

Release Notes

Published
2024-08-07

Junos[®] OS Evolved Release 21.4R1

SUPPORTED PLATFORMS

- ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700

SOFTWARE HIGHLIGHTS

- [Layer 2 VPN for FAT Label \(PTX10001-36MR, PTX10004, PTX10008, and PTX10016\)](#)
- [Export BGP Community and AS Path Information using IPFIX \(PTX Series\)](#)
- [MPLS Support for FTI tunnels \(PTX10001-36MR, PTX10004, PTX10008, and PTX10016\)](#)
- [GPS and BITS interface support \(PTX10008\)](#)
- [Layer 2 and layer 3 support for flood policers \(PTX10001-36MR, PTX10004, PTX10008, and PTX10016\)](#)

HARDWARE HIGHLIGHTS

- New ACX7509 Router

Table of Contents

Introduction | 1

Junos OS Evolved Release Notes for ACX7100-32C, ACX7100-48L, and ACX7509 Devices

What's New | 1

What's New in 21.4R1 | 2

Hardware | 2

EVPN | 14

Network Management and Monitoring | 14

Services Applications | 15

Additional Features | 15

What's Changed | 17

What's Changed in Release 21.4R1-S2 | 17

What's Changed in Release 21.4R1-S1 | 17

What's Changed in Release 21.4R1 | 18

Known Limitations | 20

Open Issues | 23

Resolved Issues | 31

Resolved Issues: 21.4R1 | 31

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Devices

What's New | 33

What's New in 21.4R1 | 34

Hardware | 35

Device Security | 38

EVPN | 39

High Availability | 39

Juniper Extension Toolkit (JET) | 39

Junos Telemetry Interface | 39

- Layer 2 VPN | 42
- Network Management and Monitoring | 43
- OpenConfig | 44
- Routing Policy and Firewall Filters | 45
- Routing Protocols | 46
- Services Applications | 47
- Software Installation and Upgrade | 47
- Additional Features | 47

What's Changed | 51

- What's Changed in Release 21.4R1-S2 | 51
- What's Changed in Release 21.4R1-S1 | 52
- What's Changed in Release 21.4R1 | 52

Known Limitations | 56

Open Issues | 58

Resolved Issues | 66

- Resolved Issues: 21.4R1 | 66

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

What's New | 75

- What's New in 21.4R1 | 75
 - What's Changed in Release 21.4R1-S2 | 76
 - What's Changed in Release 21.4R1-S1 | 76
 - Device Security | 77
 - EVPN | 77
 - Juniper Extension Toolkit (JET) | 78
 - Network Management and Monitoring | 78
 - Routing Policy and Firewall Filters | 78
 - Routing Protocols | 79
 - Software Installation and Upgrade | 79
 - Additional Features | 80

What's Changed | 81

- What's Changed in Release 21.4R1 | 81

Known Limitations | 84

Open Issues | 85

Resolved Issues | 86

| **Resolved Issues: 21.4R1 | 86**

Upgrade Your Junos OS Evolved Software | 87

Licensing | 88

Finding More Information | 88

Documentation Feedback | 89

Requesting Technical Support | 89

Revision History | 91

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 21.4R1.

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, ACX7100-48L, and ACX7509 Devices

IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 17](#)
- [Known Limitations | 20](#)
- [Open Issues | 23](#)
- [Resolved Issues | 31](#)

These release notes accompany Junos OS Evolved Release 21.4R1 for ACX7100-32C, ACX7100-48L, and ACX7509 routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- [What's New in 21.4R1 | 2](#)

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

What's New in 21.4R1

IN THIS SECTION

- [Hardware | 2](#)
- [EVPN | 14](#)
- [Network Management and Monitoring | 14](#)
- [Services Applications | 15](#)
- [Additional Features | 15](#)

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 21.4R1, 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 ACX7509 router (ACX Series)**—In Junos OS Evolved Release 21.4R1, we introduce the ACX7509 Router, a high-performance multiservice router designed to address the needs of emerging 5G, IoT, and low latency virtualized edge services. The ACX7509 routers enable a wide variety of cloud-ready port speeds, making it ideal for metro high-end aggregation, large enterprise, and data center applications.

The ACX7509 provides several capabilities that include advanced timing and synchronization, Control and User Plane Separation (CUPS), network slicing, and a comprehensive set of features that are suited for cloud metro architectures.

With a 5-U modular and energy efficient design, the ACX7509 has a 4.8-Tbps forwarding-capacity ASIC, with over-subscription support and Flexible PIC Concentrators (FPCs) with 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE, 200GbE, and 400GbE port configurations.

We ship the ACX7509 routers with redundant power supply modules (PSMs). You can order the ACX7509-PREMIUM configuration routers with AC or DC power supply modules (PSMs).

Table 1 on page 3 summarizes the descriptions of the features either supported on or added to the ACX7509 platform in Junos OS Evolved Release 21.4R1.

Table 1: Features supported on or added to the ACX7509

Feature	Description
Chassis	<ul style="list-style-type: none"> • Support for new line cards (ACX7509): <p>We introduce the new ACX7509-FPC-4CD, ACX7509-FPC-16C, and ACX7509-FPC-20Y line cards. The line cards come in 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE, 200GbE, and 400GbE port configurations and have Precision Time Protocol (PTP), SyncE (Class C), and Media Access Control Security (MACsec) support.</p> • The ACX7509-FPC-16C line card has 16 QSFP/QSFP28 ports that support 40-Gbps and 100-Gbps speeds. Default port speed is 100-Gbps. The line card has a throughput of 1.6 Tbps when installed in FPC slots 1 and 5; 800 Gbps of throughput when installed in other FPC slots. The line card supports all 16 ports when it is plugged into FPC slot 1 and 5. But when it is plugged to other FPC slots, only ports 0,1 4,5,8,9 and 12, 13 are active. <ul style="list-style-type: none"> • You can channelize the ports operating at: <ul style="list-style-type: none"> • 100-Gbps speed to 4x25-Gbps speed • 40-Gbps speed to 4x10-Gbps speed <p>You can configure four 25-Gbps speed and four 10-Gbps speed breakout cables at 0,1,4,5, 8,9 and 12,13 ports.</p>

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
	<ul style="list-style-type: none"> • The ACX7509-FPC-4CD line card has 4 QSFP56/QSFP56-DD ports that support 100-Gbps, 200-Gbps, and 400-Gbps speeds. Default port speed is 400 Gbps. The line card has a throughput of up to 1.6 Tbps when installed in FPC slots 1 and 5 and is supported only in these FPC slots. You can channelize the ports operating at 400-Gbps speed to 4x100-Gbps speed. You can configure four 100-Gbps speed and two 100-Gbps speed breakout cables at 0,1,2,3 ports. • The ACX7509-FPC-20Y line card has 20 SFP/SFP+/SFP28/SFP56 ports that support 1-Gbps, 10-Gbps, 25-Gbps, and 50-Gbps speeds. Default port speed is 25 Gbps. 1-Gbps is supported in all slots except FPC slots 1 and 5. 50-Gbps is supported only in FPC slots 1 and 5. The line card has a throughput of up to 1 Tbps when installed in FPC slots 1 and 5; 500 Gbps of throughput when installed in other slots. Channelization is not supported on this line card. <p>[See Configuring Port Speed at Port Level.]</p> <ul style="list-style-type: none"> • New ACX7509-FEB Forwarding Engine Board (FEB)—In Junos OS Evolved 21.4R1, we introduce the ACX7509-FEB, which consists of two forwarding ASIC's for packet processing and traffic management. Each forwarding ASIC is capable of supporting 2.4 Tbps with entire FEB support 4.8 Tbps throughput, with oversubscription upto 6 Tbps. You can install one or two FEBs on the ACX7509 router. The ACX7509-FEB supports 1+1 redundancy. • Support for SFP56, QSFP28, and QSFP56-DD ports on the 50-Gbps (20 ports), 100-Gbps (16 ports), and 400-Gbps (4 ports) FPCs respectively.

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
	<ul style="list-style-type: none"> • New ACX7509-RCB Routing and Control Board (RCB)— We introduce the ACX7509-RCB, which consists of Intel Xeon D-1600 processor, 64-GB DDR4 RAM, and two 100-GB SATA SSDs. The ACX7509 router supports two ACX7509-RCBs. ACX7509-RCB is an integrated board and a single field-replaceable unit (FRU) that provides Routing Engine and Control Board (CB) functionality. The Routing Engine performs all route-processing functions, whereas the CB performs chassis control and management plane functionality. The RCB provides control plane functions. You can install one or two RCBs on the ACX7509 router. • Platform software support: Platform software support includes chassis management and power management for FPC and FEB. We support the following power management functions: <ul style="list-style-type: none"> • Secure boot support for chassis. • FPC will be online based on available power, power ratings, and priority sequence configured by user. • You can restart or make an FEB online or offline using the request chassis feb slot command. • You can configure the power-on sequence for FPC using the set chassis fru-poweron-sequence command. The default priority sequence for FPCs is 0-7 where 0 is the highest priority. <p>[See request chassis feb slot, set chassis fru-poweron-sequence, and show chassis feb.]</p>

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
	<ul style="list-style-type: none"> • Support for secure boot and secure BIOS in Intel Xeon D-1600 series processors being used on the RCB. • Support for RCB infrastructure, peripherals, reachability, resiliency, and CLI commands. [See show chassis routing-engine.] • Fault detection, logging, recovery, and debugging support for platform (FEB, FPCs, PSUs, Fan Tray Controllers (FTCs)) resiliency. [See show system alarms.]
Class of Service	<ul style="list-style-type: none"> • Support for QoS classification and rewrites for all types of interfaces (MPLS, virtual private LAN service (VPLS), Layer 3 VPN, L2CKT, CCC, and IRB). [See CoS on ACX Series Routers Features Overview.]
Class of Service	<ul style="list-style-type: none"> • Support for configuring CoS-specific buffer sizes, deep buffers, and drop profiles (tail drop and weighted random early detection (WRED)). [See Shared and Dedicated Buffer Memory Pools on ACX Series Routers.] • Support for port shaping and scheduling with eight virtual output queues (VoQs) per port and two scheduling priority levels (strict-high and low). Queue management with multiple strict high-priority queues (round-robin scheduling), multiple low-priority queues (weighted fair queuing (WFQ) scheduling), and default deep buffers. [See Schedulers Overview for ACX Series Routers.] • Support for CoS classifiers, forwarding classes, and rewrite rules. [See CoS on ACX Series Routers Features Overview.]

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
Dynamic Host Configuration Protocol	<ul style="list-style-type: none"> • DHCP server, client, and relay agent for DHCPv4 and DHCPv6. [See DHCP Overview.]
Ethernet Switching and Bridging	<ul style="list-style-type: none"> • Support for Ethernet ring protection switching (ERPS) G.8032v2. [See Ethernet Switching User Guide.] • Support for STP, Multiple Spanning Tree Protocol (MSTP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), Storm Control, and No Local Switching (NLS). [See Spanning-Tree Protocols User Guide.]
IP Tunneling	<ul style="list-style-type: none"> • Support for Layer 2 protocol tunneling (L2PT) with service provider-style configuration. [See Layer 2 Protocol Tunneling.]
Layer 2 VPN	<ul style="list-style-type: none"> • Support for Layer 2 VPN and Layer 2 Circuits. [See Understanding Layer 2 VPNs and Configuring Interfaces for Layer 2 Circuits.] • L2 bridging support for E-LINE, E-LAN, E-ACCESS, and E-Transit services. [See Layer 2 Bridge Domains on ACX Series Overview.]

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
MPLS	<ul style="list-style-type: none">• Multiprotocol Label Switching (MPLS) support<ul style="list-style-type: none">• LDP Downstream-on-Demand• LDP over RSVP-TE areas• CSPF/single-area CSPF for RSVP-TE• RSVP-TE LSP over inter IGP areas• SRLG for RSVP LSPs• MPLS label stacking• LSP protection• Inter-AS RSVP-TE LSP <p>[See MPLS Overview]</p>

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
Multicast	<ul style="list-style-type: none"> • Ingress and egress multicast replication— Support for ingress and egress multicast replication using the Recycle replication model where the multicast and broadcast packets destined for their local ports are replicated. There are two passes for each multicast packet. In the first pass, each packet is replicated to each egress PFE in the system using recycle port of the egress PFE. In the second pass, packets are replicated to the ports of the egress PFE. <p>Use the set <code>ingress-only</code> statement at the system <code>packet-forwarding-options mcast-repl-type</code> hierarchy level for the ingress replication model and the set <code>ingress-egress-recycle</code> statement at the system <code>packet-forwarding-options mcast-repl-type</code> hierarchy level for ingress and egress multicast replication using the Recycle replication model. If you don't select the replication type, ingress replication (<code>ingress-only</code>) is considered as the default replication type.</p> <p>[See Ingress and Egress Multicast Replication using Recycle Replication Model, ingress-egress-recycle, and ingress-only.]</p> <ul style="list-style-type: none"> • CoS support for ingress-egress-multicast replication—By default, the recycle interface bandwidth is based on chassis oversubscription and platform-specific FPC or interface configuration. Starting in Junos OS Evolved Release 21.4R1, you can configure the bandwidth of the recycle interface to enable: <ul style="list-style-type: none"> • Equal distribution of the recycle interface bandwidth at the PFE level. • A percentage of the recycle interface bandwidth for ingress-egress-multicast replication.

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
	<p>To configure the recycle interface bandwidth, use the following CLI commands:</p> <ul style="list-style-type: none"> • set system packet-forwarding-options internal-forwarding-bandwidth <value in Gbps> • set system packet-forwarding-options internal-forwarding-bandwidth <value> reserve ingress-egress-mcast <percentage> <p>[See Multicast Ingress-Egress Replication.]</p> <ul style="list-style-type: none"> • IGMP snooping and Multicast Listener Discovery (MLD) snooping support for L2 multicast. [See Configuring IGMP and MLD.] • Support for IPv4 and IPv6 L3 multicast protocols: PIM, IGMP (v1, v2, v3), and MLD (v1, v2). [See Multicast Overview.]
Network Management and Monitoring	<ul style="list-style-type: none"> • Support for the following RFC4293 tables: <ul style="list-style-type: none"> • icmpStatsTable—Generic system-wide ICMP counters • icmpMsgStatsTable—System-wide per-version and per-message type ICMP counters <p>[See Standard MIBs for Junos OS Evolved.]</p>

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
Operations, Administration, and Maintenance	<ul style="list-style-type: none"> • EOAM support: <ul style="list-style-type: none"> • Connectivity Fault Management (CFM) and Performance Monitoring (Y.1731) protocols over Ethernet interfaces for Bridge and INET services. • CFM (Up MEPs) for Layer 2 Circuits. • Ethernet in the first mile (Down MEPs). <p>[See Ethernet OAM Connectivity Fault Management.]</p> • BFD support (single hop, multihop) and micro BFD support for base L3 Protocols (non-MPLS). [See Bidirectional Forwarding Detection (BFD).] • BFD support for MPLS: BFD for LSP (LDP/RSVP). [See Bidirectional Forwarding Detection (BFD) for MPLS.]
Protection Against DDOS Attacks	Support for control plane distributed denial-of-service (DDoS) protection. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview .]

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
Routing Protocols	<ul style="list-style-type: none"> • Support for the following Layer 3 features: <ul style="list-style-type: none"> • IP forwarding and exception packet handling • IEEE 802.1Q support (VLAN trunking) on IRB interfaces • ARP, ND, uRPF, and ECMP • IGP protocols: OSPF, RIP, and IS-IS • Support for BGP multipath, BGP-LU, and BMP. [See Understanding BGP.] • VRRP support for IPv4 and IPv6. [See VRRP and VRRP for IPv6 Overview.]
Routing Policy and Firewall Filters	<ul style="list-style-type: none"> • Firewall filters and policer support for family types inet, inet6, bridge, vpls, ccc, mpls, and any. [See Firewall Filters Overview.] • Support for configuring firewall filters with up to 16,000 terms. [See Understanding Firewall Filter Match Conditions.] • Support for configuring multifield classifiers and forwarding table filters (FTFs). • Support for loose/strict mode unicast RPF. [See Understanding Unicast RPF (Routers).] • Multiple database profile support. [See hw-db-profile.]
Software Installation and Upgrade	<ul style="list-style-type: none"> • Zero-touch provisioning (ZTP) support for WAN interfaces or management interfaces. [See Zero Touch Provisioning.]

