Release Notes

Published
2023-08-09

Junos ® OS 20.2R1 Release Notes

SUPPORTED ON

- ACX Series, EX Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, JRR Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series

SOFTWARE HIGHLIGHTS

- Retain the authentication session based on DHCP or SLAAC snooping entries (EX Series)
- Rest API support for EX2300, EX2300-MP, EX3400, EX4300, EX4300-MP, EX4600, EX4650, and EX9200
- TI-LFA SRLG protection for IS-IS (MX Series and PTX Series)
- MX Series Virtual Chassis support for the ephemeral database (MX Series)
- Change the default re-merge behavior on the P2MP LSP (MX Series)
- BGP-LU over SR-MPLS and IS-IS segment routing underlay
- Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast (MX Series and EX Series)
- Packet capture of unknown application traffic (NFX Series, SRX Series, and vSRX)
- Safe search enhancement for Web filtering (SRX Series and vSRX)
- Encrypted traffic analysis
- Support for Application Quality of Experience (AppQoE) (SRX4600)

IN FOCUS GUIDE

- Use this new guide to quickly learn about the most important Junos OS features and how you can deploy them in your network.

Release Notes: Junos® OS Release 20.2R1 for the ACX Series, EX Series, Junos Fusion, JRR Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series

20 July 2023

Contents

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, Junos Fusion, JRR Series, M Series, MX Series, NFX Series, PTX Series, QFabric systems, QFX Series, and SRX Series, T Series.

These release notes accompany Junos OS Release 20.2R1 for the ACX Series, EX Series, Junos Fusion, JRR Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

- In Focus guide—We have a document called In Focus that provides details on the most important features for the release in one place. We hope this document will quickly get you to the latest information about Junos OS features. Let us know if you find this information useful by sending an e-mail to techpubs-comments@juniper.net.
- **Important Information:**
  - Upgrading Using ISSU on page 270
  - Licensing on page 270
  - Compliance Advisor on page 271
  - Finding More Information on page 271
  - Documentation Feedback on page 271
  - Requesting Technical Support on page 273

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for ACX Series routers.

Hardware

- **New ACX710 Universal Metro Routers (ACX Series)**—In Junos OS Release 20.2R1, we introduce the ACX710 router. The ACX710 is a compact 1-U router that provides system throughput of up to 320 Gbps through the following port configurations:
  - Twenty-four 10GbE or 1GbE ports (ports 0 through 23) that operate at 10-Gbps speed when you use small form-factor pluggable plus (SFP+) transceivers or at 1-Gbps speed when you use small form-factor pluggable (SFP) optics. Ports 0 through 15 also support 1000 Mbps speeds when you use tri-rate SFP optics. Ports 16 through 23 support 100 Mbps and 1000 Mbps speeds when you use tri-rate SFP optics.

- Four 100GbE ports (ports 0 through 3) that support quad small form-factor pluggable 28 (QSFP28) transceivers. You can channelize these ports into four 25-Gbps interfaces using breakout cables and channelization configuration. These ports also support 40-Gbps speed when you use quad small form-factor pluggable plus (QSFP+) optics. You can channelize these 40-Gbps ports into four 10-Gbps interfaces using breakout cables and channelization configuration. [See Channelize Interfaces on ACX710 Routers.]

The ACX710 router is a DC-powered device that is cooled using a fan tray with five high-performance fans to cool the chassis.

To install the ACX710 router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the ACX710 Universal Metro Router Hardware Guide.

Table 1 on page 14 summarizes the ACX710 features supported in Junos OS Release 20.2R1.

Table 1: Features Supported by the ACX710 Routers

Feature	Description
Class of service (CoS)	• Standard CoS feature support, including configuring classification, rewrite, shaping, buffering, and scheduling parameters for traffic management. [See CoS on ACX Series Routers Features Overview.]
DHCP	• DHCP server, DHCP client, and DHCP relay configuration for IPv4 and IPv6 services. [See Understanding DHCP Client Operation on ACX Series.]
EVPN	• EVPN-VPWS. [See Overview of VPWS with EVPN Signaling Mechanisms EVPN-VPWS with flexible cross-connect (FXC).] • EVPN-VPWS with flexible cross-connect (FXC). [See Overview of Flexible Cross-Connect Support on VPWS with EVPN.] • EVPN with ELAN services over MPLS. [See EVPN Overview.]
Firewalls and policers	• Configure firewall filters on packets (families such as bridge domain, IPv4, IPv6, CCC, and MPLS) based on packet match conditions. Along with the match conditions, actions such as count, discard, log, syslog, policer are performed on the packets that match the filter. You can configure policers and attach them to a firewall term. [See Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview.]

Table 1: Features Supported by the ACX710 Routers *(continued)*

Feature	Description
High availability (HA) and resiliency	• VRRP protocol support with Broadcom's DNX chipset. [See Understanding VRRP Overview.] • Configure alarm input and output, manage FRUs, and monitor environment. The router also supports field-replaceable unit (FRU) management and environmental monitoring. [See alarm-port.] • Platform resiliency to handle failures and faults of the components such as fan trays, temperature sensors, and power supplies. The router also supports firmware upgrade for FPGA and U-boot. [See show chassis alarms and show system firmware.]
Layer 2 features	• Layer 2 support: bridging, bridge domain with no vlan-id, with vlan-id none, or with single vlan-id, single learning domain support,.Q-in-Q service for bridging, MAC limit feature support, no local switching support for bridge domain, and E-LINE from a bridge with no MAC learning. [See Layer 2 Bridge Domains on ACX Series Overview.] • Layer 2 support for bridge interfaces for vlan-map push operation, swap operation, pop operation, and swap-swap operation. [See Layer 2 Bridging Interfaces Overview.] • Layer 2 support for control protocols (L2CP): RSTP, MSTP, LLDP, BPDU guard/protection, loop protection, root protection, Layer 2 protocol tunneling, storm control, IRB interface, LAG support with corresponding hashing algorithm, E-LINE, E-LAN, E-ACCESS, and E-Transit service over L2/Bridge with the following AC interface types: Port, VLAN, Q-in-Q, VLAN range and VLAN list. [See Layer 2 Control Protocols on ACX Series Routers.] • Layer 2 circuit cross-connect (L2CCC) support for Layer 2 switching cross-connects. You can leverage the hardware support available for cross-connects on the ACX710 device with the Layer 2 local switching functionality using certain models. With this support, you can provide the EVP and EVPL services. [See Configuring MPLS for Switching Cross-Connects.] • Reflector function support in RFC 2544. [See RFC 2544-Based Benchmarking Tests Overview.]

Table 1: Features Supported by the ACX710 Routers *(continued)*

Feature	Description
Layer 3 features	• Layer 3 VPN and Layer 3 IPv6 VPN Provider Edge router (6VPE) support over MPLS. The router uses MPLS as a transport mechanism with support for label-switching router (LSR), label edge routers (LERs), and pseudowire services. These protocols are also supported: ECMP, OSPF, IS-IS, and BGP. [See Understanding Layer 3 VPNs.] • Basic Layer 3 services over segment routing infrastructure. The segment routing features supported are: segment routing with OSPF through MPLS, segment routing with IS-IS through MPLS, segment routing traffic engineering (SR-TE), segment routing global block (SRGB) range label used by source packet routing in networking (SPRING), anycast segment identifiers (SIDs) and prefix SIDs in SPRING, and segment routing with topology independent (TI)-loop-free alternate (LFA) provides fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. [See Segment Routing LSP Configuration.] • Enhanced timing and synchronization support using Synchronous Ethernet with ESMC and BITS-Out. [See Synchronous Ethernet Overview and synchronization (ACX Series).] • Supports full-mesh VPLS domain deployment. The router supports interworking of both BGP as well as LDP-based VPLS. BGP can be used only for auto-discovery of the VPLS PEs, while LDP signaling for VPLS connectivity. [See Introduction to VPLS.]
MPLS	• Supports the Path Computation Element Protocol (PCEP). You can configure the PCEP implementation for both RSVP-TE and segment routing label-switched paths (LSPs). [See PCEP Configuration.] • Support for MPLS fast reroute (FRR) and unicast reverse-path forwarding (uRPF). [See fast-reroute (Protocols MPLS) and Guidelines for Configuring Unicast RPF on ACX Series Routers.] • Provides MPLS ping and traceroute support. [See MPLS Connectivity Verification and Troubleshooting Methods.]
Multicast	• Multicast support for IPv4 and IPv6 PIM-SM, SSM, IGMP snooping and proxy support, IGMP, IGMPv1/v2/v3 snooping, IGMP snooping support for LAG, global multicast support, MLD, and multicast support on IRB. [See Multicast Overview.]

Table 1: Features Supported by the ACX710 Routers *(continued)*

Feature	Description
Network management and monitoring	• TWAMP support. [See Two-Way Active Measurement Protocol on ACX Series.] • NETCONF sessions over TLS. [See NETCONF Sessions over Transport Layer Security (TLS).] • Support for adding custom YANG data models to the Junos OS schema [See Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS.] • Secure boot support in U-boot phase to authenticate and verify the loaded software image while also preventing software-based attack. [See Software Installation and Upgrade Guide.]
OAM	• IEEE 802.3ah standard for operation, administration, and management (OAM) connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM. [See IEEE 802.1ag OAM Connectivity Fault Management Overview.]
System management	• Zero-touch provisioning (ZTP) can automate the provisioning of the device configuration and software image. [See Software Installation and Upgrade Guide.]
To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported across all platforms, see the Hardware Compatibility Tool.	

Class of Service (CoS)

- **Support for hierarchical class of service (HCoS) (ACX5448)**—Starting with Junos OS Release 20.2R1, ACX5448 devices support up to four levels of hierarchical scheduling (physical interfaces, logical interface sets, logical interfaces, and queues). By default, all interfaces on the ACX5448 use port-based scheduling (eight queues per physical port). To enable hierarchical scheduling, set **hierarchical-scheduler** at the **[edit interfaces *interface-name*]** hierarchy level.

  [See Hierarchical Class of Service in ACX Series Routers.]

EVPN

- **Noncolored SR-TE LSPs with EVPN-MPLS (ACX5448, EX9200, MX Series, and vMX)—**Starting in Junos OS Release 20.2R1, ACX5448, EX9200, MX Series, and vMX routers support noncolored static segment routing-traffic engineered (SR-TE) label-switched paths (LSPs) with an EVPN-MPLS core network and the following Layer 2 services running at the edges of the network:

  - E-LAN

- EVPN-ETREE
- EVPN-VPWS with E-Line

Without color, all LSPs resolve using a BGP next hop only.

The Juniper Networks routers support noncolored SR-TE LSPs in an EVPN-MPLS core network with the following configurations:

- EVPN running in a virtual switch routing instance
- Multihoming in active/active and active/standby modes

The Juniper Networks routers also support noncolored SR-TE LSPs when functioning as a Data Center Interconnect (DCI) device that handles EVPN Type 5 routes.

[See Static Segment Routing Label Switched Path.]

Interfaces and Chassis

- **Port speeds and channelization (ACX710 routers)**—Starting in Junos OS Release 20.2R1, you can configure multiple speeds and interface channelization on our new ACX710 router. The router has 28 ports, which support the following speeds:
  - Ports **0** through **23** on PIC 0 support 1-Gbps speed (with SFP transceivers) and 10-Gbps speed (with SFP+ transceivers).
  - Ports **0** through **3** on PIC 1 support the default 100-Gbps speed (with QSFP28 transceivers) or the configured 40-Gbps speed (with QSFP+ transceivers). You can use the **set chassis fpc *slot-number* pic *pic-number* port *port-number* speed *speed*** CLI command and breakout cables to channelize each:
    - 100-Gbps port into four 25-Gbps interfaces
    - 40-Gbps port into four 10-Gbps interfaces

  [See Channelize Interfaces on ACX710 Routers.]
- **Ethernet OAM and BFD support (ACX710)**—Starting in Junos OS Release 20.2R1, the ACX710 routers support IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM.

  [See Introduction to OAM Connectivity Fault Management (CFM).]
- **Alarm port configuration, FRU management, and environmental monitoring (ACX710)**—Starting in Junos OS Release 20.2R1, you can configure the alarm port on the ACX710 router. You can use the alarm input to connect the router to external alarm sources such as security sensors so that the router receives alarms from these sources and displays those alarms. You can use the alarm output to connect the router to an external alarm device that gives audible or visual alarm signals based on the configuration. You can configure three alarm inputs and one alarm output by using the **alarm-port** statement at the **[edit**

chassis] hierarchy level. You can view the alarm port details by using the **show chassis craft-interface** command.

The ACX710 also supports FRU management and environmental monitoring.

[See alarm-port.]

- **Multichassis link aggregation groups, configuration synchronization, and configuration consistency check (ACX5448 routers)**—Starting in Junos OS Release 20.2R1, multichassis link aggregation (MC-LAG) includes support of Layer 2 circuit functionality with **ether-ccc** and **vlan-ccc** encapsulations.

  MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running spanning-tree protocols (STPs).

  [See Multichassis Link Aggregation Features, Terms, and Best Practices.]

Juniper Extension Toolkit (JET)

- **JET Clang toolchain supports cross-compiling JET applications for use on ARM platforms (ACX710)**—Starting in Junos OS Release 20.2R1, you can use the Clang toolchain to compile JET applications written in C, Python, or Ruby to run on the ARM architecture as well as Junos OS with FreeBSD and upgraded FreeBSD. The Clang toolchain for ARM is included in the JET software bundle. After you have downloaded the JET software bundle, you can access the Clang toolchain at **/usr/local/junos-jet/toolchain/llvm/**. Use the **mk-arm,bsdx** command to use the Clang toolchain to compile your application.

  [See Develop On-Device JET Applications.]

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

  [See language (Scripts), Develop Off-Device JET Applications, and Develop On-Device JET Applications.]

Junos Telemetry Interface

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

  [See Mapping OpenConfig Routing Policy Commands to Junos Configuration and Mapping OpenConfig Network Instance Commands to Junos Operation.]

MPLS

- **Support for MPLS ping and traceroute for segment routing (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.2R1, we extend the MPLS ping and traceroute support for all types segment routing--traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

  We also support the following features:

  - FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
  - ECMP traceroute support for all types of SR-TE paths.

  We do not support the following:

  - Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
  - OAM for IPv6 prefix.
  - BFD

[See traceroute mpls segment-routing spring-te and ping mpls segment routing spring-te.]

Multicast

- **Support for IPv6 multicast using MLD (ACX5448)**—Starting with Junos OS Release 20.2R1, ACX5448 routers support Multicast Listener Discovery (MLD) snooping with MLDv1 and MLDv2 for both any source multicast and SSM. Support for MLD snooping in EVPN was introduced in Junos OS Release 19.4R2.

  MLD snooping for IPv6 is used to optimize Layer 2 multicast forwarding. It works by checking the MLD messages sent between hosts and multicast routers to identify which hosts are interested in receiving IPv6 multicast traffic, and then forwarding the multicast streams to only those VLAN interfaces that are connected to the interested hosts (rather than flooding the traffic to all interfaces). You can enable or disable MLD snooping per VLAN at the **[edit protocols mld-snooping vlan *vlan-ID*]** hierarchy level. Note, however, that you cannot use ACX Series routers to connect to a multicast source.

  [See Understanding MLD Snooping, Understanding MLD, and Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment.]

Network Management and Monitoring

- **NETCONF sessions over TLS (ACX710)**—Starting in Junos OS Release 20.2R1, ACX710 routers support establishing Network Configuration Protocol (NETCONF) sessions over Transport Layer Security (TLS) to manage devices running Junos OS. TLS uses mutual X.509 certificate-based authentication and provides encryption and data integrity to establish a secure and reliable connection. NETCONF sessions over TLS enable you to remotely manage devices using certificate-based authentication and to more easily manage networks on a larger scale than when using NETCONF over SSH.

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

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **Support for port mirroring (ACX5448)**—Starting in Junos OS Release 20.2R1, you can use analyzers to mirror copies of packets to a configured destination. Mirroring helps in debugging network problems and also in defending the network against attacks. You can mirror all ingress traffic to a configured port (or port list), using a protocol analyzer application that passes the input to mirror through a list of ports configured through the logical interface. You configure the analyzer at the **[edit forwarding-options analyzer]** hierarchy level.

Configuration guidelines and limitations:

- Maximum of four default analyzer sessions
- LAGs supported as mirror output; a maximum of eight child members
- Not supported:
  - Egress mirroring
  - Mirroring on IRB, Virtual Chassis, or management interfaces
  - Nondefault analyzers

[See show forwarding-options analyzer.]

Routing Policy and Firewall Filters

- **Support for firewall filters and policers (ACX710)**—Starting with Junos OS Release 20.2R1, the ACX710 router supports configuring firewall filters on packets (families such as bridge domain, IPv4, IPv6, CCC, and MPLS) based on packet match conditions. Along with the match conditions, actions such as count, discard, log, syslog, and policer are performed on the packets that match the filter. You can configure policers and attach them to a firewall term.

  [See Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview.]

