hop BFD in distributed mode with an interval of 1 second or more
	  • Single-hop BFD in centralized mode with a minimum interval of 1 second to detect IRB failures
	  • Multihop BFD with an interval of 1 second or more
	  • Micro-BFD for LAG in centralized or distributed mode with an interval of 1 second or more
	[See Bidirectional Forwarding Detection (BFD).]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

Layer 2 features	• Support for the following advanced Layer2 features:
	  • Unqualified bridging support
	  • Bridge domain without a `vlan-id` *number* statement
	  • Bridge domain with the `vlan-id` value set to none
	  • Bridge domain with a single VLAN ID
	  • Single-learning domain
	  • MAC limiting
	  • Ethernet service types:
	  • E-Line with AC interface types port, VLAN, Q-in-Q, VLAN list, and VLAN maps
	  • E-Line
	  • E-LAN
	  • E-Access
	  • E-Transit
	  • LLDP
	  • LACP
	  • IRB interface
	  • Link aggregation group (LAG) support with the following hashing algorithms:
	  • For family `multiservice`, destination and source MAC addresses
	  • For family `inet`, Layer 3 and Layer 4
	  • For family `inet6`, Layer 3 destination and source addresses

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

	• For family `inet6`, Layer 4 destination and source ports
	• Encapsulation types:
	  • `extended-vlan-bridge`
	  • `vlan-bridge`
	• Q-in-Q tunneling
	[See Understanding Layer 2 Bridge Domains and Q-in-Q Tunneling on ACX Series.]

Layer 2 VPN	• Support for VPLS. The ACX7100-32C routers support a single VLAN for each `virtual switch` routing instance type. Junos OS Evolved does not support the `family vpls` option. To configure VPLS on the ACX7100-32C routers, configure the `instance-type virtual switch` statement at the [`edit routing-instances` *`routing instance`*] hierarchy level. [See Introduction to Configuring VPLS.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

Layer 3 VPN	• Support for the following Layer 3 VPN features:
	  • IP-VPN services:
	  • Virtual routing and forwarding (VRF) and virtual-router instance type
	  • All control plane configuration options
	  • Per-prefix and per-table label signaling
	  • Layer 3 VPN support with ECMP
	  • BGP policies support for different Layer 3 VPN use cases such as full mesh VPN, hub-spoke VPN, management VPN, and leaking routes
	  • Layer 3 VPN with `vrf-table-label` mode
	  • Layer 3 VPN with `chained-composite-next-hop` mode
	  • Layer 3 VPN ping using `ping mpls l3vpn prefix prefix-name` *`l3vpn-name`* command.
	  NOTE: The ping command works only with the `vrf-table-label` configuration.
	  • Import and export of routes across non-default to non-default virtual routing and forwarding (VRF)
	  NOTE: Table next hop is not supported.
	  • Inter-autonomous system (Inter-AS) options A, B, and C.
	  NOTE: You can deploy inter-AS option B in a hierarchical network design within a single IGP autonomous system.
	  • Provider edge-to-customer edge (PE-to-CE) route using static route and routing protocols such as eBGP, IS-IS, OSPF, and RIP
	  • 6PE and 6VPE with PE-CE routing-static and PE-to-CE BGPv6
	Currently, we do not support virtual tunnel (VT) interface-based Layer 3 VPN. [See Layer 3 VPNs User Guide for Routing Devices.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

Layer 3 features	• Support for the following Layer 3 features:
	  • Longest prefix match
	  • Exception packets handling
	  • VLAN tagging modes
	  • Neighbor solicitation
	  • Unicast RPF
	  • Interface-based routing
	  The ACX7100-32C router also supports interior gateway protocols such as OSPF, IS-IS, RIP, and ECMP. [See Configure ICMP Features, Enabling VLAN Tagging, Neighbor solicitation, Understanding Unicast RPF (Routers), OSPF Overview, IS-IS Overview, RIP User Guide, and BGP Overview.]

	• Support for BGP for IPv4 and IPv6. [See BGP Overview.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

MPLS	• Support for the following MPLS features:
	  • IP/MPLS infrastructure feature set for the Layer 3 VPN service
	  • Basic BGP control plane features such as LDP-DOD, CSPF, single-area CSPF
	  • MPLS label stack
	  • MPLS protections:
	  • Fast reroute (FRR)/ Make-before-break (MBB)
	  • Link protection
	  • Node protection
	  • Label-switching router (LSR)
	  • Shared Risk Link Group (SRLG) for MPLS
	  • RSVP label-switched path (LSP) over IPv4 include refresh reduction
	  • LDP LSP over IPv4
	  • RSVP 1:1
	  • RSVP-Traffic Engineering (RSVP- TE)
	  • LDP over RSVP
	  • Inter-autonomous Systems LSP intra-area LSP
	[See MPLS Applications User Guide.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

Multicast	• Support for IPv4 multicast for Layer 3. You can configure IGMP snooping with IGMPv2 and IGMPv3, which includes support for the following:
	  • Auto-rendezvous point (auto-RP)
	  • Anycast RP
	  • IGMP filter
	  • IGMP querier
	  • Protocol Independent Multicast source-specific multicast (PIM SSM)
	  • PIM sparse mode (PIM SM)
	[See IGMP Snooping Overview.]
	NOTE: In this Junos OS Evolved release, the ACX7100-32C doesn't support IPv6 multicast and Layer3 multicast protocols (such as IGMP, MLD, or PIM) over IPv4 and IPv6 IRB interfaces.

	• Support for Layer 2 multicast-related features, including support for IGMP and Multicast Listener Discovery (MLD) snooping. You can configure IGMP snooping with IGMPv1, IGMPv2, and IGMPv3, which includes support for the following:
	  • IGMP snooping in a bridge domain
	  • IGMP snooping with IRB interface configured in a bridge domain
	  • MLD snooping in a bridge domain
	  • MLD snooping with IRB interface configured in a bridge domain
	[See IGMP Snooping Overview.]

Network management and monitoring	• Support for NETCONF event notifications. NETCONF clients can subscribe to NETCONF event notifications to receive alerts for events that might impact device operations or management activities. The events include netconf-config-change, netconf-session-start, and netconf-session-end events. [See NETCONF Event Notifications.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description
Resiliency	• Support for platform resiliency to handle failures and faults related to components such as CPU, fan trays, temperature sensors, power supply units, FPGA, and optics. Fault handling includes detecting and logging the error, raising alarms, sending SNMP traps, providing indication about the error through LEDs, self-healing, and taking components out of service. [See show system errors active.]
Routing policy	• Unicast reverse path forwarding (unicast RPF) support for IPv4 and IPv6. You can reduce the impact of denial-of-service (DoS) attacks for IPv4 and IPv6 interfaces by configuring unicast RPF. You can use unicast RPF to determine the source of attacks and reject packets from unexpected source addresses on interfaces. However, we do not support unicast RPF checking for: • Transit packets exiting a tunnel source interface • Asymmetrical routing [See Understanding Unicast RPF (Routers).]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description
Services applications	• Support for RFC 5357, Two-Way Active Measurement Protocol (TWAMP) monitoring service. You can configure the TWAMP monitoring service, which sends out probes to measure network performance. You often use TWAMP to check compliance with service-level agreements. In Junos OS Evolved, you configure TWAMP at the `[edit services monitoring twamp]` hierarchy level. The support for this service is limited to the following: • IPv4 traffic only for control sessions and test sessions • Probe statistics and history • Control and test session status • Test session probe generation and reception, as well as reflection • Timestamps set by the Routing Engine or the Packet Forwarding Engine • Error reporting through system log messages only • Unauthenticated mode only [See Understanding Two-Way Active Measurement Protocol on Routers.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description

	• Support for RFC 2544-based benchmarking tests. We support only the Layer 3 reflector function for these tests. The support is further restricted to:
	  • `family inet` only
	  • IPv4 source and destination addresses
	  • Software-based reflection, with a maximum frame rate for the traffic flows of 1000 kbps
	You use the RFC 2544 benchmarking tests to measure and demonstrate the service-level agreement (SLA) parameters before service activation. The tests measure throughput, latency, frame loss rate, and number of back-to-back frames.
	You can configure these tests at the `[edit services monitoring rfc2544]` hierarchy level. [See RFC 2544-Based Benchmarking Tests Overview.]

Software installation and upgrade	• 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.]
	• Support for secure-boot implementation based on the UEFI 2.4 standard. [See Software Installation and Upgrade Guide.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description
Timing and synchronization	• Support for enhanced Ethernet equipment clock (eEEC). Enhanced EEC enables new clocks to operate with different quality levels defined in the Synchronous Ethernet chain. To enable enhanced EEC on your router, configure the `enable-extended-ql-tlv` statement at the `[edit chassis synchronization]` hierarchy level. The ACX7100-32C supports the following new clock quality levels for enhanced EEC: • Enhanced primary reference time clock (ePRTC) • Primary reference time clock (PRTC) • Enhanced primary reference clock (ePRC) • Enhanced Ethernet equipment clock (eEEC) [See enable-extended-ql-tlv, Ethernet Synchronization Message Channel Overview, and synchronization (ACX Series).]
	• Support for frequency synchronization using Synchronous Ethernet. [See Synchronous Ethernet Overview.]

Table 1: Feature Support on the ACX7100-32C *(Continued)*

Feature	Description
	• Support for Synchronous Ethernet with Precision Time Protocol (PTP) Transparent Clock. The transparent clock measures the residence time of PTP packets as the packets pass through the router. The network load and device architecture might result in queuing or buffering delays. These delays are the main source of packet delay variation in the router. Transparent clock adds the residence time into the correction field of the PTP packet. The client or boundary clocks can determine this resident time while they receive the PTP packet from the upstream transparent clock router. The client clock can estimate and remove these delays from the offset computation and reduce the packet's jitter effects. In syntonized transparent clock, the transparent clock requires physical layer frequency based on the ITU-T G.8262/.1 standard. Synchronous Ethernet configuration is mandatory for enabling syntonized transparent clock. Use the `show protocols ptp` and `show ptp global-information` commands to verify the PTP transparent clock configuration status. To enable PTP transparent clock on your router, configure `syntonized-e2e-transparent` statement at the `[edit protocols ptp]` hierarchy level. See [Understanding Transparent Clocks in Precision Time Protocol and show ptp global-information.]

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported on ACX7100-32C, see the Hardware Compatibility Tool.

- Table 2 on page 18 summarizes the features added to the ACX7100-48L in Junos OS Evolved Release 21.2R1.

Table 2: Feature Support Added to the ACX7100-48L

Feature	Description
DHCP	• Support for DHCP server and DHCP relay configuration for IPv4 and IPv6 services. [See DHCP Overview.]

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description

Ethernet OAM	• Support for Operations, Administration, and Maintenance (OAM). You can configure connectivity fault management (CFM), and the ITU-T Y.1731 standard for Ethernet service OAM. The ACX7100-48L also supports the following link-fault management (LFM) features:
	  • Discovery
	  • Link monitoring
	  • Remote fault detection
	  [See ITU-T Y.1731 Ethernet Service OAM Overview.]

High availability	• Support for the following BFD features:
	  • BFD for IPv4 and IPv6 routes
	  • Single-hop BFD in inline mode with an interval range of 10 milliseconds to 1 second
	  • Single-hop BFD in distributed mode with an interval of 1 second or more
	  • Single-hop BFD in centralized mode with a minimum interval of 1 second to detect IRB failures
	  • Multihop BFD with an interval of 1 second or more
	  • Micro-BFD for LAG in centralized or distributed mode with an interval of 1 second or more
	  [See Bidirectional Forwarding Detection (BFD).]

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description
Layer 2 VPN	• Support for VPLS. The ACX7100-48L routers support a single VLAN for each `virtual switch` routing instance type. Junos OS Evolved does not support the `family vpls` option. To configure VPLS on the ACX7100-48L routers, configure the `instance-type virtual switch` statement at the [`edit routing-instances` *`routing instance`*] hierarchy level. If you configure normalized VLANs, either by not configuring VLAN IDs or by including the `vlan-id none` statement, then you must also include the `service-type single` statement at the [`edit routing-instances` *`routing-instance`* `protocol vpls`] hierarchy level. [See Introduction to Configuring VPLS.]

Layer 3 VPN	• Support for the following Layer 3 VPN features:
	  • IP-VPN services:
	  • Virtual routing and forwarding (VRF) and virtual-router instance type
	  • All control plane configuration options
	  • Per-prefix and per-table label signaling
	  • Layer 3 VPN support with ECMP
	  • BGP policies support for different Layer 3 VPN use cases such as full mesh VPN, hub-spoke VPN, management VPN, and leaking routes
	  • Layer 3 VPN with `vrf-table-label` mode
	  • Layer 3 VPN with `chained-composite-next-hop` mode
	  • Layer 3 VPN ping using `ping mpls l3vpn prefix prefix-name` *`l3vpn-name`* command.
	  NOTE: The ping command works only with the `vrf-table-label` configuration.
	  • 6PE and 6VPE with PE-CE routing-static and PE-to-CE BGPv6

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description
	• Import and export of routes across non-default to non-default virtual routing and forwarding (VRF) **NOTE**: Table next hop is not supported. • Inter-autonomous system (Inter-AS) options A, B, and C. **NOTE**: You can deploy inter-AS option B in a hierarchical network design within a single IGP autonomous system. • Provider edge-to-customer edge (PE-to-CE) route using static route and routing protocols such as eBGP, IS-IS, OSPF, and RIP Currently, we do not support virtual tunnel (VT) interface-based Layer 3 VPN. [See Layer 3 VPNs User Guide for Routing Devices.]

Layer 3 features	• Support for the following Layer 3 features:
	  • Longest prefix match (LPM)
	  • Exception packet handling
	  • VLAN tagging modes
	  • Neighbor solicitation
	  • Unicast reverse path forwarding (unicast RPF)
	The ACX7100-48L also supports interior gateway protocols (IGPs) such as OSPF, IS-IS, RIP, and ECMP.