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
Source Packet Routing in Networking (SPRING) or segment routing	<p>Segment routing support:</p> <ul style="list-style-type: none"> • Segment routing global block (SRGB) for OSPF and IS-IS • Fast reroute (FRR) • Metro Ethernet services over segment routing infrastructure • Static SR (node segment, prefix segment, adjacency, and anycast segments) for OSPF and IS-IS • SR services—L3VPN, 6VPE, and 6PE, L2Circuit, L2VPN, and VPLS • Topology independent loop-free alternate (TI-LFA) with segment routing for OSPF and IS-IS <p>[See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS, Understanding Source Packet Routing in Networking (SPRING) and Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING.]</p>

Table 1: Features supported on or added to the ACX7509 (Continued)

Feature	Description
VPNs	<ul style="list-style-type: none"> • L2 VPN support with VPLS over MPLS. [See Introduction to VPLS.] • Support for the following Layer 3 VPN features: <ul style="list-style-type: none"> • Per prefix labels and per VRF table labels for MPLS VPN • VRF load balancing using ECMP • Layer 3 VPN LSP protection using FRR • Export and import of routes between local VRFs (VRF to VRF) and global VRF • PE-CE route exchange using Static, eBGP, IS-IS, OSPF, and RIP protocols • IPv6 over MPLS (6PE and 6VPE) • Service provider chaining using Inter AS options A, B, and C <p>[See Understanding Virtual Routing and Forwarding Tables.]</p>

EVPN

- **DHCP relay in EVPN-VXLAN (ACX7100)**—Starting in Junos OS Evolved Release 21.4R1, we support DHCP relay in an EVPN-VXLAN fabric with edge-routed bridging overlays. Support for DHCP relay includes DHCPv4 and DHCPv6.

[See [DHCP Relay Agent over EVPN-VXLAN.](#)]

Network Management and Monitoring

- **SNMP MIB support for ICMP (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX1003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 21.4R1, we've added support for the following RFC4293 tables:
 - icmpStatsTable—Generic system-wide ICMP counters

- icmpMsgStatsTable—System-wide per-version and per-message type ICMP counters

[See [Standard MIBs for Junos OS Evolved](#).]

Services Applications

- **TWAMP IPv6 address support (ACX7100)**—Starting in Junos OS Evolved Release 21.4R1, you can specify IPv6 source and target addresses for Two-Way Active Measurement Protocol (TWAMP) clients, control connections, and test sessions.

[See [Understand Two-Way Active Measurement Protocol](#).]

Additional Features

We've extended support for the following features to these platforms.

- **Hotfix support (ACX7100-32C and ACX7100-48L)**—Support for hot fix using a delivery mechanism of full .iso images.

[See [request system software add](#) and [show system software list](#) .]

- **IPsec support for OSPFv2 and OSPFv3 (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130-48C, and QFX5220)**

[See [Overview of IPsec](#), [Configuring OSPF Authentication](#), and [Configuring IPsec Security Associations](#).]

- **Layer 3 Gateway Encapsulation support for Layer 3 over an IPv4 underlay (ACX7100-48L)**—Junos OS Evolved Release 21.4R1 provides an L3 gateway that can route packets across subnets. At the ingress VXLAN tunnel endpoint (VTEP), L3 traffic entering from a routing interface is routed to the VXLAN tunnel. At the egress VTEP, the VXLAN tunnel is terminated and the forwarding decision is determined based on the inner header. L3 VXLAN gateways are configurable as both centralized and distributed gateways.

[See [Working with VXLANs](#).]

- **Metro Ethernet services (ACX7100-32C and ACX7100-48L)**
 - IRB interface support for IPv4 and IPv6
 - IPv6 multicast protocol support (FHR, LHR, and RP roles) for any-source multicast (ASM) and source-specific multicast (SSM) modes

NOTE: Ensure that you enable L3-xl or balanced-exem profile for IPv6 multicast protocol to function.

[See [Multicast Overview](#) and [Understanding Integrated Routing and Bridging](#).]

- **Storm control support in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) overlay network (ACX7100-48L)**

[See [Understanding Storm Control](#).]

- **Support for EVPN-VXLAN Active/Active multihoming with Ethernet Switch Identifier for Link Aggregation Groups (ESI-LAG) (ACX7100-48L)**—Junos OS Evolved Release 21.4R1 introduces EVPN control plane support for BGP and VXLAN with VNI (VXLAN Network Identifier) segment ID encoding for L2 and L3 gateway unicast functionality. This feature supports Active/Active multihoming for servers and nodes connecting to multiple leaf nodes. You can configure multihomed Ethernet segments configured on the physical interface (IFD).

[See [EVPN VXLAN User Guide - Leaf Device Forwarding](#).]

- **Supported transceivers, optical interfaces, and DAC cables (ACX Series, PTX Series, and QFX Series)**—Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
-
- **VXLAN Layer 2 gateway with encapsulation support over an IPv4 underlay (ACX7100-48L)**—Junos OS Evolved Release 21.4R1 includes VXLAN L2 gateways that provide L2 connectivity between leaf devices separated by an underlay Layer 3 network.

NOTE: The CLI command `set protocols evpn default-gateway no-gateway-community` is a required configuration.

[See [Layer 2 Feature Guide for Routing Devices](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 21.4R1-S2 | 17](#)
- [What's Changed in Release 21.4R1-S1 | 17](#)
- [What's Changed in Release 21.4R1 | 18](#)

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

What's Changed in Release 21.4R1-S2

IN THIS SECTION

- [User Interface and Configuration | 17](#)

User Interface and Configuration

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after rollback timeout.

What's Changed in Release 21.4R1-S1

IN THIS SECTION

- [User Interface and Configuration | 17](#)

User Interface and Configuration

- When you configure `max-cli-sessions` at the **edit system** hierarchy level, it restricts the maximum number of cli sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new

CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login.

The `max-cli-sessions` is configured so you can control the memory usage for the CLI. You may set the `max-cli-sessions` per your requirement. However, if `max-cli-sessions` is not configured, there is no control on the number of CLIs getting invoked.

What's Changed in Release 21.4R1

IN THIS SECTION

- [Class of Service | 18](#)
- [EVPN | 18](#)
- [General Routing | 18](#)
- [Interfaces and Chassis | 19](#)
- [Junos OS API and Scripting | 19](#)
- [Network Management and Monitoring | 19](#)
- [Routing Protocols | 20](#)

Class of Service

- Junos OS Evolved now correctly displays the index for `show class-of-service` commands.

EVPN

- **Output for `show Ethernet switching flood extensive`**—The output for `show ethernet-switching flood extensive` now displays the correct next-hop type for Virtual Ethernet and WAN mesh group in an EVPN-VXLAN network as `unilist`. Previously, the output for `show ethernet-switching flood extensive` would misidentify the next-hop type as `composite`.

General Routing

- In Junos OS Release 21.4R1 and later, `ssh` is enabled by default on all the routers with VM host support.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Previously, the source and destination address pair had to be unique across all encapsulation type.

Junos OS API and Scripting

- **Limits increased for the `max-datasize` statement (ACX Series, PTX Series, and QFX Series)**—The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

Network Management and Monitoring

- **SNMP support for MIB**—Operational command `show snmp mib walk system` now shows the latest software version and does not show the build date.

[See [show snmp mib](#).]
- **Change in behavior of SNMP MIB object `ifAlias`**—SNMP MIB object `ifAlias` now shows the configured interface alias. In earlier releases, `ifAlias` used to show configured interface description.
- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, PTX Series, and QFX Series)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

[See [ssh \(NETCONF\)](#).]
- **The configuration accepts only defined identity values for nodes of type `identityref` in YANG data models (ACX Series, PTX Series, and QFX Series)**—If you configure a statement that has type `identityref` in the corresponding YANG data model, the device accepts only defined identity values (as defined by an `identity` statement) as valid input. In earlier releases, the device also accepts values that are not defined identity values.

Routing Protocols

- The RPD_OSPF_LDP_SYNC message not logged? On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the hold-time for ldp-synchronization at the **edit protocols ospf area area id interface interface name** hierarchy level less than three minutes. The message is printed after three minutes.
- To achieve consistency among resource paths, the resource path `/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter ip-addr='address'/state/countersname='name'/out-pkts/` is changed to `/mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/counters name='name'/`. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").

Known Limitations

Learn about limitations in Junos OS Evolved 21.4R1 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.

General Routing

- In some corner cases traffic is not scheduled equally between strict priority queues. This can happen in the following scenario: Priority queue configured and is completely utilizing the bandwidth and remaining queues are starved and traffic completely drops on those queues. In this state if we configure a second strict high priority queue traffic is not scheduled equally between strict priority queues. This is hardware specific issue, Q2C specific. [PR1577035](#)
- VOQ buffer in Packet Forwarding Engine has been changed from LOW to MID for some protocols. Because of this, frames do not drop as soon as ingress traffic rate crosses the configured pps rate. However, the rate towards CPU is still maintained at configured pps rate although violation counters do not increment. And this is evident from the DDoS Routing Engine-based statistics. [PR1592370](#)
- When PTP is enabled, ACX7100-32C port 31 is configured to 1G. A vendor's J2 only supports to 10G, so another vendor's gearbox converts 10G to 1G by rate matching. However, there is one limitation: it can only support frame size of 2000 bytes officially, so jumbo frame cannot be supported on port 31 when PTP mode is enabled. [PR1593015](#)

- Degradation is due to both unit programming during synchronous install calls. Therefore, programming happens to both units one after the other. [PR1597191](#)
- ACX7509, ACX7100-32C and ACX7100-48L do hardware timestamping. However, the peer device (in this case MX Series) does software timestamping, which takes more time resulting in high delays. This issue is seen in all the cases where peer setup does software time stamping. [PR1599777](#)
- 1M MPLS label scale is not possible with symmetric traffic testing (where encapsulation and decapsulation are scaled together). In this case, the encapsulation is limited by FEC hardware resource. 52000 scale is tried and tested with symmetric approach. [PR1600545](#)
- The queue statistics are not appropriate in the following scenarios:

1. When DRAM block is asserted in case of continuous congestion. DRAM block is a mechanism to stop traffic from going to DRAM when it is oversubscribed. This prevents packet stuck in the DRAM forever. (DRAM block feature is enabled by default in the ASIC). For this issue we have provided a CLI to disable the DRAM block. You might not see the issue after disabling it.
2. Packet deleted in VoQ due to non-arrival of credit (Watchdog timer flushing the VoQ).

Per the vendor, packet deleted in VoQ due to non-arrival of credit (Watchdog timer flushing the VoQ) is not counted as dropped packet. This can happen in two scenarios:

- The bandwidth is completely configured on a few queues and the other queues do not get any credit.
- When priority queue is configured and only priority traffic gets scheduled and the other queues do not get any credit and no shaper on priority queue.

In scenario b) the statistics on the SHP also fluctuates. We recommend to configure shaper on priority queue to prevent starvation of other queues. [PR1601320](#)

- ACX7509 has a mesh fabric and it is not a real fabric. There is a flow control mechanism from fabric to CGM. In case of fabric congestion there is flow control and back pressure to the VOQs and packets are dropped in VOQs. These VOQs cannot be mapped directly to fabric congestion as VOQ drops might occur for different reasons. There is no such specific dedicated registers or counters from BCM to provide values for fabric drops. As of now its accounted as queue drops as fabric is back pressuring ingress. [PR1601332](#)
- In case of scaled scenarios with route table scale reaching upto their full limit (close to 1M), a few triggers which reprogram the routing table might result in route entries getting missed in hardware. This happens in a worst-case scenarios of prefix distribution. In real-time scenario this issue must not be seen while the internet route scale is around 800,000 which must hold good. [PR1604034](#)
- This is the common issue of Junos OS Evolved platform for 400G DAC. On picd restart, the first pcsE has no Xcvr information, and the default Xcvr type is optics, so pfemand configures the port as 400G optics and turn off AN, later on. When the second pcsE arrives, it carries the correct Xcvr

information, and turns on AN again. This might flap one link. In the early stage of picd restart, the underlying link is still up, this updates to picd inevitably. Then the first PcsE event comes, and link is down. After the second PcsE event comes, link is up. [PR1604454](#)

- Route installation (/32) routes in Packet Forwarding Engine takes longer time to install in LPM, when LEM table limit is reached in the hardware. This issue is observed, only when LEM scale is lesser and LPM scale is on the higher side. when LEM reports 'Table Full', it takes around 1ms to install 1 route in hardware. Since we have around 2.3M routes to be pushed under this condition, the LPM installation will take longer time. [PR1604954](#)
- PTP to PTP noise transfer fails for frequencies 1. 0.03125 HZ 2. 0.123125 HZ. [PR1608786](#)
- syncE to PTP and syncE to 1pps noise transfer tests fails for frequencies 1. 0.00781 HZ 2. 0.01563 HZ 3. 0.03125 HZ 4. 0.06156 HZ 5. 0.12313 HZ. [PR1608866](#)
- On ACX7100-48L, enabling or disabling of PTP TC or BC causes all the interfaces to flap at the same time. [PR1609927](#)
- If more than 400 PTP clients are connecting to a single interface configured for PTPoE, PTP delay statistics might not be collected correctly. [PR1611130](#)
- Fabtoken is not used by ACX Series platforms except multiline card where similar mechanism of pfetoken is used instead of fabtoken. [PR1614670](#)
- Egress firewall ACL has limited lookup width available, so **family-any** and **family ethernet-switching/ccc/vpls/inet6 filters** cannot co-exist. [PR1617952](#)
- In a system with different type of interfaces (IRB, ccc, rest of interfaces), classifier might not work in specific scenarios. [PR1619441](#)
- Egress firewall ACL rules with count or policer actions might be skipped for BUM traffic and might not hit. This is because of ASIC limitation for BUM traffic and applies to all services like L2, L3, EVPN etc. [PR1621609](#)
- The learning rate of ACX7509 is same as ACX7100 when the host routes /128 routes are within the limit of ACX7100 scale (375K) or LEM table size limit. More than 50% degradation is observed in Fib learning rate on IPv6 host routes (/128) as compared to ACX7100, only when the scale exceeds 375K routes . [PR1624365](#)
- The error message "ACX7509: HWD_FIRMWARE_VERSION_READ_ERROR_NOTICE: Failed to get firmware version for fru: Feb, slot: 1, device: Dpll - Error message " is to be seen, when the feb is offlined and the dpll version is tried to be read. From Cli the FEB slot1 DPLL version is properly read later. FEB 1 DPLL 9 600 6022 OK. [PR1625324](#)
- CCM transmission keeps happening (for the UP MEP aggregated Ethernet case) during FPC restart. This is because, aggregated Ethernet interface might not be deleted during the FPC restart. Therefore, the session information that is already programmed in the hardware will not get deleted.

As a result, during the FPC restart case, the CCM transmission keeps happening. Due to this reason, sessions are up on the peer. Whereas in case of physical interface, the sessions will get deleted and added back during FPC restart. [PR1627369](#)

- Ingress replication is default for multicast. Packet Forwarding Engine can replicate upto 2.4T which is the chip capacity. Inter-Packet Forwarding Engine replication is limited to 2.1T due to fabric bandwidth limitations. Therefore, with ingress replication, it is not possible to achieve total egress traffic across ports more than 2.1T. [PR1629280](#)
- 2M is not achievable without errors, 1.5M routes are the appropriate value for the scaled scenarios. [PR1630088](#)
- The fields in output of the CLI `show pfe statistics traffic`, which are not supported in ACX7509, cannot be hidden, as this is a design limitation. The fields continue to show zero value and can be ignored. There is no impact on any feature. [PR1630562](#)
- Maximum support logical interface is 16000. Each local XC needs 2 logical interfaces. Therefore, achieved scale will be 8000. [PR1636947](#)
- Maximum supported logical interface is 16000. Each pseudowire require 1 logical interface. Therefore, 16000 pseudowire scale will be achieved. [PR1637066](#)
- Configd Object-info anomalies are seen after you commit the `show platform object-info anomalies app configd` configuration statement. [PR1635668](#)

Routing Protocols

- IPv6 link-local based BFD sessions work in a centralised BFD mode in ACX7509. So we can not have aggressive BFD timers supported for IPv6 link-local based BFD session. [PR1599658](#)

Open Issues

IN THIS SECTION

- [EVPN | 24](#)
- [Fault Management | 24](#)
- [General Routing | 24](#)
- [Interfaces and Chassis | 30](#)
- [Layer 2 Features | 30](#)
- [User Interface and Configuration | 30](#)

Learn about open issues 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.