SEE ALSO

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for ACX Series routers.

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

  [See commit (System).]

- **New major alarms (ACX-710)** —We have introduced the following major alarms:
  - PTP No Foreign Master—Indicates that the external Precision Time Protocol (PTP) master is not sending announce packets.
  - PTP Sync Fail—Indicates that the PTP lock-status is not in Phase Aligned state.
  - Chassis Loss of all Equipment Clock Synch References—Indicates that both the primary and secondary SyncE references have failed and the chassis PLL is in holdover.
  - Chassis Loss of Equipment Clock Synch Reference 1—Indicates that the primary SyncE reference has failed, and no secondary SyncE reference is configured or present.
  - Chassis Loss of Equipment Clock Synch Reference 2—Indicates that you have configured at least two or more SyncE sources and the secondary SyncE source has failed.

  NOTE: These alarms get cleared when the system recovers from the error condition.

  See show chassis alarms**.**

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

  [See Develop Off-Device JET Applications and Develop On-Device JET Applications.]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

  [See Juniper EngNet.]

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—You can set the verbosity of the trace log to only show error messages using the error option at the edit system services extension-service traceoptions level hierarchy.

  [See traceoptions (Services).]

Network Management and Monitoring

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

SEE ALSO

What's New | 13

Known Limitations | 25

Open Issues | 28

Resolved Issues | 31

Documentation Updates | 34

Migration, Upgrade, and Downgrade Instructions | 34

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for ACX Series routers.

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

General Routing

- If Layer 2 VPN sessions have the OAM **control-channel** option set to **router-alert-label**, the **no-control-word** option in the Layer 2 VPN should not be used for BFD sessions to come up. PR1432854
- The time consumed on 1-Gigabit performance is not equal to that on 10-Gigabit performance. Compensation is done to bring the mean value under class A but the peak-to-peak variations are high and can go beyond 100 ns. It has a latency variation with peak-to-peak variations of around 125–250 ns without any traffic (for example, 5–10 percent of the mean latency introduced by each phy which is of around 2.5us). PR1437175
- On the ACX710 router, variable amount of time is taking to reflect the TWAMP packets. Because of this, the packet latency is not uniform. PR1477329
- On the ACX710 router, as per current design and BCOM input, load balancing does not work on any packet which is injected from host path. PR1477797
- On the ACX710 router, OSPF neighbors are not learned via VPLS connections because the **vlan-tags outer *vlan-id1* inner *vlan-id2*** statement is not supported in VPLS routing instance. PR1477957
- On the ACX710 router, sequential increment of both SRC and DST MAC do not provide better load balance as per HASH result. PR1477964
- On the ACX710 router, load balancing does not happen based on inner IP address when MPLS labelled traffic is received on NNI interface. PR1478945
- On the ACX710 router, for TCP protocol as well as for non-TCP protocol, loss-priority medium-low is not supported. PR1479164
- For ethernet-vpls encapsulation, if both DST IP and SRC IP are identically varied at the same octet, then hashing might not happen and leads to undefined behavior in load balancing on the ACX710 router. PR1479767

- For bridge LB with vlan-bridge encapsulation, if both SRC IP and DST IP are incremented or decremented by the same order (such as DIP = 10.1.1.1 (increment by 1 upto 100) and SIP = 20.2.3.1 (increment by 1 upto 100), then hashing does not happen on the ACX710 router. PR1479986
- For vlan-ccc encapsulation, if both SRC IP and DST IP are incremented or decremented by the same order (such as DIP = 10.1.1.1 (increment by 1 upto 100) and SIP = 20.2.3.1 (increment by 1 upto 100), then hashing does not happen on the ACX710 router. PR1480228
- On the ACX710 router, the input packet statistics for the **show interfaces** command represents the input packets at the MAC. The error packets which get dropped by MAC and that do not reach PHY will not be accounted. PR1480413
- Fragmentation or reassembly is not supported on ACX710 platforms due to the lack of hardware support. PR1481867
- On ACX5448 and ACX710 routers, each traffic stream is measured independently per port. Storm control is initiated only if one of the streams exceeds the storm control level. For example, if you set a storm control level of 100 Megabits and the broadcast and unknown unicast streams on the port are each flowing at 80 Mbps, storm control is not triggered. PR1482005
- On the ACX710 router, RFC2544 reports high latency and throughput loss when the packet size is 64 bytes at 100 percent line rate on the ASIC. The ASIC has low threshold value due to which packets are moved to DRAM from SRAM. When packets are moved to DRAM, high latency and packet drop are observed. PR1483370
- On the ACX710 router, VRRP over aggregated Ethernet interface is not supported. PR1483594
- On the ACX710 router, traffic loss is seen for segment routing, if protection (FRR) is enabled for 128 IPv6 prefix route. PR1484234
- Counters for PCS bit errors are not supported because of hardware limitations. Hence "Bit errors" and "Errored blocks" are not supported on an ACX710. PR1484766
- If any queue is configured with high priority, it is expected that accuracy of traffic distribution might vary for normal queues because of chip limitation. PR1485405
- For Layer 3 VPN configuration, sequential increment of both SRC IP and DST IP address would not provide better load balance as per hash result on the ACX710 router. PR1486406
- On the ACX710 router, double tagged interfaces implicit normalization to VLAN ID none is not supported. PR1486515
- On the ACX710 router, double tagged interfaces implicit normalization to VLAN ID none, ingress VLAN map operation, and pop-pop are not supported. PR1486520
- On the ACX710 router, packet priority at egress is derived from the internal priority. This internal priority is derived from the outer VLAN priority at ingress. Thus, the exiting packet retains the same priority as the ingress outer VLAN priority. PR1486571
- When you add or delete a configuration or a LAG member link flaps, configuration updates happen for all other members of the LAG too. This results in transient traffic drop on the ACX710 devices. PR1486997

- On the ACX710 router, double tagged ELMI and LLDP PDUs are dropped when L2PT is enabled for these protocols on the ingress interface. These PDUs are supposed to be untagged/native VLAN tagged and hence the drop. PR1487931
- On the ACX710 router, VLAN map operations like swap/swap does not work because the **vlan-tags outer *vlan-id1* inner *vlan-id2*** statement is not supported in VPLS routing instance. PR1488084
- On the ACX710 router, whenever the 100-Gigabit Ethernet interface is disabled, the alarm is not shown in the **jnxDomMib jnxDomCurrentLaneWarnings** and **jnxDomCurrentLaneAlarms**. PR1489940
- On the ACX710 router, in case of Layer 2 circuit, load balancing does not occur based on inner MAC address when MPLS labelled traffic is received on an NNI interface. PR1490441
- On the ACX710 router, unable to scale 1000 CFM sessions at 3 ms intervals; an error message is observed. PR1495753
- On ACX5448 routers, aggregated Ethernet LACP toggles with host path traffic with MAC rewrite configuration enabled. PR1495768
- The **traceroute mpls ldp** command does not work in case **explicit-null** is configured. It does not affect data path traffic. PR1498339
- On the ACX710 router, the convergence time for the traffic to switch over from the primary to the secondary link during link flap could be expected to be around 60 to 200 ms with the basic link aggregation configuration. PR1499965
- The maximum FIB route scale supported in an ACX710 router are as below:

  FIB IPv6 route scale - 80,000

  FIB IPv4 route scale - 170,000

  If routes are added above this scale, an error indicating **lpm route add** failure is reported. PR1515545

SEE ALSO

What's New | **13**

What's Changed | **23**

Open Issues | **28**

Resolved Issues | **31**

Documentation Updates | **34**

Migration, Upgrade, and Downgrade Instructions | **34**

Open Issues

IN THIS SECTION

Learn about open issues in this release for ACX Series routers.

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

General Routing

- Policer discarded packets are marked as color black. Black color is used to discard the packets in the pipeline. These packets are not really enqueued into the queues (VoQs) in hardware. The hardware queue statistics shows this as discarded. However today, both actual-enqueued and the discarded counts are shown as queue-stats in software. This is a software queue-statistics show issue. PR1414887
- DHCP clients are not able to scale to 96,000. PR1432849
- Protocols get forwarded when using non-existing SSM map source address in IGMPv3 instead of pruning. PR1435648
- Memory leaks are expected in this release. PR1438358
- When there is a failure of the I2C daemon, core files are generated on ACX5448. PR1455928
- On ACX5048 routers, the egress queue statistics are not working for the aggregated Ethernet interfaces. PR1472467
- On ACX710 routers, VPLS OAM sessions are detected with error(remote defect indication sent by some MEPs) after changing VLANs. PR1478346
- On ACX710 routers, initial few packet drop is observed after changing ALT port cost for RSTP. PR1482566
- On ACX710 routers, VRRP over dual tagged interface is not supported. PR1483759
- Issue is seen during unified ISSU to Junos OS Releases 20.2. Unified ISSU is completed, but the Packet Forwarding Engine does not function. Because of this, forwarding is affected. PR1483959
- On ACX710 routers, FEC of channel 0 in a channelized 25-Gigabit Ethernet interface is set to **None** while channels 1, 2, and 3 have FEC74 as the default value for 100G LR4 optics. The desired FEC value can be set through the CLI command **set interfaces et-x/y/z: *channel no* gigether-options fec *fec value***. PR1488040

- Queue statistics are not as expected after you configure physical interface and logical interface shaping with **transmit-rate** and **scheduler-map**. PR1488935
- Port mirroring is not supported on ACX6360-OR. PR1491789
- On ACX710 routers, the **ping mpls l2ckt/l2vpn** command does not work if the **no-control-word** statement is configured. PR1492963
- On ACX710 routers, the **ping mpls l2circuit** command does not work if **explicit-null** is configured. It does not affect data path traffic. PR1494152
- On ACX710 routers, high convergence is seen with EVPN-ELAN service in a scaled scenario during FRR switchover. With 150 EVPN-ELAN session, the switchover can go as high as ~200 msec with aggregated Ethernet in the core. PR1497251
- When the NETCONF session is established over an outbound SSH connection, the high rate of pushing the configuration to an ephemeral database might result in flapping of the outbound SSH connection or a memory leak issue. PR1497575
- The local link speed parameter under the autonegotiation information displays configured speed instead of negotiated speed. However, the link partner speed indicates the negotiated speed. When speed is not configured, the local link speed is not displayed with reboot or Packet Forwarding Engine restart. It is displayed when the speed is configured and later deleted. PR1499012
- On ACX710 routers with an EVPN-VPWS and EVPN-FXC circuits, Layer 3 VPN destination reachable over composite next hop (this is enabled using CLI **set routing-options forwarding-table chained-composite-next-hop ingress l3vpn**) does not get HW FRR behavior (less than 50 ms convergence). The traffic convergence dependents on control plane convergence. PR1499483
- I2C errors and SFP toxic message might be seen during boot, if the port with copper SFP is disabled before reboot. These I2C errors do not flood and see during boot, then they stop. Even though SFP toxic message is seen, there is no functional impact. PR1501332
- On ACX710 routers, if we configure DHCP option 012 host-name in DHCP server and the actual base configuration file also has the host-name in it, then overwriting of the base configuration file's host-name with the DHCP option 012 host-name is happening. PR1503958
- On the ACX6360 platform, the core file **core-ripsaw-node-aftd-expr** is generated and you are unable to back trace the file. PR1504717
- On ACX710 routers, in case the following steps are done for PTP:

  1. One or two port as source for chassis synchronization and both PTP and SyncE locked.
  2. Disable both Logical Interfaces.
  3. Restart clksyncd.
  4. Rollback 1 chassis does not lock again.

  This can be recovered by deleting PTP configuration, restart clksyncd, and reconfiguring the PTP post this operation. PR1505405

- MPLS LSP check is failing while verifying basic lsp_retry_limit. Reset the src_address of the LSP to 0 (if src_address is not configured) whenever it changes its state from up to down. So when the ingress LSP goes to down state, reset it to 0. The script is failing because the script is checking for src_address to be present for the ingress ISP session. PR1505474
- In a PTP environment, some vendor devices act as slave expecting announce messages at an interval of -3 (8pps) from upstream master device. Currently, announce messages are configured in a range of 0 to 3. To support -3 requirement, a hidden CLI statement **set protocol ptp master announce-interval -3** is introduced. In a network or design where you have this requirement, you can configure the hidden CLI or the regular CLI which is in the range of 0 to 3. Both the CLI statements are mutually exclusive, commit error is expected if both are configured. This new change is applicable to all ACX platforms except ACX5000 line of routers. PR1507782
- On ACX710 routers, unexpected delay counter values are seen under **show ptp statistics detail** when upstream master stops sending the PTP packets. PR1508031
- On ACX710 routers, if the ukern is restarted with the **chassis-control restart** command, the state of the PTP lock status on the Routing Engine will transition between holdover/acquiring/phase locked. The clock data is displayed accordingly. Once the Packet Forwarding Engine is up and running after restart, clock data is stable and correct. During the time the Packet Forwarding Engine is not up, the clock display is inconsistent but eventually it becomes valid once the Packet Forwarding Engine is up and the clock is created and announce packets are being generated. PR1508385
- On ACX710 routers, the Packet Forwarding Engine might crash and the FPC might remain down. This issue occurs when the PTP is configured and removed, and then the router is rebooted. This issue happens when the DMA in QAX device goes in bad state when host bound PTP traffic is pumped and router rebooted. This causes the router to crash and it does not come up. PR1509402
- On ACX710 routers, EXP re-marking is supported only for a single MPLS label packet. PR1509627
- On ACX710 routers, local repair can be in seconds (>50 ms) during FRR convergence. If explicit NULL is configured on the PHP node and on the PHP node of the backup path, the link failure is observed at PHP node. Global repair resumes the traffic flow. PR1515512
- On ACX710 routers, whenever EVPN core link is flapped, the following errors might be seen for a few seconds **LOG: Err] dnx_nh_indr_bcm_nh_install: BCM L3 Egress create object failed for:Indirect nh 2097905 (-4:Invalid parameter)**, **LOG: Err] ACX_L2_CFG_FAILED: ACX Error (L2):dnx_l2alm_get_gport_from_ifl_index : Failed to get hw nh index for evpn ifl 270533361**, **LOG: Err] ACX_PFE_ERROR: dnx_l2alm_add_mac_table_entry_in_hw: Get port from ifl failed ifl index 270533361**. PR1515516
- On ACX710 routers, the L2ALD process might restart unexpectedly during interface flaps. PR1517074
- On ACX710 routers, CFMD memory leak is observed for scaled configurations involving IPv4 and IPv6 logical interfaces with operations like deactivate, activate logical interfaces, and restart FPCs etc. This memory leak can lead to CFMD core file generation. PR1517775
- On ACX710 routers with trirate copper SFP, the interface speed in CLI is seen as 10-Gbps intermittently when the configurations are deleted. PR1518111

- On ACX5448 and ACX5448-D routers, Packet Forward Engine memory exhaustion is reported because of continuous IPv6 neighbor flaps. PR1519372
- On ACX710 routers, delete chassis alarm-port does not delete the alarm port configuration and **show chassis craft-interface** command displays the old configuration. PR1520326

Platform and Infrastructure

- The CFM REMOTE MEP does not come up after configuration or remains in Start state. PR1460555

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for ACX Series routers.