Network management and monitoring	• Support for NETCONF event notifications. NETCONF clients can subscribe to NETCONF event notifications to receive alerts for events that might impact device operations or management activities. The events include netconf-config-change, netconf-session-start, and netconf-session-end events. [See NETCONF Event Notifications.]

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description

MPLS	• Support for the following MPLS features:
	  • IP/MPLS infrastructure feature set for the Layer 3 VPN service
	  • Basic BGP control plane features such as LDP-DOD, CSPF, single-area CSPF
	  • MPLS label stack
	  • MPLS protections:
	  • Fast reroute (FRR)/ Make-before-break (MBB)
	  • Link protection
	  • Node protection
	  • Label-switching router (LSR)
	  • Shared Risk Link Group (SRLG) for MPLS
	  • RSVP label-switched path (LSP) over IPv4 include refresh reduction
	  • LDP LSP over IPv4
	  • RSVP 1:1

	• RSVP-Traffic Engineering (RSVP- TE) • LDP over RSVP • Inter-autonomous Systems LSP intra-area LSP [See MPLS Overview.]

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description

Multicast	• Support for IPv4 multicast for Layer 3. You can configure IGMP snooping with IGMPv2 and IGMPv3, which includes support for the following:
	  • Auto-rendezvous point (auto-RP)
	  • Anycast RP
	  • IGMP filter
	  • IGMP querier
	  • Protocol Independent Multicast source-specific multicast (PIM SSM)
	  • PIM sparse mode (PIM SM)
	  [See IGMP Snooping Overview.]
	  NOTE: In this Junos OS Evolved release, the ACX7100-48L doesn't support IPv6 multicast and Layer3 multicast protocols (such as IGMP, MLD, or PIM) over IPv4 and IPv6 IRB interfaces.

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description
Services applications	• Support for RFC 5357, Two-Way Active Measurement Protocol (TWAMP) monitoring service. You can configure the TWAMP monitoring service, which sends out probes to measure network performance. You often use TWAMP to check compliance with service-level agreements. In Junos OS Evolved, you configure TWAMP at the `[edit services monitoring twamp]` hierarchy level. The support for this service is limited to the following: • IPv4 traffic only for control sessions and test sessions • Probe statistics and history • Control and test session status • Test session probe generation and reception, as well as reflection • Timestamps set by the Routing Engine or the Packet Forwarding Engine • Error reporting through system log messages only • Unauthenticated mode only [See Understanding Two-Way Active Measurement Protocol on Routers.]

Table 2: Feature Support Added to the ACX7100-48L *(Continued)*

Feature	Description
Timing and synchronization	• Support for enhanced Ethernet equipment clock (eEEC). Enhanced EEC enables new clocks to operate with different quality levels defined in the Synchronous Ethernet chain. To enable enhanced EEC on your router, configure the `enable-extended-ql-tlv` statement at the `[edit chassis synchronization]` hierarchy level. The ACX7100-48L supports the following new clock quality levels for enhanced EEC: • Enhanced primary reference time clock (ePRTC) • Primary reference time clock (PRTC) • Enhanced primary reference clock (ePRC) • Enhanced Ethernet equipment clock (eEEC) [See enable-extended-ql-tlv, Ethernet Synchronization Message Channel Overview, and synchronization (ACX Series).]
	• Support for frequency synchronization using the Synchronous Ethernet protocol in accordance with the ITU-T G.8262 and G.8262.1 standards. [See Synchronous Ethernet Overview.]

What's Changed

IN THIS SECTION

Learn about what changed in this release for ACX Series Routers.

What's Changed in Release 21.2R1

IN THIS SECTION

Authentication and Access Control

- **Command to automate SSH key-based authentication (ACX Series, PTX Series, and QFX Series)**—You can set up SSH-key based authentication between the network device and a remote host by issuing the `request security ssh password-less-authentication` operational mode command. When you execute the command with the appropriate options, the device generates SSH keys for the current user, provided the user does not already have existing keys, and transfers the user's public key to the **authorized_keys** file of the specified user on the remote host.

  [See request security ssh password-less-authentication.]

General Routing

- **Enhancement to the show chassis pic command (Junos OS Evolved)**— You can now view additional information about the optics when you run the `show chassis pic` command. The output now displays the following additional field:

  MSA Version: Multi-source Agreements (MSA) version that the specified optics is compliant to. Values supported are: SFP+/SFP28 — SFF-8472 (versions 9.3 - 12.3), QSFP+/QSFP28 — SFF 8363 (versions 1.3 - 2.10), and QSFP-DD — CMIS 3.0, 4.0, 5.0.

  Previously, the `show chassis pic` command did not display this additional field.

  [See show chassis pic.]

- **Enhancement to the default remnant-holdtime (Junos OS Evolved platforms: ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**— Starting this release, the default `remnant-holdtime` has been increased from 180 seconds to 300 seconds. This provides sufficient time for protocols to start and sync routes from neighbors in a scaled

environment, during rpd restart. You can configure `remnant-holdtime` at the `edit routing-options forwarding-table` hierarchy level.

[See forwarding-table.]

- **Mozilla certification authority (CA) certificates removed (ACX Series, PTX Series, and QFX Series)**—To minimize security risks, Junos OS Evolved no longer includes Mozilla's set of root certificates from various CA operators by default. To use Docker container images from a registry that requires TLS authentication, you must first save the image as a tar archive on a remote device and then import the contents of the archive on the device running Junos OS Evolved.

  [See Running Third-Party Applications in Containers..]

Junos XML API and Scripting

- **Changes to how command-line arguments are passed to Python op scripts (ACX Series, PTX Series, and QFX Series)**—When the device passes command-line arguments to a Python op script, it prefixes a hyphen (-) to single-character argument names, and it prefixes two hyphens (--) to multi-character argument names. The prefix enables you to use standard command-line parsing libraries to handle the arguments. In earlier releases, the device prefixes a single hyphen (-) to all argument names.

  [See Declaring and Using Command-Line Arguments in Op Scripts.]

- **The `language python` statement is enabled by default (ACX Series, PTX Series, and QFX Series)**—The `language python` statement is configured by default in the `junos-defaults` configuration group on devices running Junos OS Evolved. Thus, you can execute unsigned Python scripts using the default Python version without explicitly configuring the statement on the device.

  [See Requirements for Executing Python Automation Scripts on Devices Running Junos OS.]

Network Management and Monitoring

- **Changes to how command-line arguments are passed to Python action scripts (ACX Series, PTX Series, and QFX Series)**—When a custom YANG RPC invokes a Python action script and passes command-line arguments to the script, the device prefixes a hyphen (-) to single-character argument names, and it prefixes two hyphens (--) to multi-character argument names. The prefix enables you to use standard command-line parsing libraries to handle the arguments. In earlier releases, the device passes the unmodified argument names to the script.

  [See Creating Action Scripts for YANG RPCs on Devices Running Junos OS and Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules.]

- **Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, PTX Series, and QFX Series)**—When you configure the `rfc-compliant` statement at the `[edit system services netconf]`

hierarchy level, the NETCONF server's response for <commit> operations includes the following changes:

- If a successful <commit> operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
- The NETCONF server response emits the <source-daemon> element as a child of the <error-info> element instead of the <rpc-error> element.
- If you also configure the flatten-commit-results statement at the [edit system services netconf] hierarchy level, the NETCONF server suppresses any <commit-results> XML subtree in the response and only emits an <ok/> or <rpc-error> element.

[See Configuring RFC-Compliant NETCONF Sessions.]

- **Changes in contextEngineID for SNMPv3 INFORMS (ACX Series, PTX Series, and QFX Series)**— Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos OS devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  [See SNMP MIBs and Traps Supported by Junos OS.]

Platform and Infrastructure

- **The jcs:invoke() function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX 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).]

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Release 21.2R1 for ACX7 Series routers.

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

General Routing

- The internal switch does not support hot swap. Removing internal switch PCI device or power down the main board does not trigger any signal to evo-pfemand. The user space application cannot detect this fault. PR1554516
- Egress IP MTU exception and fragmentation is not supported. Outgoing IP packets bigger than the configured interface MTU does not get fragmented. PR1558327
- For ACX7100-48L, on configuring or unconfiguring PTP TC, interface needs reinitialization and therefore it is expected to see interfaces go down and come up. PR1558603
- ACX7100-48L :: Multicast : IPv6 multicast traffic loss on downstream receiver along with inconsistent forwarding S,G route entry on the ACX7100-48L device. The following limitation is applicable only for IPv6 multicast.

  The hardware limitation for IPv6 multicast is, the route entry (In-Intf, S1, G1) programmed as (In-Intf, *, G1).

  - It means when multiple sources for the same group are in the same incoming interface and the OIF sets are different, there could be potential traffic loss.
  - When multiple sources for the same group are in the same incoming interface and the OIF sets are same, then there is no potential traffic loss expected.
  - Where can we place the ACX7100-48L device in the network only for IPv6 multicast.
  - ACX7100-48L acts as RP or FHR:
    - When acting as RP, there is a possibility that RPT prune and SPT join happens, which could result in above limitation mentioned. So ACX7100-48L cannot be placed as RP or FHR.
  - ACX7100-48L acts as LHR:
    - The OIF set is not different for different (S1, G1) entries, hence overwriting the same (In-Intf, *, G1) in hardware can result in no traffic loss.

      When moving from RPT tree to SPT tree, the (S1,G1) entry moves to a different IIF interface, hence we program different (In-Intf, *, G1) entry in hardware and there is no expectation of traffic loss.

PR1564654

- ACX7100-48L and ACX7100-32C does not support micro bfd with `no-dedicated-mac-for-micro-bfd`. When packets received with no-dedicated mac, ACX7100-48L and ACX7100-32C does not accept and punt to the host path, micro bfd session does not come up. So configuring the configuration statement `no-dedicated-mac-for-micro-bfd` on the peer router if the device is ACX7100-48L and ACX7100-32C is not supported. PR1566672
- The firewall filter can only be implemented by adding the micro bfd dedicated MAC to VRRP TCAM table. But VRRP TCAM is limited in size only 11 entries in balanced profile), reserving one entry for storing the micro bfd reduces the VRRP TCAM size by 1 and might impact the scale. PR1569841
- When a high scale of routes are dependent on a ECMP group or unilist, few streams might observe 8-9 seconds traffic drop during ECMP member link flap. PR1573295

Open Issues

IN THIS SECTION

Learn about open issues in this release for the ACX7100-32C and ACX7100-48L.

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

General Routing

- Monitor interface traffic statistics is not flushed out after interface goes down. This is a common behavior across Junos OS Evolved platforms. Even on re-enabling it is still not flushed.

  PR1564672

- Errors related to pfe_portE, loopback on system reboot. **[Interface: can't get pfe_portE for ifd:ae0, Interface: PIL set loopback failed for IFD:ae0 error:-1]** seen. There is no functional impact. PR1565831

- If you configure unsupported actions in egress firewall filter, commit errors are thrown. If the unsupported actions are removed from the configuration and committed, the filter might not work in the Packet Forwarding Engine even then. This behavior depends on the order the BIND and MODIFY of filter comes to the Packet Forwarding Engine from firewalld. Syslog errors and pfe-cli commands can be used to check the filter install status. PR1568525

- PFC feature is not supported on ACX platforms. Therefore `show interfaces` *`interface-name`* `extensive` command always displays the **MAC Priority Flow Control Statistics table** with 0 value. PR1576317

- Each routing instance has an associated {vrf, jtdv} device pair, and these devices are managed as a service by systemd. As such, whenever a routing instance is created or deleted there is work that systemd needs to do, and when using scaled numbers of routing instances the

  The amount of CPU consumed by systemd can be high whilst these devices are being created or deleted.

  The time to take down or boot the system is longer. PR1581283

- If evo-pfemand process is restarted and at the same time if ethernet-connectivity-fault-management has also restarted, there is a possibility that a ppmd core occurs. As evo-pfemand is restarted, the services are also interrupted. Once the process is back online, all the CCM sessions recover post the mentioned triggers. PR1585178

- In scaled cfm config scenarios, ppmd or evo-pfemand process core file might be hit after a router reboot. PR/1591534

- After a router reboot, in scaled cfm config scenario, the sessions (some or all) might not come up. Deactivating and reactivating cfm config must resolve the issue. PR/1591978

- On ACX7100 evo-pfemand core file can be seen rarely with link flaps when there is VPLS configuration with high number of MAC entries learnt (> 10,000). PR/1596587

- Few LSI MACs are not properly learnt in the software with 8000 vpls instance scale: The issue is seen above 4000 VPLS instances, hardware LSI MAC learning is fine. However, we observe some difference in LSI MAC learned in the software. PR1597125

- No MAC address present in the Ethernet table or Junos OS Evolved Packet Forwarding Engine, however arp is present: The issue is seen only while upgrading or downgrading an image on the router. The issue is only applicable to aggregated Ethernet interfaces, physical interfaces are fine. The hardware learns the MACs properly, only from the software point of view MACs are missing. PR1597277

- In scaled scenarios (4k BDs, IRBs), with restart of l2ald and pfe daemons, arpd and ndp daemon crash is observed and it recovers by itself and no functionality impact expected. PR1598217

- Egress firewall ACL rules with count or policer actions are skipped for BUM traffic and does not hit. This is because of ASIC limitation for BUM traffic and applies to all services like Layer 2, Layer 3, EVPN etc. PR1598489
- In ACX7100-32C, it causes the traffic to drop while sending continuous traffic against the ACX7100-32C 400G ports. PR1601151
- For ACX7100-32C and ACX7100-48L, the voltage thresholds are not correct which might cause the **Voltage Threshold Crossed** alarm sometimes, however it does not impact the functionality. PR1601493
- On ACX7100-32C or ACX7100-48L, `show system firmware` some times might show Current Firmware version for FPC 0 as blank. This is only show CLI issue due to random read failure issue of FPC 0's firmware version and this does not impact any normal operations of FPC 0. PR1618949
- When IPv6 NDP is configured on IRB interface with one of the l2-interfaces from BD, it programs the NDP correctly. However, when the Layer 2 interface is changed for that NDP with other Layer 2 interface from same BD, the configuration does not taking effect. NDP still points to the old Layer 2 interface. PR1602894

Resolved Issues

IN THIS SECTION

Learn about the issues fixed in this release for ACX Series Routers.