EVPN

- On Junos OS Evolved platforms, if MAC-VRF instances with VXLAN encapsulation associated with IRB Logical interfaces in primary instance (and not in a separate VRF) and ospf is enabled on all interfaces, there could be a next-hop loop where remote PE loopback IP first resolves over IP underlay, remote IRB IP is reachable over the VXLAN tunnel to remote PE loopback and then remote PE loopback is resolved due to lower OSPF metric over IRB. Configure `protocols ospf interface irb.x passive statement remote PE loopback isn't learnt over IRB adjacency`. Typically, IRB logical interfaces are added to a VRF, which prevents this loop. [PR1626574](#)

Fault Management

- On FPC log, when the system is up, the Packet Forwarding Engine error message `Jexpr: JexprHandleDdos faile to update plct pfe:1f proto:0x8400` will be seen. This error message is just a debugging error message, there is no functional impact with this log error messages. [PR1610764](#)

General Routing

- There is no impact on these traces. Here are some RPC calls which throws error when there is no handler. This does not have any functionality impact. [PR1589410](#)
- 4x400G FPC is supported on FPC slot 1 and FPC slot 5 only as per design. 4x400G FPC plugged into any slot might bring FPC online but links come up only in FPC1 or FPC5. When 4x400G FPC is plugged into other slots, it does not allow the FPC to come online and an alarm is raised. [PR1582183](#)
- ACX7509: some of the interfaces from 16x100G and 20XSFP56 will not go down after evo-pfemand restart. [PR1592388](#)
- Error messages are observed while accessing the supercon scratch pad registers on bootup and there is no functional impact. [PR1594136](#)
- On Junos OS Evolved platforms, if a firewall filter has both Source Class Usage (SCU) or Destination Class Usage (DCU) matches and forwarding-class or loss-priority matches used on the same term in

the 'from' action, the traffic supposed to be subjected to the firewall filter might not get subjected to it. [PR1595788](#)

- When you restart FPC with scale configuration or routes or MAC entries, picd might generate a core file. PICD backup automatically and user intervention is not required. [PR1595861](#) and [PR1602352](#)
- The timingd application cannot be successfully restarted for ACX7100-32C and ACX7100-48L in the first G.8275.1 software release. [PR1597120](#)
- Currently, support for load balancing is on ECMP and not on ECMP FRR. Behaviour is the same as in ACX7100. [PR1599752](#)
- RPF fail counters reset on family modification is expected due to product limitations. [PR1598135](#)
- ACX7509 will do hardware timestamping, but peer device (in this case MX) does the software timestamping, which takes more time resulting in high delays. This issue will be seen in all the cases where peer setup does software timestamping. [PR1599777](#)
- With IGMPv3 reports received at a higher rate more than 1600 pps, packets are dropped due to control plane rate limit. Therefore, it is not possible to form 2,56,000 IGMP groups. Need to tweak DDOS configuration for reaching 2,56,000 IGMP groups. [PR1599998](#)
- In scaled environment, the interfaces do not come up during FEB offline - online since the Packet Forwarding Engine state is not transitioning to online though FEB has transitioned to online state. [PR1601158](#)
- ACX7509 has a mesh fabric and it is not a real fabric. There is a flow control mechanism from fabric to CGM. In case of fabric congestion there is flow control and back pressure to the VOQs and packets are dropped in VOQs. These VOQs cannot be mapped directly to fabric congestion as VOQ drops can happen due to different reasons. There is no such specific dedicated registers/counters from BCM to provide values for fabric drops. As of now its accounted as queue drops as fabric is back pressuring ingress. [PR1601332](#)
- The output of the `show system processes extensive` command shows high short term CPU utilization. The values ranges from 50 percent or higher for evo-pfemand. This is a single CPU view. As the ACX7509 system is a multi-core CPU, this has no impact on performance. [PR1603899](#)
- When dual Routing Engines are inserted in the system and these logs come from the backup Routing Engine. Supercon fpga has two end point devices (for RE0 and RE1) in which the primary Routing Engine end point device is enabled and the backup Routing Engine is disabled in the hierarchy. During boot up, the backup Routing Engine tries to access [supercon fpga in pci] hierarchy that is disabled. Therefore, supercon scratch pad register read failure logs and pci uncorrected errors are seen. [PR1605797](#) and [PR1615157](#)
- [PR 1606585](#)

The error message **Failed to attach ACPI GPIO chip** seen on console in early bootup is not an issue and it does not have any known functionality issue. It might be ignored. [PR1606585](#)

- The evo-pfemand process that programs the Packet Forwarding Engine ASIC restarts. As the evo-pfemand restarts, it re-programs the system ports. The hardware linkscan thread that monitors the link status records a momentary link down or link up transition. Each time the linkscan thread observes a change, it sends out a notification. Due to state compression in DDS, consumer applications like Picd do not receive intermediate link-down status notification from evo-pfemand and the interface does not toggle. This is because the notifications are produced rapidly. The remote peer interfaces observe an interface down while some of the local interfaces remain up. [PR1608215](#)



CAUTION: For 20xSFP FPC, following speed combination are possible

Slot	PortGroups	Speed
0	0-3, 4-7, 8-11, 12-15 and 16-19.	25g/10g/1g
1	0-7, 8-16, and 16-19	25g/10g
2	0-3, 4-7,8-11,12-15 and 16-19.	25g/10g/1g
3	0-3, 4-7,8-11,12-15 and 16-19.	25g/10g/1g
4	0-3, 4-7,8-11,12-15 and 16-19.	25g/10g/1g
5	0-7, 8-16, and 16-19	25g/10g
6	0-3, 4-7,8-11,12-15 and 16-19	25g/10g/1g
7	0-3, 4-7,8-11,12-15 and 16-19.	25g/10g/1g

Please note that:

1. There can be a combination of speed as initial configuration, however, a link flap is observed within the ports of the port groups if the speeds of any port are reconfigured to other speeds.
2. To avoid such a situation all the 1614286 ports of the port group can have single-speed or do not reconfigure the speed within the port group.

[PR1608223](#)

- While rebooting the system, the difference between time in which the interface comes up and the time in which RE1 comes up is approximately 3 minutes when compared to RE0 interface up time. [PR1608527](#)
- The syncE to PTP and syncE to 1pps transient response marginally fails. This happens when the servo get the initial 100 nano seconds jump in one measurement window and the next 100 nano seconds in the next measurement window adjusting less initially. [PR1608934](#) and [PR1611848](#)
- The maximum number of ECMP paths supported is 128. [PR1609063](#) and [PR1632055](#)

- A restart of DHCP takes more time because of the internal issues with the SIGTERM event. [PR1610229](#)
- While sending traffic from both the core files or units, all (16) VOQs/connectors need credits from same port. As a result, credits are distributed per the weight ratio and flows that need higher credits get higher bandwidth than the others. [PR1611028](#)
- PTP to PTP noise transfer fails for frequency 0.03125 HZ [PR1611838](#)
- The syncE to PTP and syncE to 1pps noise transfer tests fail for 1. 0.00781 HZ 2. 0.01563 HZ 3. 0.03125 HZ 4. 0.06156 HZ 5. 0.12313 HZ frequencies. [PR1611911](#)
- [PR 1612105](#)
it is day1 issue in ACX card. this is not problem only for multicast packet field in mac statistics . it is problem for all fields in mac statistics when queried for AE interface alone. problem is not there for actual ifd interface which is under AE interface. [PR1612105](#)
- On an FTC FRU hot removal on ACX7509, the following error is seen: **FAN -Error: value = 0xffff880e ret = 0xffffffffb** and there could be a kernel call stack trace core file in journal or dmesg. These are harmless messages and can be ignored. [PR1613389](#)
- "clear mpls lsp" operation is a destructive operation where it wipes off all existing routes and next-hops in the system and does a fresh reinstallation, the 10 seconds delay in traffic restoration for 16000 l3vpn routes might be attributed to programming delay in the hardware units combined with software model and the CPU capacity. [PR1614413](#)
- On ACX7509, 1GE interface does not come up with copper 1G SFP-T optics and this issue is specific to copper 1G cables. [PR1614286](#)
- Transient zl30642 PLL alarms are seen for the Forwarding Engine Board (FEB) or Flexible PIC Concentrators (FPC) seen when FEB or FPC is brought online on ACX7509. These alarms occur during the period of powerup and will clear in a matter of seconds when the PLL locks are achieved. Therefore, it can be safely ignored if they occur and clear during the onlining of these field-replaceable units (FRUs) - **Major FEB 0 zl30642 PLL Input Failure Major FEB 0 zl30642 PLL LOCK Failure Major FEB 0 zl30642 PLL OCXO Failure Major FPC 1 zl30642 PLL Input Failure Major FPC 1 zl30642 PLL LOCK Failure Major FPC 1 zl30642 PLL OCXO Failure.** [PR1615688](#)
- A race condition between overlay and kernel deregister is seen. Workaround is before doing jackout, run the command `request chassis fpc offline` and then jack-out. [PR1618133](#)
- Following scenarios warrant reboot or power cycle of the box.
 - Ungraceful removal of FRU without turning it offline.
 - Power failure in FRU.
 - PCI link failures in the system.

These issues can not be recovered in a live box and require reboot/power cycle. [PR1619368](#)

- The issue is specific to slot-7, where slot-7, port-13 is not supported. As part of channelization constraints, both adjacent ports have to configure or both to non channelization. Since slot-7 port is unsupported, need to ignore channelization constraints for port the two ports. [PR1620425](#)
- PE1 ---- P ---- PE2 ping mpls l2circuit does not work in case of explicit-null is configured in PE routers. It does not affect data path traffic. [PR1621111](#)
- In certain Segment Routing topologies having routers with varying CPU load, higher convergence time might be seen for a switchover event because of the micro-loops. This is a known behaviour when one node converges faster than the neighbouring node. [PR1621263](#)
- In case of multicast replication mode as ingress-egress-recycle, load balancing in the aggregate Ethernet member interfaces does not occur. [PR1621377](#)
- Due to major code changes and more regression impact, no details of auto negotiation settings is available on 1GE interfaces. [PR1621991](#)
- The ACX7509 has multiple FPC slots, statistics are maintained with single slot (slot 0) since Packet Forwarding Engine is centralized. ACX7509 is the only platform which has multiple FPC slot among ACX Series card. The pfstatsd is querying statistics based upon FPC slot for clear command which is causing issue since Junos OS Evolved infrastructure is not maintaining statistics based upon FPC slot for ACX7509. This needs change in pfstatsd which requires testing in other multi FPC slot platform also . As a workaround, other traffic counters like interface statistics instead of Packet Forwarding Engine statistics can be used. Packet Forwarding Engine statistics is rarely used to debug. [PR1622515](#)
- On upgrading Junos OS Evolved, FPCs transition PowerOff -> PowerOnWait -> PowerOnPermission -> PowerOnStart. But there is a glitch while it was trying to PowerOn and state went to Empty state from PowerOnStart. This leads to Fpc::OnDelete where we set SetFrueDeleted(true); to true. However, it does not go for FpcFrueDeleteCleanup(fpc_name); fruedelete cleanup where we reset the SetFrueDeleted flag back to false as FPC is still in not ready state (FPC did not come online as dmf init not yet completed) and will defer the termination. [PR1623455](#)
- The user can configure a family ethernet-switching filter to match the dmac(01:80:c2:00:0e) and etype(0x88f7) and drop such packets. Configuration gives us the flexibility to install filter only for the services where the drop is needed and others can forward as required. [PR1623756](#)
- RPC key failures error messages "Invalid RPC request key: 0x00110000" and "Invalid RPC request key: 0x000b0009" are observed while loading multid configuration. This does not have any functionality impact. [PR1624635](#)
- AOC cables go down on multiple restart of FPC. [PR1624992](#)

- The error message is expected while creation of IRB intermittently, this has no functional impact. [PR1625782](#)
- The Routing Engine can panic on back-to-back routing restarts due to timer_list corruption. 22.x images use newer Linux kernel (version 5.2+) where timer subsystem and callers have been updated to prevent this. [PR1626361](#)
- 400G-ZR optic transceiver firmware upgrade fails. It is a generic issue seen across all single RU platforms supporting QDD-400G-ZR. [PR1626882](#)
- After picd or rpdagent application restart multipile object-info anomalies for evo-pfemand, below are the types of anomalies seen Type : net::juniper::rtnh::Route Type : net::juniper::rtnh::NHOpaqueTlv Type : net::juniper::rtnh::Nexthop Type : net::juniper::rtnh::Unilist Type : net::juniper::rtnh::BfdSessionId. [PR1628843](#)
- If a system is fully scaled across features and firewall is also scaled, CPU consumption might be more for a small window of around 5 seconds after every 18 seconds or so. Evo-pfemand might be busy collecting the scaled firewall statistics for that 5 second window and any other applications like 'pfe-cli' trying to execute commands might fail during it. [PR1629342](#)
- In the Layer 2 performance test we have observed that for some of the packet sizes we are not able to achieve 100 percent line rate across the mesh fabric between the Q2C s. We use 42, 100G ports and expecting 4.2Tbps. However, we are able to achieve 4.17T instead of 4.2T. This happens for some specific packet sizes after 352 bytes and before 406 bytes. The vendor already clarified that for some of the packet sizes we won't be able to achieve 4.2T because of inefficiency of packing into cells. [PR1631948](#)
- When a fan tray fails an alarm is raised. After rebooting the alarm gets cleared. However, the fan failure condition will be logged in the log file. [PR1633353](#)
- In a working and non-working logs, l2d index is different for vrrp group number 187. This is the same group for which packet is getting dropped out of 400 groups, other groups are working as expected. So there is some fix which went between working and NKWR related to l2dld which has exposed VRRP issue. Both VRRP MAC and interface MAC gets stored in SLU my_mac_hash table. For finding hash index for vrrp mac we use l2dld, protocol type and vrrp group number as a key. In a non-working scenario there is a collision between interface mac and vrrp mac on same hash index. Ideally hash movement should have happened to address collision but somehow it is not properly done. Going further we need to debug why hash movement is not happening and fix that code area. This code is very sensitive and requires a lot of testing before doing any changes. So we should commit it in DCB first before committing to any other release. [PR1633986](#)
- Interfaces in 4x400g linecard are flapping when unrelated FPC is offlined in multiple FPC offline or online test. [PR1635115](#)
- In scaled setup, while FPC restart, PICD is not releasing fruHwStatus status and causing PIC struck in online. [PR1635941](#)

- FEB jackout and jackin might generate a hwd core file while rebooting the system (but not at jackout or jackin time). There will not be any functionality issue as the core file is generated only during reboot sequence. System is expected to come up fine. [PR1636243](#)

Interfaces and Chassis

- On ACX7509, there is a limitation in adding more than 64 member links in 1 AE, whereas from ASIC, there is no limitation. [PR1627951](#)
- On ACX7509, ACX7100-48L, and ACX7100-32C, when a large number of CFM sessions with 10ms or 3.3ms continuity-check interval are configured and evo-pfemand process is restarted or gets crashed, then the CFM sessions might not come up even after evo-pfemand is up. [PR1634721](#)

Layer 2 Features

- On Junos OS Evolved 21.4-based software, with VPLS instances at some scale, if the RPD process is killed, it might lead to inconsistency for some instances. In LSI information between control and forwarding planes, specifically show vpls connections instance and show ethernet-switching table instance and traffic drops due to incorrect VPLS label imposed at ingress. As a workaround, restart the rpd and/or l2ald gracefully. [PR1627593](#)

User Interface and Configuration

- In a rare event after configuring and deleting the DDL configuration statements such as set chassis aggregated-devices and set interface ae1 aggregated-ether-options, the configuration object gets out of scope with no functional impact. [PR1617667](#)
- File delete with regex might fail, if using filename without regex it works. [PR1624562](#)
- After picd or rpdagent application restarts, multiple object-info anomalies for evo-pfemand are observed. Below are the types of anomalies seen Type : net::juniper::rtnh::Route Type : net::juniper::rtnh::NHOpaqueTlv Type : net::juniper::rtnh::NextHop Type : net::juniper::rtnh::Unilist Type : net::juniper::rtnh::BfdSessionId. [PR1628843](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 21.4R1 | 31](#)

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: 21.4R1

IN THIS SECTION

- [General Routing | 31](#)
- [Infrastructure | 32](#)
- [User Interface and Configuration | 32](#)

General Routing

- Egress IP MTU exception and fragmentation are not supported. [PR1558327](#)
- High CPU seen mostly with systemd with 4000 mac-vrf instances activate or deactivate. [PR1581283](#)
- PICD restart or crash might result in junk statistics for carrier transition. [PR1594253](#)
- Default CoS profiles applied to Layer 2 VPN and Layer 2 circuit logical interfaces are not shown in the output of `run show class-of-service-interface IF`. [PR1596342](#)
- Few label switched interfaces (LSI) MACs do not get properly learnt in the software with 8000 VPLS instance scale. [PR1597125](#)
- 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 is seen. [PR1598217](#)

- The egress access control list (ACL) actions are skipped for Broadcast, Unknown Unicast, and Multicast (BUM) traffic and does not hit. [PR1598489](#)
- The Address Resolution Protocols (ARPs) might not be resolved on the integrated routing and bridging (IRB) interface which is replaced by another IRB interface. [PR1600209](#)
- For ACX7100-32C and ACX7100-48L routers, the **Voltage Threshold Crossed** alarm might be observed sometimes. [PR1601493](#)
- Traffic loss might be observed when evo-pfemamd restarts. [PR1608004](#)
- Default dscp-ipv6-compatibility classification does not work when deactivating and then activating dscp-ipv6 classification scenario. [PR1614249](#)
- PCI device missing FPC[0] FPC supercon FPGA alarms are observed. [PR1615519](#)
- On performing request `system snapshot`, the snapshot message is not captured in the `/etc/motd` file. [PR1618946](#)
- On ACX7100-32C and ACX7100-48L, the `show system firmware` command some times might show current firmware version for FPC 0 as blank. [PR1618949](#)

Infrastructure

- Memory leak is observed in packet send path. [PR1620610](#)
- Lookups command might not work. [PR1608401](#)

User Interface and Configuration

- The `file copy` command does not accept HTTPS URLs. [PR1596881](#)
- The file copy failure is seen via netconf or operation script. [PR1597550](#)

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Devices

IN THIS SECTION

- [What's New | 33](#)
- [What's Changed | 51](#)
- [Known Limitations | 56](#)
- [Open Issues | 58](#)
- [Resolved Issues | 66](#)

These release notes accompany Junos OS Evolved Release 21.4R1 for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Packet Transport Routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

NOTE: The PTX10016 is supported in Junos OS Evolved Release 21.4R1. See [Feature Explorer](#) for the full list of supported features.