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

General Routing

- Drift messages in ACX2200, which is a PTP hybrid (PTP + Synchronous Ethernet) device. PR1426910
- ACX5448-D interfaces support: The input bytes value for the **show interfaces extensive** command is not at par with older ACX Series or MX Series devices. PR1430108
- On an ACX5448 device, DHCP packets are not transparent over Layer 2 circuit. PR1439518
- On an ACX5048 device, SNMP polling stops after the link is flapped or the SFP transceiver is replaced, and **ACX_COS_HALP(acx_cos_gport_sched_set_strict_priority:987): Failed to detach** logs might be seen. PR1455722
- ACX5448-D and ACX5448-M devices do not display airflow information and temperature sensors as expected. PR1456593
- Unable to get shared buffer count as expected. PR1468618
- Loss of manageability on ACX6360-OX platform when its disk gets full. PR1470217
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. PR1473610
- On an ACX710 device, MPLS packet load balancing is done without hashing enabled. PR1475363
- FPC might continuously crash after deactivating or activating loopback filter or reboot the system after configuring the loopback filter. PR1477740
- The dcpfe core file is generated when disabling or enabling MACsec through Toby scripts. PR1479710
- Link does not come up when a 100-Gigabit Ethernet port is channelized into four port 25-Gigabit Ethernet interfaces. PR1479733
- Memory utilization enhancement on ACX platforms to reduce the memory foot print. PR1481151
- On ACX5448 devices, **dnx_nh_mpls_tunnel_install** logs are seen. PR1482529
- ACX AUTHD process memory usage is 15 percent. PR1482598
- FPC crash is seen on ACX5448 platform. PR1485315
- On an ACX5448 device, Layer 2 VPN with interface ethernet-ccc **input-vlan-map/output-vlan-map** can cause traffic to be discarded silently. PR1485444
- On the ACX710 router, VPLS flood group results in IPv4 traffic drop after core interface flap. PR1491261
- On the ACX710 routers, LSP (primary and standby) does not Act/Up after routing or rpd restart. PR1494210
- During speed mismatch, QSFP28/QSFp+ optics/cables might or might not work. PR1494600
- ACX710 BFD sessions are in initialization state with CFM scale of 1000 on reboot or chassis control restart. PR1503429
- On an ACX500-i router, SFW sessions are not getting updated on ms- interfaces. PR1505089
- On an ACX710 router, wavelength changed from CLI does not take effect in tunable optics. PR1506647

- PIC slot might be shut down in less than 240 seconds due to the over-temperature start time is handled incorrectly. PR1506938
- BFD flaps with the error **ACX_OAM_CFG_FAILED: ACX Error (oam):dnx_bfd_l3_egress_create : Unable to create egress object** after random time interval. PR1513644

Interfaces and Chassis

- The status of the MC-AE interface might be shown as unknown when you add the subinterface as part of the VLAN on the peer MC-AE node. PR1479012

Layer 2 Ethernet Services

- Member links state might be asychronized on a connection between a PE device and a CE device in an EVPN active/active scenario. PR1463791

MPLS

- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. PR1493431

Routing Protocols

- The BGP route target family might prevent route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743

VPNs

- The Layer 2 circuit neighbor might be stuck in RD state at one end of the MG-LAG peer. PR1498040
- The rpd core files are generated while disabling Layer 2 circuit with connection protection, backup neighbor configuration, and Layer 2 circuit trace logs enabled. PR1502003

SEE ALSO

What's New | **13**

What's Changed | **23**

Known Limitations | **25**

Open Issues | **28**

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for ACX Series routers.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for ACX Series routers. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

SEE ALSO

What's New | 13

What's Changed | 23

Known Limitations | 25

Open Issues | 28

Resolved Issues | 31

Documentation Updates | 34

Junos OS Release Notes for EX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

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

NOTE: The following EX Series switches are supported in Release 20.2R1: EX2300, EX2300-C, EX3400, EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in Release 20.2R1-S1

Software Installation and Upgrade

- **Zero touch provisioning (ZTP) with IPv6 support (EX3400, EX4300, QFX5100 and QFX5200 switches, MX-Series routers)**—Starting in Junos OS Release 20.2R1-S1, ZTP supports the DHCPv6 client. 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 will continue to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device will check for DHCPv6 bindings and follow the same process as for DHCPv4 until the device can be provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

  The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

  > **NOTE:** Only HTTP and HTTPS transport protocols are supported EX3400, EX4300, QFX5100, and QFX5200 devices.

  [See Zero Touch Provisioning.]

What's New in Release 20.2R1

Authentication, Authorization, and Accounting

- **Retain the authentication session based on DHCP or SLAAC snooping entries (EX2300, EX3400, and EX4300)**—Starting in Junos OS Release 20.2R1, you can configure the authenticator to check for a DHCP, DHCPv6, or SLAAC snooping entry before terminating the authentication session when the MAC address ages out. If a snooping entry is present, the authentication session for the end device with that MAC address remains active. This ensures that the end device will be reachable even if the MAC address ages out.

  [See Authentication Session Timeouts.]

EVPN

- **802.1X authentication with EVPN-VXLAN (EX4300-48MP and EX4300-48MP Virtual Chassis)**—Starting in Junos OS Release 20.2R1, EX4300-48MP switches that act as access switches can use 802.1X authentication to protect an EVPN-VXLAN network from unauthorized end devices. EX4300-48MP switches support the following 802.1X authentication features on access and trunk ports:
  - Access ports: single, single-secure, and multiple supplicant modes
  - Trunk ports: single and single-secure supplicant modes
  - Guest VLAN

- Server fail
- Server reject
- Dynamic VLAN
- Dynamic firewall filters
- RADIUS accounting
- Port bounce with Change of Authorization (CoA) requests
- MAC RADIUS client authentication
- Central Web Authentication (CWA) with redirect URL
- Captive portal client authentication
- Flexible authentication with fallback scenarios

[See 802.1X Authentication.]

- **Support for firewall filtering on EVPN-VXLAN traffic (EX4300-MP)**—Starting with Junos OS Release 20.2R1, you can configure firewall filters and policers on the VXLAN traffic in an EVPN network (EVPN-VXLAN traffic). You set the rules that the devices uses to accept or discard packets by defining the terms for a firewall filter. For filters that you would apply to a port or VLAN, configure firewall filters at the **[edit firewall family ethernet-switching]** hierarchy level. For filters that you would apply to an IRB interface, configure firewall filters at the **[edit firewall family inet]** hierarchy level. After a firewall filter is defined, you can then apply it at an interface.

  [See Firewall Filtering and Policing Support for EVPN-VXLAN.]

- **Noncolored SR-TE LSPs with EVPN-MPLS (ACX5448, EX9200, MX Series, and vMX)—**Starting in Junos OS Release 20.2R1, ACX5448, EX9200, MX Series, and vMX routers support noncolored static segment routing-traffic engineered (SR-TE) label-switched paths (LSPs) with an EVPN-MPLS core network and the following Layer 2 services running at the edges of the network:

  - E-LAN
  - EVPN-ETREE
  - EVPN-VPWS with E-Line

  Without color, all LSPs resolve using a BGP next hop only.

  The Juniper Networks routers support noncolored SR-TE LSPs in an EVPN-MPLS core network with the following configurations:

  - EVPN running in a virtual switch routing instance
  - Multihoming in active/active and active/standby modes

  The Juniper Networks routers also support noncolored SR-TE LSPs when functioning as a Data Center Interconnect (DCI) device that handles EVPN Type 5 routes.

[See Static Segment Routing Label Switched Path.]

- **MAC filtering, storm control, and port mirroring support in EVPN-VXLAN overlay networks (EX4300-48MP)**—Starting with Junos OS Release 20.2R1, EX4300-48MP switches support the following features in an EVPN-VXLAN overlay network:
  - MAC filtering
  - Storm control
  - Port mirroring and analyzers

  [See MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.]

- **Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface (EX4600)**—Starting in Junos OS Release 20.2R1, you can configure and successfully commit the following on a physical interface of an EX4600 switch in an EVPN-VXLAN environment:
  - Layer 2 bridging (**family ethernet-switching**) on any logical interface unit number (unit 0 and any nonzero unit number).
  - VXLAN on any logical interface unit number (unit 0 and any nonzero unit number).
  - Layer 2 bridging (**family ethernet-switching** and **encapsulation vlan-bridge**) on different logical interfaces (unit 0 and any nonzero unit number).
  - Layer 3 IPv4 routing (**family inet**) and VXLAN on different logical interfaces (unit 0 and any nonzero unit number).

  For these configurations to be successfully committed and work properly, you must specify the **encapsulation flexible-ethernet-services** configuration statements at the physical interface level—for example, **set interfaces xe-0 /0/5 encapsulation flexible-ethernet-services**.

  [See Understanding Flexible Ethernet Services Support With EVPN-VXLAN.]

High Availability (HA) and Resiliency

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes mastership. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

  [See Understanding the Ephemeral Configuration Database.]

Juniper Extension Toolkit (JET)

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

  [See language (Scripts), Develop Off-Device JET Applications, and Develop On-Device JET Applications.]

Junos OS XML, API, and Scripting

- **Support for Rest API (EX2300, EX2300-MP, EX3400, EX4300, EX4300-MP, EX4600, EX4650, and EX9200)**—Starting in Release 20.2R1, Junos OS supports the REST API on EX2300, EX2300-MP, EX3400, EX4300, EX4300-MP, EX4600, EX4650, and EX9200 switches. The REST API enables you to securely connect to the Junos OS devices, execute remote procedure calls (RPC) commands, use REST API explorer GUI to conveniently experiment with any of the REST APIs, and use a variety of formatting and display options including JavaScript Object Notation (JSON).

  [See REST API Guide.]

Junos Telemetry Interface

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

  [See Mapping OpenConfig Routing Policy Commands to Junos Configuration and Mapping OpenConfig Network Instance Commands to Junos Operation.]

- **Support for OpenConfig configuration model version 4.0.1 for BGP with JTI (EX2300, EX3400, EX4300, EX4600, and EX9200)**— Junos OS Release 20.2R1 provides support for the OpenConfig version 4.0.1 data models **openconfig-bgp-neighbor.yang** and **openconfig-bgp-policy.yang** using Junos telemetry

interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream telemetry statistics to an outside collector.

The following major resource paths are supported with gRPC and JTI:

- **/network-instances/network-instance/protocols/protocol/bgp/global/**
- **/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/**
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/**
- **/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/**

[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface and OpenConfig Data Model Version.]

- **Support for OpenConfig configuration model version 1.0.0 for local routing with JTI (EX2300, EX3400, EX4300, EX4600, and EX9200)**— Junos OS Release 20.2R1 provides support for the OpenConfig version 1.0.0 data model **openconfig-local-routing.yang** using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream telemetry statistics to an outside collector.

  The following major resource paths are supported with gRPC and JTI:

  - **/local-routes/static-routes/static/**
  - **/local-routes/local-aggregates/aggregate/**

  [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface and OpenConfig Data Model Version.]

- **Packet Forwarding Engine and Routing Engine sensor support with JTI (EX2300, EX2300-MP, and EX3400)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) with remote procedure call (gRPC) services to export Packet Forwarding Engine statistics and Routing Engine statistics from EX2300, EX2300-MP, and EX3400 switches to an outside collector. These statistics can also be exported through UDP (native) sensors.

  Supported Packet Forwarding Engine sensors are:

  - Sensor for CPU (ukernel) memory (resource path **/junos/system/linecard/cpu/memory/**)
  - Sensor for firewall filter statistics (resource path **/junos/system/linecard/firewall/**)
  - Sensor for physical interface traffic (resource path **/junos/system/linecard/interface/**)
  - Sensor for logical interface traffic (resource path **/junos/system/linecard/interface/logical/usage/**). Not supported on EX2300 or 2300-MP switches.
  - Sensor for software-polled queue-monitoring statistics (resource path **/junos/system/linecard/qmon-sw/**). Not supported on EX2300 or 2300-MP switches.

  Supported Routing Engine sensors are:

- Sensor for LACP state export (resource path **/lacp/**)
- Sensor for chassis environmentals export (resource path **/junos/system/components/component/**)
- Sensor for chassis components export (resource path **/components/**)
- Sensor for LLDP statistics export (resource path **/lldp/interfaces/interface[name='*name*']/**)
- Sensor for BGP peer information export (resource path **/network-instances/network-instance/protocols/protocol/bgp/**). Not supported on EX2300 or 2300-MP switches.
- Sensor for RPD task memory utilization export (resource path **/junos/task-memory-information/**)
- Sensor network discovery ARP table state (resource path **/arp-information/**)
- Sensor for network discovery NDP table state (resource path **/nd6-information/**)

[See Understanding OpenConfig and gRPC and gNMI on Junos Telemetry Interface, sensor (Junos Telemetry Interface), and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface.]

Multicast

- **Static multicast route leaking for VRF and virtual router instances (EX4650 and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can configure the switch to statically share (leak) IPv4 multicast routes for IGMPv3 (S,G) traffic among different virtual router or virtual routing and forwarding (VRF) instances. You can only leak static multicast routes per group, not per source and group. The destination prefix length must be 32.

  To configure multicast route leaking to the VRF or virtual router instance *routing-instance-name*, configure the **next-table *routing-instance-name*.inet.0** statement at the [**edit routing-instances *routing-instance-name* routing-options static route destination-prefix/32]** hierarchy level.

  [See Understanding Multicast Route Leaking for VRF and Virtual Router Instances.]

- **Multicast-only fast reroute (MoFRR) (EX4650 and QFX5120-48Y)**—Starting in Junos OS Release 20.2R1, you can configure MoFRR to minimize multicast packet loss in PIM domains when link failures occur. With MoFRR enabled, the switch maintains primary and backup traffic paths, forwarding traffic from the primary path and dropping traffic from the backup path. If the primary path fails, the switch can quickly start forwarding the backup path stream (which becomes the primary path). The switch creates a new backup path if it detects available alternative paths. MoFRR applies to all multicast (S,G) streams by default, or you can configure a policy for the (S,G) entries where you want MoFRR to apply.

  [See Understanding Multicast-Only Fast Reroute.]

Network Management and Monitoring

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

  [See NETCONF Sessions over Outbound HTTPS.]

Routing Policy and Firewall Filters

- **Support for MPLS firewall filter on loopback interface (EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can apply an MPLS firewall filter to a loopback interface on a Label switching router (LSR). For example, you can configure an MPLS packet with **ttl=1** along with MPLS qualifiers such as **label**, **exp**, and Layer 4 **tcp/udp** port numbers. Supported actions include **accept**, **discard**, and **count**.

  You configure this feature at the **[edit firewall family mpls]** hierarchy level. You can only apply a loopback filters on **family mpls** in the ingress direction.

  [See Overview of MPLS Firewall Filters on Loopback Interface.]

Routing Protocols

- **Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast (MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices)**—Starting with Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices support BGP PIC Edge protection for Layer 2 circuit, Layer 2 VPN, and VPLS (BGP VPLS, LDP VPLS and FEC 129 VPLS) services with BGP labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

  A prerequisite for BGP PIC Edge protection is to program the Packet Forwarding Engine (PFE) with expanded next-hop hierarchy.

  To enable BGP PIC Edge protection, use the following CLI configuration statements:

  - Expand next-hop hierarchy for BGP labeled unicast family:

    ```
    [edit protocols]
    user@host#set bgp group group-name family inet labeled-unicast nexthop-resolution
       preserve-nexthop-hierarchy;
    ```

- BGP PIC for MPLS load balance nexthops:

```
[edit routing-options]
user@host#set rib routing-table-name protect core;
```

- Fast convergence for Layer 2 circuit and LDP VPLS:

```
[edit protocols]
user@host#set l2circuit resolution preserve-nexthop-heirarchy;
```

- Fast convergence for Layer 2 VPN, BGP VPLS, and FEC129:

```
[edit protocols]
user@host#set l2vpn resolution preserve-nexthop-heirarchy;
```

[See Load Balancing for a BGP Session.]

SEE ALSO

What's Changed

IN THIS SECTION

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

Class of Service (CoS)

- We've corrected the output of the "show class-of-service interface | display xml" command. Output of the following sort: <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> to will now appear correctly as: <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container>

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

  [See commit (System).]

Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

  [See Develop Off-Device JET Applications and Develop On-Device JET Applications.]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

  [See Juniper EngNet.]