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: 21.2R1-S1

IN THIS SECTION

General Routing

- arpd and ndp daemons crash in scale setups. PR1598217

Resolved Issues: 21.2R1

IN THIS SECTION

General Routing

- ACX7100-32C :: **Invalid RPC request key: 0x000a0000** error messages seen in**/var/log/messages**. PR1552436
- Incorrect **Active defect** alarm generated when local or remote fault is inserted. PR1554852
- The device NMI watchdog starts after USB scratch install completes and we wait for the user action to reboot, resulting in a system exception. PR1555142
- ACX7100-32C : Interface : 400G DAC link does not come up between ACX7100-32C and ACX7100-48L. PR1560431
- [timing] [ptp] - ACX7100-48L TC - T1 time error and 2-way CTE does not meet Class C metrics for links for 100G-in and 40G-out combination. PR1562699
- Continuous interface MAC change on the neighbor switch results in evo-pfemand running at high CPU and never returns to the normal state. PR1564137
- ACX7100 : In multicast route extensive output does not show proper packet rate in pps for any multicast routes. PR1566475
- IRB logical interface does not get created after a sequence of events. PR1565842
- On ACX7100-48L and ACX7100-32C platforms, when a large amount of ARP resolutions happen on IRBs in a very short time, the ARPD process usage can shoot to 100%. PR1568206

- ACX7100-48L supports **Host 0 RTC battery error alarm** in case of RTC failure. PR1568440
- ACX7100-48L :IPv6 ping does not work, when strict uRPF is enabled. PR1568938
- ACX7100-48L: Router should not boot up with USB installation again after selecting the second option **Type 'reboot' and hit <return>** to complete the installation. PR1571930
- ACX7100-32C : :: mismatch in the snapshot recovery steps displays message in Junos OS Evolved 21.1R1. PR1578556
- ACX7100 - Evo-pfemand can crash on deactivating interfaces, VLANs and protocols in scaled setup with ECMP routes. PR1580565
- PICD restart or crash results in junks stats for carrier transition. PR1594253
- The evo-pfemand might crash after restarting app evo-pfemand. PR1594331
- SyncE feature does not work in 4x10G in ACX7100-32C. PR1595487

Infrastructure

- ToS of self-initiated packets might get changed unexpectedly. PR1578247

Routing Policy and Firewall Filters

- Syslog as an action of filter by default dump logs in syslog in Junos OS Evolved which is different from Junos OS. PR1564088

Routing Protocols

- Multipath information displayed for BGP route even after disabling the interface for one path. PR1557604

User Interface and Configuration

- The Junos OS Evolved operational state is incorrect on the system and CoS schedulers configuration change might not take effect.PR1536615

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

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 21.2R1 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.

What's New

IN THIS SECTION

Learn about new features introduced in this release for PTX Series Routers.

What's New in 21.2R1

IN THIS SECTION

To view features supported on the PTX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 21.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- PTX10001-36MR
- PTX10003
- PTX10004
- PTX10008

The following sections highlight the key features in this release.

Ethernet Switching and Bridging

- **Support for Proxy ARP on IRB interfaces (PTX10008)**—Starting in Junos OS Evolved Release 21.2R1, you can configure Proxy ARP on IRB interfaces.

  [See Proxy ARP.]

- **Data Center Bridging Capability Exchange protocol (DCBX) (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 support is added for DCBX. DCBX

is a discovery and exchange protocol for communicating configuration and capabilities among neighbors to ensure consistent configuration across the data center bridging network.

[See Understanding DCBX and Understanding CoS Flow Control (Ethernet PAUSE and PFC).]

High Availability

- **Support for VRRP on IRB interfaces (PTX10001-36MR, PTX10003)**—Starting in Junos OS Evolved Release 21.2R1, PTX10001-36MR and PTX10003 devices support VRRP on IRB interfaces. The PTX10001-36MR and PTX10003 do not support ISSU, proxy ARP, and MC-LAG with VRRP in this release.

  [See Understanding VRRP.]

- **NSR support for controller-initiated RSVP LSPs (PTX10003, PTX10004, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 21.2R1, we support NSR for controller-initiated RSVP-based point-to-point (P2P) and point-to-multipoint (P2MP) LSPs. The primary Routing Engine synchronizes all RSVP LSPs initiated by Path Computational Elements (PCEs) with the backup Routing Engine. It also synchronizes multicast flow specifications for any PCE-initiated P2MP LSPs with the backup Routing Engine. This ensures zero traffic loss for the traffic carried over PCE-initiated RSVP LSPs during Routing Engine switchovers. You enable this feature when you configure NSR.

  [See PCEP Configuration.]

- **Segment routing–traffic engineering (SR-TE) NSR support (PTX10001-36MR, PTX10003, PTX10008)**—Starting in Junos OS Evolved Release 21.2R1, we support NSR for SR-TE, enabling hitless traffic flow during Routing Engine switchover. Routes using next hops from SR-TE policies that don't support NSR may experience traffic loss on switchover. The SR-TE policies that don't support NSR are DCSPF and Path Computation Element (PCE).

  [See Segment Routing for Traffic Engineering.]

- **Warm standby mode for routing protocols process (rpd) on backup Routing Engine (PTX10003, PTX10004, PTX10008)**—Starting in Junos OS Evolved Release 21.2R1, you can set the rpd mode to `warm-standby` by using the `set routing-options warm-standby` command. Warm standby mode enables the backup Routing Engine to remain synchronized with the primary Routing Engine, allowing for faster Routing Engine switchover during GRES.

  [See warm-standby.]

- **Per PFE-Restart (PTX10008)**—Starting in Junos OS Evolved release 21.2R1 you can restart each PFE instance on the device.

  [See fpc.]

- **Non-Stop Routing (NSR) (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 support is added for NSR using the LDP and BGP protocols.

[See Nonstop Active Routing Concepts.]

- **Routing Engine Switchover (PTX10004, PTX10008)**—Starting in Junos OS Evolved release 21.2R1 you can configure Routing Engine switchover for failure scenario such as Disk Read Threshold Failover, Component Failover, and Disk Failure.

  [See Understanding Routing Engine Redundancy on Juniper Networks Routers.]

Interfaces

- **Support for performance monitoring and TCA (PTX10001-36MR)**—The PTX10001-36MR router supports performance monitoring and threshold crossing alert (TCA) information for the QSFP56-DD 400GE ZR optical transceiver modules (shown in CLI outputs as QSFP56-DD-400GE-ZR). TCAs provide the management system an early indication of the deteriorating health of an optical network connection when a monitored performance parameter crosses a certain threshold. You can view the current and historical performance monitoring metrics which are accumulated into 15-minute and 1-day interval bins by using the `show interfaces transport pm`command for optical transport link management.

  [See show interfaces transport pm.]

- **Support for local loopback on 400G coherent ZR DWDM (PTX10001-36MR)**— Starting in Junos OS Evolved Release 21.2R1, we support local loopback on 400G coherent ZR DWDM optics. To enable local loopback, use `loopback local` at the `[edit interfaces optics-options]` hierarchy level. We also support 75 grid spacing for the 400G ZR optics.

  [See optics-options, show interfaces, show interfaces diagnostics optics, and show chassis pic.]

- **Proxy ARP on IRB (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 you can configure Proxy ARP for an integrated routing and bridging (IRB) interface.

  [See Proxy ARP.]

- **IEEE 802.3ah Link Fault Management (LFM) (PTX10004, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 we support the IEEE 802.3ah LFM standard for Operation, Administration, and Maintenance (OAM) of Ethernet networks.

  [See Introduction to OAM Link Fault Management (LFM).]

Junos Telemetry Interface

> **NOTE**: For Routing Engine telemetry sensors supported by this platform, see Telemetry Sensor Explorer. If any Platform Forwarding Engine sensors have been added for this release, they are listed below.

- **JTI: Infrastructure: UDP streaming of PFE/line card sensors over management port (PTX10003, PTX 10004, PTX10001-36MR (PTX10K-LC-1202036MR), PTX10008)**—Starting in Junos OS Evolved Release 21.2R1, Junos telemetry interface (JTI) provides an infrastructure for exporting AFT-based Packet Forwarding Engine sensors using UDP tranport over a management interface.

  The following sensors are supported:

  - **/junos/system/linecard/interface** (logical interface)
  - **/junos/system/linecard/interface/logical/usage/** (physical interface)
  - **/junos/services/label-switched-path/usage/**(LSP)
  - **/junos/system/linecard/npu/memory/** (NPU memory)
  - **/junos/system/linecard/npu/utilization/**NPU utilization
  - **/junos/system/linecard/firewall/** (firewall)

  [See sensor (Junos Telemetry Interface).]

- **gRPC mutual authentication (PTX10001-36MR, PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 21.2R1, Junos telemetry interface (JTI), you can configure Junos telemetry interface (JTI) and Remote Procedure Calls (gRPC) to require client authentication as well as server authentication. Previously, only the client initiating an RPC request was able to authenticate the server; that is, a Juniper Networks device, using SSL certificates.

  [See Configuring Bidirectional Authentication for gRPC for Junos Telemetry Interface.]

MPLS

- **RSVP-TE supports preempting secondary LSPs that are signaled but not active (PTX Series and QFX Series)**—Starting in Junos OS Evolved Release 21.2R1, you can configure the hold priority of the secondary standby label-switched path (LSP) for RSVP-Traffic Engineering (RSVP-TE). The hold priority will be used to determine if the standby non-active LSP can be preempted. This will help to bring up non-standby secondary path LSPs with higher setup priority which are not able to come-up because of bandwidth crunch. To configure the non-active hold priority value for a secondary standby path, use the non-active-hold-priority statement at the [`edit protocols mpls label-switched-path`

<*lsp-name*>] hierarchy level. You can set the priority from 0 through 7, where 0 is the highest priority and 7 is the lowest.

[See RSVP Overview.]

Multicast

- **Multicast-only fast reroute (MoFRR) (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 we support MoFRR - Draft-Karan-mofrr-02. MoFRR minimizes the packet loss in PIM and multipoint LDP domains.

  [See Understanding Multicast-Only Fast Reroute.]

Network Security

- **Public Key Infrastructure (PKI) (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 we support PKI-based authentication.

  [See Public Key Infrastructure (PKI) and PKI.]

Routing Options

- **BFD (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 you can configure centralized and distributed BFD sessions. The centralized BFD runs on the Routing Engine, while the distributed BFD runs on the Packet Forwarding Engine.
  [See Bidirectional Forwarding Detection (BFD).]

Routing Protocols

- **Flexible algorithm inter-level leaking support for SR-MPLS in IS-IS (PTX10001-36MR, PTX10003, and PTX10008)**—Starting in Junos OS Evolved Release 21.2R1, we support flexible algorithm inter-level leaking for SR-MPLS in IS-IS. Also, we partially support TE-App draft ietf-isis-te-app in its basic form and extends 6PE support for flexible algorithm.

- **RPKI Sharding (PTX10003, PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 we support origin validation over BGP RIB shards. BGP sharding allows concurrent processing of BGP routes over different threads.

  [See rib-sharding.]

Routing Policy and Firewall Filters

- **Class-Based Forwarding Class Usage (PTX10004, PTX10008, PTX10001-36MR)**—Starting in Junos OS Evolved release 21.2R1 you can configure Source Class Usage (SCU) and Destination Class Usage

(DCU). SCU enables you to monitor the amount of traffic originating from a specific prefix and DCU enables you to track how much traffic is sent to a specific prefix in the core of the network originating from a specific interface.

[See Configuring the Filter Profile.]

- **Support for route's next-hop weight in policy match condition (PTX10001, PTX10003, PTX10008)**—Starting in Junos OS Evolved Release 21.2R1, a route with multiple next-hop paths can use weight associated with the paths to identify primary and backup paths. The path with the lowest weight is used as the primary path, and any paths with higher weights are treated as backup paths. You can use the next-hop weight as a match condition in export policies to redistribute IGP and BGP routes based on whether the primary or backup paths are active. You can also use this match condition in conjunction with the programmed match condition to avoid exporting routes programmed by APIs that don't have a primary path configured.

  Configure this match condition at the `[edit policy-options policy-statement policy-name term term-name from]` hierarchy level.

  [See Basic BGP Routing Policies.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Static route resolution over SR-TE tunnel (PTX10001-36MR, PTX10003, and PTX10008)**—Starting in Junos OS Release 21.2R1, we support static route resolution over segment routing–traffic engineered (SR-TE) colored and uncolored label-switched paths (LSPs). To enable this feature, configure the `spring-te-lsp-next-hop` statement at the `[edit routing-options static` *`destination`*`]` and `[edit routing-options rib` *`rib name`* `static` *`destination`*`]` hierarchy levels. The feature support extends towards static, DTM, BGP-SR-TE, and PCEP source types that are currently supported by Source Packet Routing in Networking–Traffic Engineering (SPRING-TE). If a source is not configured, by default, it takes the next hop as static.

  You must configure the `tunnel-tracking` statement at the `[edit protocols source-packet-routing]` hierarchy level to enable this feature. This feature enhances the accuracy of first-hop label-based tunnel status for SR-TE tunnels according to their route resolution.

  [See spring-te-lsp-next-hop and source-packet-routing.]

- **Distributed Constrained Shortest Path First (CSPF) for Segment Routing Traffic Engineering (SR-TE) paths (PTX10004, PTX10008)**—Starting in Junos OS Evolved release 21.2R1 we support distributed CSPF for segment routing LSPs. With distributed CSPF for segment routing, LSPs are optimized by the ingress device based on the configured constraints and metric type (traffic-engineering or IGP). The LSPs are computed to utilize the available ECMP paths to the destination with segment routing label stack compression enabled or disabled.