What's New

IN THIS SECTION

- [What's New in 21.4R1 | 34](#)

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

What's New in 21.4R1

IN THIS SECTION

- [Hardware | 35](#)
- [Device Security | 38](#)
- [EVPN | 39](#)
- [High Availability | 39](#)
- [Juniper Extension Toolkit \(JET\) | 39](#)
- [Junos Telemetry Interface | 39](#)
- [Layer 2 VPN | 42](#)
- [Network Management and Monitoring | 43](#)
- [OpenConfig | 44](#)
- [Routing Policy and Firewall Filters | 45](#)
- [Routing Protocols | 46](#)
- [Services Applications | 47](#)
- [Software Installation and Upgrade | 47](#)
- [Additional Features | 47](#)

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 21.4R1, click the Group by Release link. You can collapse and expand the list as needed.

- [PTX10001-36MR](#)
- [PTX10003](#)
- [PTX10004](#)
- [PTX10008](#)
- [PTX10016](#)

The following sections highlight the key features in this release.

Hardware

- The PTX10016 is supported on Junos OS Evolved Release 21.4R1. See [Feature Explorer](#) for the full list of supported features.

Support for New PTX10016 Router Configurations—In Junos OS Evolved Release 21.4R1, we introduce the following new configurations of the PTX10016 router that support JNP10016-SF3 switch interface board (SIB):

- PTX10016-BASE3
- PTX10016-PREM2
- PTX10016-PREM3

The new SIB provides higher routing capacity and support for the PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards.

We also introduce a new front door (JNP10016-FRNT-PNL) that supports side EMI deflectors and a new rack mount kit (JNP10004-RMK-4POST).

To install the PTX10016 router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see [PTX10016 Packet Transport Router Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform. See [PTX10016 Switch Fabric](#) for more information about the SIBs. See [PTX10016 Components and Configurations](#) for more information about the router configurations.

Table 2: Features Supported on PTX10016 Devices

Feature	Description
Class of Service (CoS)	Support for CoS features includes classifiers (behavior aggregate (BA), fixed, and multifield (MF)) rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, and drop-profiles. [See CoS Features and Limitations on PTX Series Routers .]

Table 2: Features Supported on PTX10016 Devices *(Continued)*

Feature	Description
Interfaces and Chassis	<ul style="list-style-type: none"> • Support for redundant power supply includes environment monitoring. Use the request chassis fpc command to override the default power budget for an FPC line card. The SFBs have a higher priority over the line cards for power allocation. [See Managing Power and fpc command (PTX Series).] • Fabric manager support for: <ul style="list-style-type: none"> • PFE • PTX10K-LC1201-36CD line card • CCL link training and error handling • fabric statistics collection • RE switchover handling for SIB, and reduced SIB initialization time. [See PTX-10016 System Overview.] • Support for chassis to run with one or two RCBs. [See PTX10016 Routing and Control Board Components and Descriptions.]
Licensing	<p>You need a license to use the software features on the PTX10016 devices. [To find out the features supported on this device, see PTX10016 Routers Support for the Juniper Flex Program . To add, delete, and manage licenses, see Managing Licenses.]</p>

Table 2: Features Supported on PTX10016 Devices *(Continued)*

Feature	Description
Services and Applications	<p>Support for the following functionalities:</p> <ul style="list-style-type: none"> • IPv4 and IPv6 support for gateway [See Protocol-Independent Routing Properties User Guide and BGP User Guide.] • IPv4 and IPv6 inline flow monitoring. [See Understanding Inline Active Flow Monitoring.] • Tunneling for MPLS over UDP packets over RSVP LSPs and CSC, MPLS over Dynamic GRE with dynamic scaling option for number of tunnels up to 256K. [See Configuring a Next-Hop-Based Dynamic GRE Tunnels.] • Dynamic next-hop tunnel localization and dynamic GRE Tunnel creation based on new IPv4-mapped-IPv6 next-hops. [See Next-Hop-Based Dynamic Tunnel Localization Overview.] • Dynamic next-hop tunnel statistics reporting. [See Overview of Junos Telemetry Interface.] • BGP Signaling for MPLS over UDP and MPLS over GRE tunnels. [See Configuring MPLS over GRE.] • Mapping single IPv6 anycast address on multiple PFEs. [See get-route-range and IPv4 over IPv6.]

Table 2: Features Supported on PTX10016 Devices (*Continued*)

Feature	Description
Resiliency	<ul style="list-style-type: none"> • Platform resiliency support with logs, actions, faults, component recovery, and enhanced debugging capability for hardware failures. [See Understanding Chassis Alarms.] • Resiliency support for RCBs with fault-handling actions. [See PTX-10016 RCB Components.] • Fabric resiliency support for the PTX10K-LC1202-36MR line card (JNP10K-LC1202) for the links between FPCs. [See PTX-10016 Line Card Components.]

- **Express4 ASIC temperature-based performance throttling (PTX Series)**—PTX devices have five Express4 ASICs, each capable of operating at 3.2Tbps for a total of 14.4Tbps per line card with various temperature sensors for software monitoring and alarm issuance. If the Express4 ASICs get too hot, the line card shuts down. Reducing the Express4 ASIC performance reduces power consumption and reduces the ASIC shutdown temperature. To avoid a shutdown on the line card, performance throttling is enabled by default. Use the command `set chassis fpc slot pfe pfe-id temp-perf-throttle disable` to disable performance throttling. `show chassis fpc slot pfe-instance detail` displays the real-time status of all the available PFEs and whether the temperature throttle is on. [See [fpc](#), [show chassis fpc](#), and [No Link Title.](#)]

Device Security

- **Juniper Malware Removal Tool (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220)**—Starting in Junos OS Evolved Release 21.4R1, you can use the Juniper Malware Removal Tool (JMRT) to scan for and remove malware running on Junos OS Evolved devices. You can run two types of scans:
 - Quick scan—To scan each running program file
 - Integrity check—To check whether integrity mechanisms are enabled
 [See [Juniper Malware Removal Tool](#) and [request system malware-scan.](#)]

EVPN

- **VXLAN tunnel UDP source port entropy to support ECMP load balancing in EVPN-VXLAN fabrics (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, these routers use a software-supported method to provide source port entropy for ECMP load balancing on a VXLAN tunnel in an EVPN-VXLAN fabric. To support ECMP load balancing across a VXLAN tunnel, the device selects source port values based on a hash of fields in the VXLAN UDP packet inner header.

Before this release, these routers use the default hardware method that assigns only one UDP source port per VXLAN tunnel (based on the VXLAN egress port and the packet forwarding class).

[See [Load Balancing VXLAN Traffic.](#)]

High Availability

- **NSR support for up to 2000 TCP connections (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, PTX10004, PTX10008, and PTX10016 devices support NSR for up to 2000 TCP sessions. A maximum of 2000 BGP sessions are supported, and a maximum of 500 LDP or mLDP sessions, or a combination of both, are supported. Only node-level switchover is supported, application-level switchover is not supported.

Juniper Extension Toolkit (JET)

- **Support for programming FTIs using JET APIs (PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can use the Interfaces Service API to configure flexible tunnel interfaces (FTIs) in Junos OS. You can change the attributes of the tunnel configurations for the unit under an existing FTI but cannot change the existing tunnel encapsulation type using the APIs. For the following families, you can configure only the listed attributes when you use Juniper Extension Toolkit (JET) APIs:

- INET and INET6: address and destination-udp-port
- MPLS and ISO: destination-udp-port

[See [Overview of JET APIs](#) and [Configure Flexible Tunnel Interfaces.](#)]

Junos Telemetry Interface

- **BGP enhanced state sensors (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 supports Junos Telemetry Interface (JTI) streaming new BGP resource paths.

[See [Telemetry Sensor Explorer.](#)]

- **Chassis management configuration and counter support (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 supports JTI streaming of chassis management error (cmerror) configuration and counters to an outside collector using the UDP (native) sensors:

- /junos/chassis/cmerror/configuration
- /junos/chassis/cmerror/counters

[See [sensor \(Junos Telemetry Interface\)](#).]

- **CoS interface telemetry support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 supports gRPC Network Management Interface (gNMI) streaming of CoS interface queue statistics. To stream statistics, use the resource path **/qos/interfaces/interface/output/queues/queue/state/**.

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

- **INITIAL_SYNC telemetry support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 extends INITIAL_SYNC support for OpenConfig data model openconfig-interfaces.yang to new platforms. This feature lets the collector have a complete view of the current state of every sensor it is subscribed to. INITIAL_SYNC requires that at least one copy of all the sensors be sent to the collector. This feature has gRPC and gNMI support.

We support these resource paths:

/interfaces/interface/

/interfaces/interface/ethernet

/interface/holdtime/state

/interfaces/interface/subinterfaces/subinterface/

- /interfaces/interface/
- /interfaces/interface/ethernet
- /interface/holdtime/state
- /interfaces/interface/subinterfaces/subinterface/

[See [Telemetry Sensor Explorer](#).]

- **INITIAL_SYNC telemetry support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 extends INITIAL_SYNC support for OpenConfig data model openconfig-platform.yang and openconfig-alarms.yang to new platforms. This feature lets the collector have a complete view of the current state of every sensor it is subscribed to. INITIAL_SYNC requires that at least one copy of all the sensors be sent to the collector. This feature is supported with gRPC and gNMI.

We support the following resource paths:

- `/components/component`
- `/system/alarms/alarm`

[See [Telemetry Sensor Explorer](#).]

- **Pause frame receive and transmit counter support (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 includes JTI support for the sensor `/interfaces/interface/ethernet/state/counters/jnx-pfc-statistics` to stream pause frame counters from a router and transmit them to an outside collector.

Transport statistics using remote procedure call (gRPC), gNMI, or UDP (native).

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

- **Packet Forwarding Engine sensor support (PTX10001-36MR, PTX10004, and PTX10008)**—Junos OS Evolved Release 21.4R1 has JTI support for enhanced counters in the following areas:
 - Telemetry for device monitoring
 - Hardware components that require support NPU (using the INTEGRATED_CIRCUIT type)
 - Hardware drop reason counters
 - IP-in-IP encapsulation and de-encapsulation counters
 - Forwarding table occupancy

Transport statistics using gRPC or gNMI using the following resource paths:

- `/junos/system/linecard/packet/usage/`
- `/junos/system/linecard/npu/memory/`
- `/junos/system/linecard/npu/utilization/`

[See [Telemetry Sensor Explorer](#).]

- **Packet Forwarding Engine sensor support for non-viable aggregated interface member links (PTX10008 and PTX10016)**—Junos OS Evolved Release 21.4R1 has JTI support for non-viable aggregated interface member links. This feature does not support non-LAG link members. Use the resource path `/interfaces/interface/state/forwarding-viable` to stream statistics from a router to an outside collector by means of gRPC or gNMI.

[See [Telemetry Sensor Explorer](#).]

- **Platform sensors (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 includes JTI UDP streaming support for the sensors `/junos/system/cmerror/configuration/` and `/junos/system/cmerror/counters/`.

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

- **SR-TE colored policy RIB5 and SR-TE colored telemetry sensor support (PTX10008)**—Junos OS Evolved Release 21.4R1 introduces JTI streaming and ON-CHANGE sensors that deliver operational state statistics for Segment Routing–Traffic Engineering (SR-TE) colored policy RIB5 and SR-TE colored telemetry sensors. Statistics are delivered to an outside collector using gRPC or gNMI. Included are new OpenConfig resource paths for existing and new SR-TE policy (tunnel) and SR-TE per-LSP colored statistics.

[See [Telemetry Sensor Explorer](#).]

- **SR telemetry sensor enhancements (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 supports segment routing sensor enhancements for per SID level and interface level traffic counts. These changes comply with the current supported sensors in the OpenConfig models `openconfig-segment-routing.yang` and `openconfig-mpls.yang`.

We've included new resource paths with the enhancements:

- `/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/interface-ref/state/interface ['et-1/0/1']`
- `/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/interface-ref/state/subinterface [0]`

[See [Telemetry Sensor Explorer](#).]

Layer 2 VPN

- **Support for Layer 2 services on SR-TE tunnels using transport class (PTX10008)**—Starting with Junos OS Evolved Release 21.4R1, you can configure the following L2 services on colored and noncolored segment routing–traffic engineering (SR-TE) tunnels using transport classes:
 - L2 circuits
 - L2 VPN

[See [Understanding BGP Classful Transport Planes](#), [Configuring Interfaces for Layer 2 Circuits](#), and [Understanding Layer 2 VPNs](#).]

- **Layer 2 VPN for FAT label (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting with Junos OS Evolved Release 21.4R1, we support Flow Aware Transport for Pseudowires (FAT) labels on ingress routers, with parsing that includes all the payload fields in the hash calculation.

We support the following labels:

- L2 circuit, LDP-signaled pseudowires
- L2 VPN, BGP-signaled pseudowires
- L2 VPN with FEC129 (BGP autodiscovery)

Enable flow labels using the `flow-label-transmit` and `flow-label-receive` statements at the [edit protocols l2circuit neighbor neighbor-id interface interface-name] hierarchy level. Issue the `show ldp database extensive` command to display database information received from and advertised to the peer. Issue the `show l2circuit connections` command to see the status of the flow label after an exchange of signaling information.

[See [flow-label-receive](#) and [flow-label-transmit](#).]

Network Management and Monitoring

- **DHCP stateless relay MIB support (PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130, and QFX5220)**—Starting in Junos OS Evolved Release 21.4R1, you can use the Juniper Networks enterprise-specific DHCPv6 and DHCPv6 MIBs to retrieve statistics for DHCP stateless relay. We provide support only for the following MIB objects:

- `jnxJdhcpRelayStatistics`
- `jnxJdhcpRelayIfcStats`
- `jnxJdhcpv6RelayStatistics`
- `jnxJdhcpv6RelayIfcStats`

[See [Enterprise-Specific MIBs for Junos OS Evolved](#).]

- **NETCONF sessions over TLS (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)** – Starting in Junos OS Evolved Release 21.4R1, you can manage devices by running Network Configuration Protocol (NETCONF) sessions over Transport Layer Security (TLS). TLS uses mutual exchange of X.509 certificates for authentication to ensure the data integrity of exchanged messages.

[See [NETCONF Sessions over Transport Layer Security \(TLS\)](#).]

- **SNMP support for SCU and DCU counter (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, the SNMP MIB tables, `jnxDcuStatsTable`, `jnxDCUsTable`, and `jnxScuStatsTable` displays the source class usage (SCU) and destination class usage (DCU) statistics .

[See [Enterprise-Specific MIBs for Junos OS Evolved](#).]

- **SNMP MIB support for ICMP (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX1003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 21.4R1, we've added support for the following RFC4293 tables:
 - icmpStatsTable—Generic system-wide ICMP counters
 - icmpMsgStatsTable—System-wide per-version and per-message type ICMP counters

[See [Standard MIBs for Junos OS Evolved](#).]

OpenConfig

- **OpenConfig LACP configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 21.4R1 introduces LACP OpenConfig configuration support based on OpenConfig data model openconfig-lacp.yang.

[See [OpenConfig User Guide](#).]

- **OpenConfig LLDP configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008 and PTX10016)**—Junos OS Evolved Release 21.4R1 introduces LLDP OpenConfig configuration support based on OpenConfig data model openconfig-lldp.yang.

[See [OpenConfig User Guide](#).]

- **OpenConfig routing policies and options configuration support (PTX10008 and PTX10016)**—Junos OS Evolved Release 21.4R1 introduces OpenConfig routing policies and options configuration support to new platforms. OpenConfig configuration support is based on the OpenConfig data models:

- openconfig-network-instance.yang
- openconfig-routing-policy.yang
- openconfig-local-routing.yang

Additional OpenConfig augmentation for Junos CLI configurations that are not provided through these OpenConfig models is provided through the following data models:

- jnx-aug-openconfig-ni.yang
- jnx-aug-openconfig-routing-policy.yang
- jnx-aug-openconfig-local-routing.yang

[See [OpenConfig User Guide](#).]