Network Management and Monitoring

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

SEE ALSO

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for EX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

EVPN

- When only one link is present between the leaf devices and it goes down resulting in a silent drop in traffic. PR1480847

Infrastructure

- Depending on the actual traffic pattern and the order in which the MACs are learned, the actual MAC DB scale may vary. This is due to the way the MACs are internally stored in the hardware. PR1485319
- On EX-4300MP, 9000 IPv6 MC routes can be installed. If you try to add more IPv6 MC routes, error messages will be seen. PR1493671

SEE ALSO

What's New | 36

What's Changed | 44

Open Issues | 47

Resolved Issues | 51

Documentation Updates | 56

Migration, Upgrade, and Downgrade Instructions | 56

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.2R1 for EX Series switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Authentication and Access Control

- When a 802.1X session terminates, an event denoting the same was not logged in single supplicant mode. As fix, a new event **DOT1XD_USR_SESSION_DISCONNECTED** is logged consistently whenever a session terminates irrespective of supplicant mode. **DOT1XD_AUTH_SESSION_DELETED** events still get generated too but only for multiple and single-secure supplicant modes (as per design). PR1512724
- On EX2300/EX3400/EX4300 that supports Private-Vlan and dot1x platforms, the authenticated dot1x client in the isolated (secondary vlan) is not cleared when the authenticated PVLAN is deleted. PR1516341

EVPN

- In all platforms with VXLAN Static VTEP tunnels scenario (including Static VXLAN without EVPN), after Routing Engine switchover or restart of l2-learning, if you create a new VTEP interface, the interface may not work. PR1520078

Infrastructure

- qmon-sw sensor is not supported in EX3400. PR1506710
- The IP communication between directly connected interfaces on EX4600 TVP platforms would fail. This issue only might occur in this special scenario and it might have traffic/service impact. PR1515689

Interfaces and Chassis

- On GRES, VSTP port cost on aggregated Ethernet interfaces might get changed, leading to topology change. PR1174213
- The same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating incorrect configuration. PR1221993

Layer 2 Ethernet Services

- If **forward-only** is set within **dhcp-reply** in a Juniper Networks device as a DHCP relay agent, the DHCP DECLINE packets that are broadcast from the DHCP client are dropped and not forwarded to the DHCP server. PR1429456

- In a DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case, the subscriber might need more time to get IP address assigned. The subscriber might also remain in this state until its lease expires if it has previously bound with the address in the option 50. PR1435039
- Sometimes image upgrade through ZTP may fail due to not having enough space on EX3400. Below kb article talks about how to free up the space : KB31198. PR1515013

Layer 2 Features

- GARPs were being sent whenever there was a MAC (fdb) operation (add or delete). This is now updated to send GARP when the interface is up and l3 interface attached to the VLAN. PR1192520
- On QFX5000/EX46xx, if **forwarding-options enhanced-hash-key hash-params** is not configured and if the hash function and pre-process for LAG is the same on ingress nodes and QFX5K/EX46xx, egress traffic imbalance might be observed when ECMP or LAG is used. It might cause traffic congestion unexpectedly. PR1514793

Platform and Infrastructure

- EX3400/EX2300 upgrade may fail due to space and the system generates the following messages: **/usr/libexec/ui/package: /var/tmp/mchassis-install.tgz: no such file**PR1440122
- On EX9208 switches, 33 percent degradation in MAC learning rate is seen in Junos OS Release 19.3R1 while comparing with Junos OS Release 18.4R1. PR1450729
- On EX4300 platforms configured with ERP, after multiple devices reboot/restart at the same time, ERP might not revert back to the IDLE state. This issue might be seen in situations where the ERP node-id is not configured manually and after the restart, the default node-id (switch base MAC address) might get reset to 00:00:00:00:00:00, effectively causing multiple devices to have the same node-id. PR1461434
- On MX series platforms, when a route's next-hop is an IRB interface with lt- as the underlying L2 interface, it is not getting programmed on PFE, resulting in packet drop. PR1494594
- Chassis connection dropped often in AD-2 while when dot1x clients connect/disconnect. The issue is seen when dot1x clients connect/disconnect. PR1513274
- After GRES, interfaces may flap due to which DHCP bindings may be lost. PR1515234
- Craftd messages are generated on MX10003 and MX204 platforms. These platforms do not have a craft interface. Hence these errors are expected, and can safely be ignored. When Craftd daemon tries to open the device, it fails with a junk char in the fatal error message because the error no is not mapped to a string in the kernel code. The following messages are seen: **Feb 20 01:49:38 MX craftd[xxxx]: craftd detected platform mx10002 Feb 20 01:49:38 MX craftd[xxxx]: LIBJSNMP_SA_IPC_REG_ROWS:**

ns_subagent_register_mibs: registering 1 rows Feb 20 01:49:38 MX craftd[xxxx]: fatal error, failed to open smb device: ‚JÎÈ"" PR1359929

- On an EX9208 switch, a few xe- interfaces are going down with the error **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error**. PR1377840
- Unicast RPF check in strict mode might not work properly. PR1417546
- On the EX9214 device, if the MACsec-enabled link flaps after reboot, the error **errorlib_set_error_log(): err_id(-1718026239)** is observed. PR1448368
- In overall commit time, the evaluation of mustd constraints is taking 2 seconds more than usual. This is because the persist-group-inheritance feature has been made a default feature in the latest Junos OS releases. Eventually, this feature helps improve the subsequent commit times for scaled configurations significantly. The persist-group-inheritance feature is useful in customer scenarios where groups and nested groups are used extensively. In those scenarios, the group inheritance paths are not built every time, thus subsequent commits are faster. This issue is seen only with a QFX platform or other low end devices. PR1457939
- On EX4300 switches, when packets entering a port exceed a size of 144 bytes, they might get dropped in very few cases. PR1464365
- While verifying Last-change op-state value through XML, rpc-reply message is inappropriate. PR1492449
- When the NETCONF session is established over outbound ssh, the high rate of pushing the configuration to the ephemeral DB might result in flapping of the outbound SSH connection or a or memory leak issue. PR1497575
- EX4300-48MP-EX4300-VC: This issue is very rarely seen and is Virtual Chassis specific. For the issue to get triggered, the Lag IRB interface where OSPF is stuck should be present in the Standby switch. The problem state is recovered by rebooting the master and switch is not seen again. PR1498903
- On EX4300/EX3400/EX2300 Virtual-Chassis with NSB and xSTP enabled, the continuous traffic loss might be observed while doing GRES. PR1500783
- LLDP packets are not acquired when **native-vlan** configured is same as tagged **vlan-id**. PR1504354
- On EX/QFX virtual-chassis setup, when LLDP is configured along with the PVLAN and the interface is connected to the backup or linecard member port, LLDP might not work on the other end of Virtual Chassis. PR1511073
- 35 seconds delay is added in reboot time from Junos OS Release 20.2R1 release compared to Release 19.4R2. PR1514364
- Memory leak is seen in 'dot1xd' daemon when no 'dot1x' is configured. Memory leak is seen for the allocation while creating socket from 'dot1xd' daemon to 'authd' daemon. If 'authd' is not running , 'dot1xd' daemon tries to connect to 'authd' periodically and every time it was allocating memory for string "/var/run/authd_control" for socket creation. The memory does not free in this scenario and we see memory leak for string "/var/run/authd_control". There will be no service impact to other services/daemons other than dot1x. PR1515972

Routing Protocols

- ECDSA256+SHA256 is not used for software integrity checking. PR1504211
- On 48 port (2 units), partial packet drops may be seen when unicast stream is configured. PR1520059

SEE ALSO

What's New | 36

What's Changed | 44

Known Limitations | 46

Resolved Issues | 51

Documentation Updates | 56

Migration, Upgrade, and Downgrade Instructions | 56

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in Junos OS Release 20.2R1 for EX Series switches.

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

Authentication and Access Control

- EX2300-48MP: Client did not receive captive-portal success page by downloading the ACL parameter as Authentication failed. PR1504818

EVPN

- The ESI of IRB interfaces does not get updated after an autonomous-system number change if the interface is down. PR1482790
- The VXLAN function might be broken due to a timing issue after the change in PR 1495098. PR1502357

High Availability (HA) and Resiliency

- Kernel core files on backup Routing Engine might cause traffic drop if multicast-mac is configured on IRB interface. PR1467847

Infrastructure

- Memory leak leads to kernel crash (vmcore) due to SNMP polling. PR1482379
- Kernel core files might be observed if you deactivate the daemon on EX2300/EX3400 platforms. PR1483644
- The fxpc process might crash when configuring scaled configuration with 4093 VLANs. PR1493121

Interfaces and Chassis

- The following FRU with no connection arguments is observed after MX Series Virtual Chassis local or global switchover: **fru_send_msg Global FPC x**. PR1428254
- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. PR1459201
- Executing commit might hang up due to a stuck dcd process. PR1470622
- A stale IP address might be seen after a specific order of configuration changes under a logical-systems scenario. PR1477084
- Traffic might get dropped as the next hop points to ICL even though the local MC-LAG is up. PR1486919

Junos Fusion Enterprise

- SDPD core files found: **vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry**. PR1454335
- Loop detection might not work on extended ports in Junos fusion scenarios. PR1460209

Junos Fusion Satellite Software

- Temperature sensor alarm is seen on EX4300 in a Junos fusion scenario. PR1466324

Layer 2 Ethernet Services

- Member links state might be asychronized on a connection between PE and CE devices in an EVPN active/active scenario. PR1463791
- Issues with DHCPv6 relay processing Confirm and Reply packets. PR1496220

Layer 2 Features

- The LLDP function might fail when a Juniper device connects to a non-Juniper one. PR1462171
- EX4650/QFX5120: QinQ: The third VLAN tag is not pushed onto the stack and SWAP is being done instead. PR1469149
- Traffic might be affected if composite next hop is enabled. PR1474142

MPLS

- BGP session might keep flapping between two directly connected BGP peers because of the wrong TCP-MSS in use. PR1493431

Platform and Infrastructure

- The IRB traffic might get dropped after mastership switchover. PR1453025
- The switch might not be able to learn MAC addresses with **dot1x** and **interface-mac-limit** configured. PR1470424
- EX4300: Input firewall filter attached to isolated or community VLANs not matching 802.1p bits on the VLAN header. PR1478240
- MAC learning under bridge-domain stops after an MC-LAG interface flap. PR1488251
- The NSSU upgrade might fail on EX4300 switches due to a storage issue in the **/var/tmp** directory. PR1494963

- On the EX4300 device, high CPU load due to receipt of specific IPv4 packets is observed. PR1495129
- Traffic loss might be seen with framing errors or runts if MACsec is configured on EX4300. PR1502726
- The MAC Pause frames will be incrementing in the Receive direction if half-duplex mode on 10-Mbps or 100-Mbps speed is configured. PR1452209
- Link up delay and traffic drop might be seen on mixed SP L2/L3 and EP L2 type configurations. PR1456336
- MAC addresses learned on RTG may not be aged out after the aging time. PR1461293
- RTG link faces nearly 20 seconds down during backup node rebooting. PR1461554
- The jdhcpd process might consume high CPU and no further subscribers can be brought up if there are more than 4000 DHCP relay clients in the MAC move scenario. PR1465277
- FPCs might get disconnected from the EX3400 Virtual Chassis briefly after a reboot or an upgrade. PR1467707
- Traffic loss might be seen with framing errors or runts if MACsec is configured on EX4600 or QFX5100 platforms. PR1469663
- SSH session closes while checking for the **show configuration | display set** command for both local and nonlocal users. PR1470695
- The shaping of CoS does not work after reboot. PR1472223
- CoS 802.1p bits rewrite might not happen in Q-in-Q mode. PR1472350
- DSCP marking might not work as expected if the fixed classifiers are applied to interfaces on QFX5000 or EX4600 platforms. PR1472771
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. PR1473610
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. PR1473685
- On EX4300, the output of **show security macsec statisitics** shows high values incorrectly. PR1476719
- EX3400 me0 interface might remain down. PR1477165
- The dhcpd process may crash in a Junos fusion environment. PR1478375
- Trio based linecard might crash when there is bulk route update failure in a corner case. PR1478392
- TFTP installation from loader prompt may not succeed on the EX Series devices. PR1480348
- ARP request packets for an unknown host might get dropped in remote PE in EVPN-VXLAN scenario. PR1480776
- On EX2300 switches, SNMP traps are not generated when the MAC addresses limit threshold is reached. PR1482709
- Incorrect frame length of 132 bytes might be displayed in the packet header. PR1487876
- Virtual Chassis ports might go down in a mixed Virtual Chassis setup of QFX5100-24Q-2P/EX4300 and EX4600/EX4300. PR1489985

- DHCP binding fails while you verify DHCPv4 snooping functionality with P-VLAN with a firewall to block or allow certain IPv4 packets. PR1490689
- On the EX2300 device, high CPU load due to the receipt of specific multicast packets on Layer 2 interface is observed. PR1491905
- Traffic loss could be observed in a mixed-Virtual Chassis setup of QFX5100 and EX4300. PR1493258
- Traffic loss could be seen in an MC-LAG scenario on QFX5120 and EX4650. PR1494507
- The fxpc process might crash when renumbering the primary member ID value of the EX2300 or EX3400 Virtual Chassis. PR1497523
- Traffic might get dropped if AE member interface is deleted/added or a SFP of the AE member interface is unplugged/plugged. PR1497993
- On EX4650 switches, the firewall filter might not get applied. PR1499647
- The isolated VLAN from radius is not deleted when the interface flaps. PR1506427

Routing Protocols

- BGP IPv4/IPv6 convergence and RIB install and delete time is degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. PR1414121
- MUX State in LACP interface does not go to **collecting and distributing** and remains **attached** after enabling the ae interface. PR1484523
- FPC might go to "NotPrsnt" state after upgrading with non-TVP image in VC/VCF setup. PR1485612
- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- On the EX4300-MP and EX4600 devices, high CPU load due to receipt of specific Layer 2 frames in EVPN-VXLAN deployment. PR1495890
- Firewall filter could not work in certain conditions in an Virtual Chassis setup. PR1497133
- Packets drop might be seen when multicast MAC with static ARP is configured on one IRB interface. PR1489374

User Interface and Configuration

- **umount: unmount of /.mount/var/val/chroot/packages/mnt/jweb-ex32-d2cf6f6b failed: Device busy** message is seen when Junos OS is upgraded with the validate option. PR1478291
- On the EX2300 and EX3400 devices, installing J-Web application package might fail. PR1513612

SEE ALSO

What's New | 36
What's Changed | 44
Known Limitations | 46
Open Issues | 47
Documentation Updates | 56
Migration, Upgrade, and Downgrade Instructions | 56

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for EX Series switches.

SEE ALSO

What's New | 36
What's Changed | 44
Known Limitations | 46
Open Issues | 47
Resolved Issues | 51
Migration, Upgrade, and Downgrade Instructions | 56

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 57

This section contains the upgrade and downgrade support policy for Junos OS for EX Series switches. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network. For information about software installation and upgrade, see the Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

SEE ALSO

What's New | 36

What's Changed | 44

Known Limitations | 46

Open Issues | 47

Resolved Issues | 51

Documentation Updates | 56

Junos OS Release Notes for JRR Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the JRR Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in Junos OS Release 20.2R1 for JRR Series Route Reflectors.

Layer 2 Features

- **Support for Link Layer Discovery Protocol (JRR200)**—Starting in Junos OS Release 20.2R1, JRR Series devices support Link Layer Discovery Protocol (LLDP) is supported both on the management port em0 and on the WAN ports em2 through em9. LLDP is a link-layer protocol defined in IEEE 802.1AB that allows network devices to advertise their identity, capabilities, and configuration to other devices on the LAN.

  [See Understanding LLDP.]

SEE ALSO

- What's Changed | 59
- Known Limitations | 60
- Open Issues | 60
- Resolved Issues | 60
- Documentation Updates | 61
- Migration, Upgrade, and Downgrade Instructions | 62

What's Changed

There are no changes in behavior and syntax in Junos OS Release 20.2R1 for JRR Series Route Reflectors.

SEE ALSO

- What's New | 58
- Known Limitations | 60
- Open Issues | 60
- Resolved Issues | 60
- Documentation Updates | 61
- Migration, Upgrade, and Downgrade Instructions | 62

Known Limitations

There are no known limitations JRR Series in Junos OS Release 20.2R1 for JRR Series Route Reflectors.

SEE ALSO

- What's New | 58
- What's Changed | 59
- Open Issues | 60
- Resolved Issues | 60
- Documentation Updates | 61
- Migration, Upgrade, and Downgrade Instructions | 62

Open Issues

There are no open issues in Junos OS 20.2R1 Release for JRR Series Route Reflectors.

SEE ALSO

- What's New | 58
- What's Changed | 59
- Known Limitations | 60
- Resolved Issues | 60
- Documentation Updates | 61
- Migration, Upgrade, and Downgrade Instructions | 62

Resolved Issues

IN THIS SECTION

- General Routing | 61

Learn about resolved issues for JRR Series in Junos OS 20.2R1 Release.

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

General Routing

- USB install image is not working for JRR200 platform. PR1471986
- Link state of virtual em interfaces in Junos OS might not reflect the true link status of corresponding physical interfaces in the Linux host. PR1492087

SEE ALSO

- What's New | 58
- What's Changed | 59
- Known Limitations | 60
- Open Issues | 60
- Documentation Updates | 61
- Migration, Upgrade, and Downgrade Instructions | 62

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for JRR200 Route Reflectors.

SEE ALSO

- What's New | 58
- What's Changed | 59
- Known Limitations | 60
- Open Issues | 60
- Resolved Issues | 60
- Migration, Upgrade, and Downgrade Instructions | 62

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for the JRR Series Route Reflector. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

For information about software installation and upgrade, see the JRR200 Route Reflector Quick Start and the Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No

Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

SEE ALSO

Junos OS Release Notes for Junos Fusion for Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the Junos fusion for enterprise. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

There are no new features or enhancements to existing features in Junos OS Release 20.2R1 for Junos fusion for enterprise.

> **NOTE:** For more information about the Junos fusion for enterprise features, see the Junos fusion for enterprise User Guide.

SEE ALSO

- What's Changed | 64
- Known Limitations | 65
- Open Issues | 65
- Resolved Issues | 66
- Documentation Updates | 67
- Migration, Upgrade, and Downgrade Instructions | 67

What's Changed

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.2R1 for Junos fusion for enterprise.

SEE ALSO

- What's New | 64
- Known Limitations | 65
- Open Issues | 65
- Resolved Issues | 66

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.2R1 for Junos fusion for enterprise.

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

SEE ALSO

Open Issues

There are no known issues in hardware and software in Junos OS Release for 20.2R1 Junos fusion for enterprise.

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

SEE ALSO

Migration, Upgrade, and Downgrade Instructions | 67

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for Junos fusion for enterprise.

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

Resolved Issues: Release 20.2R1

- Observing duplicate ECID values for cluster and extended ports on member ports of same cluster. PR1408947
- The SDPD process generates a core file at **vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry**. PR1454335
- Loop detection might not work on extended ports in a Junos fusion scenario. PR1460209
- The temperature sensor alarm is seen on EX4300 in a Junos fusion scenario. PR1466324

SEE ALSO

What's New | 64
What's Changed | 64
Known Limitations | 65
Open Issues | 65
Documentation Updates | 67
Migration, Upgrade, and Downgrade Instructions | 67

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 for documentation for Junos fusion for enterprise.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade or downgrade Junos OS and satellite software for a Junos fusion for enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos fusion for enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the **junos-install** package. Use other packages (such as the **jbundle** package) only when so instructed by a Juniper Networks support

representative. For information about the contents of the **junos-install** package and details of the installation process, see the Installation and Upgrade Guide.