  [See Enabling Distributed CSPF for Segment Routing LSPs.]

What's Changed

IN THIS SECTION

Learn about what changed in this release for PTX Series Routers.

What's Changed in Release 21.2R1-S2

IN THIS SECTION

Network Management and Monitoring

- **SNMP support for MIB**—Operational command `show snmp mib walk system` now shows the latest software version and doesn?t show the build date.

  [See show snmp mib.]

What's Changed in Release 21.2R1

IN THIS SECTION

Authentication and Access Control

- **Command to automate SSH key-based authentication (ACX Series, PTX Series, and QFX Series)**—You can set up SSH-key based authentication between the network device and a remote host by issuing the `request security ssh password-less-authentication` operational mode command. When you execute the command with the appropriate options, the device generates SSH keys for the current user, provided the user does not already have existing keys, and transfers the user's public key to the **authorized_keys** file of the specified user on the remote host.

  [See request security ssh password-less-authentication.]

General Routing

- **SSH session connection limit and rate limit per connection (PTX Series and QFX Series)**—We have introduced SSH `connection-limit` and `rate-limit` options at the `edit system services ssh` hierarchy levels to enable SSH connection limit and rate limit per connection. The default connection limit value is 75 connections and there is no default value associated with rate limit.

- **Secure boot disabled alarm is raised (PTX10008)**—The `Secure boot disabled` alarm is raised when the system boots with secure boot disabled in bios.

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

- **Enhancement to the show chassis pic command**— You can now view additional information about the optics when you run the `show chassis pic` command. The output now displays the following additional field:

  MSA Version: Multi-source Agreements (MSA) version that the specified optics is compliant to. Values supported are: SFP+/SFP28 — SFF-8472 (versions 9.3 - 12.3), QSFP+/QSFP28 — SFF 8363 (versions 1.3 - 2.10), and QSFP-DD — CMIS 3.0, 4.0, 5.0.

  Previously, the `show chassis pic` command did not display this additional field.

  [See show chassis pic.]

- **Maximum Transmission Unit (MTU) Support (PTX10001-36MR, PTX10008 and PTX10004)**—MTU 16KB is only for transiting traffic of WAN interfaces. MTU is 9500B for protocols and 16KB for transit traffic.

  [See mtu.]

- **Enhancement to the show interfaces (Aggregated Ethernet) command (PTX Series and QFX Series)**—When you run the `show interfaces extensive` command for aggregated Ethernet interfaces, you can now view following additional fields for MAC statistics : Receive, Transmit, Broadcast and Multicast packets.
  [See show chassis pic.]

- **Enhancement to the default `remnant-holdtime` (Junos OS Evolved platforms: ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)** —Starting this release, the default `remnant-holdtime` has been increased from 180 seconds to 300 seconds. This provides sufficient time for protocols to start and sync routes from neighbors in a scaled environment, during rpd restart. You can configure `remnant-holdtime` at the `edit routing-options forwarding-table` hierarchy level.

  [See forwarding-table.]

- You can enter **zero suppression no-zero-suppression** at the **edit services analytics** hierarchy level.to disable zero suppression for gRPC-based sensors. When this feature is enabled, data for a sensor is sent to the collector if the sensor value is zero. All key value pair updates will be streamed to a collector without performing any zero suppression. To enable zero suppression again (the default), delete the configuration by entering #delete services analytics zero-suppression no-zero-suppression. Whenever this feature is set or deleted, any existing collector is disconnected

- **Mozilla certification authority (CA) certificates removed (ACX Series, PTX Series, and QFX Series)**—To minimize security risks, Junos OS Evolved no longer includes Mozilla's set of root certificates from various CA operators by default. To use Docker container images from a registry that requires TLS authentication, you must first save the image as a tar archive on a remote device and then import the contents of the archive on the device running Junos OS Evolved.

  [See Running Third-Party Applications in Containers..]

Interfaces and Chassis

- **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.]

Junos XML API and Scripting

- **Changes to how command-line arguments are passed to Python op scripts (ACX Series, PTX Series, and QFX Series)**—When the device passes command-line arguments to a Python op script, it prefixes a hyphen (-) to single-character argument names, and it prefixes two hyphens (--) to multi-character argument names. The prefix enables you to use standard command-line parsing libraries to handle the arguments. In earlier releases, the device prefixes a single hyphen (-) to all argument names.

  [See Declaring and Using Command-Line Arguments in Op Scripts.]

- **The `language python` statement is enabled by default (ACX Series, PTX Series, and QFX Series)**—The `language python` statement is configured by default in the `junos-defaults` configuration group on devices running Junos OS Evolved. Thus, you can execute unsigned Python scripts using the default Python version without explicitly configuring the statement on the device.

  [See Requirements for Executing Python Automation Scripts on Devices Running Junos OS.]

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

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

Network Management and Monitoring

- **Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, PTX Series, and QFX Series)**—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:

  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.

  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.

  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and only emits an `<ok/>` or `<rpc-error>` element.

[See Configuring RFC-Compliant NETCONF Sessions.]

- **Changes to how command-line arguments are passed to Python action scripts (ACX Series, PTX Series, and QFX Series)**—When a custom YANG RPC invokes a Python action script and passes command-line arguments to the script, the device prefixes a hyphen (-) to single-character argument names, and it prefixes two hyphens (--) to multi-character argument names. The prefix enables you to use standard command-line parsing libraries to handle the arguments. In earlier releases, the device passes the unmodified argument names to the script.

  [See Creating Action Scripts for YANG RPCs on Devices Running Junos OS and Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules.]

- **Changes in contextEngineID for SNMPv3 INFORMS (ACX Series, PTX Series, and QFX Series)**— Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos OS devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  [See SNMP MIBs and Traps Supported by Junos OS.]

Platform and Infrastructure

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX 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).]

Services Applications

- Changes to inline active flow monitoring (PTX Series)—When you use inline active flow monitoring, no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before. The IPFIX and version 9 Options Template Data Record now contains 0 in the `Flow Active Timeout` (IE 36) and `Flow Inactive Timeout` (IE 37) fields. Therefore, the Options Template Data Record is not compliant with IPFIX RFC 7011. We do not recommend that you configure the `nexthop-learning` statement at the `[edit services flow-monitoring version` *`version`* `template` *`template-name`*`]` hierarchy level, as

it reduces the number of packets that can be processed. The `show services accounting flow inline-jflow fpc-slot` *slot* operational mode command now displays 0 for all of the `Active Flows` and `Timed Out` fields. The various `Total Flows` fields are now equal to their respective `Flow Packets` fields. The various `Flows Inactive Timed Out` fields are now equal to their respective `Flow Packets` fields.

[See Understanding Inline Active Flow Monitoring.]

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Evolved 21.2R1 Release for PTX Series Routers.

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 unknown inner ether-type is received at the device over an EVPN-MPLS tunnel, then such packet is dropped. PR1564431

General Routing

- Restarting the fabspoked-pfe application for the line card restarts the line card. PR1486023
- For a faulty link, after autoheal is triggered and before the action is completed, if fabspoked-fchip restarts, the link stays in fault state. PR1508915

- Excess-rate configuration in port schedulers might not be completely honored in certain scenarios. In such scenarios, with explicit excess-rate configuration, the actual excess-rate achieved might still be more in proportion to the configured transmit-rate.. PR1528124

- Double-fault scenarios are not handled by link auto-heal feature and fabric links remain down if the Routing Engine switchover is attempted while auto-heal recovery is in progress. PR1529599

- 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 does not be apply to this interface and all the "interface queue" counters for this interface shows statistics as 0. PR1539052

- PTX10008: By default IPv6 addressing is configured at the /64 subnet by default irrespective of the subnet configured on the DHCP server side. PR1539839

- On all Junos OS Evolved platforms when the next-hop is added or changed to the Packet Forwarding Engine and the same next-hop is also forwarding nexthop of an indirect route, if ingress Packet Forwarding Engine is fast and egress Packet Forwarding Engine is slow (MPC and DPC cards for example or MPC cards with different versions), then this results in packet loss as ingress Packet Forwarding Engine being faster would have seen new FNH and also the indirect change. But egress Packet Forwarding Engine being slower would not have consumed indirect change yet. PR1547432

- On Junos OS Evolved PTX10008 platforms, if multiple SIBs are in offline state and GRES is performed immediately, SIBs may be stuck in offline state for sometime. PR1554423

- During fabric link bringup, fabspoked-fchip[0-5] app restart, sib offline, or sib online, /re0/fabspoked-fchip[0-5] might not respond to CLI commands (`show chassis sibs`). This is indicated by the following output on CLI

  error: communication failure with /re0/fabspoked-fchip[0-5]/.

  The producer app is either down or unresponsive. Run the `show system processes node node | grep app` command to check if the app is running. Run the `show system application` app to check the state of the app. If there was a change in primary role recently, wait for the switchover to complete for the app to be online. `Run request chassis routing-engine master switch check` to check the status of the switchover. If the command is reissued after training or ifs detraining reaches a steady state, the valid output shows up in the CLI. PR1570605

- The PTP FPGA is kept in reset during BIOS boot. During Linux boot, the PTP FPGA is taken out of reset and pcie-tree is re-enumerated. Hence you would be seeing the Link-up/down during this sequence. PR1572061

- Sometimes during ZTP auto-speed detection, ports in 4x25G mode might not come up if there is difference in FEC across the nodes. PR1572666

- Ungraceful FPC restart followed by an immediate master switchover is a double fault scenario. The final state of system is unpredictable. The user has to wait until the system stabilizes after ungraceful FPC restart before triggering events like primary switchover. PR1576549
- Double-fault scenarios are not handled by link auto heal feature and fabric links remain down if the Routing Engine switchover is attempted while auto-heal recovery is in progress. PR1578615
- During NSR switchover, there is a potential of loss of a small number of packets addressed to the host or Routing Engine. This may cause a LDP sessions to flap since BFD sends only a single keepalive and loss of this keepalive causes the BFD to be declared down. PR1580355
- UDP encapsulated MPLS packets with `explicit null label` received on FTI tunnel gets dropped after UDP decapsulation. After UDP tunnel header decapsulation MPLS payload with explicit null label cannot be forwarded as it requires popping MPLS explicit null label and lookup of MPLS inner payload which is not supported in BT ASIC based products without looping back the MPLS payload for additional lookup. We support only scenarios where we decapsulate the tunnel header and forward the packets based on the exposed MPLS Label. PR1580641
- PTX10003-60C or PTX10003-160C interface queue and voq does not report drops when the low priority queue is slightly oversubscribed (seeing page timeouts). PR1581490
- If we offline multiple sibs and halt the primary Routing Engine, the SIBs can be stuck in the offline state for 15 minutes, before it goes to offline state. PR1584712
- Change in flow label configuration is handled by deleting and re-adding the Layer 2 VPN route. Therefore, traffic loss is seen during the configuration change. PR1585059
- Ungraceful power off or removal of SIBs can cause PCIE errors for SIB devices and the system might report the alarm **PCI Uncorrected error on dev 0000:00:03.0**.PR1590286
- Additional triggers like fabricHub or Fabspoked-fchip daemon restarts while FPC offline or online is in transition, results in traffic loss after the FPC is online. PR1596818

Routing Policy and Firewall Filters

- On Junos OS Evolved PTX Series platforms, the firewall `prefix-action-stats` is not supported, so if a user tries to use the `show firewall prefix-action-stats` CLI then the respective XML tags `<firewall-prefix-action-information>` is not be visible. PR1443325

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Evolved 21.2R1 for PTX Series Routers.

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

General Routing

- On all the platforms with BGP SR-TE (Spring-TE), in segment routing topology the transit v4 traffic might have missing labels and might get dropped in first hop, when ingress is forwarding traffic. It might miss out all the labels except the last hop in the v4 traffic forwarded by NH interface. PR1505592
- VSTP vport scale crosses 3500 leading to VSTP convergence issues due to BPDU drops in packetio of the Packet Forwarding Engine. This issue can be solved by adding DDOS configuration to support the increased scale. PR1509685
- BGP-SRTE binding-sid with more than one label stack needs enhancement for PTX10003-80C PTX10003-160C. For PTX Series does not support more than 1 chain-composite-next-hop, it can

configure this CLI configuration statement to disable chain-composite-next-hop as a workaround. With this configuration statement configured, all labels are pushed from egress.