- **System OpenConfig configuration support and gNMI mixed-mode support (PTX10008 and PTX10016)**—Junos OS Evolved Release 21.4R1 introduces OpenConfig configuration support for OpenConfig data models:
 - System Model `openconfig-system.yang`
 - System Logging Model `openconfig-system-logging.yang`
 - System Management Model `openconfig-system-management.yang`
 - System Terminal Model `openconfig-system-terminal.yang`

[See [OpenConfig User Guide](#)].

Routing Policy and Firewall Filters

- **Layer 2 and layer 3 support for flood policers (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can configure firewall filters for flood policers on L2 (family CCC) and L3 (family any) traffic, in both the ingress and egress directions. Most match conditions (except Packet-length) and most actions are supported.

See Firewall Filter Match Conditions and Actions.

- **Support for forwarding class and PLP as policer actions (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can use forwarding class, and both forwarding class and packet loss priority (PLP) together, as policer actions in policer policy configurations. This includes both the ingress and egress directions.

See Forwarding Class and Loss Priority.

- **Support for input-chain and output-chain CLI filters (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting from Junos OS Evolved Release 21.4R1, you can use multiple levels of CLI filters. The filter chain helps in logically grouping filters with a specific pattern of rules, instead of evaluating all the filter terms in one filter and deciding at the last term of it. The feature provides you the flexibility in modeling the filters as and when it is applicable in the solution. You can configure up to 8 filters in both input-chain and output-chains.

You can apply the filter chain as follows:

```
set interfaces interface-name unit unit family inet filter input-chain [filter1 filter2 filter3];
```

```
set interfaces interface-name unit unit family inet filter output-chain [filter1 filter2 filter3];
```

[See [input-chain](#), [output-chain](#), and [Example: Using Firewall Filter Chains](#).]

Routing Protocols

- **Enhanced support to handle S flag, D flag and A flags in IS-IS (PTX10001, PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.4R1, you can set the S flag to allow the label binding type, length and values (TLV) to leak through the IS-IS level (Level 1 or Level 2). You can set the A flag to program the penultimate-hop popping (PHP). You can set the D flag to prevent the leaking of the label binding TLV from Level 2 back to Level 1. Use the `no-binding` configuration statement at the `[edit protocols isis source-packet-routing no-binding-sid-leaking]` hierarchy level to disable label binding TLV leaks.

[See [Handling of the IS-IS Binding SID 'S' Flag and RFC 7794 Prefix Attribute Flags.](#)]

- **Service to correlate FPCs with P2MP sub-LSPs (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can correlate FPCs with point-to-multipoint (P2MP) sub-LSPs. An FPC can act as an ingress, egress, or transit point for more than one sub-LSP of the same P2MP LSP. If an FPC fails, then all the sub-LSPs that it serves are affected.

You can configure a service that enables you to monitor the correlation between FPCs and the P2MP sub-LSPs—branch paths—that are on an LSR. This information helps you evaluate the effect a failed FPC has on the correlated sub-LSPs. When you enable tracing, the service also provides system log messages (during an FPC outage) that provide detailed information about the sub-LSPs affected.

[See [Configuring a Service to Correlate Point-to-Multipoint sub-LSPs with FPCs.](#)]

- **Support for FAD and FAPM on the traffic engineering database and BGP-LS (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Junos OS Evolved Release 21.4R1, we support Flex ALgo Definition (FAD) and Flex Algo Prefix Metric (FAPM) on traffic engineering database and BGP-LS. You can store FADs and FAPMs into TED and BGP Link-State and store multiple prefix-sids for a prefix into BGP Link-State. You can import the FAD and FAPM entries from the traffic engineering database to BGP-LS and export the FAD entries from BGP-LS to the traffic engineering database.

[See [What is Flexible Algorithm Definition \(FAD\).](#)]

- **Support for IPv6 ping, IPv6 traceroute, and ECMP traceroute for Labeled-IS-IS segment routing paths (PTX10001-36MR, PTX10003, PTX10004, and PTX10008)**—Starting in Release 21.4R1, Junos OS Evolved supports IPv6 ping, IPv6 traceroute, and ECMP traceroute for Labeled-IS-IS segment routing paths.

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

- **Support for link delay measurement and advertising in IS-IS (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can measure and advertise various performance metrics in IP networks with scalability, by using several IS-IS probe messages. These metrics can then be used to make path-selection decisions based on network performance.

[See [How to Enable Link Delay Measurement and Advertising in IS-IS.](#)]

Services Applications

- **Export BGP community and AS path information using IPFIX (PTX Series)**—

Starting in Junos OS Evolved Release 21.4R1, you can export BGP community and AS path information using IP Flow Information Export (IPFIX) information elements 483 through 491, 16, and 17, per RFCs 8549 and 6313. Content providers can use this information to identify a transit service provider degrading the quality of the service. You configure these elements with the statement `data-record-fields` at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level.

[See [data-record-fields.](#)]

- **MPLS support for FTI tunnels (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can configure MPLS protocols over FTI (flexible tunnel interface), thereby transporting MPLS packets over IP networks which do not support MPLS. GRE and UDP tunnels support MPLS protocol for both IPv4 and IPv6 traffic. You can configure encapsulation and de-encapsulation for the GRE and UDP tunnels.

To allow the MPLS traffic on the UDP tunnels, include the `mpls port-number` statement at the `[edit forwarding-options tunnels udp port-profile profile-name]` hierarchy level. To allow the MPLS traffic on the GRE tunnels, include the `mpls` statement at the `[edit interfaces fti0 unit unit family]` hierarchy.

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

Software Installation and Upgrade

- **Support for outbound SSH service (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220-32CD, and QFX5220-128C)**—Starting in Junos OS Evolved Release 21.4R1, we support outbound SSH service.

You can use the `restart service-deployment` command to restart the Service Deployment System (SDX) process.

[See [outbound-ssh.](#)]

Additional Features

We've extended support for the following features to these platforms.

- **Avoid microloops in IS-IS segment routing MPLS networks (PTX Series)**—Starting in Junos OS Evolved Release 21.4R1, you can enable post-convergence path calculation on a device to avoid microloops between network devices.

[See [Understanding Microloop Avoidance.](#)]

- **BGP and statically configured SR-TE traffic statistics sensor support for JTI (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**

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

- **Forwarding information base (FIB) sensor support on JTI (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**

[See [Telemetry Sensor Explorer](#).]

- **Forwarding information base (FIB) sensor support on JTI (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Use JTI and gRPC services to stream or export ON_CHANGE FIB (also known as *forwarding table*) statistics from routers to outside collectors. This feature supports the OpenConfig YANG model OC-AFT.

[See [Telemetry Sensor Explorer](#).]

- **GPS and BITS interface support (PTX10008)**—

- 10MHz clock input and output support
- BITS input and output on T1/E1 framed and 2.048MHz unframed clock input

[See [Centralized Clocking Overview](#).]

- **ICMPv6 and ARP statistics support for Junos OS Evolved platforms (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700)**—The existing ARP and ICMPv6 statistics reports are enhanced to support the following statistics.

- ARP
 - Proxy requests
 - Requests dropped on entry
 - Number of arp packets dropped as nexthop allocation failed
- ICMPv6
 - Output histogram—unreach, echo request, echo reply, neighbor solicitation, neighbor advertisement, router solicitation, and router advertisement details
 - Input histogram—unreach, echo request, echo reply, neighbor solicitation, neighbor advertisement, router solicitation, and router advertisement details
 - messages with bad code fields, messages < minimum length, bad checksums, messages with bad length

- message responses generated, messages with too many ND options, current public ND nexthops present
 - [See [show system statistics arp](#) and [show system statistics icmp6](#)]
- **IPsec support for OSPFv2 and OSPFv3 (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130-48C, and QFX5220)**

[See [Overview of IPsec](#), [Configuring OSPF Authentication](#), and [Configuring IPsec Security Associations](#).]
- **LACP minimum link with sync reset and micro BFD configuration support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**

[See [Configuration Guidelines for Micro-BFD Sessions](#).]
- **MACsec fallback preshared key (PSK) and timer-based SAK refresh (PTX10001-36MR and PTX100016)**

[See [Configuring MACsec with Fallback PSK](#).]
- **New firewall filter feature: shared-bandwidth and percentage policers on the BT chip (PTX10003)**

[See [Configuring the Filter Profile](#).]
- **Packet Forwarding Engine restart (PTX10003-80C and PTX10003-160C):**
 - GRE tunnel encapsulation and de-encapsulation
 - Unicast RPF

[See [Powering Off Packet Forwarding Engines](#) .]
- **Physical Ethernet Interface sensor support on JTI (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**

[See [Telemetry Sensor Explorer](#).]
- **SNMP support for LACP and LAG MIB objects (PTX10001-36MR, PTX10008, and PTX10016)**

[See [Enterprise-Specific MIBs for Junos OS Evolved](#) and [SNMP MIB Explorer](#).]
- **Support for Application Specific Link Attributes (ASLA) for flexible algorithms (PTX10001-36MR, PTX10008, PTX10016):** IS-IS supports advertising different te-metric and admin-groups for RSVP and flexible algorithm on the same link using flexible-algorithm specific ASLA as defined in RFC 8919.

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

- **Support for disable, offline and restart of individual PFEs with error log and action, using `request chassis pfe` and `set chassis fpc` commands for specific line cards (PTX10K-LC1202-36MR) on PTX series devices.**

[See [request chassis fabric pfe](#) and [fpc \(PTX Series Routers\)](#).]

- **Support for EVPN Type 5 routes for Layer 2 and Layer 3 VXLAN gateways with overlay load balancing in an EVPN-VXLAN network (PTX10001-36MR, PTX10004, and PTX10008)**

[See [Understanding EVPN Pure Type-5 Routes](#) and [EVPN Type-5 Route with VXLAN encapsulation for EVPN-VXLAN](#).]

- **Support for G.8275.1 profile, Precision Time Protocol over Ethernet encapsulation, and Hybrid mode (PTX10008)**

[See [G.8275.1 Telecom Profile](#), [Guidelines for Configure PTP over Ethernet](#), and [Hybrid Mode Overview](#).]

- **Support for inner VLAN transparency (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**

We support the pop/push/swap/pop-pop/pop-swap/swap-push/push-push/swap-swap operations on port-based and VLAN-based Metro Ethernet Forum (MEF) L2 services. VLAN transparency refers to preserving inner VLANs in the packet that are not subject to manipulation and are not used for forwarding. Based on the scenarios, VLAN transparency works on up to four VLAN tags on Express4 ASICs.

[See [Understanding VLAN Manipulation \(Normalization and VLAN Mapping\) on Ethernet Services](#).]

- **Support for MIB accounting profiles (PTX10001-36MR, PTX10004, PTX10008, PTX10016, and QFX5220)**

[See [Accounting Profiles: An Alternative to SNMP Statistics](#).]

- **Support for PIM multicast features on Layer 2 IRB interface (PTX10001-36MR and PTX10004)**

[See [Multicast Routing Protocols](#).]

- **Supported transceivers, optical interfaces, and DAC cables (ACX Series, PTX Series, and QFX Series)** –Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

- **Telemetry stream path resolution by MPLS and RSVP interfaces (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**

[See [Telemetry Sensor Explorer](#).]

- **VRRP support (PTX10001-36MR, PTX10003, PTX10004, and PTX10008):** Support for advertisement interval at 1 second and interworking of VRRP version 2 and VRRP version 3.

- VRRP active inheritance support
- VRRP tracked routes and tracked interfaces
- VRRP support for IPv4 and IPv6
- VRRP support at physical, aggregated Ethernet, and IRB IFL-level interfaces
- VRRP support over the single tagged interface
- VRRP support in centralized and delegated mode

[See [Understanding VRRP](#) .]

What's Changed

IN THIS SECTION

- [What's Changed in Release 21.4R1-S2 | 51](#)
- [What's Changed in Release 21.4R1-S1 | 52](#)
- [What's Changed in Release 21.4R1 | 52](#)

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

What's Changed in Release 21.4R1-S2

IN THIS SECTION

- [User Interface and Configuration | 51](#)

User Interface and Configuration

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after rollback timeout.

What's Changed in Release 21.4R1-S1

IN THIS SECTION

- [General Routing | 52](#)
- [User Interface and Configuration | 52](#)

General Routing

- JNP10K-PWR-DC2 power supplies installed in PTX10008 and PTX10016 routers display as online when the power supplies are switched off - JNP10K-PWR-DC2 power supplies installed in PTX10008 and PTX10016 routers in which Junos OS Evolved Release 21.4R1.13-EVO is installed display as online in the output of the command `show chassis environment psm` when the input power feeds are connected, but the power switch on the power supplies are switched off.

User Interface and Configuration

- When you configure `max-cli-sessions` at the **edit system** hierarchy level, it restricts the maximum number of cli sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login.

The `max-cli-sessions` is configured so you can control the memory usage for the CLI. You may set the `max-cli-sessions` per your requirement. However, if `max-cli-sessions` is not configured, there is no control on the number of CLIs getting invoked.

What's Changed in Release 21.4R1

IN THIS SECTION

- [Class of Service | 53](#)
- [EVPN | 53](#)
- [General Routing | 53](#)
- [Interfaces and Chassis | 54](#)
- [Junos OS API and Scripting | 54](#)
- [Layer 2 Features | 54](#)

- Network Management and Monitoring | 54
- Routing Protocols | 55
- Security | 55
- Software Licensing | 56

Class of Service

- Junos OS Evolved now correctly displays the index for `show class-of-service` commands.

EVPN

- **Output for the `show Ethernet switching flood extensive` command**—The output for the `show ethernet-switching flood extensive` command now displays the correct next-hop type for Virtual Ethernet and WAN mesh group in an EVPN-VXLAN network as unilist. Previously, the output for the `show ethernet-switching flood extensive` command would misidentify the next-hop type as composite.

General Routing

- In Junos OS Release 21.4R1 and later, `ssh` is enabled by default on all the routers with VM host support.
- **Validation of TCA threshold values (PTX10008)**—We've implemented immediate validation of threshold values configured in the `tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number)` statement under the `[edit interface <interface name> optics-optics tca]` hierarchy level to ensure the threshold value entered is valid.

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

- **Enhancement to the `request system license add terminal` command (PTX10001-36MR)**—When you run the `request system license add terminal` command, you can now view following additional fields for information: JUNOS564022985: Ignoring unknown feature.

[See [Managing vMX Licenses.](#)]

- **A major alarm is raised (PTX10008)**—A major alarm is raised when a fan tray controller is removed from the chassis.
- **SNMP MIB support for field-replaceable unit (FRU) LEDs (PTX10008)**—SNMP MIB object `jnxLEDEntry` now indicates multiple LED indexes and color values for FRUs.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Previously, the source and destination address pair had to be unique across all encapsulation type.
- **Enhancement to snmp mib command behavior (PTX10008)**—Starting in Junos OS Evolved, when you execute `show snmp mib walk decimal` command, the output parameter `jnxRedundancySwitchoverReason` is not working as expected, which always show the value 0 instead of expected values. Now, `jnxRedundancySwitchoverReason` output parameter is corrected to expected behavior with the following expected values. `jnxRedundancySwitchoverReason OBJECT-TYPE SYNTAX INTEGER href=' other(1), - others neverSwitched(2), - never switched userSwitched(3), - user-initiated switchover autoSwitched(4) - automatic switchover ' format="html" scope="external">`

[See [show snmp mib](#).]

Junos OS API and Scripting

- **Limits increased for the max-datasize statement (ACX Series, PTX Series, and QFX Series)**—The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

Layer 2 Features

- **New Commit check for Layer 2 Interfaces (PTX10003)**—We've introduced a commit check to prevent you from misconfiguring ethernet encapsulation on Layer 2 interfaces. Ethernet encapsulation is not supported on Layer 2 interfaces.

[See [encapsulation \(Logical Interface\)](#) and [Layer 2 Address Learning and Forwarding Overview](#).]

Network Management and Monitoring

- **SNMP support for MIB**—Operational command `show snmp mib walk systemnow` shows the latest software version and does not show the build date.

[See [show snmp mib](#).]

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

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

- **The configuration accepts only defined identity values for nodes of type `identityref` in YANG data models (ACX Series, PTX Series, and QFX Series)**—If you configure a statement that has type `identityref` in the corresponding YANG data model, the device accepts only defined identity values (as defined by an identity statement) as valid input. In earlier releases, the device also accepts values that are not defined identity values.

Routing Protocols

- **The RPD_OSPF_LDP_SYNC message not logged?**On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the `hold-time` for `ldp-synchronization` at the **`edit protocols ospf area area id interface interface name`** hierarchy level less than three minutes. The message is printed after three minutes.
- To achieve consistency among resource paths, the resource path `//mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter ip-addr='address'/state/countersname='name'/out-pkts/` is changed to `/mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/counters name='name'/`. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").