> **NOTE:** Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:
>
> ```
> user@host> request system snapshot
> ```
>
> The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the Junos OS Administration Library.

To download and install Junos OS:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos fusion** to find the software that you want to download.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list on the right of the page.

5. Select the **Software** tab.

6. Select the software package for the release.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new **junos-install** package on the aggregation 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.

Customers in the United States and Canada, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot source/package-name.n.tgz
```

All other customers, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot source/package-name.n-limited.tgz
```

Replace ***source*** with one of the following values:

- ***/pathname***—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname/pathname***
  - **http://*hostname/pathname***
  - **scp://*hostname/pathname*** (available only for Canada and U.S. version)

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to minimize disrupting network operations as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.

3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the *Installation and Upgrade Guide*.

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos fusion for enterprise. See Configuring or Expanding a Junos fusion for enterprise.

For satellite device hardware and software requirements, see Understanding Junos fusion for enterprise Software and Hardware Requirements.

Use the following command to install Junos OS on a switch before converting it into a satellite device:

```
user@host> request system software add validate reboot source/package-name
```

> **NOTE:** The following conditions must be met before a Junos switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:
>
> - The switch running Junos OS can be converted only to SNOS 3.1 and later.
> - Either the switch must be set to factory-default configuration by using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

When the interim installation has completed and the switch is running a version of Junos OS that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device using the console port.
2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

> **NOTE:** The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port *port-number***

   For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 0**
   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 1**
   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 2**
   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 3**

   This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, or preconfiguration. See Configuring or Expanding a Junos fusion for enterprise for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Switch

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove it from the Junos fusion topology. For more information, see Converting a Satellite Device to a Standalone Device.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

Downgrading Junos OS

Junos fusion for enterprise is first supported in Junos OS Release 16.1, although you can downgrade a standalone EX9200 switch to earlier Junos OS releases.

> **NOTE:** You cannot downgrade more than three releases.
>
> For more information, see the Installation and Upgrade Guide.

To downgrade a Junos fusion for enterprise from Junos OS Release 20.2, follow the procedure for upgrading, but replace the 20.2 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

- What's New | 64
- What's Changed | 64
- Known Limitations | 65
- Open Issues | 65

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for fusion for provider edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in this release for Junos fusion for provider edge.

Hardware

- **Support for QFX5110 as a satellite device in a Junos fusion for provider edge on a GNF(MX480 and MX960)**—With Junos Node Slicing, you can create guest network functions (GNFs), partitions where an aggregation device can be configured. The aggregation device on a GNF supports a maximum of 10 satellite devices. Starting in Junos OS Release 20.2R1, Junos OS supports QFX5110 switches as satellite devices in Junos fusion for provider edge on a GNF.

  [See Understanding Junos Fusion Provider Edge Software and Hardware Requirements and Junos Node Slicing Overview.]

Junos Fusion

- **MPC10E and MPC11E interoperability with Junos fusion for provider edge (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 20.2R1, Junos OS supports using the MPC10E and MPC11E alongside other MPC line cards in the same MX Series router chassis that has been configured with Junos fusion for provider edge. The line cards can coexist in the same router chassis, and the router passes traffic between the devices connected to the MPC10E/MPC11E and the satellite devices that are connected to other MPC line cards through the switch fabric. You cannot use MPC10E/MPC11E in Junos fusion, which means you cannot connect satellite devices to ports on the MPC10E/MPC11E line cards.

  Junos fusion does not support hyper mode. To support Junos fusion in an MX Series router where MPC10E/MPC11E coexists with other MPC line cards, use the **set forwarding-options no-hyper-mode** statement. In addition, you must also use an FPC slot ID in the range of 160—252 for the satellite device interfaces. To configure the FPC slot ID, use the **set chassis satellite-management fpc *slot-id*** statement.

  [See Junos Fusion Provider Edge Overview.]

SEE ALSO

What's Changed

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in this release for Junos fusion for provider edge.

SEE ALSO

- What's New | 73
- Known Limitations | 75
- Open Issues | 75
- Resolved Issues | 76
- Documentation Updates | 77
- Migration, Upgrade, and Downgrade Instructions | 77

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.2R1 for Junos fusion for provider edge.

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

SEE ALSO

- What's New | 73
- What's Changed | 75
- Open Issues | 75
- Resolved Issues | 76
- Documentation Updates | 77
- Migration, Upgrade, and Downgrade Instructions | 77

Open Issues

There are no known issues in the Junos OS Release 20.2R1 for Junos fusion for provider edge.

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

SEE ALSO

- What's New | 73
- What's Changed | 75
- Known Limitations | 75
- Resolved Issues | 76
- Documentation Updates | 77
- Migration, Upgrade, and Downgrade Instructions | 77

Resolved Issues

IN THIS SECTION

- Fusion for Provider Edge | 76

This section lists the issues fixed in the Junos OS Release 20.2R1 for Junos fusion for provider edge.

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

Fusion for Provider Edge

- On the EX4300 devices in the Junos fusion scenario, the temperature sensor alarm is observed. PR1466324

SEE ALSO

- What's New | 73
- What's Changed | 75
- Known Limitations | 75

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for Junos fusion for provider edge.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for Junos fusion for provider edge. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

Basic Procedure for Upgrading an Aggregation Device

When upgrading or downgrading Junos OS, always use the **jinstall** package. Use other packages (such as the **jbundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **jinstall** package and details of the installation process, see the Installation and Upgrade Guide.

> **NOTE:** Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:
>
> ```
> user@host> request system snapshot
> ```
>
> The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the Installation and Upgrade Guide.

The download and installation process for Junos OS Release 20.2R1 is different from that for earlier Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos fusion** to find the software that you want to download.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the page.

5. Select the **Software** tab.

6. Select the software package for the release.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new **jinstall** package on the aggregation 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.

Customers in the United States and Canada, use the following commands.

- For 64-bit software:

> **NOTE:** We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot
source/jinstall64-20.2R1.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.2R1.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

> **NOTE:** We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot
source/jinstall64-20.2R1.SPIN-export-signed.tgz
```

- For 32-bit software:

  ```
  user@host> request system software add validate reboot
  source/jinstall-20.2R1.SPIN-export-signed.tgz
  ```

  Replace ***source*** with one of the following values:

- ***/pathname***—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname/pathname***
  - **http://*hostname/pathname***
  - **scp://*hostname/pathname*** (available only for the Canada and U.S. version)

The **validate** option validates the software package against the current configuration as a prerequisite for adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is for a different release.

Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 20.2R1 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately as follows to minimize disrupting network operations:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.

3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the *Installation and Upgrade Guide*.

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see Understanding Junos fusion Software and Hardware Requirements

> **NOTE:** The following conditions must be met before a standalone switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:
>
> - The switch can be converted to only SNOS 3.1 and later.
> - Either the switch must be set to factory-default configuration by using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

Customers with EX4300 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-ex-4300-14.1X53-D43.3-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add reboot
source/jinstall-qfx-5-14.1X53-D43.3-domestic-signed.tgz
```

When the interim installation has completed and the switch is running a version of Junos and OS on one line that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device by using the console port.

2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

> **NOTE:** The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device by using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose your connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

   ```
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port port-number
   ```

   For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

   ```
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 0
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 1
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 2
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 3
   ```

   This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See Configuring Junos fusion for provider edge for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Device

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove the satellite device from the Junos fusion topology.

NOTE: If the satellite device is a QFX5100 switch, you need to install a PXE version of Junos OS. The PXE version of Junos OS is software that includes *pxe* in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D43 is named install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz . If the satellite device is an EX4300 switch, you install a standard jinstall-ex-4300 version of Junos OS.

The following steps explain how to download software, remove the satellite device from Junos fusion, and install the Junos OS software image on the satellite device so that the device can operate as a standalone device.

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads

2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos fusion** from the drop-down list and select the switch platform series and model for your satellite device.

4. Select the Junos OS Release 14.1X53-D30 software image for your platform.

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

7. Copy the software to the routing platform or to your internal software distribution site.

8. Remove the satellite device from the automatic satellite conversion configuration.

   If automatic satellite conversion is enabled for the satellite device's member number, remove the member number from the automatic satellite conversion configuration. The satellite device's member number is the same as the FPC slot ID.

   ```
   [edit]
   user@aggregation-device# delete chassis satellite-management auto-satellite-conversion satellite member-number
   ```

   For example, to remove member number 101 from Junos fusion:

   ```
   [edit]
   ```

```
user@aggregation-device# delete chassis satellite-management auto-satellite-conversion
satellite 101
```

You can check the automatic satellite conversion configuration by entering the **show** command at the **[edit chassis satellite-management auto-satellite-conversion]** hierarchy level.

9. Commit the configuration.

   To commit the configuration to both Routing Engines:

   ```
   [edit]
   user@aggregation-device# commit synchronize
   ```

   Otherwise, commit the configuration to a single Routing Engine:

   ```
   [edit]
   user@aggregation-device# commit
   ```

10. Install the Junos OS software on the satellite device to convert the device to a standalone device.

    ```
    [edit]
    user@aggregation-device> request chassis satellite install URL-to-software-package fpc-slot
    member-number
    ```

    For example, to install a PXE software package stored in the **/var/tmp** directory on the aggregation device onto a QFX5100 switch acting as the satellite device using FPC slot 101:

    ```
    [edit]
    user@aggregation-device> request chassis satellite install
    /var/tmp/install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz fpc-slot 101
    ```

    For example, to install a software package stored in the **var/tmp** directory on the aggregation device onto an EX4300 switch acting as the satellite device using FPC slot 101:

    ```
    [edit]
    user@aggregation-device> request chassis satellite install
    /var/tmp/jinstall-ex-4300-14.1X53-D30.3-domestic-signed.tgz fpc-slot 101
    ```

    The satellite device stops participating in the Junos fusion topology after the software installation starts. The software upgrade starts after this command is entered.

11. Wait for the reboot that accompanies the software installation to complete.

12. When you are prompted to log back into your device, uncable the device from the Junos fusion topology. See *Removing a Transceiver from a QFX Series Device* or *Remove a Transceiver*, as needed. Your device has been removed from Junos fusion.

> **NOTE:** The device uses a factory-default configuration after the Junos OS installation is complete.

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 20.2R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No

Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

Downgrading from Junos OS Release 20.1

To downgrade from Release 20.1 to another supported release, follow the procedure for upgrading, but replace the 20.1 **jinstall** package with one that corresponds to the appropriate release.

> **NOTE:** You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

SEE ALSO

Junos OS Release Notes for MX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for MX Series routers.

What's New in Release 20.2R1-S1

Software Installation and Upgrade

- **Zero touch provisioning (ZTP) with IPv6 support (EX3400, EX4300, QFX5100 and QFX5200 switches, MX-Series routers)**—Starting in Junos OS Release 20.2R1-S1, ZTP supports the DHCPv6 client. 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 will continue to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device will check for DHCPv6 bindings and follow the same process as for DHCPv4 until the device can be provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

  The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

  > **NOTE:** Only HTTP and HTTPS transport protocols are supported EX3400, EX4300, QFX5100, and QFX5200 devices.

  [See Zero Touch Provisioning.]

What's New in Release 20.2R1

Class of Service (CoS)

- **Support for rewrite rules on a per-customer basis on MPC10 and MPC11 (MX Series)**—Starting in Junos OS Release 20.2R1, we support creating rewrite rules on a per-customer basis on MPC10 and MPC11 cards. You can create rewrite rules on a per-customer basis through a policy map. You define policy maps at the **[edit class-of-service policy-map]** hierarchy level, and assign the policy map to a customer through a firewall action, an ingress interface, or a routing policy.

  [See Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps Overview.]

EVPN

- **IPv4 unicast VXLAN encapsulation optimization (MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX10016)—**Starting in Junos OS Release 20.2R1, by default, the listed MX Series routers optimize the IPv4 unicast VXLAN encapsulation process for the following tunnel types:
  - PIM-based VXLAN
  - EVPN-VXLAN
  - Static VXLAN

The optimized encapsulation process results in an increased throughput rate for IPv4 unicast packets between 512 to 1500 bytes in size.

The optimization feature does not support the following:

- EVPN Type-5 tunnels, which are already optimized
- Forwarding table filters

[See Understanding VXLANs.]

- **EVPN on MPLS-over-UDP tunnels (MX Series and vMX)**—Starting in Junos OS Release 20.2R1, Junos OS supports an EVPN network with MPLS-over-UDP tunnels. EVPN uses indirect next hop while MPLS-over-UDP tunnels use tunnel composite next hop (TCNH) in resolving routes in the routing table. In Junos OS releases before Release 20.2R1, indirect next hops for EVPN traffic on MPLS-over-UDP tunnels resolve into unicast next hops. With this release, the indirect next hops for EVPN traffic on MPLS-over-UDP tunnels will resolve into TCNH.

  [See EVPN Overview and Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels.]

- **Support for inline performance monitoring services on EVPN (MX Series)**—Starting in Junos OS Release 20.2R1, you can enable inline performance monitoring services on an EVPN network. With inline performance monitoring, you can configure a greater number of performance monitoring sessions. Inline performance monitoring applies only to delay measurements and synthetic loss measurements. You must also enable both enhanced IP network services and enhanced CFM mode in the device.

  To enable inline performance monitoring, include the following statements:

  - **hardware-assisted-pm** and **hardware-assisted-keepalives enable** statements at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level.
  - **enhanced-ip** statement at the **[edit chassis network-services]** hierarchy level.
  - **enhanced-cfm-mode** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

  [See Connectivity Fault Management Support for EVPN and Layer 2 VPN Overview.]

- **Noncolored SR-TE LSPs with EVPN-MPLS (ACX5448, EX9200, MX Series, and vMX)—**Starting in Junos OS Release 20.2R1, ACX5448, EX9200, MX Series, and vMX routers support noncolored static segment routing-traffic engineered (SR-TE) label-switched paths (LSPs) with an EVPN-MPLS core network and the following Layer 2 services running at the edges of the network:

  - E-LAN
  - EVPN-ETREE
  - EVPN-VPWS with E-Line

  Without color, all LSPs resolve using a BGP next hop only.

The Juniper Networks routers support noncolored SR-TE LSPs in an EVPN-MPLS core network with the following configurations:

- EVPN running in a virtual switch routing instance
- Multihoming in active/active and active/standby modes

The Juniper Networks routers also support noncolored SR-TE LSPs when functioning as a Data Center Interconnect (DCI) device that handles EVPN Type 5 routes.

[See Static Segment Routing Label Switched Path.]

- **Layer 3 gateway in an EVPN-MPLS environment (MPC10 and MPC11 line cards with MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, the supported MX Series routers with MPC10 and MPC11 line cards can act as a default Layer 3 gateway for an EVPN instance (EVI), which can span a set of routers. In this role, the MX Series routers can perform inter-subnet forwarding. With inter-subnet forwarding, each subnet represents a distinct broadcast domain.

  The Layer 3 gateway supports the following features:

  - IRB interfaces through which the default gateway routes IPv4 and IPv6 traffic from one bridge domain to another [See Example: Configuring EVPN with IRB Solution.]
  - Dynamic list next hop [See Configuring Dynamic List Next Hop.]
  - EVPN proxy ARP and ARP suppression, and proxy NDP and NDP suppression on IRB interfaces [See EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression.]
  - The substitution of a source MAC address with a proxy MAC address in an ARP or NDP reply [See ARP and NDP Request with a Proxy MAC Address.]
  - Data center interconnectivity using EVPN Type 5 routes [See EVPN Type-5 Route with MPLS encapsulation for EVPN-MPLS.]

- **Multihoming in an EVPN-MPLS environment (MPC10 and MPC11 line cards with MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, you can multihome a customer edge (CE) device to two or more provider edge (PE) devices (the supported MX Series routers with MPC10 and MPC11 line cards) in an EVPN-MPLS network. We support the following multihoming features:

  - Single-active and all-active modes
  - The configuration of an Ethernet segment identifier (ESI) per interface
  - Preference-based designated forwarder election

  [See EVPN Multihoming Overview.]