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

PR1512213

- 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

- GRES is not supported when FPCs are restarted. PR1539685
- Issue the CLI command `show system fabric app-status` to see list of fabric applications that are running in the system and fetch the application status using the right application name. PR1540706
- On PTX10008 with inherit VRRP configuration some of the sessions might take longer time to converge. And during this period we see traffic loss for those sessions. So if system has inherit sessions configured then it is recommended to reduce the inherit advertisement interval timer to 6 seconds to avoid traffic loss. Default value of the inherit time is 120 seconds. PR1571339
- For input subscription paths containing a ":" character, the extension header in case of GNMI and certain fields for the `show network-agent statistics` CLI has incorrect values. PR1581659
- sflow ingress sampling does not work for user IPv6 traffic with aggregated Ethernet ECMP case at last hop router with UHP LSP. PR1582960
- IPv6 based PKID enrollment operations are not supported over revenue ports on Junos OS Evolved platform.PR1584378
- Used-power and allocated-power leaves for telemetry under /components path are currently shown as "0" for all FRUs. PR1587184

- On all Junos OS Evolved platforms configured with PIM (Protocol Independent Multicast) and MoFRR (Multicast-only Fast Reroute), a brief multicast traffic drop might be observed towards the downstream receiver after performing Routing Engine switchover or restart of rpd with GRES (Graceful Routing Engine Switchover) and warm-standby enabled.PR1593810
- Explicit fault trigger while Packet Forwarding Engine is in transition is not supported. Once Packet Forwarding Engine is in FAULT, Packet Forwarding Engine restart does not recover the Packet Forwarding Engine. PR1596476
- The **[Error]Jexpr: cannot find ifToken for counterType:12** has no impact on the operation of the system and can be safely ignored. Nothing needs to be done if this error is seen. PR1597355
- With QFX5220-32CD interconnected to PTX10008 chassis line card, intermittently issue is seen where 100G SR4 optics link does not come up or flaps, and continuously increasing FEC errors are seen on the PTX Series end. PR1597528
- Major host 13 Ethernet interface link down false alarm is seen after Routing Engine 1 replacement manually. PR1597763
- When all Packet Forwarding Engines go offline and are brought online with multicast route being active during this cycle, multicast traffic is permanently lost due to the absence of mcast route. PR1598894
- Due to issue in AGEOUT notification for inline sessions, sessions remain up till peer sends BFD down packet or BFD client brings it down. PR1599257
- Do not restart fabric hub and fabric spoke applications during the FPC offline or online event. Wait for the FPC to be in stable state, before restarting the application. PR1600094
- On PTX10008 routers, `the set chassis redundancy routing-engine 1 master` command does not change the default Routing Engine election priority. PR1601430
- In PTX10008, at times, MVRP enabled trunk ports can go into blocked or designated state if the peer connected to the interface has no VLANs configured in its trunk port. PR1601915
- If multiple Packet Forwarding Engines are brought online rapidly using the `request chassis fpc x pfe-instance y online | restart`, it might cause evo-aftmand to terminate and the FPC to reboot. PR1602035
- On PTX10003, IRB ping fails post powering off and on underlying Packet Forwarding Engines for Aggregate Ethernet child member. PR1602181
- Powering off Packet Forwarding Engines gives error messages **Jexpr: deleteFdbEntry: Null**. This happens when any of FPC has all the Packet Forwarding Engines go offline but the FPC is online and there are some routes deleted from the control plane. It does not affect any functionality. PR1602670

- Slow counters (off-chip) used for ingress nexthops might not work. PR1603518
- Due to the product limitation of the PTX10008 and PTX10001-36MR having to communicate between COSIM and CDA running in different VMs, periodic counter reads induce the large delays in a scale configuration. This scale scenario needs to be executed on a hardware setup. PR1603975

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 alarm **Host 0 Active Disk Usage Exceeded** might be generated due to large files which were already marked as deleted. PR1601251

Interfaces and Chassis

- Change to IFD properties in 25G Mode brings down the adjacent channels. PR1594740
- The licenses used value for **Port Bandwidth Usage (PAYG)** field under **show system license** does not update when new interfaces are created. In order to correct the value, the IFMAND process needs to be restarted to invoke a refresh. PR1595179
- The `show platform object-info anomalies summary` CLI command may time out after the system is up for some time. This does not affect normal operation of the system but this debug command will not work.
- Should the marvd daemon restart for any reason, the device becomes unreachable because of PCI uncorrectable non-fatal AER errors getting generated on marvd restart and system goes for reboot without any core. This is a rare event that occurs if there marvd crashes. PR1600870

MPLS

- The rpd generates a core file most of the times in standby LDP module when static LDP P2MP configuration is deleted. PR1594405

- Multipoint LDP (mLDP) transit statistics reset after Routing Engine switchover. PR1596395

Routing Protocols

- This scenario happens when we enable IS-IS authentication key-chain having multiple keys between routers. TheIS-IS adjacency is up since both routers have the same key active. When we manually change the system time in such a way that routers have different keys active in the key chain, the IS-IS adjacency must go down. However, that does not happen. PR1572441

Network Management and Monitoring

- Reported Junos OS User Interface Scripting Environment (JUISE) generates a core file while using request-system-storage-cleanup RPC with POST method in HTTP. The root cause is that this RPC sends a long list of file details. Other RPCs do not have this issue. PR1587337

Routing Protocols

- This scenario happens when we enable IS-IS authentication key-chain having multiple keys between routers. The IS-IS adjacency is up because both the routers have the same key active. When we manually change the system time in such a way that routers have different keys active in the key chain, the IS-IS adjacency should go down. However, that is not happening. PR1572441

User Interface and Configuration

- If an inet filter is configured with or without `family` statement then changing its configuration to without or with `family` statement might cause the filter process (firewalld) to restart unexpectedly. PR1556426
- Post `request system zeroize` operation, sshd service is not enabled by default due to race condition on PTX10008. Hence it is recommended to enable `edit system services ssh` as part of the first commit on the device so that user 'ssh' credentials are created properly on the device and there is no need for password to log in to the backup Routing Engine. PR1594258

Resolved Issues

IN THIS SECTION

Learn about the issues fixed in this release for PTX Series Routers.

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: 21.2R1-S1

IN THIS SECTION

General Routing

- GRE keepalive packet with recursion control bit set gets dropped on PTX10003. PR1602353

Infrastructure

- Malformed packets might be sent out on egress interfaces in Junos OS Evolved platform. PR1603783

Interfaces and Chassis

- PTX10003:`show platform object-info anomalies summary` CLI command times out. PR1598337

Routing Policy and Firewall Filters

- The firewalld might crash if you configure `fragment-offset` out of the range (fragment-offset 1-900000000000). PR1605805

Resolved Issues: 21.2R1

IN THIS SECTION

Class of Service (CoS)

- Applying non-configured traffic-control-profiles on the interface for the first time causes cosd crash. PR1571742
- Cosd core can be seen on issuing `run show class-of-service`. PR1580573
- The user-defined CoS might not get applied on the interface when you configure class-of-service stanza with `interface all`. PR1592900

EVPN

- Sometimes Broadcast, Unknown Unicast, and Multicast (BUM) traffic that comes 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. PR1578314
- EVPN missing option under routing-instances <> protocols. PR1581821

General Routing

- Ungraceful SIB failures results in transient loss of traffic. PR1497212
- VCCV type 1 connectivity verification is not supported. PR1503724
- Junos OS Evolved - The aggregated Ethernet interfaces do not display member links' statistic. PR1505596
- [Junos OS Evolved PTX10008] : DCDCEdge-VPNTunnelMulticastL3L2:serviceability :core file creation failure, aftmand core file is stuck at **/var/lib/ftp/in/**]. PR1522404
- The user might not be able login to PTX10001-36MR after multiple abrupt power cycles or reboots. PR1523238
- Traffic loss might be observed after rpd restart or GRES in a scaled MPLS scenario. PR1525710
- Set of info level no passwd entry cron logs displayed every 1 minute. PR1527266
- Global Port-Mirroring applied with deactivation does not display XML correctly for PTX10003. PR1529413
- FPC vmcore files can be stored at **/var/lib/ftp/in/fpc_slot/** on RE0 or RE1. PR1531214
- PTX10008: Need support for `show chassis fabric summary` output.PR1532163
- The interfaces might take longer to come up after loading baseline and rollback configuration. PR1534996
- The `show chassis alarms` should be redirected to `show system alarm`. PR1536020
- The port mirroring stops working for the FTI interface when GRE source is changed. PR1536223
- EXP rewrite might not take effect on the exposed label while performing PHP. PR1538918
- PTX10001-36MR : :: IP-in-IP: Routing Engine initiated traceroute packets do not use IP-in-IP encapsulation. PR1545049
- [ddos] [ddostag] PTX10001-36MR : `show ddos-protection protocols bgp statistics brief` throws error - **communication failure with /re0/evo-aftmand-bt/**. PR1547491
- The segment routing-traffic engineering (SR-TE) might stay in the Up state when the routes are deleted through policy. PR1547933
- The process aftmand might crash if you restart FPC or delete configuration. PR1548866
- The firewalld crash might be seen if deactivating or activating the firewall during back to back switchovers. PR1549856

- The backup Routing Engine should not clear the primary Routing Engine ARP entry. PR1550959
- Junos Evolved Ethernet MAC Address allocation is incorrectly distributed for FPCs. PR1552955
- jnxFruLastPowerOn value is incorrect for FPCs. PR1553924
- **CoS WRED Curve: Create Expr Curve: No curve data points!!** errors are seen when interpolate is configured under drop profile. PR1554220
- [pfe] [cos_installation] : PTX10008 PFC: Global Ethernet flow-control should be disabled when priority-based flow control (PFC) congestion notification profile (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
- [ddos] [ddostag] Junos OS Evolved PTX10003 : :: Violations observed even after clearing the states and statistics. PR1554515
- [pfe] [pfe] PTX10008 : : The following error messages are seen **@Error] Jexpr: CoS Scheduler Express Handle: Destructor: Interface Physical Handle is NULL. streamIndex:1147 schedNodeToken:4508** while executing 2000 static SR-TE scale in PTX10008. PR1558328
- Junos OS Evolved:JDI_FT_REGRESSION:FIREWALL:PTX10003:FIREWALL:[firewall] [filter_installation] :: firewalld: anomolies are seen in firewalld app for publish publish-deleted. PR1559046
- Major alarms might be seen when a large class-of-service buffer-size is configured. PR1559459
- The `request system software sync all-versions` command might cause the CLI to hang. PR1560315
- PTX10008:NSR Support for LDP/RSVP/BGP: BGP NH_index (indirect and unilist) change after GRES +NSR trigger causes a momentary (unexpected) traffic loss. PR1560323
- The `request system software sync all-versions` command might cause the CLI to hang. PR1560315
- The class-of-service RED feature might work unexpectedly and cause traffic drop. PR1560495
- The FPC might reboot in a high-scale configuration scenario on Junos OS Evolved PTX10008 platforms. PR1560757
- Traffic drop might be seen after Packet Forwarding Engine restart. PR1560901
- Support switchover on routing-crash configuration statement during abnormal termination of rpd. PR1561059
- The pfe-disable action being taken due to 'viqm_intr_viqm_gs_deq_dry_err' interrupt. PR1561265

- Timingd-lc errors, **CdaExprClient: grpc api call ExprServerInfoGet failed" and "CdaExprClient: Failed to fetch server info error:5** seen on all FPCs after restarting the router or FPC restart. PR1561362
- Junos OS Evolved System - after recovering from restart routing immediately, object-info anomalies is observed on rpd agent. PR1561812
- PTX10008: After sync all followed by rollback and then reboot, RE1 booted on snapshot. PR1562189
- CPU utilization of evo-aftman process goes to 100% in a certain scenario on Junos OS Evolved PTX Series devices. PR1562328
- Complete ingress multicast traffic loss might be seen on interfaces that are flapped using Packet Forwarding Engine offine or online command. PR1562452
- The interface loopback might not work if there is no optics connected to the port on PTX10008. PR1562471
- For topologies involving high ingress and transit LSP scale on PTX10008 Junos OS Evolved platforms, error messages can be seen in journalctl when tearing down the ingress and transit LSPs. This leads to slow hardware resource leaks for the ASIC onchip memory associated with nexthops. PR1562503
- The aftmand crashes when sFlow is enabled on the loopback interface. PR1562869
- FPC is not be powered on using `request node power-on fpc`. PR1562981
- The ARP might not resolve and traffic might be dropped on Junos Evolved platforms. PR1563684
- PTX10008: RE0 goes into reboot loop continuously during validate restart (INDB unsupported). PR1563742
- It might take a long time to create physical interfaces after restarting the FPC. PR1564156
- The evo-cda-bt might crash in large scaled configuration scenario. PR1565427
- Improve `request system software delete` CLI command to add new option `archived` to delete all old software versions except current and rollback. PR1566173
- MACsec-Encrypted packets counter displays 0 under **Secure Association transmitted** of MACsec statistics output when AN rollovers with `sak-rekey-interval` configured. PR1566665
- Observed license-check core on RE-1 during runtime removal of CB[0] SAM FPGA from PCIe device. PR1567066
- The `request node halt re` does not work as expected when the Routing Engine is a primary Routing Engine. PR1567558

- Drop counts in `show interfaces voq ae0` might not match with `show interfaces queue` when `clear interface` command is issued while traffic flows. PR1567598
- User folders are not created when snapshot taken. PR1567880
- Junos OS Evolved:JDI_FT_REGRESSION:BGP:PTX10003 :: State is not established for `show bgp bmp < Station name>` post `authentication-key bmp-auth` configuration. PR1568046
- More memory usage may occurs in ndpd (NDP daemon). PR1568370
- Routes learned through IRB interface might not be reachable in IBGP setup. PR1568566
- The firewalld crash might be seen if GRES is executed as soon as the firewall is activated (for example, commit is done).PR1569427
- PTX10008: User script output should be logged during ZTP execution for determining failure in the logs. PR1570167
- The interface hold-time down feature might not work in some conditions. PR1570204
- The ZTP state machine might be stuck on the management interface for about 12 minutes. PR1570598
- Traffic loss time more than link failover time might be seen on PTX10008. PR1570665
- Silent switchover might be triggered on executing restart routing. PR1570993
- Certain leaves in /components/component[name='FPC1:CPU']/properties/property/cpu-utilization-total is not in Junos OS Evolved 20.4R2.PR1571502
- The grpcd process might crash and telemetry subscription might retry until grpcd restarts. PR1572107
- The log and syslog action does not work along with port-mirror or sample in PTX10003 and PTX10008.PR1572239
- FPCs get restarted automatically after ungraceful removal of SIBs. PR1572431
- [Junos Telemetry Interface] PTX10008:: NPU Memory KHT DLU IDB value. PR1572704
- The rpd agent crashes during interface flapping. PR1572940
- Junos OS Evolved: Specially crafted packets may cause the AFT manager process to crash and restart (CVE-2021-0286). PR1572969
- The `hash-key label-1-exp` CLI configuration statement does not take effect. PR1573109
- Junos OS Evolved:JDI_FT_REGRESSION: PTX10008 [jflow][Firewall]: Counter value from sampling firewall fails while validating IPv4 and IPv6 egress sampling with static routes. PR1573969