Security

Renamed `verixec-check` option—We have changed the `verixec-check` option of the `request system malware-scan` command to `integrity-check`. This update does not include any functional changes. You can use the `integrity-check` option to check whether integrity mechanisms are enabled for the Juniper Malware Removal Tool.

Software Licensing

- **Juniper Agile Licensing (PTX10003, PTX10016, QFX5130-32CD, and QFX5220)**—Starting from this release onwards, the Juniper Agile License Manager is deprecated. You can use the Juniper Agile Licensing Portal to activate, install, manage, and monitor licenses on Juniper Networks devices.

[See [Juniper Agile Licensing Guide](#).]

Known Limitations

IN THIS SECTION

- [General Routing | 56](#)
- [Network Management and Monitoring | 57](#)
- [Routing Protocols | 58](#)
- [VPNs | 58](#)

Learn about known limitations in Junos OS Evolved Release 21.4R1 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

- Restarting the fabspoked-pfe application for the line card, restarts the line card. [PR1486023](#)
- By default, the IPv6 addressing gets configured with 64 subnet irrespective of the subnet configured on the DHCP server side. [PR1539839](#)
- The PTP FPGA is in the reset mode during BIOS boot. During Linux boot, the PTP FPGA is taken out of the reset mode and the pcie-tree is re-enumerated. Therefore, the link goes up or comes down when this sequence occurs. [PR1572061](#)
- On PTX10001-36MR routers, which is a fixed form factory system with a single CPU, the FPC is a logical one. There is no dedicated CPU to handle the FPC processes and there are three different processes that handle the Packet Forwarding Engine functionality. Also, there is no dedicated

memory reserved for the Packet Forwarding Engine processes, due to which it is not feasible to figure out the percentage of CPU or memory utilization by the FPC. Hence, the `show chassis fpc` output does not display a percentage for neither the CPU or memory computed. [PR1586907](#)

- Configuring a 400G ZR link with 4x100GE channelized mode on one end and 1x400GE on the other end results in link status mismatch. The 1x400GE port gets displayed with the port in the Up state even when the port is in the Down state. [PR1597707](#)
- Ungraceful FRU removal results in CRC errors on the peer FRUs, that might potentially wedge the Packet Forwarding Engines. [PR1615297](#)
- Traffic convergences, interface deactivates/disables, and deletes one of the SRTE path in a policy in two different scenarios. During these periods, the hardware counters gets deleted and added again for new paths where traffic flows. During this transient period of time counting, the statistics do not increment and as per the design SR statistics counters, reading from hardware occurs every 4 seconds. By any chance, if the convergence and interfaces triggers right after reading the hardware counters, then no statistics increment for that 4 seconds of time period. [PR1630517](#)
- The request `system node reboot<>` command restarts the FRU ungracefully. On an ungraceful FRU restart, transient link errors on peer FRU appears. [PR1630787](#)
- On PTX10001-36MR routers, PTP TC over IPv6 feature does not meet the T1/T4 error for 4x25G mode when connected between two retimer ports (retimer-----retimer), that is port 4 to port 7 in each PIC. This does not impact the basic PTP TC over IPv6 functionality and other performance metrics such as two-way time error. cTE and dTE meets class B for all the port combinations. [PR1609110](#)
- On all Junos OS Evolved platforms, when nexthop gets added or changed to the Packet Forwarding Engine and the same next-hop is also forwarding nexthop of an indirect route, if the ingress Packet Forwarding Engine is faster than the egress Packet Forwarding Engine, this causes packet loss as the ingress Packet Forwarding Engine being faster would have seen the new forwarding nexthop and also the indirect change. However, the egress Packet Forwarding Engine being slower would not have consumed indirect change yet. [PR1547432](#)

Network Management and Monitoring

- Configuring the `set system no-hidden-commands` command must not block the netconf sessions. As a workaround, you must disable the `no-hidden-commands` command. [PR1590350](#)

Routing Protocols

- Protocols might not restart gracefully on a scaled setup with the `restart routing gracefully` command due to which long time is consumed in synchronization the routing states during restart. (It take around 4 minutes 20 seconds for protocols to start back again after restart with 2.8m routes.) [PR1575626](#)

VPNs

- ASIC does not have a way to drop the packets with invalid configuration while packets go out. Hence, the unsupported VLAN rewrite configuration might send packets out with wrong VLAN encapsulation. [PR1627484](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 59](#)
- [General Routing | 59](#)
- [Infrastructure | 64](#)
- [Interfaces and Chassis | 64](#)
- [Juniper Extension Toolkit \(JET\) | 65](#)
- [MPLS | 65](#)
- [Network Management and Monitoring | 65](#)
- [Routing Protocols | 65](#)
- [User Interface and Configuration | 65](#)

Learn about open issues in Junos OS Evolved Release 21.4R1 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.

Class of Service (CoS)

- Cosd generates a core file only with the `show` command consistently while the ifd scheduler map bindings do not get updated correctly or in a `Transition` state during the booting time, and initiates the `show interface <ifd-name> name` command. The cosd application recovers after the crash and does not affect any other programming. As a workaround, wait until all the interfaces are in the `Up` state and complete the interface binding configurations commit before triggering the `show interface <if-name> extensive` command. [PR1604169](#)

General Routing

- When a aggregated Ethernet link is brought down, a transient error message, **[Error] Nexthop: EalNhHandler: failed to add Nh: xxxx, type: composite, as pil add failed** might be seen. There is no functional impact due to these errors. [PR1570710](#)
- The PTP FPGA is reset during BIOS boot. During Linux boot, the PTP FPGA is taken out of reset and the pcie-tree is reenumerated. Hence you see the link going up or coming down during this sequence. [PR1572061](#)
- A vulnerability in the handling of exceptional conditions in Juniper Networks Junos OS Evolved allows an attacker to send specially crafted packets to the device, causing the Advanced Forwarding Toolkit manager (evo-aftmand-bt or evo-aftmand-zx) process to crash and restart, impacting all traffic going through the FPC, resulting in a Denial of Service (DoS). [PR1572969](#)
- Changing decapsulation of only the tunnel destination address configuration after tunnel is up might not work and end up using the previously configured tunnel destination address for the decapsulation. After the system enters this state, any further configuration changes to tunnel configuration cannot be handled. [PR1575724](#)
- IPv6 based PKID enrollment operations are not supported over revenue ports on all Junos OS Evolved platforms. [PR1584378](#)
- When all the Packet Forwarding Engines are in the `Offline` state and being in the `Online` state with multicast route being active during this cycle, the multicast traffic gets lost permanently due to the absence of multicast route. [PR1598894](#)
- Sflow ingress sampling reports incorrect OIF and NH with MPLS traffic(MPLS-IPV4/MPLS-IPV6 with ECMP and Ultimate Hop Pop(UHP)) on the egress node. [PR1602448](#)
- On PTX10016 routers with a fully loaded chassis and all the FPCs being in the `Online` state, if you take a SIB offline or online, the fabric link training takes around 10 minutes. [PR1611753](#)

- On PTX10008 routers, few input PTP & SyncE-1pps noise transfer test frequencies might fail. The following result is the expected performance :

```
Profile Freq (Hz) NoiseTransfer_0_00391_Results PASS PASS
NoiseTransfer_0_00781_Results FAIL FAIL NoiseTransfer_0_01563_Results FAIL FAIL
NoiseTransfer_0_03125_Results FAIL FAIL NoiseTransfer_0_06156_Results FAIL FAIL
NoiseTransfer_0_12313_Results FAIL FAIL NoiseTransfer_0_24625_Results PASS PASS
NoiseTransfer_0_4925_Results PASS PASS NoiseTransfer_0_985_Results PASS PASS
NoiseTransfer_1_985_Results PASS PASS NoiseTransfer_3_985_Results PASS PASS
NoiseTransfer_7_985_Results PASS PASS
```

[PR1624502](#)

- On PTX Evolved platforms with a mac-vrf instance configured with VLAN aware service, BUM(Broadcast Unicast Multicast) flooding issue might occur on the ESI (Ethernet Segment Identifier) interfaces and packets might get dropped on the egress PE devices. [PR1624677](#)
- On PTX10008 routers, the transit CCM sessions comes up but transit oopback(LB) ping or LinkTrace(LT) PDUs does not go through. [PR1632255](#)
- In a scaled environment, it might take up to five minutes to gather and display the p2mp LSP correlation information. [PR1629771](#)
- On PTX10008 routers, after testing the laser-off link member to bring the aggregated Ethernet interface down, the following error message gets generated:

```
[Error] Nexthop: fail to get target nh facad for indr nh 20209, [Error] Nexthop:
EalIndirectNh::add: failed to get target nh for nh id: 20209, [Error] Nexthop: NH add during
change major failed for NH 20209, status:1000
```

[PR1617388](#)

- After configuring the warm-standby option, you must wait for three minutes before Routing Engine switchover. [PR1623601](#)
- With back-to-back P2MP branch joins/leaves, out-of-order error might be generated in the Packet Forwarding Engine while processing the P2MP Multicast Nexthop CHG. It is a transient issue of Packet Forwarding Engine being not able to process MCNH CHG, which should be corrected by an immediate the Multicast Nexthop CHG and Make-before-break (MBB) sequence. [PR1628049](#)
- Due to limitation in CBC FPGA, interrupts do not get generated if the Routing Engine mastership daemon stalls and a hardware timer-expiry mastership switchover gets initiated. Post switchover, you must reboot the backup routing-engine before another switchover. [PR1633552](#)

- System reboot or boot up with traffic could result in init time fabric link crc errors and small percentage of traffic drop. [PR1635178](#)
- Traceroute in MPLS OAM on SR over IPv6 might fail in ECMP case if the Junos OS Evolved device is in topology. This is because Linux kernel in Junos OS Evolved puts an autoflowlabel on every IPv6 packet. This flow label is transparent to daemon process, which uses a null value for it and calculates the nexthop details. The Packet Forwarding Engine however takes the flow label into account and calculates the nexthop details. This difference in calculation of the nexthop details leads to a mismatch in the path the packet takes to the destination and might cause traceroute to fail. [PR1618406](#)
- In some reload scenarios, the config-sync service might fail due to xinetd limiting the SSH service, leading to the Major Application config-sync fail on node Re1 alarm status. This might cause configuration difference between the primary and backup Routing Engine. As a workaround, restart the service using the restart config-sync command. [PR1629952](#)
- The evo-aftmand-bt process does not fully support request node halt <node> command. The command might cause cm-errors on other FPCs as the fabric links are not gracefully shutdown. Use the request chassis fpc slot <x> offline command to shutdown an FPC instead. [PR1636271/xref](#)
- On PTX10008 routers, FPC JNP10K-LC1201 frequently generates the following error message:

```
zephyr_clock_get_tod_ext_sync_sample(xxx): READ BT-X tod_sec: xxxxxxxxxx, tod_ns: xxxxxxxxxx
```

[PR1635771](#)

- The switch from 10M to 5M (or 1M), when both output and source is configured for GPS ports (i.e output and source configuration), the REF status does not get updated in the table. This causes incorrect status to be displayed for the Clock Status and Clock Event. When you change the signal type to 10M from other frequencies, the DPLL remains in the Holdover state. As a workaround to restore the normal behavior, deactivate and activate the chassis synchronization command. [PR1631689](#)
- On PTX10008 routers, transient JSR replication error messages 113/115 gets generated when you disable or enable OSPF. No functionality impact as NSR gets enabled again quickly on all protocols after the error messages. The error messages also gets generated when rpd-agent crashes. [PR1627625](#)
- The CDA application crashes as soon as FPC boots up that results in fabricPfeE object not being published. Hence, the fabricPfeE objects do not get created or deleted. SIB endpoints cleanup are dependent on the deletion of the fabricPfeE objects. Stale SIB fabric endpoints prevent fabric link to bring up in further FPC reboots. [PR1624765](#)

- On PTX10008 routers, few frequencies might fail to meet the mask. The following result is the expected performance for PTP-PTP and PTP-1PPS:

```
Profile Freq (Hz) PTP 1PPS NoiseTransfer_0_00391_Results PASS PASS NoiseTransfer_0_00781_Results PASS PASS
NoiseTransfer_0_01563_Results FAIL FAIL NoiseTransfer_0_03125_Results FAIL FAIL NoiseTransfer_0_06156_Results
FAIL FAIL NoiseTransfer_0_12313_Results PASS PASS NoiseTransfer_0_24625_Results PASS PASS
NoiseTransfer_0_4925_Results PASS PASS NoiseTransfer_0_985_Results PASS PASS NoiseTransfer_1_985_Results PASS
PASS PR1624478
```

- Errors do not appear until there is a composite nexthop with two labels in the nexthop. This scenario must not appear and there is no impact in behavior.

[PR1621689](#)

- On scaling system, load override and commit of baseline configuration causes the rpd to spin high on CPU. As GRPC configuration gets removed, sensors must be uninstalled. But, rpd does not respond to these telemetry sensors uninstallation requests failing the sensors uninstallation. Later, when you enable GRPC back on box and request the same sensor profile (cfg .json file used with jtimon) from collector, the rpd sends packets with higher sequence numbers. (Because sensor was not removed from the rpd earlier and is considered as drops by collectors, which rely on sequence-numbers.)

[PR1621347](#)

- With the scale configuration and NSR switched on, the cosd (Class of service control plane) application generates core files with multiple Routing Engine switchover. Most of the time, the cosd application recovers automatically. The status of the cosd application after the Routing Engine switchover displays the following message:

```
root@ptx10008-pfe-03-re0> show system applications app cosd
Applications Information:
Application : cosd
Node : re0 App
App State : online ready
App Weight : 1
App Zookeeper Session : 50000024b200036
Object Producer details
Producer ID : 50331769
Epoch ID : 70
Production Topic : /Root/cosd/5443209913770603356
Producer State : active
Command to restart cosd (Manual)
root@escapa-pfe-03-re0> request system application node re0 app cosd restart
This may affect traffic in the system. Proceed ? [yes,no] (no) yes
```

App cosd on node re0 restart request is submitted

Note: use re0 or re1 based on the mastership

PR1620758

- When you use 400G or 4x100G interfaces using 400G-ZRoptics, after the data path of the optics gets powered up and Tx laser enabled, sometimes the actual transmitting power might appear very low. Inside 400G-ZR software driver, a check-and-recovery mechanism has been implemented. If the actual Tx power appears very low after powering up and enabling Tx laser, a data path powering-down and re-init sequence gets exercised. The Tx power level gets recovered back to normal. After such hardware issue appears, the actual link up time appears longer than expected.

PR1616445

- On PTX10008 routers, the rpd process generates a core file at rt_evo_free_sesid_info in the backup Routing Engine.