- **EVPN-VXLAN (MPC10 and MPC11 line cards with MX2010, MX2020)**—Starting in Junos OS Release 20.2R1, the MX2010 and MX2020 routers with MPC10 and MPC11 line cards installed support the following EVPN-VXLAN features:

  - Layer 2 VXLAN

    - Multihoming with active/active and active/standby modes, an Ethernet segment identifier (ESI) per interface, and preference-based designated forwarder (DF) election
    - MAC pinning, MAC move, MAC limiting, and MAC aging
    - QoS
    - DHCP and DHCP relay
    - Prevention of broadcast, unknown unicast, and multicast (BUM) traffic loops when a leaf device is multihomed to more than one spine device
- Layer 3 VXLAN
    - IRB interfaces
    - IPv6 over IRB interfaces
    - Support for OSPF, IS-IS, BGP, and static routing over IRB interfaces
    - Proxy ARP and ARP suppression, and proxy NDP and NDP suppression with and without IRB interfaces
    - IPv6 underlay
    - Virtual machine traffic optimization (VMTO) for ingress traffic
- Data Center Interconnect (DCI)
    - Nonpure and pure EVPN Type-5 routes
- High availability
    - Nonstop active routing (NSR)
    - Graceful Routing Engine switchover (GRES)
    - Graceful restart from a routing process restart or Routing Engine switchover without NSR enabled
- Operations and management
    - Core isolation feature
    - Ping over EVPN Type-5 tunnel
- Static VXLAN
    - Overlay ping and traceroute

[See EVPN User Guide.]

High Availability (HA) and Resiliency

- **Support for VRRP on the MPC11 (MX2010 and MX2020)**—Starting in Junos OS Release 20.2R1, VRRP is supported on the MPC11 line card. All VRRP features are supported.

  [See Understanding VRRP.]

- **LACP inline support during unified ISSU for multivendor networks (MX104, MX240, MX480, MX960, and MX10003)**—Starting with Junos OS Release 20.2R1, unified ISSU supports LACP interoperability with other vendor devices for fast periodic interval sessions. LACP sessions in full-scale scenarios with interoperability will no longer experience timeouts during unified ISSU.

  Use the **set protocols lacp ppm inline** command to enable LACP inline support.

  [See Getting Started with Unified In-Service Software Upgrade.]

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes mastership. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

  [See Understanding the Ephemeral Configuration Database.]

- **Support for VRRP on the MPC10 and MPC11 (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, VRRP is supported on the MPC11 and MPC10 line cards. All VRRP features are supported.

  [See Understanding VRRP.]

- **Unsupported hardware for unified ISSU (MX240, MX480, MX960, MX10003, and PTX3000)**—The following cards do not support unified ISSU upgrading to Junos OS Release 20.2R1:

  - MPC7E-MRATE
  - MPC8E with MRATE MIC
  - MPC9E with MRATE MIC
  - MPC10E-10C-MRATE
  - MPC10E-15C-MRATE
  - PTX5000 with 24-Port 10-Gigabit Ethernet, 40-Gigabit Ethernet PIC with QSFP+ or 15-Port 10-Gigabit, 40-Gigabit Ethernet, 100-Gigabit Ethernet PIC with QSFP28
  - MX10003 with QSFP28 Ethernet TIC

Interfaces and Chassis

- **Transparent forwarding of CFM packets over VPLS (MX Series)**—In Junos OS Release 20.2R1 and later, MX Series router supports VLAN transparency for connectivity fault management (CFM) packets over Virtual private LAN service (VPLS). If the incoming CFM packets have more **vlan-tags** than the configured interface **vlan-tags**, then CFM PDU is treated transparent. In the earlier Junos OS releases, CFM frame filtering was applied on all CFM PDU including on CFM PDU that had more number of tags than the interface configuration.

  We do not support the following on MX Series routers:

  - Transparency for tagged CFM PDU incoming on untagged interface.
  - Transparency for untagged CFM PDU on interface with native VLAN configuration.

  [See Example: Configuring Ethernet CFM over VPLS.]

- **Support for 400-Gbps port speed (MX240, MX480, and MX960)**—In Junos OS Release 20.2R1, you can configure port speed of 400-Gbps for MPC10E (MPC10E-10C-MRATE and MPC10E-15C-MRATE) on MX240, MX480, and MX960 routers. Use the QSFP56-DD optics to configure 400-Gbps port speed on:

  - MPC10E-10C-MRATE: Port 4 of the MPC
  - MPC10E-15C-MRATE: Port 4 of the MPC

  [See Port Speed.]

- **Support for monitoring link degradation (MX Series routers with MPC10E)**—Starting in Junos OS Release 20.2R1, you can monitor link degradation of the 10-Gigabit Ethernet interfaces, 40-Gigabit Ethernet interfaces, and 100-Gigabit Ethernet interfaces on the MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line cards. Link degradation monitoring enables you to monitor the quality of physical links on interfaces and take corrective action when the link quality degrades beyond a certain value.

  To enable your device to monitor the links, use the **link-degrade-monitor** statement at the **[edit interfaces *interface-name*]** hierarchy level.

  [See Link Degrade Monitoring Overview.]

- **Targeted broadcast support (MPC10E and MX2K-MPC11E)**—Starting in Junos OS Release 20.2R1, you can configure targeted broadcast on broadcast interfaces on the MPC10E and MX2K-MPC11E line cards. Targeted broadcast enables a broadcast packet, destined for a remote network, to transit across networks until the destination network is reached. In the destination network, the packet is broadcast as a normal broadcast packet. This feature is useful when the Routing Engine is flooded with packets to process. You can configure targeted broadcast to forward the packets to :

  - Both the egress interface and the Routing Engine.
  - Egress interface only.

To configure targeted broadcast on an interface, include the **targeted-broadcast** statement at the [**edit interfaces *interface-name* unit *logical-unit-number* family inet**] hierarchy level.

[See Understanding Targeted Broadcast.]

Juniper Extension Toolkit (JET)

- **RIB service APIs support dynamic next-hop interface binding (MX Series, PTX Series, and vMX)**—Starting in Junos OS Release 20.2R1, programmed RIB routes react to Up, Down, Add, and Delete events for direct next-hop interfaces. When all direct next-hop interfaces are unusable, the route becomes inactive. This prevents traffic from being dropped and keeps inactive routes from being propagated through the network.

  This feature applies to all routes programmed using the rib_service JET API where an interface is configured as a direct next hop, including interfaces that are part of a flexible tunnel. It also applies to tunnels configured with the flexible_tunnel_service JET API.

  To disable this feature, use **edit routing-options programmable-rpd rib-service dynamic-next-hop-interface disable**.

  [See rib-service (programmable-rpd), Juniper Extension Toolkit Developer Guide, and Juniper Engineering Network website.]

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

  [See language (Scripts), Develop Off-Device JET Applications, and Develop On-Device JET Applications.]

Junos Telemetry Interface

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

  [See Mapping OpenConfig Routing Policy Commands to Junos Configuration and Mapping OpenConfig Network Instance Commands to Junos Operation.]

- **ON-CHANGE BGP peer information statistics support for JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 provides BGP peer sensor support using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON_CHANGE statistics are sent to an outside collector.

  The following resource paths are supported:

- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/active** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/received** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/sent** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/rejected** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/admin-state** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/established-transitions** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/last-established** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/received/notification** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/messages/received/update** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/notification** (stream
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/update** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/session-state** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/supported-capabilities** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/local-address** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-address** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-port** (ON_CHANGE)

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

- **Telemetry support for LDP and MLDP traffic statistics (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, the following LDP and multipoint LDP native sensors are added for the Junos telemetry interface:
  - /junos/services/ldp/label-switched-path/ingress/usage/
  - /junos/services/ldp/label-switched-path/transit/usage/
  - /junos/services/ldp/p2mp/interface/receive/usage/
  - /junos/services/ldp/p2mp/interface/transmit/usage/
  - /junos/services/ldp/p2mp/label-switched-path/usage/

  You must enable telemetry streaming with the **sensor-based-stats** option at the **[edit protocols ldp traffic-statistics]** hierarchy level.

  The **show ldp traffic-statistics** command is enhanced to display upstream LDP traffic statistics and to display multipoint LDP traffic statistics per interface.

  On PTX Series routers, this feature is not supported for the following variants:
  - PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine
  - PTX10003
  - PTX10008 with the PTX10K-LC1201-36CD line card
  - FPC2 line cards do not support ingress multipoint LDP statistics.

  [See sensor (Junos Telemetry Interface).]

- **gRPC telemetry support for LDP and MLDP traffic statistics (MX Series)**—Starting in Junos OS Release 20.2R1, gRPC support is available to export LDP and multipoint LDP traffic statistics. You can use the following resource paths to export sensor data:
  - LDP LSP transit traffic—**/mpls/signaling-protocols/ldp/lsp-transit-policies/lsp-transit-policy/state/counters**
  - LDP LSP ingress traffic—**/mpls/signaling-protocols/ldp/lsp-ingress-policies/lsp-ingress-policy/state/counters**
  - Multipoint LDP traffic—**/mpls/signaling-protocols/ldp/p2mp-lsps/p2mp-lsp/state/counters**
  - Multipoint LDP egress traffic per-interface—**/mpls/signalling-protocols/ldp/p2mp-interfaces/p2mp-interface/state/counters**
  - Multipoint LDP ingress traffic per-interface—**/mpls/signalling-protocols/ldp/p2mp-interfaces/p2mp-interface/**

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

- **JTI sensor support for Packet Forwarding Engine and Routing Engine sensors (MX Series Virtual Chassis and MX Series routers with dual Routing Engines)**—Junos OS Release 20.2R1 extends Junos telemetry interface (JTI) sensor support for all Packet Forwarding Engine and Routing Engine sensors currently

supported on MX Series routers to include MX routers with dual Routing Engines or MX Series Virtual Chassis. The level of sensor support currently available for MX Series routers applies, whether through streaming or ON_CHANGE statistics export, using UDP, remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. Additionally, JTI operational mode commands will provide details for all Routing Engines and MX Series Virtual Chassis, too.

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

- **JTI sensor support for standby Routing Engine statistics (MX480, MX960, MX10003, MX2010, and MX2020)**—Junos OS Release 20.2R1 provides Junos telemetry interface (JTI) sensor support for standby Routing Engine statistics using remote procedure call (gRPC) services. This feature is supported on both single chassis and virtual chassis unless otherwise indicated. Use this feature to better track the state of software components running on a standby Routing Engine. Statistics exported to an outside collector through the following sensors (primarily under subscriber management) provide a more complete view of the system health and resiliency state:

  - Chassis role (backup or master) sensor **/junos/system/subscriber-management/chassis** and **/junos/system/subscriber-management/chassis[chassis-index=*chassis-index*]** (for specifying an index for an MX Series Virtual Chassis)
  - Routing Engine status and GRES notification sensor **/junos/system/subscriber-management/chassis/routing-engines/routing-engine** and **/junos/system/subscriber-management/chassis/routing-engines/routing-engine[re-index=*RoutingEngineIndex*]** (to specify an index number for a specific Routing Engine)
  - Subscriber management process sensor **/junos/system/subscriber-management/chassis/routing-engines/process-status/subscriber-management-processes/subscriber-management-process** and **/junos/system/subscriber-management/chassis/routing-engines/process-status/subscriber-management-processes/subscriber-management-process[pid=*ProcessIdentifier*]** (to specify a PID for a specific process)
  - Per Routing Engine DHCP binding statistics for server or relay sensor **/junos/system/subscriber-management/chassis/routing-engines/routing-engine/dhcp-bindings/dhcp-element[dhcp-type-name=*RelayOrServer*/v4]** and **/junos/system/subscriber-management/chassis/routing-engines/routing-engine/dhcp-bindings/dhcp-element[dhcp-type-name=*RelayOrServer*/v6]**
  - Virtual Chassis port counter sensor **/junos/system/subscriber-management/chassis/virtual-chassis-ports/virtual-chassis-port** and **/junos/system/subscriber-management/chassis/virtual-chassis-ports/virtual-chassis-port[vcp-interface-name=*vcp-interface-port-string*] (to specify the interface name)**. This resource path is only supported on a virtual chassis.

  [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface) and Understanding OpenConfig and gRPC on Junos Telemetry Interface.]

- **CPU statistics support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 supports streaming various CPU statistics and process parameters using remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) services and Junos telemetry interface (JTI). You can stream CPU usage per process (statistics are similar to output

from the **show system process detail** operational mode command), as well as CPU usage per Routing Engine core.

This feature supports the private data model openconfig-procmon.yang.

To stream statistics to an outside collector, include the following resource paths in a gRPC or gNMI subscription:

- Individual process level information (resource path **/system/processes/process/**)
- Individual Routing Engine core information (resource path **/components/component/cpu/**)

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

- **TARGET_DEFINED subscription mode support with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX10016)**—Junos OS Release 20.2R1 adds support for TARGET-DEFINED mode for subscriptions made using gRPC Network Management Interface (gNMI) services.

  Using a gNMI subscription, an external collector stipulates how sensor data should be delivered:

  - STREAMING mode periodically streams sensor data from the DUT at a specified interval.
  - ON_CHANGE mode sends updates for sensor data from the DUT only when data values change.
  - Newly supported TARGET_DEFINED mode (submode 0) instructs the DUT to select the relevant mode (STREAMING or ON_CHANGE) to deliver each element (leaf) of sensor data to the external collector. When a subscription for a sensor with submode 0 is sent from the external collector to the DUT, the DUT responds, activating the sensor subscription so that periodic streaming does not include any of the ON_CHANGE updates. However, the DUT will notify the collector whenever qualifying ON_CHANGE events occur.

  [See Understanding OpenConfig and gRPC and gNMI on Junos Telemetry Interface.]

- **Packet Forwarding Engine sensor support with INITIAL_SYNC on JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000 line of routers, QFX5100, and QFX5200)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export Packet Forwarding Engine statistics from devices to an outside collector using gNMI submode INITIAL_SYNC. When an external collector sends a subscription request for a sensor with INITIAL_SYNC (gnmi-submode 2), the host sends all supported target leaves (fields) under that resource path at least once to the collector with the current value. This is valuable because:

  - The collector has a complete view of the current state of every field on the device for that sensor path.
  - Event-driven data (ON_CHANGE) is received by the collector at least once before the next event is seen. In this way, the collector is aware of the data state before the next event happens.
  - Packet Forwarding Engine sensors that contain zero counter values (zero-suppressed) that normally do not show up in streamed data are sent, ensuring that all fields from each line card (also referred to as source) are known to the collector.

> **NOTE:** ON_CHANGE data is not available for native (UDP) Packet Forwarding Engine Sensors.

INITIAL_SYNC submode requires that at least one copy to be sent to the collector; however, sending more than one is acceptable.

INITIAL_SYNC submode is supported for the following sensors:

- Sensor for CPU (ukernel) memory (resource path **/junos/system/linecard/cpu/memory/**)
- Sensor for firewall filter statistics (resource path **/junos/system/linecard/firewall/**)
- Sensor for physical interface traffic (resource path **/junos/system/linecard/interface/**)
- Sensor for logical interface traffic (resource path **/junos/system/linecard/interface/logical/usage/**)
- Sensor for physical interface queue traffic (resource path **/junos/system/linecard/interface/queue/**)
- Sensor for physical interface traffic except queue statistics (resource path **/junos/system/linecard/interface/traffic/**)
- Sensor for NPU memory (resource path **/junos/system/linecard/npu/memory/**)
- Sensor for NPU utilization (resource path **/junos/system/linecard/npu/utilization/**)
- Sensor for packet statistics (resource path **/junos/system/linecard/packet/usage/**)
- Sensor for software-polled queue-monitoring statistics (resource path **/junos/system/linecard/qmon-sw/**)

[See Understanding OpenConfig and gRPC and gNMI on Junos Telemetry Interface and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

- **Export data using JSON encoding format with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX10016)**—Junos OS Release 20.2R1 adds support for JavaScript Object Notation (JSON) encoding to export telemetry data using gRPC network management interface (gNMI) services and Junos telemetry interface (JTI). JSON is an open standard file format and data interchange format that provides a good balance of usability and performance. It uses human-readable text to store and transmit data objects consisting of attribute–value pairs and array data types.

  To export telemetry data using JSON encoding, include **format json-gnmi** at the [**edit services analytics export-profile *profile-name***] hierarchy level. This is part of the export profile CLI configuration used to configure collector and sensor details in Junos OS.

  [See export-profile (Junos Telemetry Interface).]

- **SR-TE statistics for uncolored SR-TE policies streaming on JTI (MX240. MX480, MX960, MX2010, and MX2020 with MPC-10E or MPC-11E)**—Junos OS Release 20.2R1 provides segment routing-traffic engineering (SR-TE) per label-switched path (LSP) route statistics using Junos telemetry interface (JTI)

and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream SR-TE telemetry statistics for uncolored SR-TE policies to an outside collector.

Ingress statistics include statistics for all traffic steered by means of an SR-TE LSP. Transit statistics include statistics for traffic to the binding SID (BSID) of the SR-TE policy.

To enable these statistics, include the **per-source per-segment-list** statement at the [**edit protocols source-packet-routing telemetry statistics**] hierarchy level.

If you issue the **set protocols source-packet-routing telemetry statistics no-ingress** command, ingress sensors are not created.

If you issue the **set protocols source-packet-routing telemetry statistics no-transit** command, transit sensors are not created. Otherwise, if BSID is configured for a tunnel, transit statistics are created.

The following resource paths (sensors) are supported:

- **/junos/services/segment-routing/traffic-engineering/tunnel/lsp/ingress/usage/**
- **/junos/services/segment-routing/traffic-engineering/tunnel/lsp/transit/usage/**

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface), source-packet-routing, and show spring-traffic-engineering lsp detail name *name*.]