- Traffic might not get load balanced after setting and deleting the hash-seed value. PR1574108
- All queues are not getting correct rate as per the commit when more than one queue are configured with **transmit-rate remainder**. PR1574121
- The Packet Forwarding Engines might crash accidentally on FPC if using the service provider style layer-2 interfaces on Junos OS Evolved platforms. PR1574146
- Junos OS Evolved-PTX10008: CLI timout error **communication failure with /fpc0/evo-aftmand-bt/** and traffic loss seen. PR1574513
- The rpd might continuously crash if deleting forwarding-class policy with discard action. PR1575177
- Some error messages might be seen when performing continuous aggregated Ethernet deactivate or activate on PTX Series.PR1574714
- Huge Invalid statistics shown in `show interface statistics` when an interface is removed and added from the aggregated Ethernet bundle. PR1575623
- [Junos OS Evolved_NSR_longivety] sr_master_kkcm_thread runs at 100% CPU for long time. PR1575661
- The distribution of buffer with buffer-size remainder is not correct on Junos OS Evolved PTX Series. PR1575798
- Loss on IPv6 traffic streams might be observed after NSR SWO. PR1576369
- Clock status holdover when configured for free-run. PR1576487
- [PTX10008 Junos OS Evolved] Incorrect capacity value is shown on JNP10K-PWR-AC2 /JNP10K-PWR-DC2 PSM. PR1578682
- The kernel might hang if multiple Routine Engine primary switchovers are performed in a short span of few seconds. PR1578693
- FPC Status LED do not turn RED with power fault. PR1579466
- The Packet Forwarding Engine function might break down on all FPCs after performing Routing Engine switchover on Junos OS Evolved platform. PR1579683
- FPC is stuck in online state and seen continuously rebooting during unified ISSU. PR1580374
- The l2cpd process might crash on Junos OS Evolved platforms with dual Routing Engines. PR1580479
- Junos Telemtry Interface properties missing after HwD app restart. PR1580735
- In certain scenarios, shapers applied on a 10G interface might drop the traffic more than the configured max-rate. PR1580795

- Streaming over IPv6 fails in Junos OS Evolved. PR1581341
- [mpls] PTX10004 ::After disabling the active path, forcing FRR, we see large traffic loss, also we see that the irp.core.trapcode.cfg_err counter increased. PR1582170
- The broadcast traffic is matched by multicast filter. PR1582217
- The CLI `show chassis craft-interface` does not show correct PSM LED status on PTX10008 Junos OS Evolved.PR1582444
- Node locked license addition fails in Junos OS Evolved. PR1582704
- There can be a failure of the config-sync service and a major system alarm is raised after upgrade. PR1582717
- Junos Telemetry Interface: Interfaces: Missing Leaves - Transceiver/state. PR1583076
- New primary might be struck with **Switchover is in transition, Please wait** after primary reboot test case if the switchover happens back-to-back within 2-3 seconds.PR1583347
- The system may crash if configuring IPv6 FBF with prefix < /88 on all Junos OS Evolved platforms. PR1583374
- The FRR convergence number is high with ALB enabled on aggregated Ethernet bundle. PR1583866
- The ospf-hello ddos statistics pktCnt is listed as 0. PR1584458
- After PIC offline and online, `show interfaces queue <intf>` shows large values for cumulative tail-drop and RED-drop packets and bytes. PR1585552
- Packet loss might be seen during global repair of FRR. PR1586122
- PTX10008: NPU HBM statistics. PR1586148
- Npu Sensor: "components-memory" vs "components" PR1588242
- NPU Sensor: "components-memory" vs "components"
- Traffic loss observed on global repair after disabling of active path forcing FRR. PR1589803
- [platform_re] [re_generic] PTX10008 : The Packet Forwarding Engine stuck in READY state with anomalies type net::juniper::fabric::fabricPfeE. PR1590319
- If a system has power shortage, then post switchover we see unexpected FPCs or SIBs go down on the new primary. FPCs that were down on the previous primary might be online if they are discovered earlier in the powerManager on the new primary. PR1592004
- PTX10004: Power off re1 message is very vague and might panic users. PR1592145

- PTX10008 Serviceability: picd log floods when there is **Optics does not support configured speed** system alarm.

  PR1592165

- ZTP occasionally fails to apply user configuration after the system upgrade.PR1592281
- Duplicate Junos Telemetry interface leaf "oper-status" tag for physical interface index 16386 has mismatch value. PR1592468
- Port related component sensor does not get exported when subscribed to **/components/component/state/** path. PR1593031
- Port mirroring instance down with mirrored output as tagged interface. PR1593276
- [rpda] PTX10004 : :: PDT: rpdagent crash seen in the primary Routing Engine @NHTable::insert , comp_nh_rts_handler after fourth GRES with GR enabled. PR1593104
- Load-balance might not take effect for the Layer 2 VPN traffic on the PTX10008 platforms. PR1593548
- On Junos OS Evolved platforms, "type" leaf value for "FPC3:PIC0:PORT0:XcvrX" displays XCVR as opposed to TRANSCEIVER displayed in Junos OS. PR1595103
- Some TCP sessions might not be established after performing the `request system snapshot` command. PR1595470
- The ifmand core files detected during the NSR switchovers with BT @0x00007fd98dc609d8 in bq_hash_elinks_equal (ds2=0x7fd83b74f600, ds1=0x7fd83be4d2d0) at ../../src/cbq.c:771. PR1595846
- The aftmand core might be observed on all Junos EVO platforms. PR1597649
- Continuous FPC restart might be observed on Junos Evolved products with firewall policer configuration. PR1602446

Infrastructure

- Junos OS Evolved Platforms: The IPv6 BGP session might flap continuously due to transport error (CVE-2021-0226). PR1544978
- The TCP session might fail on Junos OS Evolved platforms with dual Routing Engines. PR1555441
- Default multicast ff00::/8 route not available while verifying IPv6 multicast routes. PR1563940
- There might be continuous memory increase noticed for systemd daemon.PR1566717
- ZTP over IPv6 on a management interface is not functional over 21.1R1-EVO. PR1567967

- The backup router might get stuck in the idle state during the NSR replication for IBGP single hop peers. PR1569696
- In Junos OS Evolved, BGP and LDP sessions with TCP MD5 authentication established with peers not configured for authentication (CVE-2021-0297). PR1569843
- The configuration statement `default-address-selection` might not work on all Junos EVO platforms. PR1570552
- Next-hop incorrectly associated with lo0 in the forwarding table when the interface is configured as unnumbered. PR1570918
- ToS of self-initiated packets might get changed unexpectedly. PR1578247
- The FTP IPv6 server function might be failure on all Junos OS Evolved platforms. PR1591733
- The TCP-based protocol sessions might remain down after multiple Routing Engine switchovers. PR1593580

Interfaces and Chassis

- [hostpath] [hostpathtag] Junos OS Evolved-PTX10003 : :: "picd" Publish-deleted anomolies seen for the type "net::juniper::hwd::serdesDfeTuneStatusE" .PR1547484
- [lacp] [lacptag] PTX10004 lacpd core is observed after router reboot .PR1553196
- PTX10001-36MR: Control physical interface might not be present for ports et-0/0/11 and et-0/2/11. PR1566752
- There might be increase in memory for the fabspoked process .PR1574391
- The `show interface description` display order is different from Junos OS and Junos OS Evolved.PR1576224
- PTX10008: ifmand core file seen at IFAManager::handleIFACModify.PR1583681
- When changing the Micro BFD session's address from IPv4 to IPv6 or vice versa, the BFD session and aggregated Ethernet interfaces go down.PR1584853
- Some interface units description are missing from the output of `show interfaces description` on certain PTX Series platforms running Junos OS Evolved .PR1591340
- PTX10003-160C Junos OS Evolved, interface is not programed in routing-instance. PR1596768

Network Management and Monitoring

- The SNMP hostname does not match the configured hostname on Junos OS Evolved based device. PR1567835
- The snmpd core might be generated on Junos OS Evolved platforms. PR1572236

Routing Policy and Firewall Filters

- Syslog as an action of filter by default generates logs in syslog in Junos OS Evolved which is different from Junos OS. PR1564088
- Error while applying filter as output-list with DSCP action. PR1569691
- Toggle of the interface-specific field of filter already bound to interface is not allowed. PR1571654

Routing Protocols

- The rpd might crash on the backup Routing Engine after rpd restart is triggered on the primary Routing Engine. PR1563350
- The rpd process might crash when there is BGP session re-establishing or flapping. PR1567182
- There might be 10 seconds delay to upload the LSP on the point-to-point interface if rpd is restarted on its direct neighbor. PR1571395
- The ppmd might crash when enabling MD5 authentication on OSPF with BFD flapping. PR1576893

User Interface and Configuration

- The configuration under groups stanza is not inherited properly. PR1529989
- The Junos OS Evolved operational state is incorrect on the system and CoS schedulers and configuration change might not take effect. PR1536615
- shell-init: error retrieving current directory: getcwd: cannot access parent directories: No such file or directory. PR1549479
- The LACP might stop working after disabling lacp sync-reset. PR1576146

Junos OS Evolved Release Notes for QFX5130-32CD, QFX5220, and QFX5700 Devices

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 21.2R1 for QFX5130-32CD, QFX5220-32CD, QFX5220-128C, and QFX5700 switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

Learn about new features introduced in this release for QFX Series Switches.

What's New in 21.2R1

IN THIS SECTION

Hardware

- **New QFX5700 Switch (QFX Series)**—In Junos OS Evolved Release 21.2R1, we introduce QFX5700 switch as the first modular chassis from Juniper that uses Broadcom's Trident 4 chipset to support and deliver a diverse set of use cases. Featuring a 5-U form factor, QFX5700 supports very large, dense, and fast 400GbE IP fabrics with a fully redundant 12.8Tb capacity that benefits large public cloud providers. It also offers 10/40/100/400GE high port density and delivers high-performance, scale, and flexibility to support IP services and functions for service provider, web, and enterprise networks.

  NOTE: Upgrade to Junos OS Evolved Release 21.2R2 for the latest software features and bug fixes.

  "Table 3" on page 67 summarizes the QFX5700 features supported for Junos OS Evolved Release 21.2R1.

  Table 3: Features Supported by the QFX5700 Switches

Class of service	• Support for class of service (CoS) configuration with these limitations:
	  • 802.3X Ethernet PAUSE is not supported.
	  • CoS flexible hierarchical scheduling, also known as enhanced transmission selection (ETS), is not supported.
	  • Neither MPLS EXP ingress packet classification nor egress rewrite rules are supported.
	  • Classifiers and rewrite rules are applied to logical interfaces instead of physical interfaces.
	  [See CoS Support on QFX Series Switches, EX4600 Line of Switches, and QFabric Systems.]
	• CoS support on EVPN-Virtual Extensivle LAN (VXLAN). Support is provided for defining classifiers and rewrite rules on leaf (initiation and terminations) and spine nodes for EXPN VXLANs. Support is also provided for defining schedulers, interpolated drop profiles, explicit congestion notification (ECN), and priority based flow control (PFC).
	  [See CoS Support on EVPN VXLANs.]

DHCP	• Support for DHCPv4 and DHCPv6 stateless relay over Layer 3 (L3) interfaces.
	  Support includes:
	  • Option-82 for DHCPv4.
	  • Option-18 and Option-37 for DHCPv6.
	  • Virtual router aware DHCPv4/v6 stateless relay.
	  NOTE: Support does not include IRB interfaces.
	  [See DHCP Relay Agent.]

EVPN	• EVPN-VXLAN support for unicast traffic using media access control (MAC) virtual routing and forwarding (VRF) routing instances. The switch can be a leaf device or a spine device in an edge-routed bridging or centrally routed bridging overlay design. You can configure multiple EVPN instances (EVIs) of type `mac-vrf`. Each EVI can support a different EVPN service type (`vlan-based`, `vlan-aware`, or `vlan-bundle`). The switch creates and uses one virtual tunnel endpoint (VTEP) logical interface per remote provider edge (PE) device by default, regardless of the number of routing instances. This implementation improves VXLAN VTEP scaling with multiple routing instances. EVPN-VXLAN support on this switch includes:
	  • External BGP (EBGP) and internal BGP (IBGP) overlays.
	  • Single-homing and all-active multihoming.
	  • Layer 2 and Layer 3 unicast for IPv4 and IPv6 with ARP suppression.
	  • Proxy ARP and ARP suppression for Layer 2 and Layer 3 IRB traffic.
	  • Proxy Neighbor Discovery Protocol (NDP) and NDP suppression for Layer 2 and Layer 3 IRB traffic.
	  • Service provider style configuration with Q-in-Q support, supported for Layer 2 gateways only and only with MAC VRF `vlan-bundle` service type.
	  • Storm control.
	  • MAC mobility, MAC limiting, MAC move limiting, and duplicate MAC detection.
	  • Core isolation.
	  • Proxy re-advertisement of EVPN Type 2 MAC+IP routes by all provider edge devices on the same Ethernet segment identifier (ESI).
	  • Virtual machine traffic optimization (VMTO).
	  • OSPF and BGP routing protocols on IRB interfaces.
	  [See EVPN User Guide.]

	• Support for sFlow monitoring technology in EVPN-VXLAN fabrics. [See Overview of sFlow Technology.] • Support for port mirroring in EVPN-VXLAN fabrics. [See How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics.] • Support for EVPN Type-5 routes in EVPN-VXLAN fabrics. [See Understanding EVPN Pure Type-5 Routes.] • Support for EVPN-VXLAN firewall filtering and policing. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]
Hardware	• The QFX5700 switch supports up to 8 line cards and each line card supports 16 QSFP28 ports that operate at 100GbE speed or 4 QSFP56-DD ports that operate at 400GbE speed. The QFX5700 switch has AC or DC power supplies and front-to-back airflow. [See QFX5700 Hardware Guide.]
High availability (HA) and resiliency	• Resiliency support for Routing and Control Board (RCB) includes CPU memory and dual in-line memory module. You can configure fault-handling actions such as logging the error, raising alarms, sending SNMP traps, and indicating error conditions by using LEDs. [See routing-engine (Chassis).] • Resiliency support for the chassis, line card (JNP-FPC-4CD), and Forwarding Engine Board (FEB)—QFX5130-FEB—includes handling the faults related to the links between components; for example, between a line card (such as JNP-FPC16C) and the switch fabric. [See QFX Hardware Components.]