PR1607553

- Vulnerability during the handling of exceptional conditions in Juniper Networks Junos OS Evolved allows an attacker to send specially crafted packets to the device, causing the Advanced Forwarding Toolkit manager (evo-aftmand-bt or evo-aftmand-zx) process to crash and restart. This impacts all the traffic going through the FPC, resulting in a Denial of Service (DoS). For for information, please refer to <https://kb.juniper.net/JSA11188>. [PR1572969](#)
- When an aggregated Ethernet interface link is brought down, a transient error message [Error] Nexthop: EalNhHandler: failed to add Nh: xxxx, type: composite , as pil add failed might be generated. There is no functional impact due to these errors. [PR1570710](#)
- On Junos OS Evolved platforms, if a firewall filter has both SCU (Source Class Usage) or DCU (Destination Class Usage) matches and forwarding-class/loss-priority matches used on the same term in the from action, the traffic supposed to be subjected to the firewall filter might not get subjected to it. [PR1595788](#)
- Others FPCs gets rebooted if you restart FPC 2. [PR1625065](#)
- The show snmp mib walk jnxBoxAnatomy command lists partial output. [PR1628788](#)
- You must disable the inline BFD processing and only allow the delegate processing. [PR1630797](#)
- During system boot up, the following PCIe Bus Error error message indicating on pcieport 0000:00:02.2 gets generated:

```
pcieport 0000:00:02.2: PCIe Bus Error: severity=Corrected, type=Physical Layer,
id=0012(Receiver ID)
```

PR1631300

- At some corner cases, the MPLS family filter matches fails in the latest Junos OS Evolved Release 21.4R1 images, which was due to uninitialized variable during the filter programming. Deactivating and activating the MPLS filter resumes the functionality. [PR1634644](#)
- When the Packet Filtering Engine triggers reset, there is MPLS LSP transit traffic that gets wrong encapsulation and traffic gets dropped in the down stream. [PR1635130](#)
- The UDP export from FPC does not work for the `/junos/system/cmerror/counters/` and `/junos/system/cmerror/configuration/` sensors files. The UDP export of these sensors works on the fixed-form-factor systems like PTX10001 and PTX10003. However, the UDP export of these sensors does not work on the chassis-based (modular) systems like PTX10008, PTX10016, and PTX10004. GRPC and GNMI export of these sensors do not get affected. [PR1638262](#)

Infrastructure

- When you use a source IP address as the management interface with RPF check set to strict on the interface, the response for the ICMP ping from the peer on the management interface gets dropped by the Linux kernel. The Linux kernel expects the path to the peer to be on the WAN side. [PR1498255](#)
- GRES gets triggered through the request chassis routing-engine master switch command that starts displaying the following connector driver overlay messages:

```
{master} root@ptx10004-05-re0> [ 1185.081257] gpio-jnx-i2cs gpio-jnx-i2cs.50: Asserting
power_status irq 59
[ 1185.125182] OF: overlay: overlay_is_topmost: #9 clashes #10 @/ftc0/i2c-bus/i2cs@54/
fan_hwmon
[ 1185.125183] OF: overlay: overlay #9 is not topmost
```

PR1539232

- Rebooting the PTX10003 router during a broadcast storm on the management port might cause a fault on the router. [PR1423216](#)

Interfaces and Chassis

- The local switching traffic sequence number does not reset. [PR1560111](#)

- On PTX10003 routers, SNMP mib walk on jnxLEDTTable fails. [PR1620398](#)

Juniper Extension Toolkit (JET)

- From Junos OS Evolved 20.4 Release, there are two different grpc python files for each JAPI file. The names of the files are `*pb2_grpc.py` and `*pb2.py`. The stub creation functions are present in the `*pb2_grpc.py` file. [PR1580789](#)

MPLS

- In Junos OS Evolved platforms, the LDP session does the authentication key-chain configuration based on the session remote-id on initiator stop from the session establishment even though you configure the authentication key-chain of the responder for the responder remote-id. [PR1592431](#)

Network Management and Monitoring

- The mgd process might crash when you configure an invalid value though JSON and XML load for the identityref type leafs or leaf-lists while configuring Openconfig or any other third-party YANG. This issue occurs with JSON and XML loads. [PR1615773](#)

Routing Protocols

- This scenario occurs when you enable the IS-IS authentication key-chain that have multiple keys between the routers. The IS-IS adjacency is up since both routers have the same key active. When you manually change the system time in such a way that the routers have different keys active in the key chain, the IS-IS adjacency must go down, which does not happen. [PR1572441](#)

User Interface and Configuration

- For the list definition `list bar { key foo; leaf foo { type uint8; } leaf baz { type string; } }`, Junos cannot parse the `{ "bar": [{ "baz": "zig", "foo": 123 }, { "baz": "zag", "foo": 0 }] }` configuration payload where leaf foo (key as per Yang model) does not appear as a first element. Junos understands

only { "bar": [{ "foo": 123, "baz": "zig" }, { "foo": 0, "baz": "zag" }] } kind of configuration payload as shown below where leaf foo (key as per Yang model) appears as a first element. As per Junos, the list keys are always required (must always) to precede any other siblings within a list entry and appear in the order specified by the data model. [PR1616216](#)

- In a rare event after configuring and deleting the DDL commands, the configuration object correctly gets out of scope. However, the configuration object does not get released due to a shared reference set chassis aggregated-devices set interface ae1 aggregated-ether-options. There is no functional impact and the situation can be cleared by restarting configuration daemon from CLI--> configd > restart configd. [PR1617667](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 21.4R1 | 66](#)

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: 21.4R1

IN THIS SECTION

- [Authentication and Access Control | 67](#)
- [Class of Service \(CoS\) | 67](#)
- [Flow-based and Packet-based Processing | 67](#)
- [General Routing | 67](#)
- [Infrastructure | 72](#)
- [Interfaces and Chassis | 73](#)
- [MPLS | 73](#)
- [Network Management and Monitoring | 73](#)

- [Routing Policy and Firewall Filters | 74](#)
- [Routing Protocols | 74](#)
- [User Interface and Configuration | 74](#)

Authentication and Access Control

- The ssh process crashes on Junos OS Evolved platforms. [PR1601150](#)
- The root password might not be accepted under su on the Junos OS Evolved platforms. [PR1607861](#)

Class of Service (CoS)

- The MPLS fixed classifiers might not work on the Junos OS Evolved platforms. [PR1616492](#)

Flow-based and Packet-based Processing

- The msvcsd process might get crashed when you enable the `nexthop-learning` command with the jFlow service. [PR1620569](#)

General Routing

- New alarm is reported for three consecutive failures on a particular fan tray (this can be power or fan failure). [PR1500920](#)
- The device might run out of service post GRES or unified ISSU. [PR1558958](#)
- `request system zeroize` command does not delete snapshot images from the backup hard disk drive. [PR1569294](#)
- The static MACs configured over the aggregated Ethernet interface might not get programmed while forwarding after the FPC restarts. [PR1581325](#)
- Certain fields in the GNMI extension header and `show network-agent statistics` command display incorrect values if the input subscription path contains the `:` character. [PR1581659](#)
- Packet might be lost during global repair of FRR. [PR1586122](#)
- The `RPD_KRT_KERNEL_BAD_ROUTE` error message might be generated in certain scenarios when the rpd restarts or GRES when you enable NSR, which has no functional impact. [PR1586466](#)
- Removing SIB without switching it to the `Offline` mode might impact traffic. [PR1586820](#)

- Telemetry leaves used-power and allocated-power under the `/components` file do not display correct value. [PR1587184](#)
- On PTX10008 routers, the error message that gets generated when you issue the request `chassis cb slot 1 offline` command needs to be corrected before the node goes offline. [PR1589433](#)
- Duplicate Junos Telemetry Interface leaf `oper-status` tag for IFD index 16386 displays mismatched value. [PR1592468](#)
- The `I2cpd-agent` process might become unresponsive after starting the telemetry service. [PR1592473](#)
- After the Routing Engine switchover, the error message `JexprSlowCntrRead - Unable to get the plct Inst for pfeIdx: 255, User-type: OVFM_OFFCHIP_NEXTHOP_CNTR` gets generated. [PR1593079](#)
- The TCP connections to the telemetry server might become unresponsive in the `CLOSE_WAIT` state. [PR1593113](#)
- The BFD session for MPLS LSP goes down after enabling ultimate-hop-popping. [PR1594621](#)
- PTX10008 routers:: Inconsistent component name for FPC CPU. [PR1595109](#)
- On PTX10008 routers, the application error alarms gets generated and the trace-writer process generates core file due to defunct rcp zombie. [PR1595409](#)
- Layer 2 VPN stops forwarding when you change the interface encapsulation to VLAN circuit cross-connect (VLAN CCC) from ethernet- circuit cross-connect (Ethernet CCC) and back. [PR1595455](#)
- Some TCP sessions might not be established after you issue the request `system snapshot` command. [PR1595470](#)
- `[Error]Jexpr: cannot find ifToken for counterType:12`. [PR1597355](#)
- The `aftmand` process might generate core file on Junos OS Evolved platforms. [PR1597649](#)
- On PTX10016 routers, the `Major Host # Ethernet Interface Link Down false` alarm gets generated. [PR1597763](#)
- On PTX10001-36MR routers, there is inconsistency in the platform name that gets used in multiple places, version, and snmp mibs. [PR1597999](#)
- Master-only IP address in the old primary Routing Engine (new backup) and the device become inaccessible after the Routing Engine switchover. [PR1598173](#)
- On PTX10008 routers, the BFD sessions do not go down after adding the input or output filter to block the BFD IPv4 or IPv6 packets. [PR1599257](#)
- On PTX10008 routers, the FTC status LED and SIB power LED are unlit or off. [PR1600178](#)

- The config interface `ip remove` command does not work correctly. [PR1600932](#)
- On PTX10008 routers, the `set chassis redundancy routing-engine 1 master` command does not change the default Routing Engine election priority. [PR1601430](#)
- After you delete the old VLAN and then add new VLAN, the lag interface is not learnt by the static VLAN. [PR1601915](#)
- On PTX10008 routers, the `aftman` process generates core files at `jexpr_if_logical_l2d_alloc` while you power off or on all the Packet Forwarding Engines across all FPCs. [PR1602035](#)
- On PTX10003 routers, the IRB ping fails after you power off or on the underlying Packet Forwarding Engine for the aggregated Ethernet child member. [PR1602181](#)
- On PTX10003 routers, the GRE keepalive packet with recursion control bit set get dropped. [PR1602353](#)
- On PTX10008 routers, powering off the Packet Forwarding Engine generates the `Jexpr: deleteFdbEntry: Null` error messages. [PR1602670](#)
- The `show system errors fru detail` command does not display `reset-pfe` as the `cmerror` configured action. [PR1602726](#)
- The `evo-aftmand-bt` process generates the `[Error] IfStats:map entry not present for ifl:1039` error message. [PR1604334](#)
- The channel 0 IFDs does not come up after adding the correct speed configuration. [PR1604810](#)
- Remote aggregated Ethernet interface member failure (through disable or laser-off) might cause high tail drop that results in a high traffic loss. [PR1604823](#)
- The host loopback wedge might be detected in the Packet Forwarding Engine when you delete the aggregated Ethernet interface bundle configuration. [PR1605599](#)
- On PTX10008 routers, the Fan Tray Controller removal/absence alarm gets generated. [PR1605987](#)
- Segment Routing License issue might occur when you use the default chained-composite-next-hop configuration. [PR1606377](#)
- On PTX10008 routers, around 500ms to 800ms of traffic might be lost with one of the aggregated Ethernet interface member links of p2mp LSP branches down. [PR1606839](#)
- Memory might leak on the `l2cpd` process when you perform certain LLDP operations. [PR1608699](#)
- On PTX10000 routers, defunct `rcp` processes increase that might cause the primary Routing Engine to reboot. [PR1608776](#)

- On PTX10008 routers, the `evo-aftmand-bt.fpc_x86_64` process generates core file at `jexpr_pile_malloc` with LSR core profile configuration. [PR1608999](#)
- High priority queue might not get the expected bandwidth on Junos OS Evolved platforms. [PR1609823](#)
- The `show pfe statistics traffic` command does not display the host bound traffic. [PR1611115](#)
- After PICD restarts, the interface goes into the `Down` state in the channelized 100G link. [PR1611379](#)
- The IS-IS session might not come up when the network type is `p2p` for IRB interface. [PR1612606](#)
- Some of the fabric links might go into the `Faulty` state after swapping the FPC LC1201 with LC1202. [PR1612624](#)
- Mitigate false wrap gets drops statistics when IFDs move into or out-of an aggregated Ethernet interface while IFD drops excess traffic. [PR1613889](#)
- On PTX10016 routers, SIBs and FPCs go into the `Active` or `Online` state twenty-two minutes after loading the Junos OS Evolved Release image 21.2R1-S1-202108130158.0-EVO. [PR1614489](#)
- Line cards might become unstable due to the continuous memory usage of the `evo-cda-bt` application. [PR1614952](#)
- Twenty-seven percent traffic loss appears at the 221B packet size in Junos OS Evolved Release 21.4 as compared to Junos OS Evolved Release 21.2R1. [PR1615524](#)
- On PTX10008 routers, after FPC becomes offline, minor `cm-errors` might occur for ZFI block on other FPCs. [PR1616179](#)
- Inconsistent error counts appears in the `show interfaces brief` and `show interfaces extensive` command. [PR1616765](#)
- The `aftmand` process generates core files at `RtIfaHandler::notifyCommand,EalIfaHandler::registryClientCommand ,EalIfaHandler::OnAdd (this=0x7f2ffe40e9a0 EalIfaHandler::instance()::handler, ifah=...)` at `../../src/EalIfaHandler.cpp:222`. [PR1616909](#)
- While migration from Junos OS to Junos OS Evolved, you must delete the chassis redundancy failover or disable the chassis redundancy failover. [PR1617720](#)
- Multicast traffic to a mixed high-priority and low-priority OQs of a busy port might suffer drops of higher-priority multicast. [PR1618026](#)
- On PTX10001-36MR routers, issuing `request system snapshot` command does not capture the snapshot message in the `/etc/motd` file. [PR1618946](#)
- `InputIntf` is reported incorrectly for the MPLS-IPv4 and MPLS-IPv6 ingress sampling in case of Layer 3 VPN. [PR1619052](#)

- The hwdre process might crash when an FPC gets pulled out or some FPC power failure occurs. [PR1619102](#)
- The `/interfaces/interface/subinterfaces/subinterface/state/counters` are not exported during the initial synchronization for an on-change. [PR1620160](#)
- PIC becomes unresponsive in the Offline state when you issue the Offline command right after the transceiver plugin. [PR1621694](#)
- ZTP does not work properly on all the PTX Series routers if you use an EX Series switch as a DHCP server. [PR1621987](#)
- On PTX10008 routers, interface goes into the Down state while performing Custom Optics Profile validation for low power mode in the non-channelized mode. [PR1624228](#)
- On PTX10008 routers, the continuous information level `evo-aftmand-bt:Pfe:controller add for fru :controller modify for fru:fru power-on (block mode) for fryu:initiating online (block mode) for fru` syslog message might be generated. [PR1624375](#)
- The `show pfe route ip` command times out when the route table size is large. [PR1624629](#)
- Need to add the `show cda pipestats` command to more Junos OS Evolved Releases for serviceability. [PR1626687](#)
- Junos OS Evolved-based PTX platforms observes incomplete objects anomalies with scale beyond 16,000 logical interfaces in the system. [PR1573994](#)
- On PTX10008 routers, CB 1 goes in to the Fault Standby state after issuing the `request node power-off re1` command. [PR1581476](#)
- On PTX10008 routers, the EVPN-VXLAN shared tunnel commands must be removed. [PR1598142](#)
- GNMI Set RPC does not work with multiple operations. [PR1609436](#)
- The CDA-BT process generates a core file when turn FPC offline. [PR1615343](#)
- Device does not respond to traceroute while checking traceroute over VPN. [PR1615677](#)
- The Strict-Priority-Scheduler (SPS) might not work accurately across the port queues. [PR1616772](#)
- Limitation with fast-lookup-filter (FFT) usage in the ISF mode exist. [PR1616804](#)
- SNMP get for MIB value for `jnxRedundancyConfig` does not work as expected. [PR1621101](#)
- SNMP get for MID ID for `jnxRedundancySwitchoverReason` does not work as expected. [PR1621103](#)
- Telnet service might be enabled when you disable the service. [PR1596411](#)
- The `evo-aftmand` process might crash in a scaled configuration scenario. [PR1597988](#)

- Telemetry data collection does not work on the FPC:PIC:PORT component if you configure telemetry on Junos OS Evolved platforms. [PR1599665](#)
- The following error message gets generated after the FPC restarts:

```
NextHop: Comp_nh: 181429 Child nh count zero, do not create state181429
```

[PR1600642](#)

- Software validation or upgrade might fail on the Junos OS Evolved platforms. [PR1603479](#)
- The telnet or ssh session might not be established on the PTX Series platforms running Junos OS Evolved. [PR1605753](#)
- The aggregated Ethernet interfaces with 32 or more member-links might observe larger load-balance deviation between the member-links. [PR1607708](#)
- Power consumption report does not display correct values. [PR1608607](#)
- IPv6 reachability fails in management-instance. [PR1608619](#)
- Session connections might reset if the backup Routing Engine with NSR configured on Junos OS Evolved platforms. [PR1611520](#)
- The hwdfpc process might crash on all Junos OS Evolved platforms. [PR1624841](#)
- The IP not-ECN-capable traffic does not get RED-dropped in an ECN-enabled congested queue. [PR1627496](#)
- Indirect next-hop (INH) version ID higher than 255 might cause INH NH FRR session in to the Down state and drop the transit traffic. [PR1630215](#)

Infrastructure

- Egress TCP RST might not correctly populate the DSCP field. [PR1612208](#)
- The Host 0 Active Disk Usage Exceeded alarm might be generated due to a large number of files under the `/var/log/journal` file, which were held by the eventd daemon. [PR1601251](#)
- ICMP tunneling might not work. [PR1605465](#)
- The default-address-selection command might not work on all Junos OS Evolved platforms. [PR1608877](#)

Interfaces and Chassis

- The sensor of the Junos Telemetry Interface optics alarm data type changes from `bool_val` to `str_val`. [PR1580113](#)
- SIB might become unresponsive at the `Offlining` state after performing offline and online operations. [PR1591076](#)
- The 25G interfaces with FEC91 goes in the `Down` state on a few configurations. [PR1594740](#)
- On PTX10003 routers, the `show platform object-info anomalies summary` command times out. [PR1598337](#)
- The LACP system priority might take a value of 0 and cause an LACP interoperability issue. [PR1602724](#)
- A few links on the channelized interface goes in the `Down` state after `oir_enable` and `oir_disable` in 4X25G. [PR1606644](#)
- IRB VGA MAC IP is missed if you configure the IRB logical interface first with a `mac-address`. [PR1611559](#)