Layer 2 VPN

- **Support for Layer 2 interworking (iw0) interface on the MPC10E and MPC11E line cards (MX Series)**—Starting in Junos OS Release 20.2R1, you can connect Layer 2 networks together by configuring a Layer 2 interworking (iw0) route with iw0 interfaces. This feature supports the following interconnections:
  - Layer 2 circuit to Layer 2 circuit
  - Layer 2 circuit to Layer 2 VPN
  - Layer 2 VPN to Layer 2 circuit
  - Layer 2 VPN to Layer 2 VPN

[See Using the Layer 2 Interworking Interface to Interconnect a Layer 2 Circuit to a Layer 2 VPN and Layer 2 VPN to Layer 2 VPN Connections.]

Layer 3 Features

- **MPC10E interoperates with MS-MPC/MS-MICs for Layer 3 Services (MX240,MX480, and MX960)**—Starting in Junos OS Release 20.2, the MPC10E interoperates with MS-MPC/MS-MICs for Layer 3 Services such as active flow monitoring, IPSec, NAT, RPM, and stateful firewall. [See Layer 2 and Layer 3 Features on MX Series Routers.]

Management

- **Error recovery, fault handling, and resiliency support for MX2K-MPC11E (MX2010 and MX2020)**—Starting in Junos OS Release 20.2R1, the MX2010 and MX2020 routers with the MX2K-MPC11E line card support error recovery, fault handling, and software resiliency. The MX2K-MPC11E line cards support detecting errors, reporting them through alarms, and triggering resultant actions. To view application-level errors, use the **show trace node fpc<#> application fabspoked-pfe** command. To check the status of the card, use the **show chassis fpc pic-status** command. Use the **show chassis errors active** command to view the fault details and the **show system alarm** command to view the alarm details.

  [See show chassis fpc pic-status and clear chassis fpc errors.]

MPLS

- **Support to change the default re-merge behavior on the P2MP LSP (MX Series)**—Starting with Junos OS Release 20.2R1, you can change the default re-merge behavior on RSVP P2MP LSP. The term re-merge refers to the case of an ingress (headend) or transit node (re-merge node) that creates a re-merge branch intersecting the P2MP LSP at another node in the network. This may occur due to events such as an error in path calculation, an error in manual configuration, or network topology changes during the establishment of the P2MP LSP.

  You can configure the no re-merge behavior on P2MP LSPs by enabling the newly introduced **no-re-merge** and **no-p2mp-re-merge** CLI commands at the ingress (headend) and transit devices (re-merge nodes), respectively.

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

- **Support for MPLS ping and traceroute for segment routing (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.2R1, we extend the MPLS ping and traceroute support for all types segment routing--traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

  We also support the following features:

  - FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
  - ECMP traceroute support for all types of SR-TE paths.

We do not support the following:

- Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
- OAM for IPv6 prefix.
- BFD

[See traceroute mpls segment-routing spring-te and ping mpls segment routing spring-te.]

- **MPLS support (MX Series routers with MPC10E and MPC11E)**—Starting in Junos OS Release 20.2R1, some of the MPLS features are supported on MX Series routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MX2K-MPC11E line cards.

  [See Protocols and Applications Supported by the MPC10E and Protocols and Applications Supported by the MX2K-MPC11E.]

Multicast

- **Fast failover according to flow rate (MX Series with MPC10E or MPC11E line cards)**—Starting in Junos OS Release 20.2R1, for routers operating in Enhanced IP Network Services mode, you can configure a threshold that triggers fast failover in next-generation MVPNs with hot-root standby on the basis of aggregate flow rate. For example, fast failover (as defined in *Draft Morin L3VPN Fast Failover 05*) is triggered if the flow rate of monitored multicast traffic from the provider tunnel drops below the set threshold.

  [See min-rate.]

Network Management and Monitoring

- **SNMP support for multicast LDP MIB objects (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS SNMP extends support for the following multicast LDP MIB tables and objects:
  - mplsMldpInterfaceStatsTable
  - mplsMldpFecUpstreamSessPackets
  - mplsMldpFecUpstreamSessBytes
  - mplsMldpFecUpstreamSessDiscontinuityTime

  The multicast LDP standard MIB builds on the objects and tables that are defined in RFC3815, which only supports LDP point-to-point label-switched paths (LSPs). This multicast LDP MIB provides support for managing multicast LDP point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs.

  [See Standard SNMP MIBs Supported by Junos OS and SNMP MIB Explorer.]

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

  [See NETCONF Sessions over Outbound HTTPS.]

- **Enhanced on-box monitoring support on the control plane (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure traceoptions to track all events related to system-level and process-level memory monitoring. You can also view the history of the actions taken for system-level and process-level memory monitoring by using the **show system monitor memory actions** command.

Next Gen Services

- **Support for Dual Stack Lite (DS-Lite) Softwires**—Starting in Junos OS Release 20.2R1, Dual Stack Lite (DS-Lite) softwires are supported for CGNAT Next Gen Services. DS-Lite allows service providers to migrate to an IPv6 network while continuing to support IPv4 services; even after the exhaustion of the IPv4 address space. You can natively allocate IPv6 addresses to customers while legacy end-user devices accessing the IPv4 Internet remain same. Thus, IPv4 devices continue to access the IPv4 Internet with minimum disruption on their home networks. DS-Lite also de-couples IPv6 deployment in the service provider network from the rest of the Internet, making incremental deployment easier.

  [See DS-Lite Softwires—IPv4 over IPv6 for Next Gen Services.]

- **Support for HTTP Content Manager (HCM)**—Starting in Junos OS Release 20.2R1, HTTP Content Manager (HCM) is supported under Next Gen Services. HCM is an application that inspects the HTTP traffic transmitted through port 80 (default) or any other port you use to transmit HTTP traffic. HCM inspects HTTP traffic even if the default port 80 is not used for HTTP traffic and is interoperable with ms, rms, and ams interface types. It supports fragmented HTTP request packets and GET, PUT, and POST requests.

  [See HTTP Content Manager (HCM).]

- **Support for Mapping of Address and Port with Encapsulation (MAP-E) Softwires for CGNAT Next Gen Services**—Starting in Junos OS Release 20.2R1, Mapping of Address and Port with Encapsulation (MAP-E) softwires are supported for CGNAT Next Gen Services. MAP-E is an automatic tunneling mechanism tailored for deployment of IPv4 to end users via a service provider's IPv6 network infrastructure. Using MAP-E technology, islands of v4 networks can be connected via v6 tunnels. The IPV4 packets are carried in IPV4-over-IPV6 tunnels from the MAP-E Customer Edge (CE) routers to the MAP-E Border Relay(s) (BR) (through IPV6 routing topology), where they are de-tunneled for further processing. MAP-E can be used by Service Providers to provide IPv4 connectivity to their subscribers over the ISP's IPv6 access network.

[See Mapping of Address and Port with Encapsulation (MAP-E) for Next Gen Services.]

- **Support for Network Address Translation and Protocol Translation for CGNAT Next Gen Services**—Starting in Junos OS Release 20.2R1, Network Address Translation and Protocol Translation (NAT-PT) [RFC2766] are supported for CGNAT Next Gen Services. NAT-PT is a IPv4-to-IPv6 transition mechanism that provides a way for end-nodes in IPv6 realm to communicate with end-nodes in IPv4 realm and vice versa. This is achieved using a combination of Network Address Translation and Protocol Translation.

  [See NAT46 Next Gen Services Configuration Examples.]

- **Support for Port Control Protocol Support (PCP) for DS-Lite for CGNAT Next Gen Services**—Starting in Junos OS Release 20.2R1, Port Control Protocol Support (PCP) for DS-Lite is supported for CGNAT Next Gen Services. DS-Lite is a technology which enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT).

  Typically, the home gateway embeds a Basic Bridging BroadBand (B4) capability that encapsulates IPv4 traffic into a IPv6 tunnel to the CGNAT, named the Address Family Transition Router (AFTR). AFTRs are run by service providers.

  PCP allows customer applications to create mappings in a NAT for new inbound communications destined to machines located behind a NAT. In a DS-Lite environment, PCP servers control AFTR devices.

  [See Port Control Protocol Overview.]

Operation, Administration, and Maintenance (OAM)

- **Support for connectivity fault management (CFM) on MPC10E and MX2K-MPC11E**—Starting in Junos OS Release 20.2R1, you can configure the IEEE 802.1ag OAM CFM Down maintenance association end points (MEPs) on MPC10E and MX2K-MPC11E to monitor Ethernet networks for connectivity faults.

  Junos OS supports the continuity check messages (CCM) and loopback messages as defined in IEEE 802.1ag.

  [See Configuring Connectivity Fault Management.]

Routing Policy and Firewall Filters

- **ARP policer support on pseudowire interfaces (MX Series)**—Starting in Junos OS Release 20.2R1, you can create policers for ARP traffic on pseudowire interfaces. Configure rate limiting for the policer by specifying the bandwidth and the burst-size limit of a firewall policer and attaching the policy to a pseudowire interface, just like you would any other interface. Traffic that exceeds the specified rate limits can be dropped or marked as low priority and delivered when congestion permits.

  In the case of denial of service (DoS) or ARP broadcast storms, ARP policers protect the Routing Engine against malicious traffic intended to degrade the network.

  Apply the ARP policer to a pseudowire interface at the [**edit interfaces *interface-name* unit *unit-number* family inet policer arp *policy-name***] level of the hierarchy.

[See ARP Policer Overview.]

- **Support for P2MP and P2P automatic LSP policers (MX Series)**—Starting in Junos OS Release 20.2R1, support for automatic policers on point-to-multipoint (P2MP) label-switched paths (LSPs) is available on MX240, MX480, MX960, MX2010, and MX2020 routers with MPC10E and MPC11E line cards.

  P2MP MPLS LSP is either an LDP-signaled, or RSVP-signaled, LSP with a single source and multiple destinations that can optimize packet replication at the ingress router. With it, packet replication only occurs for packets being forwarded to two or more different destinations requiring different network paths. Automatic LSP policing lets you provide strict service guarantees for network traffic in accordance with the bandwidth configured for the LSPs.

  Also supported with this release are the following features:

  - Graceful Routing Engine switchover (GRES) at the ingress and egress
  - Load balancing over aggregated links
  - P2MP statistics
  - Multiprotocol BGP-based multicast VPNs (or Layer 3 VPN multicast)

  [See Configuring Automatic Policers.]

- **Support for firewall forwarding (MX Series)**—Starting in Junos OS Release 20.2R1, the following traffic policers are supported on MX240, MX480, MX960, MX2010, and MX2020 routers with MPC10E or MPC11E line cards:

  - GRE tunnels, including encapsulation (**family any**), de-encapsulation, GRE-in-UDP over IPv6, and the following sub-options: sample, forwarding class, interface group, and no-ttl-decrement
  - Input and output filter chains
  - Actions, including policy-map filters, do-not-fragment, and prefix
  - Layer 2 policers
  - Policer overhead adjustment
  - Hierarchical policers
  - Shared bandwidth
  - Percentages
  - Logical interfaces

  [See Traffic Policer Types.]

Routing Protocols

- **TI-LFA SRLG protection for IS-IS (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure Shared Risk Link Group (SRLG) protection for segment routing to choose a fast reroute path that does not include SRLG links in the topology-independent loop-free alternate (TI-LFA) backup paths. This is in addition to existing fast reroute options such as **link-protection**, **node protection**, and **fate-sharing protection** for segment routing. IS-IS computes the fast reroute path that is aligned with the post-convergence path and excludes the SRLG of the protected link. All local and remote links that are from the same SRLG as the protected link are excluded from the TI-LFA back up path. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface.

  To enable TI-LFA SRLG protection with segment routing for IS-IS, include the **srlg-protection** statement at the **[edit protocols isis interface *name* level *number* post-convergence-lfa]** hierarchy level.

  [See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS.]

- **Support for BGP-LU over SR-TE for color-based mapping of VPN Services (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, we are extending support to BGP labeled unicast service for color-based mapping of VPN services over Segment Routing-Traffic Engineering (SR-TE). This enables you to advertise BGP-LU IPv6 and IPv4 prefixes with an IPv6 next-hop address in IPv6-only networks where routers do not have any IPv4 addresses configured. With this feature, BGP-LU can now resolve IPv4 and IPv6 routes over SR-TE core. BGP-LU constructs a colored protocol next hop, which is resolved on a colored SR-TE tunnel in the **inetcolor.0** or **inet6color.0** table. Currently we support BGP IPv6 LU over SR-TE with IS-IS underlay.

  See [Understanding Static Segment Routing LSP in MPLS Networks.]

- **Support for AIGP metric to MED translation (MX2010 and MX2020)**—Starting in Release 20.2R1, Junos OS supports the translation of AIGP metric to MED. You can enable this feature when you want the end to end effective AIGP metric in order to choose the best path. Effective AIGP is the AIGP value advertised with the route plus the IGP cost to reach the nexthop. This is especially useful in Inter-AS MPLS VPNs solution, where customer sites are connected via two different service providers, and customer edge routers want to take IGP metric based decision. You can configure a minimum-aigp to prevent unnecessary update of route when effective-aigp changes past the previously known lowest value.

  The following configuration statements are introduced at the **[edit protocols bgp group *<group-name>* metric-out]** hierarchy level:

  - **effective-aigp** to track the effective AIGP metric
  - **minimum-effective-aigp** to track the minimum effective AIGP metric.

  [See effective-aigp and minimum-effective-aigp.]

- **Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast (MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices)**—Starting with Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices support BGP PIC Edge protection for Layer 2 circuit, Layer 2 VPN, and VPLS (BGP VPLS, LDP VPLS and FEC 129 VPLS) services with BGP

labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

A prerequisite for BGP PIC Edge protection is to program the Packet Forwarding Engine (PFE) with expanded next-hop hierarchy.

To enable BGP PIC Edge protection, use the following CLI configuration statements:

- Expand next-hop hierarchy for BGP labeled unicast family:

```
[edit protocols]
user@host#set bgp group group-name family inet labeled-unicast nexthop-resolution
    preserve-nexthop-hierarchy;
```

- BGP PIC for MPLS load balance nexthops:

```
[edit routing-options]
user@host#set rib routing-table-name protect core;
```

- Fast convergence for Layer 2 circuit and LDP VPLS:

```
[edit protocols]
user@host#set l2circuit resolution preserve-nexthop-heirarchy;
```

- Fast convergence for Layer 2 VPN, BGP VPLS, and FEC129:

```
[edit protocols]
user@host#set l2vpn resolution preserve-nexthop-heirarchy;
```

[See Load Balancing for a BGP Session.]

- **Support for dynamic peer AS range for BGP groups (ACX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, you can configure acceptable autonomous system (AS) ranges for EBGP groups that can be used for bringing up BGP peers while establishing a BGP session. BGP accepts a peer request based on the configured AS range and rejects a peer request if the AS does not fall into the specified range. This allows you to control BGP peering when the neighbor's exact IP address is not known.

  To define peer AS range for BGP groups through policy, you can include the **as-list** statement at the **[edit policy-options]** hierarchy level. To include the specified peer AS list, include the **peer-as-list *peer-as-list*** statement at the **[edit protocols bgp group *group-name*]** hierarchy level.

  See [peer-as-list and as-list.]

- **Support for BGP-SR-TE rearchitecture (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS provides support for controller-based BGP segment routing--traffic engineering (SR-TE) routes

to be installed as source packet routing traffic-engineered (SPRING-TE) routes. BGP installs the SR-TE policy in the routing tables bgp.inetcolor.0 and bgp.inet6color.0, and these routes are subsequently installed in the routing tables inetcolor.0 or inet6color.0 by SPRING-TE.

In releases before Junos OS Release 20.2R1, controller-based BGP SR-TE routes are installed as BGP routes in the routing table. To maintain consistency and for easy maintenance, all SR-TE based routes appear as SPRING-TE routes irrespective of the source.

You need to enable **source-packet-routing** at the **[edit protocols]** hierarchy level to see the routes installed in inetcolor.0 or inet6color.0. A new option **detail** is introduced under **traceoptions (Protocols Spring-TE)** to trace the detailed information.

See [Segment Routing Traffic Engineering at BGP Ingress Peer Overview.]