Interfaces and chassis	• Support for GRE tunneling.
	[See Generic Routing Encapsulation.]
	• Support for qualifying optics on the 100GbE and the 400GbE FPC line cards includes software monitoring, sensor diagnostics, FPC interfaces node level failure or restoration, events, and error logging.
	[See QFX Hardware Components.]
	• Support for two new FPCs introduced for QFX5700 switches:
	  • JNP-FPC-16C: The line card contains a total of 16 QSFP28 ports that support 100Gbps and 40Gbps speeds. You can channelize the ports operating at:
	  • 100Gbps to four 25Gbps channels.
	  • 40Gbps to four 10Gbps channels.
	  • JNP-FPC-4CD: The line card contains a total of four QSFP56-DD ports that support 400Gbps, 100Gbps, and 40Gbps speeds. You can channelize the ports operating at:
	  • 400Gbps to four 100Gbps channels.
	  • 100Gbps to four 25Gbps channels.
	  • 40Gbps to four 10Gbps channels.
	  [See Port Settings.]
	• Support for one or two RCBs.
	[See QFX5700 Routing and Control Board.]
	• Support for the FEB with field-replaceable unit (FRU) management that includes:
	  • Health monitoring
	  • Fault handling
	  • Systems alarms
	  • Notification by LEDs
	  • Power budgeting

	• Cooling • Management [See request chassis feb.]
IP tunneling	• Support for IPv4 and IPv6 unicast IP-over-IP (IPv4) filter based decapsulation. [See Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation.]
Juniper extension toolkit	• Support for Firewall Service and Interfaces Service JET APIs. [See Juniper Engineering Network website.]
L2 features	• Support for L2 control protocols: xSTP, LACP, and LLDP. [See Ethernet Switching User Guide.] • Support for these L2 features: • Enhanced L2 Software (ELS) • 802.1D • 802.1Q VLAN tagging • 802.1Q VLAN trunking • 802.1p • Routed VLAN interface (RVI) • MAC address aging configuration • Static MAC address assignment for an interface • Disable MAC learning [See Ethernet Switching User Guide.]

MPLS	RSVP-Traffic Engineering (TE) supports preempting secondary label-switched paths (LSPs) that are signaled but not active. [See RSVP Overview.]

Multicast	• Support for these multicast forwarding features:
	  • IPv4 and IPv6 multicast
	  • IGMP
	  • Multicast Listener Discovery (MLD) protocol
	  • Protocol Independent Multicast source-specific multicast (PIM SSM)
	  • Protocol Independent Multicast sparse mode (PIM SM)
	  NOTE: Support does not extend to features beyond those listed above. In this release, IGMP snooping, MLD snooping, multicast virtual private network (MVPN) Multicast, PIM multicast-only fast reroute (MoFRR), PIM first hop router (FHR), rendezvous point (RP), and last hop router (LHR) are not supported. In addition, the IRB interface is not supported as either a source or as a receiver. Make-before-break (MBB) is not supported for existing L3 aggregated Ethernet (AE) or LAG receivers (member addition/deletions or up/downs).
	  [See Multicast Overview.]
	• Support for these IGMP snooping with IRB features:
	  • IGMP V1, V2, and V3 with plain L2 snooping with IRB (only if device acts as an LHR)
	  • Proxy mode
	  • Enterprise-style CLI only
	  • Any-source multicast (ASM) and source-specific multicast (SSM) modes
	  Limitations include:
	  • IGMP group-specific queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
	  • MBB on existing L2aggregated Ethernet or LAG interfaces, including for member additions, deletions, and link up or down events.

	• All unregistered IPv4 and IPv6 multicast packets are forwarded to the multicast router interfaces in the VLAN, even if an interface is configured as a multicast router interface only for IGMP snooping. • Neither MLD snooping, PIM snooping, snooping with VPLS, EVPN-VXLAN,nor EVPN-MPLS is supported. [See IGMP Snooping Overview and Integrated Routing and Bridging.]
Network management and monitoring	• Support for sFlow monitoring technology. [See Understanding How to Use sFlow Technology for Network Monitoring.] • Support for port mirroring. [See Understanding Port Mirroring and Analyzers.]
Protection against distributed denial of service (DDoS) attacks	• Support for DDoS protection, which is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]
Routing policy and firewall filters	• L3 support for firewall filters and policers. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]

Routing protocols	• Support for Unified Forwarding Table (UFT) and L3 sub-interface features—The QFX5700 line of switches supports these UFT and L3 sub-interface features: The UFT feature enables you to allocate forwarding table resources to optimize the memory available for different address types based on the needs of your network. The UFT stores both the L2 and L3 entries that enable you to set default sizes for different entries. These applications share UFT search banks: • L2 MAC addresses. • IPv4 and IPv6 host routes. • L2 and L3 multicast routes. • IPv4 and IPv6 longest prefix match (LPM) or prefix entries that are supported in UFT through algorithmic longest prefix match (ALPM). [See forwarding-options.] • Support for redistribution of IPv4 routes with IPv6 next hop into BGP. [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]
Software installation and upgrade	• Support for zero-touch provisioning (ZTP) on the management and the WAN interfaces. [See Zero Touch Provisioning Overview.] • Support for secure boot—The implementation is based on the UEFI 2.4 standard. [See Junos OS Software Installation and Upgrade Guide.]

- **Support for the QSFP-100G-DR and QSFP-100G-FR transceivers (QFX5220-32CD and QFX5220-128C)**—Starting in Junos OS Evolved Release 21.2R1, the QFX5220-32CD and QFX5220-128C switches support the QSFP-100G-DR and QSFP-100G-FR transceivers.

  [See Hardware Compatibility Tool.]

Class of Service

- **Support for CoS features on EVPN VXLAN (QFX5130-32CD)—**Starting in Junos OS Evolved Release 21.2R1, QFX5130-32CD routers support defining classifiers and rewrite rules on leaf (initiation and terminations) and spine nodes for EXPN VXLANs. Support is also provided for defining schedulers, interpolated drop profiles, explicit congestion notification (ECN), and priority based flow control (PFC).

  [See CoS Support on EVPN VXLANs.]

EVPN

- **EVPN-VXLAN support with MAC VRF routing instances (QFX5130-32CD)**—Starting in Junos OS Evolved Release 21.2R1, QFX5130-32CD switches support EVPN-VXLAN for unicast traffic using MAC virtual routing and forwarding (VRF) routing instances. The switch can be a leaf device or a spine device in an edge-routed bridging or centrally routed bridging overlay design. You can configure multiple EVPN instances (EVIs) of type `mac-vrf`. Each EVI can support a different EVPN service type (`vlan-based`, `vlan-aware`, or `vlan-bundle`). The switch creates and uses one VTEP logical interface per remote provider edge (PE) device by default, regardless of the number of routing instances. This implementation improves VXLAN VTEP scaling with multiple routing instances.

  EVPN-VXLAN support on this switch includes:

  - EBGP and IBGP overlays
  - Single-homing and all-active multihoming
  - Layer 2 and Layer 3 unicast for IPv4 and IPv6 with ARP suppression
  - Proxy ARP and ARP suppression for Layer 2 and Layer 3 IRB traffic
  - Proxy NDP and NDP suppression for Layer 2 and Layer 3 IRB traffic
  - Service provider style configuration with Q-in-Q support, supported for Layer 2 gateways only and only with MAC VRF `vlan-bundle` service type

    NOTE: We support only the Q-in-Q use cases described in Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network.

  - Storm control
  - MAC mobility
  - Core isolation

- Proxy re-advertisement of EVPN Type 2 MAC+IP routes by all provider edge devices on the same ESI
- Virtual machine traffic optimization (VMTO)
- OSPF and BGP routing protocols on IRB interfaces

[See EVPN User Guide.]

- **Support for EVPN Type 5 routes (QFX5130-32CD)**—Starting in Junos Evolved OS 21.2R1, the QFX5130-32CD switches support EVPN type 5 routes. EVPN Type 5 routes support inter-subnet routing across datacenters by using IP prefix advertisement.

  [See Understanding EVPN with VXLAN Data Plane Encapsulation.]

- **Support for Port Mirroring in EVPN-VXLAN fabrics (QFX5130-32CD)**—Starting in Junos Evolved OS 21.2R1, the QFX5130-32CD switches support port mirroring in EVPN-VXLAN fabrics.

  [See How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics.]

- **Firewall filtering and policing (QFX5130-32CD)**—Starting in Junos Evolved OS 21.2R1, QFX5130-32CD switches support firewall filtering and policing in EVPN-VXLAN fabrics.

  [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]

High Availability

- **NSR support for controller-initiated RSVP LSPs (PTX10003, PTX10004, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 21.2R1, we support NSR for controller-initiated RSVP-based point-to-point (P2P) and point-to-multipoint (P2MP) LSPs. The primary Routing Engine synchronizes all RSVP LSPs initiated by Path Computational Elements (PCEs) with the backup Routing Engine. It also synchronizes multicast flow specifications for any PCE-initiated P2MP LSPs with the backup Routing Engine. This ensures zero traffic loss for the traffic carried over PCE-initiated RSVP LSPs during Routing Engine switchovers. You enable this feature when you configure NSR.

  [See PCEP Configuration.]

IP Tunneling

- **Support for IPv4 and IPv6 unicast IP-over-IP (IPv4) filter based decapsulation (QFX5130-32CD and QFX5220)**—Starting in Junos OS Evolved Release 21.2R1, we provide support for filter-based decapsulation for IP-over-IP and GRE tunnels. By using the feature, you can decapsulate any source IP address, destination IP addresses, and any range of prefixes. You can use the feature in in deployments where the transit router has to decapsulate the IP-over-IP packets and forward them based on their inner destination IP address. You can configure the feature for IP-over-IP and GRE tunnels by using the following configurations at the [edit] hierarchy level:

```
set firewall family inet filter <filter name> term <ipip-decap-term-name> from source-address <address> from destination address <destination address> from protocol ipip then decapsulate ipip routing-instance <name of the routing instance>

set firewall family inet filter <filter name> term <ipip-decap-term-name> from source-address <address> from destination address <destination address> from protocol gre then decapsulate gre routing-instance <name of the routing instance>
```

NOTE: Note: On QFX5220 both source and destinations IP addresses must be configured with /32 mask.

[See Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation and Example: Configuring Next-Hop-Based IP-Over-IP Dynamic Tunnels.]

Junos Telemetry Interface

NOTE: For Routing Engine telemetry sensors supported by this platform, see Telemetry Sensor Explorer. If any Platform Forwarding Engine sensors have been added for this release, they are listed below.

- **gRPC mutual authentication (PTX10001-36MR, PTX10003, PTX10008, and QFX5220)**—Starting in Junos OS Evolved Release 21.2R1, Junos telemetry interface (JTI), you can configure Junos telemetry interface (JTI) and Remote Procedure Calls (gRPC) to require client authentication as well as server authentication. Previously, only the client initiating an RPC request was able to authenticate the server; that is, a Juniper Networks device, using SSL certificates.

  [See Configuring Bidirectional Authentication for gRPC for Junos Telemetry Interface.]

Licensing

- **Juniper Agile Licensing (QFX5130-32CD)**—Starting in Junos OS Release Evolved 21.2R1, the QFX5130-32CD 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.

[See Flex Software License for QFX Switches, Juniper Agile Licensing Guide, and Configuring Licenses in Junos OS.]

MPLS

- **RSVP-TE supports preempting secondary LSPs that are signaled but not active (PTX Series and QFX Series)**—Starting in Junos OS Evolved Release 21.2R1, you can configure the hold priority of the secondary standby label-switched path (LSP) for RSVP-Traffic Engineering (RSVP-TE). The hold priority will be used to determine if the standby non-active LSP can be preempted. This will help to bring up non-standby secondary path LSPs with higher setup priority which are not able to come-up because of bandwidth crunch. To configure the non-active hold priority value for a secondary standby path, use the non-active-hold-priority statement at the [`edit protocols mpls label-switched-path <lsp-name>`] hierarchy level. You can set the priority from 0 through 7, where 0 is the highest priority and 7 is the lowest.

  [See RSVP Overview.]

Network Management and Monitoring

- **Support for sFlow technology with VXLAN (QFX5130-32CD)**—Starting in Junos OS Evolved Release 21.2R1, we support sFlow technology with VXLAN. sFlow is a monitoring technology for high-speed switched or routed networks. The sFlow agent performs packet sampling and gathers interface statistics, and then combines the information into UDP datagrams that are sent to sFlow collectors.

  [See sFlow Monitoring Technology.]

Network Security

- **Protect your network against DDoS attacks (QFX5130-32CD)**—Distributed denial-of-service (DDoS) attacks typically use a large number of network control packets to trigger exceptions in the network, consuming resources and crippling network operations. Starting in Junos OS Evolved Release 21.2R1, QFX5130-32CD switches have DDoS protection enabled by default. DDoS protection uses firewall filters and policers to discard or rate-limit control plane traffic. Thus, malicious traffic cannot cause device failure. You can disable DDoS protection or change default policer parameters for a protocol group or supported packet types in a protocol group.

  [See Distributed Denial-of-Service (DDoS) Protection Overview.]

Routing Protocols

- **Redistribution of IPv4 routes with IPv6 next hop into BGP (QFX Series)**—Starting in Junos OS Evolved Release 21.2R1, devices running Junos OS can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, IPv4 traffic is tunneled from CPE devices to IPv4-over-IPv6gateways. These gateways are announced to CPE

devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers, which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop.