MPLS

- The `rpd` process generates core file in the backup Routing Engine at `mirror_process_recvd_data_queue` with the MLDP NSR configuration. [PR1594405](#)
- Static LDP P2MP might fail after NSR switchover. [PR1598344](#)
- SNMP traps does not get generated for the second trap group. [PR1623201](#)
- Traffic loss might occur in some SR-TE scenario. [PR1602127](#)

Network Management and Monitoring

- The syslog archival transfer fails if you configure the archive site URL with an IPv6 address. [PR1603342](#)
- Incorrect `IF-MIB::ifHCInUcastPkts` and `ifHCInBroadcastPkts` statistics gets displayed. [PR1621606](#)
- Syslog messages might be lost partially in case of excess messages generated to the `eventd` process. [PR1612535](#)

Routing Policy and Firewall Filters

- The dfwd-junos-relay process might generate core file during the Routing Engine switchover. [PR1597853](#)
- The configuration check fails if you configure more than eight FCs and enable CBF. [PR1600544](#)
- The dfwd process crashes when the no-decrement-ttl filter match condition and action is sent from the control plane to the Packet Forwarding Engine. [PR1602645](#)
- The firewall process might crash if you configure the fragment-offset out of the range (fragment-offset 1-900000000000). [PR1605805](#)
- Filters in openconfig acl execute terms in the order of their definition and not based on sequence-ids. [PR1621620](#)
- Services might not work after committing the firewall filter counter configuration with similar name of two terms. [PR1625168](#)

Routing Protocols

- The IPv4 static route might still forward traffic unexpectedly even when the static route configuration has already been deleted. [PR1599084](#)

User Interface and Configuration

- Updates to the system login configuration might not reflect after a commit. [PR1589858](#)
- The file copy command does not accept the HTTPS URIs. [PR1596881](#)
- Configuration transfer-on-commit does not work if you commit through netconf. [PR1602331](#)
- Authentication fails while logging on Junos OS Evolved platform through SSH. [PR1594327](#)
- The file copy failure occurs through the Netconf or operation script. [PR1597550](#)

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

IN THIS SECTION

- [What's New | 75](#)
- [What's Changed | 81](#)
- [Known Limitations | 84](#)
- [Open Issues | 85](#)
- [Resolved Issues | 86](#)

These release notes accompany Junos OS Evolved Release 21.4R1 for QFX Series switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- [What's New in 21.4R1 | 75](#)

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

What's New in 21.4R1

IN THIS SECTION

- [What's Changed in Release 21.4R1-S2 | 76](#)

- [What's Changed in Release 21.4R1-S1 | 76](#)
- [Device Security | 77](#)
- [EVPN | 77](#)
- [Juniper Extension Toolkit \(JET\) | 78](#)
- [Network Management and Monitoring | 78](#)
- [Routing Policy and Firewall Filters | 78](#)
- [Routing Protocols | 79](#)
- [Software Installation and Upgrade | 79](#)
- [Additional Features | 80](#)

What's Changed in Release 21.4R1-S2

IN THIS SECTION

- [User Interface and Configuration | 76](#)

User Interface and Configuration

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after rollback timeout.

What's Changed in Release 21.4R1-S1

IN THIS SECTION

- [User Interface and Configuration | 77](#)

User Interface and Configuration

- When you configure `max-cli-sessions` at the **edit system** hierarchy level, it restricts the maximum number of CLI sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login.

Device Security

- **Juniper Malware Removal Tool (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220)**—Starting in Junos OS Evolved Release 21.4R1, you can use the Juniper Malware Removal Tool (JMRT) to scan for and remove malware running on Junos OS Evolved devices. You can run two types of scans:
 - Quick scan—To scan each running program file
 - Integrity check—To check whether integrity mechanisms are enabled

[See [Juniper Malware Removal Tool](#) and [request system malware-scan](#).]

EVPN

- **EVPN Type 2 and Type 5 route coexistence (QFX5130-32CD, QFX5130-48C, and QFX5700)**—Starting in Junos OS Evolved Release 21.4R1, we support the coexistence of EVPN Type 2 and Type 5 routes in EVPN-VXLAN edge-routed bridging (ERB) overlay fabrics. This feature enables more efficient traffic flow and better usage of Packet Forwarding Engine resources. The switch applies a preference algorithm when you enable Type 5 routes. For any destinations for which the switch has no Type 5 route, the switch uses Type 2 routes by default. Otherwise, the switch gives preference to:
 - Type 2 routes for local ESI interfaces (locally learned routes)
 - Type 5 routes for all other destinations within the data center or across data centers

You can refine these preferences by configuring routing policies in the EVPN routing instance to control the Type 5 routes that the switch imports and exports.

[See [EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN](#).]

- **DHCP relay in EVPN-VXLAN (QFX5130-32CD)**—Starting in Junos OS Evolved Release 21.4R1, we support DHCP relay in an EVPN-VXLAN fabric. You can configure DHCP relay in centrally-routed and edge-routed bridging overlays. Support for DHCP relay includes DHCPv4 and DHCPv6. We introduced this feature in Junos OS Evolved Release 21.2R2..

[See [DHCP Relay Agent over EVPN-VXLAN](#).]

Juniper Extension Toolkit (JET)

- **Support for programming FTIs using JET APIs (PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 21.4R1, you can use the Interfaces Service API to configure flexible tunnel interfaces (FTIs) in Junos OS. You can change the attributes of the tunnel configurations for the unit under an existing FTI but cannot change the existing tunnel encapsulation type using the APIs. For the following families, you can configure only the listed attributes when you use Juniper Extension Toolkit (JET) APIs:

- INET and INET6: address and destination-udp-port
- MPLS and ISO: destination-udp-port

[See [Overview of JET APIs](#) and [Configure Flexible Tunnel Interfaces](#).]

Network Management and Monitoring

- **DHCP stateless relay MIB support (PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130, and QFX5220)**—Starting in Junos OS Evolved Release 21.4R1, you can use the Juniper Networks enterprise-specific DHCPv6 and DHCPv6 MIBs to retrieve statistics for DHCP stateless relay. We provide support only for the following MIB objects:

- jnxJdhcpRelayStatistics
- jnxJdhcpRelayIfcStats
- jnxJdhcpv6RelayStatistics
- jnxJdhcpv6RelayIfcStats

[See [Enterprise-Specific MIBs for Junos OS Evolved](#).]

- **SNMP MIB support for ICMP (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX1003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 21.4R1, we've added support for the following RFC4293 tables:

- icmpStatsTable—Generic system-wide ICMP counters
- icmpMsgStatsTable—System-wide per-version and per-message type ICMP counters

[See [Standard MIBs for Junos OS Evolved](#).]

Routing Policy and Firewall Filters

- **Support for profiles to improve the firewall filter scale (QFX5130-32CD, QFX5700, and QFX5220)**—Starting in Junos OS Evolved Release 21.4R1, you can apply firewall filters for inet and Ethernet-based switching using firewall filter profiles. You can use the profiles configuration statement at the

[edit system packet-forwarding-options firewall] hierarchy level to configure firewall filter profiles. The firewall filter profiles are mapped to a subset of match conditions. This helps you to plan and apply firewall filter profiles to achieve maximum scale.

You can use the following CLI commands to display the profile information and the pipe that each physical interface is mapped to:

- show pfe filter hw profile-info
- show pfe filter hw port-pipe-info

[See [How to Increase the Scale of Firewall Filters Using Profiles.](#)]

Routing Protocols

- **Support for ICMP extension (QFX5100)**—Starting in Junos OS Evolved Release 21.4R1, for both numbered and unnumbered aggregated Ethernet interfaces, we've implemented RFC5837 to enable us to append additional fields to the following ICMP (IPv4 and IPv6) messages:
 - ICMPv4 Time Exceeded
 - ICMPv4 Destination Unreachable
 - ICMPv6 Time Exceeded
 - ICMPv6 Destination Unreachable

Use the set system allow-icmp4-extension command to enable ICMP extension.

[See [Configure ICMP Features.](#)]

Software Installation and Upgrade

- **Support for DHCPv6 on ZTP (QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 21.4R1, zero-touch provisioning (ZTP) supports the DHCPv6 client on the management interface. During the bootstrap process, the device first uses the DHCPv4 client to request for information regarding image and configuration file from the DHCP server. The device checks the DHCPv4 bindings sequentially. If there is a failure with one of the DHCPv4 bindings, the device continues to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device checks for DHCPv6 bindings and follows the same process as for DHCPv4 until the device can be provisioned successfully. The DHCP server uses DHCPv6 options 59 and 17 and applicable sub-options to exchange ZTP-related information between itself and the DHCP client.

[See [Zero Touch Provisioning.](#)]

- **Support for outbound SSH service (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220-32CD, and QFX5220-128C)**—Starting in Junos OS Evolved Release 21.4R1, we support outbound SSH service.

You can use the `restart service-deployment` command to restart the Service Deployment System (SDX) process.

[See [outbound-ssh](#).]

Additional Features

We've extended support for the following features to these platforms.

- **ICMPv6 and ARP statistics support for Junos OS Evolved platforms (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700)**—The existing ARP and ICMPv6 statistics reports are enhanced to support the following statistics.
 - ARP
 - Proxy requests
 - Requests dropped on entry
 - Number of arp packets dropped as nexthop allocation failed
 - ICMPv6
 - Output histogram—unreach, echo request, echo reply, neighbor solicitation, neighbor advertisement, router solicitation, and router advertisement details
 - Input histogram—unreach, echo request, echo reply, neighbor solicitation, neighbor advertisement, router solicitation, and router advertisement details
 - messages with bad code fields, messages < minimum length, bad checksums, messages with bad length
 - message responses generated, messages with too many ND options, current public ND nexthops present
 - [See [show system statistics arp](#) and [show system statistics icmp6](#)]
- **IPsec support for OSPFv2 and OSPFv3 (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10008, PTX10016, QFX5130-48C, and QFX5220)**

[See [Overview of IPsec](#), [Configuring OSPF Authentication](#), and [Configuring IPsec Security Associations](#).]
- **Support for MIB accounting profiles (PTX10001-36MR, PTX10004, PTX10008, PTX10016, and QFX5220)**

[See [Accounting Profiles: An Alternative to SNMP Statistics](#).]

- **Support for resilient hashing on LAG groups (QFX5130-32CD and QFX5700)**

[See [Resilient Hashing on LAGs and ECMP groups](#).]

- **Support for 400GbE QSFP-DD DAC cables (QFX5700)**—Starting in Junos OS Evolved Release 21.4R1, QFX5700 switches support 400Gbps speed using QDD-400G-DAC-1M and QDD-400G-DAC-2P5M DAC cables.

[See [Hardware Compatibility Tool](#).]

- **Supported transceivers, optical interfaces, and DAC cables (ACX Series, PTX Series, and QFX Series)**—Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

What's Changed

IN THIS SECTION

- [What's Changed in Release 21.4R1 | 81](#)

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

What's Changed in Release 21.4R1

IN THIS SECTION

- [Class of Service | 82](#)
- [EVPN | 82](#)
- [General Routing | 82](#)
- [Interfaces and Chassis | 82](#)
- [Junos OS API and Scripting | 82](#)
- [Network Management and Monitoring | 83](#)
- [Routing Protocols | 83](#)

- Software Installation and Upgrade | 84
- Software Licensing | 84

Class of Service

- Junos OS Evolved now correctly displays the index for `show class-of-service` commands.

EVPN

- **Community information no longer included in VRF routing table**—The QFX series switches will no longer include the inherited advertised route target communities, EVPN extended communities, or vxlan encapsulation communities for EVPN Type 2 and EVPN Type 5 routes when an IP host is added in the VRF routing table.
- **Output for show Ethernet switching flood extensive**—The output for `show ethernet-switching flood extensive` now displays the correct next-hop type for Virtual Ethernet and WAN mesh group in an EVPN-VXLAN network as unicast. Previously, the output for `show ethernet-switching flood extensive` would misidentify the next-hop type as composite.

General Routing

- In Junos OS Release 21.4R1 and later, `ssh` is enabled by default on all the routers with VM host support.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Previously, the source and destination address pair had to be unique across all encapsulation type.

Junos OS API and Scripting

- **Limits increased for the `max-datasize` statement (ACX Series, PTX Series, and QFX Series)**—The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data

segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

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](#).]

- **Change in behavior of SNMP MIB object ifAlias**—SNMP MIB object `ifAlias` now shows the configured interface alias. In earlier releases, `ifAlias` used to show configured interface description.
- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, PTX Series, and QFX Series)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

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

- **The configuration accepts only defined identity values for nodes of type identityref in YANG data models (ACX Series, PTX Series, and QFX Series)**—If you configure a statement that has type `identityref` in the corresponding YANG data model, the device accepts only defined identity values (as defined by an identity statement) as valid input. In earlier releases, the device also accepts values that are not defined identity values.

Routing Protocols

- **The RPD_OSPF_LDP_SYNC message not logged?**On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the `hold-time` for `ldp-synchronization` at the **edit protocols ospf area area id interface interface name** hierarchy level less than three minutes. The message is printed after three minutes.
- To achieve consistency among resource paths, the resource path `/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter ip-addr='address'/state/countersname='name'/out-pkts/` is changed to `/mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/counters name='name'/`. The

leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").

Software Installation and Upgrade

- **Series profile1 image deprecation**—Starting in Junos OS Evolved 21.4R1, we've deprecated the Series Profile1 software image (also referred to as the lean rpd profile image) for QFX5130 and QFX5220 series devices. All other software image packages continue to be available for those devices. [See [Junos OS Evolved Installation Packages](#).]

Software Licensing

- **Juniper Agile Licensing (PTX10003, PTX10016, QFX5130-32CD, and QFX5220)**—Starting from this release onwards, the Juniper Agile License Manager is deprecated. You can use the Juniper Agile Licensing Portal to activate, install, manage, and monitor licenses on Juniper Networks devices. [See [Juniper Agile Licensing Guide](#).]

Known Limitations

IN THIS SECTION

- [General Routing | 84](#)

Learn about known limitations in Junos OS Evolved Release 21.4R1 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 QFX5700 platforms, the system reboots when faulty FEB is off-lined for recovery. [PR1578090](#)
- 2x200G(QSFP56-DD-2x200GBASE-CR4-CU) channel that is connected with the server through vendor connectX-6 NIC might take extra time for the link to come up. [PR1617676](#)

Open Issues

IN THIS SECTION

- [General Routing | 85](#)
- [Routing Policy and Firewall Filters | 85](#)

Learn about open issues in Junos OS Evolved Release 21.4R1 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

- 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 QFX5700 switches, few interfaces are not coming up after removing channelization through single commit that is by using **delete interfaces**. [PR1592238](#)
- On QFX5700, auto-negotiation is not supported for 100G DAC interfaces. [PR1594135](#)
- In rare situations, a 40G interface might not come up on QFX5700. A device reboot will be required to come out of this situation. [PR1637876](#)

Routing Policy and Firewall Filters

- On QFX5700 when we issue `clear firewall all` command, the policer statistics are not cleared. User defined policer statistics are removed. [PR1581867](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 21.4R1 | 86](#)

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: 21.4R1

IN THIS SECTION

- [General Routing | 86](#)
- [Infrastructure | 87](#)
- [User Interface and Configuration | 87](#)

General Routing

- If more than 512 OSPF neighbors are configured, all the OSPF routes might not be learned unless the MTU of the IRB interface is modified. [PR1570498](#)
- 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](#)
- Traffic loss might be observed when `evo-pfemand` is restarted. [PR1608004](#)
- In QFX5700 FPC, QSFP port reset signal to be inverted in the software. [PR1608849](#)
- The egress traffic is dropped due to the egress queue buffer stuck on 400G interfaces. [PR1618147](#)
- Junos OS Evolved:JDI_FT_REGRESSION:ui:ui_change :: On performing request `system snapshot`, the snapshot message is not captured in `/etc/motd` file. [PR1618946](#)

Infrastructure

- The alarm **Host 0 Active Disk Usage Exceeded** might be generated due to large files which were already marked as deleted. [PR1601251](#)

User Interface and Configuration

- The file copy command is not accepting HTTPS URLs. [PR1596881](#)

Upgrade Your Junos OS Evolved Software

Products impacted: ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, 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.
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

- [Self-Help Online Tools and Resources | 90](#)
- [Creating a Service Request with JTAC | 91](#)

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

7 August 2024—Revision 11, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

14 December 2023—Revision 10, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

10 August 2023—Revision 9, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

20 July 2023—Revision 8, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

24 November 2022—Revision 7, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

29 July 2022—Revision 6, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

16 May 2022—Revision 5, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

14 March 2022—Revision 4, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

17 February 2022—Revision 3, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

20 January 2022—Revision 3, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

28 December 2021—Revision 2, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

16 December 2021—Revision 1, Junos OS Release 21.4R1 for the ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, QFX5220, and QFX5700 Devices.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2024 Juniper Networks, Inc. All rights reserved.