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

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

System Management

- **Support for Precision Time Protocol (PTP) enterprise profile (QFX5130-32CD)**—Starting in Junos OS Evolved Release 21.2R1, the QFX5130-32CD switch supports the PTP enterprise profile, which is based on PTP version 2 (PTPv2). The PTP enterprise profile enables the enterprise and financial markets to add a timestamp to the operations of different systems and to handle a range of latencies and delays.

  Limitations

  - Because of hardware limitations, the interface et-0/0/32 is not used when the Precision Time Protocol (PTP) application is run.
  - Primary ports do not support two-step PTP.
  - The QFX5130-32CD does not support double-tagging over Layer 3.
  - The minimum packet rate for received sync packets is 8 pps. The client port recovers the primary clock only when the sync and delay-request packets are sent at a rate of not less than 8 pps.

  [See Understanding the Precision Time Protocol Enterprise Profile.]

What's Changed

IN THIS SECTION

Learn about what changed in this release for QFX Series Switches.

What's Changed in Release 21.2R1

IN THIS SECTION

Authentication and Access Control

- **Command to automate SSH key-based authentication (ACX Series, PTX Series, and QFX Series)**—You can set up SSH-key based authentication between the network device and a remote host by issuing the `request security ssh password-less-authentication` operational mode command. When you execute the command with the appropriate options, the device generates SSH keys for the current user, provided the user does not already have existing keys, and transfers the user's public key to the **authorized_keys** file of the specified user on the remote host.

  [See request security ssh password-less-authentication.]

General Routing

- **SSH session connection limit and rate limit per connection (PTX Series and QFX Series)**—We have introduced SSH `connection-limit` and `rate-limit` options at the **edit system services ssh** hierarchy levels to enable SSH connection limit and rate limit per connection. The default connection limit value is 75 connections and there is no default value associated with rate limit.

- **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.

- **Enhancement to the default `remnant-holdtime` (Junos OS Evolved platforms: ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**— Starting this release, the default `remnant-holdtime` has been increased from 180 seconds to 300 seconds. This provides sufficient time for protocols to start and sync routes from neighbors in a scaled

environment, during rpd restart. You can configure `remnant-holdtime` at the `edit routing-options forwarding-table` hierarchy level.

[See forwarding-table.]

- **Default FEC Settings (QFX5130-32CD, QFX5220-32CD, and QFX5220-128C)**—The default FEC mode for 4x25 optics is changed to FEC91 instead of FEC74. For 4x25G Direct Attach Copper Breakout Cables (DACBO), the default FEC mode remains as FEC74.

  [See show interfaces extensive.]

- **Enhancement to the show chassis pic command (Junos OS Evolved)**—You can now view additional information about the optics when you run the `show chassis pic` command. The output now displays the following additional field: MSA Version: Multi-source Agreements (MSA) version that the specified optics is compliant to. Values supported are: SFP+/SFP28 — SFF-8472 (versions 9.3 - 12.3), QSFP+/QSFP28 — SFF 8363 (versions 1.3 - 2.10), and QSFP-DD — CMIS 3.0, 4.0, 5.0. Previously, the show chassis pic command did not display this additional field.

  [See show chassis pic.]

- **Enhancement to the show interfaces (Aggregated Ethernet) command (PTX Series and QFX Series)**—When you run the `show interfaces extensive` command for aggregated Ethernet interfaces, you can now view following additional fields for MAC statistics : Receive, Transmit, Broadcast and Multicast packets.

  [See show chassis pic.]

- **Mozilla certification authority (CA) certificates removed (ACX Series, PTX Series, and QFX Series)**—To minimize security risks, Junos OS Evolved no longer includes Mozilla's set of root certificates from various CA operators by default. To use Docker container images from a registry that requires TLS authentication, you must first save the image as a tar archive on a remote device and then import the contents of the archive on the device running Junos OS Evolved.

  [See Running Third-Party Applications in Containers..]

Junos XML API and Scripting

- **Changes to how command-line arguments are passed to Python op scripts (ACX Series, PTX Series, and QFX Series)**—When the device passes command-line arguments to a Python op script, it prefixes a hyphen (-) to single-character argument names, and it prefixes two hyphens (--) to multi-character argument names. The prefix enables you to use standard command-line parsing libraries to handle the arguments. In earlier releases, the device prefixes a single hyphen (-) to all argument names.

  [See Declaring and Using Command-Line Arguments in Op Scripts.]

- **The `language python` statement is enabled by default (ACX Series, PTX Series, and QFX Series)**—The `language python` statement is configured by default in the `junos-defaults` configuration group on devices

running Junos OS Evolved. Thus, you can execute unsigned Python scripts using the default Python version without explicitly configuring the statement on the device.

[See Requirements for Executing Python Automation Scripts on Devices Running Junos OS.]

Layer 2 Features

- **Link selection support for DHCP (QFX Series)**—We've introduced `link-selection` statement at the `edit forwarding-options dhcp-relay relay-option-82` hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Earlier to this release, the DHCP relay drops packets during the renewal DHCP process as the DHCP Server uses the leaf's address as a destination to acknowledge DHCP renewal message.

  [See relay-option-82..]

Network Management and Monitoring

- **Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, PTX Series, and QFX Series)**—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:

  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.
  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and only emits an `<ok/>` or `<rpc-error>` element.

  [See Configuring RFC-Compliant NETCONF Sessions.]

- **Changes to how command-line arguments are passed to Python action scripts (ACX Series, PTX Series, and QFX Series)**—When a custom YANG RPC invokes a Python action script and passes command-line arguments to the script, the device prefixes a hyphen (-) to single-character argument names, and it prefixes two hyphens (--) to multi-character argument names. The prefix enables you to use standard command-line parsing libraries to handle the arguments. In earlier releases, the device passes the unmodified argument names to the script.

  [See Creating Action Scripts for YANG RPCs on Devices Running Junos OS and Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules.]

- **Changes in contextEngineID for SNMPv3 INFORMS (ACX Series, PTX Series, and QFX Series)**—Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  [See SNMP MIBs and Traps Supported by Junos OS.]

Platform and Infrastructure

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX 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).]

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Evolved Release 21.2R1 for QFX Series switches.

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

General Routing

- On the QFX5130-32CD and QFX5700 platforms, 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

- All stats in `show chassis fpc` will be missing as in following root@glenrothes-04> 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
- On QFX5700 seeing **BrcmPlusPfe: SI Set Failed , Port Multi Set Failed error** messages after de-channelization of 4x10G and 4X25G channelized interface. PR1569703
- On QFX5700 platforms local interface traffic statistics are not displayed for layer 2 interfaces or logical interfaces. PR1573443
- On QFX5700 platforms when configuration of multiple interfaces are modified in a single commit, any CLI command executed immediately to display the hardware/interface status will give an error message **error: communication failure with /re0/picd/** . There is however, no functional impact and output will be displayed after sometime. PR1574569
- On QFX5700 when the FEB card has a power failure the `show system alarms` CLI command will display the below error **PCI Uncorrected error on dev 0000:00:03.2** along with FEB alarm. PR1578066
- If the input voltage drops or surges, there will not be any alarm generated indicating the voltage is low or high within operating range. However, there is a PSM `psm number` Input Under Voltage Failure alarm generated when the PSM goes offline due to low voltage.PR1579105
- On QFX5130 and QFX5700, if packet hits both IPACL and IVACL MF classifiers, due to HW limitation classification will not happen properly. Classification behaviour is undefined in this case. PR1581917
- During FPGA firmware upgrade, if there is power fluctuation, the entire system goes for a reboot and the FPGA will fallback to the golden image. PR1586282

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 21.2R1 for QFX Series switches.

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

General Routing

- On QFX5220-128C, an interface might stay down for around 3 seconds after the interface is enabled or fiber is inserted, there is traffic loss due to this issue. PR1480112
- When the evo-pfemand process is restarted due to any error conditions, the channelized interfaces (if configured) will flap two times (once during the initialisation) and the second time after the configuration is programmed. PR1526984
- On all QFX Junos OS Evolved platforms, if scaled OSPF (Open Shortest Path First) neighbors (e.g 512 neighbors) are formed over IRB (Integrated Routing and Bridging) interfaces with layer 2 interfaces in the interface-mode trunk, the OSPF routes might be not learned from the neighbors. The issue results in traffic loss. PR1570498
- When optic is removed and FEB is not present, internal serdes are not powered off to reduce FPC power consumption in 21.4. PR1574330
- On QFX5700s, DAC connection at peer end will not go down while RCB is rebooted. PR1574342
- On QFX5700 platforms, system goes for reboot when faulty FEB is off-lined for recovery. PR1578090
- In QFX5220-32CD, rarely a picd daemon core seen after bootup. This doesn't affect the functionality as the device recovers itself automatically by restarting picd daemon.PR1583164
- On QFX5700 platforms few interfaces are not coming up after removing channelization through single commit that is by using **delete interfaces**. PR1592238
- On QFX5700 platforms, auto-negotiation is not supported for 100G DAC interfaces. PR1594135
- With QFX5220-32CD interconnected to PTX-10008 chassis line card, intermittently issue is seen where 100G SR4 optics link does not come up or flaps, and continuously increasing FEC errors are seen on PTX end. PR1597528

Infrastructure

- The alarm **Host 0 Active Disk Usage Exceeded** might be generated due to large files under /var/log/journal which were already marked as deleted in the `lsof` command. PR1601251

Routing Policy and Firewall Filters

- On QFX5700 when we issue `clear firewall all` command, policer statistics are not cleared. User defined policer statistics are removed. PR1581867

Resolved Issues

IN THIS SECTION

Learn about the issues fixed in this release for QFX Series switches.

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: 21.2R1

IN THIS SECTION

General Routing

- The **aggregate member links field** in the `show interfaces extensive` command output for an aggregated Ethernet interface does not reflect the number of member-links. PR1517841
- Ingress policer scale is limited to 128 due to a known issue in the Junos OS Evolved Release 20.3R1. PR1525525
- DHCPv6 relay might malfunction on QFX Series with Junos OS Evolved platforms. PR1545754

- Changing the mru and cable lengths to very low values can cause packets to drop. PR1547736
- SRTE might stay in the Up state when the routes are deleted through policy. PR1547933
- Traffic loss might happen if the flushing issue happens on the scale of ARP entries for the logical interface of the IRB interface. PR1554151
- **ICMP destination unreachable** message is not sent from the QFX5130 when a firewall filter action rejects the packet. PR1563404
- Routes learned through IRB interface might not be reachable in IBGP setup. PR1568566
- The BGP sessions might intermittently flap if the egress sFlow sampling is enabled at a high sampling rate. PR1571636
- ZTP must overwrite configuration derived from DHCP options with configuration from downloaded configuration file. PR1577004
- QFX5130 and QFX5220 object anomalies seen with the PTP TC configuration. PR1577375
- If you change the dynamic mode from DLB to non-DLB, the ECMP creation API does not work.PR1579245
- The traffic related to native VLAN might be dropped.PR1581075
- Port mirroring instance might be down on Junos OS Evolved based platforms. PR1593276
- The interface might not learn `mac-address` if it is configured with `vlan-id-list starting` with `VLAN id 1` and `native-vlan-id`. PR1597013
- The ndp process might reach to 100% and might result in traffic drop. PR1551644

Infrastructure

- ToS of self-initiated packets might change unexpectedly. PR1578247

Interfaces and Chassis

- Commit fails with error **VLAN-ID can only be specified on tagged ethernet interfaces**, while deleting Layer 3 sub-interface. PR1564703

User Interface and Configuration

- The port_speed configuration details not present in the picd configuration for ports et-0/0/128 and et-0/0/129. PR1510486

- The Junos OS Evolved operational state might be incorrect on the system and CoS scheduler configuration change might not take effect. PR1536615

Upgrade Your Junos OS Evolved Software

Products impacted: ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220-32CD, QFX5220-128C, QFX5700.

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

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

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

For more information about software installation and upgrade, see Software Installation and Upgrade Overview (Junos OS Evolved). For more information about EOL releases and to review a list of EOL releases, see https://support.juniper.net/support/eol/software/junosevo/.

Licensing

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:

  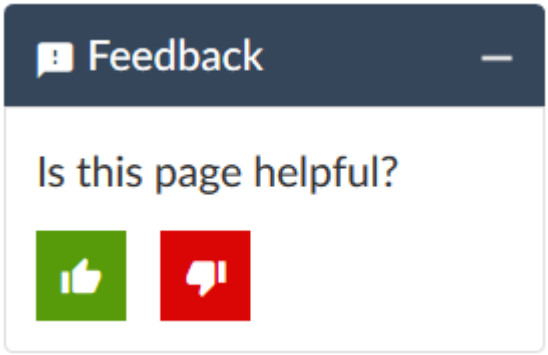

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- 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

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/content/dam/www/assets/resource-guides/us/en/jtac-user-guide.pdf.
- Product warranties—For product warranty information, visit https://support.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://support.juniper.net/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://supportportal.juniper.net/s/knowledge
- Download the latest versions of software and review release notes: https://support.juniper.net/support/downloads/
- Search technical bulletins for relevant hardware and software notifications: https://supportportal.juniper.net/s/knowledge
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://supportportal.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://support.juniper.net/support/requesting-support/
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

10 August 2023—Revision 14, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

6 July 2023—Revision 13, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

24 November 2022—Revision 12, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

29 July 2022—Revision 11, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

28 January 2022—Revision 10, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

20 January 2022—Revision 9, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

29 November 2021—Revision 8, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

13 October 2021—Revision 7, Junos OS Release 21.2R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

7 October 2021—Revision 6, Junos OS Release 21.2R1 for the ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

30 September 2021—Revision 5, Junos OS Release 21.2R1 for the ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

29 July 2021—Revision 4, Junos OS Release 21.2R1 for the ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

19 July 2021—Revision 3, Junos OS Release 21.2R1 for the ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

7 July 2021—Revision 2, Junos OS Release 21.2R1 for the ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.

30 June 2021—Revision 1, Junos OS Release 21.2R1 for the ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220 Devices.