- **Support for egress protection and BGP PIC features (MX Series Routers with MPC10E and MPC11E)**—Starting in Junos OS Release 20.2R1, you can configure the following egress link protection and BGP Prefix Independent Convergence (PIC) features on MX Series devices with MPC10E and MPC11E.
  - **Egress protection for BGP labeled unicast** —Fast protection for egress nodes is available to services in which BGP labeled unicast interconnects IGP areas, levels, or autonomous systems (ASs). If a provider router detects that an egress router (AS or area border router) is down, it immediately forwards the traffic destined to that router to a protector router that forwards the traffic downstream to the destination.
  - **Provider edge link protection for BGP labeled unicast paths**— You can configure a precomputed protection path in a Layer 3 VPN such that if a BGP labeled-unicast path between an edge router in one AS and an edge router in another AS goes down, you can use the protection path (also known as the backup path) between alternate edge routers in the two ASs. This is useful in a carrier-of-carriers deployments, where a carrier can have multiple labeled-unicast paths to another carrier. In this case, the protection path avoids disruption of service if one of the labeled-unicast paths goes down.
  - **BGP PIC for inet** —We've extended the BGP Prefix Independent Convergence (PIC) support to BGP with multiple routes in the global tables such as inet and inet6 unicast, and inet and inet6 labeled unicast. When you enable the BGP PIC feature on a router, BGP installs to the Packet Forwarding Engine the second best path in addition to the calculated best path to a destination. When an IGP loses reachability to a prefix, the router uses this backup path to reduce traffic loss until the global convergence through BGP is resolved, thereby drastically reducing the outage duration.
  - **BGP (PIC Edge for RSVP** —With BGP PIC Edge in an MPLS VPN network, IGP failure triggers a repair of the failing entries and causes the Packet Forwarding Engine to use the prepopulated protection path until global convergence has re-resolved the VPN routes. The convergence time is no longer dependent on the number of prefixes. When RSVP receives a tunnel down notification at the ingress PE router, it sends a notification to the Packet Forwarding Engine to start making use of the tunnel to the alternate egress PE router.

  [See Egress Protection for BGP Labeled Unicast ,Understanding Provider Edge Link Protection for BGP Labeled Unicast Paths, Use Case for BGP PIC for Inet, and show rsvp version.]

Services Applications

- **Interoperability of MPC10E with MS-MPC and MS-MIC for Layer 3 Services (MX240, MX480,and MX960)**—Starting in Junos OS Release 20.2R1, the MPC10E-15C-MRATE interoperates with MS-MPC and MS-MIC-16G to support the following Layer 3 Services:
  - Stateful firewall
  - NAT
  - IPSec
  - RPM
  - MS-MPC/MS-MIC based Inline flow monitoring services
- **Support for RFC 2544-based benchmarking tests (MX Series routers with MPC10E and MX2K-MPC11E)**—Junos OS Release 20.2 extends support for the reflector function and the corresponding RFC 2544-based benchmarking tests on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MX2010 and MX2020 routers with MX2K-MPC11E. The RFC 2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before activation of the service. The tests measure throughput, latency, frame loss rate, and back-to-back frames.

  RFC 2544-based benchmarking tests on MX Series routers support the following reflection functions:
  - Ethernet pseudowire reflection (ingress and egress direction) (ELINE service—supported for family **ccc**)
  - Layer 2 reflection (egress direction) (ELAN service—supported for family **bridge, vpls**)
  - Layer 3 IPv4 reflection (limited support)

  To run the benchmarking tests on the MX Series routers, you must configure reflection (Layer 2 or pseudowire) on the supported MPC. To configure the reflector function on the MPC, use the **fpc *fpc-slot-no* slamon-services rfc2544** statement at the [**edit chassis**] hierarchy level.

  [See Understanding RFC2544-Based Benchmarking Tests on MX Series Routers].
- **Support for random load balancing (MX Series routers with MPC10E and MX2K-MPC11E)**—Starting in Junos OS Release 20.2R1, you can configure per packet random load balancing on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MX2010 and MX2020 routers with MX2K-MPC11E. Per-packet random spray load balancing ensures that the members of ECMP are equally loaded without taking bandwidth into consideration. Random load balancing also eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure random load balancing on the MPC, include the **load-balance random** statement at the [**edit policy-options policy-statement *policy-name* term *term-name* then**] hierarchy level.

  [See Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers].
- **Support for static IP tunnels (MX Series routers with MPC10E and MX2K-MPC11E)**—Starting in Junos OS Release 20.2R1, MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and

MPC10E-10C-MRATE) and MX2010 and MX2020 routers with MX2K-MPC11E support static IP tunnels with:

- Encapsulation support of the following types:
  - IPv4-over IPv4
  - IPv6-over-IPv4
  - IPv4-over-IPv6
  - IPv6-over-IPv6
- Scaling upto 4000 tunnels per PIC
- Graceful Routing Engine switchover (GRES)

Software-Defined Networking (SDN)

- **Manual (PIM-based) VXLAN support (MPC10 and MPC11 line cards with MX2010 and MX2020)—**Starting in Junos OS Release 20.2R1, the MX2010 and MX2020 routers with MPC10 and MPC11 line cards installed support manual (PIM-based) VXLAN.

  [See Understanding VXLANs.]

- **GNFs with MX-SPC3 support carrier-grade NAT services over abstracted fabric interfaces (MX480 and MX960)**—Starting in Junos OS Release 20.2R1, guest network functions running Next Gen Services with the MX-SPC3 card support carrier-grade NAT services.

  The support includes the following:

  - NAT translation types—dnat-44, dynamic-nat44, basic-nat44, basic-nat66, twice-basic-nat-44, twice-dynamic-nat44, deterministic NAT. Support for interface and next-hop style service sets, EIM/EIF, PBA, XLAT464, and port forwarding are available. Support for basic-nat44, basic-nat66 over layer 3 VPN is also available.
  - SIP and RTSP Application Layer Gateways
  - carrier-grade events logging, using the Junos Traffic Vision (J-Flow).
  - Class of service (CoS)

  > **NOTE:** To support the services traffic over abstracted fabric interfaces, a GNF that has an MX-SPC3 card assigned to it must also have a line card linked to it.

  [See Junos OS Carrier-Grade NAT Implementation Overview.]

- **GNFs with MX-SPC3 support various services over abstracted fabric interfaces (MX480 and MX960)**—Starting in Junos OS Release 20.2R1, guest network functions (GNFs) running Next Gen Services with the MX-SPC3 card support the following services over abstracted fabric interfaces:

- DNS filtering to identify DNS requests for blacklisted website domains.
- URL filtering to determine which Web content is not accessible to users.

To support the services traffic over abstracted fabric interfaces, a GNF that has an MX-SPC3 card assigned to it must also have a line card linked to it.

[See DNS Request Filtering for Blacklisted Website Domains and Configuring URL Filtering]

Subscriber Management and Services

- **RADIUS-sourced connection status updates to CPE devices (MX Series)**—Starting in Junos OS Release 20.2R1, you can use RADIUS-sourced messages to convey information, such as upstream bandwidth or connection rates, that the BNG transparently forwards to CPE devices. Configure RADIUS to send the router the Juniper Networks Connection-Status-Message VSA (26-4874–218) in Access-Accept or CoA messages. Include the **lcp-connection-update** PPP option in the client dynamic profile to enable PPP to send the VSA contents to the CPE device in the Connection-Status-Message option of an LCP Connection-Update-Request message.

  [See RADIUS-Sourced Connection Status Updates to CPE Devices.]

- **Identifying dynamic profile versions with version aliases (MX Series)**—Starting in Junos OS Release 20.2R1, you can use the **versioning-alias** statement to configure a text description that identifies a particular variation of a dynamic client profile. The version alias is conveyed to the RADIUS server in the Access-Accept message in the Juniper Networks Client-Profile-Name VSA (26–4874–174).

  [See Versioning for Dynamic Profiles.]

- **IPFIX support for per-subscriber queue statistics (MX Series)**—Starting in Junos OS Release 20.2R1, you can configure the input-jti-ipfix plug-in to collect per–subscriber interface queue statistics. The output ipfix-plugin can then export the statistics as IPFIX template and data records.

  [See Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector.]

- **Junos Multi-Access User Plane support (MX204, MX10003)**—Starting with Junos OS Release 20.2R1, you can configure Junos Multi-Access User Plane on MX204 and MX10003 routers. Junos Multi-Access User Plane is a software solution that turns your MX Series router into a high-capacity user plane function called a System Architecture Evolution Gateway-User Plane (SAEGW-U). This MX Series SAEGW-U interoperates with a third-party SAEGW-C (control plane function), according to the 3GPP Release 14 Control User Plane Separation (CUPS) architecture, to provide high-throughput 4G fixed-wireless access service. CUPS enables independent scaling of the user and control planes, network architecture flexibility, operational flexibility, and an easier migration path from 4G to 5G services. The CUPS architecture is optional for 4G but inherent in 5G architecture.

  [See Junos Multi-Access User Plane User Guide.]

System Logging

- **Support to track the maximum number of routing and forwarding (RIB/FIB) routes and VRFs (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can track and display the high-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system (RPD) using the **show route summary** CLI command. High-water mark refers to the maximum number of routing and forwarding (RIB/FIB) table routes and VRFs that was present in the RPD system. The high-water mark data can also be viewed in the syslog at the **LOG_NOTICE** level.

  You can configure the interval of the high-water mark data using the **highwatermark-log-interval** CLI configuration statement at the [**edit routing-options**] hierarchy level. The minimum time gap at which the high-water mark data logged in the syslog is 30 seconds. You can configure the value for **highwatermark-log-interval** CLI configuration statement between 5 to 1200 seconds.

  [See routing-options and show route summary.]

System Management

- **Support for the G.8275.1 Profile (MX10008 and MX10016 with line card JNP10K-LC2101)**—Starting in Junos OS Release 20.2R1, ITU-T G.8275.1 Full path Timing Support (FTS) Profile and G.8273.2 Telecom Boundary Clock supported. It is phase profile and it operates with PTP based packet exchange for Phase and Time recovery, and Sync-E based frequency recovery. Also called as Sync-E assisted PTP mode of operation. This profile is required in TDD application deployment in both 4G and 5G.

  The PTP operation must be two-way in this profile in order to transport phase/time synchronization because propagation delay must be measured. Hybrid mode must be enabled for the G.8275.1 profile.

  [See profile-type.]

Virtual Chassis

- **MX Series Virtual Chassis support for the ephemeral database (MX480 and MX960)**—Starting in Junos OS Release 20.2R1, MX Series Virtual Chassis support configuring the ephemeral database. The ephemeral database is an alternate configuration database that provides a fast programmatic interface for performing configuration updates on devices running Junos OS.

  [See Understanding the Ephemeral Configuration Database.]

SEE ALSO

What's Changed | **113**

Known Limitations | **116**

Open Issues | **119**

Resolved Issues | **128**

Documentation Updates | **145**

Migration, Upgrade, and Downgrade Instructions | **146**

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for MX Series routers.

Class of Service (CoS)

- We've corrected the output of the **show class-of-service interface | display xml** command. Output of the following sort: <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> will now appear correctly as <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3></container> <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container>.

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

  [See commit (System).]

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

- **Command to view summary information for resource monitor (MX Series routers and EX9200 line of switches)**—You can use the **show system resource-monitor** command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays

information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

See show system resource-monitor and Resource Monitoring for Subscriber Management and Services.

Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

  [See Develop Off-Device JET Applications and Develop On-Device JET Applications.]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

  [See Juniper EngNet.]

Network Management and Monitoring

- **Support for new SNMP object for the ifJnx MIB**—Starting in Junos OS Release 20.2R1, we introduce a new SNMP object, **ifJnxInputErrors**, that tracks all input errors except the L3 incomplete errors. The **ifJnxInErrors** object continues to track the L3 incomplete errors.

- **Support for Clearing the Event at MEP Level (MX Series)**—In Junos OS 20.2R1, you can define an action profile for connectivity fault management at the local MEP level or at the remote MEP level. You define an action profile to monitor events and thresholds and specify an action that the device performs when the configured event occurs. When you define the action profile at the local MEP level, you can clear the event for the configured action profile at the local MEP level by specifying only the local MEP numeric identifier. When you define the action profile at the remote MEP level, you can clear the event for the configured action profile at the remote MEP level by specifying the local MEP numeric identifier as well as the remote MEP numeric identifier.

  See [clear oam ethernet connectivity-fault-management event.]

- **Request support information for IPsec function (MX Series)**—Starting in Release 20.2R1, Junos OS introduces **ipsec-vpn** option to the existing **request support information** command. The **request support information ipsec-vpn** command displays all the configurations, states, and statistics at Routing Engine and Service Card level. This new option helps in debugging IPsec-VPN related issues. The information collection is streamlined and reduces the output file size.

See [Request support information.]

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

Services Applications

- **New option for configuring delay in IPSec SA installation**—In Junos OS Releases 20.2R1 and 20.2R2, you can configure the **natt-install-interval *seconds*** option under the **[edit services ipsec-vpn rule *rule-name* term *term-name* then dynamic]** hierarchy to specify the duration of delay in installing IPSec SA in a NAT-T scenario soon after the IPSec SA negotiation is complete. The default value is 0 seconds.

Software-Defined Networking (SDN)

- **JDM install and configuration do not impact host SNMP**—Starting in Junos OS Release 20.2R1, JDM does not write any configuration to the host SNMP configuration file (**/etc/snmp/snmpd.conf**). Hence, JDM installation and subsequent configuration do not have any impact on the host SNMP. The SNMP configuration CLI command in JDM is used only to configure JDM's **snmpd.conf** file, which is present within the container.

  [See SNMP Trap Support: Configuring NMS Server (External Server Model).]

SEE ALSO

What's New | 87

Known Limitations | 116

Open Issues | 119

Resolved Issues | 128

Documentation Updates | 145

Migration, Upgrade, and Downgrade Instructions | 146

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for MX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- The **number-of-sub-ports** configuration is not supported on the MPC11E line-card interfaces. PR1442439
- On MX2010 and MX2020, the MPC11E takes longer to appear online after being moved from one GNF to another. PR1469729
- Dynamic SR-TE tunnels do not get automatically re-created at the new master Routing Engine after the Routing Engine switchover. PR1474397
- Packet Forwarding Engine lookup loop happens when firewall based re-direction under the **forwarding-options** is used to perform route lookup in a nondefault routing instance for destinations reachable over MPLS-over-UDP tunnels. PR1478000
- If the frequency of messages is too high for the number of sessions, the kernel might be overloaded and causes riftd to quit and generates the core file. PR1481169
- **SNMP_TRAP_LINK_DOWN** messages are observed in the log when configuring MTU interface. PR1486542
- The rpd core files might be generated in the absence of an explicit **route-distinguisher** configuration. PR1486922
- On MX Series with MPCs and MICs, IP-IP tunneled traffic goes out through the BGP path when IS-IS is contributing route. PR1487173
- Junos Traffic Vision gets the interface values (for example, state, counters, and in-unicast-pkts) from the Packet Forwarding Engine and sends them to the remote client (collectors). This value will be different

in the output of the **show interfaces** command after the **clear interfaces statistics all** command is run because this command does not clear up counters on the Packet Forwarding Engine. PR1488758

- NSR is not supported for BGP multipath, because of which all multipath routes are re-resolved after switchover. In scaled setups, this takes a good amount of time. After all routes are resolved, tunnel statistics can be fetched from the Packet Forwarding Engine. PR1489067
- On MX Series with MPCs and MICs, packets do not get fragmented based on the FTI MTU in data path. PR1489526
- When the number of next-hop selectors to be repaired is very high, then time to repair them during FRR would go up and could increase packet losses. This would be observed specially when there are many unilist next hops with different next-hop selectors and each has a member next hop with a logical interface over the same physical interface, which goes down. PR1490070
- Sequence numbers (initial-sync and regular streaming) are in incorrect order when multiple collectors are present. The initial-sync sequence number (2097152) might appear after the regular streaming sequence number. PR1490798
- BSID scaling limits for IPv6 policies are 16,000 per ECMP. It might vary additionally based on the underlying ECMP path and depending on topology details. PR1495330
- The restart daemon causes session to be uninstalled in the FPC, but session remains active in the Routing Engine. As a workaround, clear the stale sessions with the **clear services rpm rfc2544-benchmarking active-tests** command. PR1499285
- Active reflection sessions are not aborted when the **deleted interfaces** and **deleted services** configuration are committed. As a workaround, clear the stale sessions with the **clear services rpm rfc2544-benchmarking active-tests** command. PR1499628
- Packet reassembly is not supported for IPIP tunnel (ip- interface) on all MPC cards. Any fragmented packets received from IP tunnel source drop at the destination end point. PR1505209

Infrastructure

- On Juniper Routing Engines with the Hagiwara compact flash card installed, after upgrade to Junos OS Release 15.1 and later releases, the failure message **smartd[xxxx]: Device: /dev/ada1, failed to read SMART Attribute Data** might appear in the messages log. PR1333855

Interfaces and Chassis

- The Packet Forwarding Engine cannot identify CFM packets with more than two tags as CFM PDUs because the current design treats the packet as an unknown EtherType. PR1487351

MPLS

- If any sub-LSP in up state for the current instance cannot be brought up for the new instance, then switching to the new instance will be prevented. PR1486813
- The **p2mp-lsp no-re-merge** statement enables an ingress router to check and avoid the path computed among sub-LSPs to form a remerge condition. This is a post computation check, so there is no consideration to avoid remerging during CSPF computation. Currently, this issue is considered as a minor issue. PR1487007
- On an MX240 platform, branches do not select the common ASBR from the available list with the **single-asb** statement enabled after common ASBR failure. PR1490637

Platform and Infrastructure

- Unknown unicast filter applied in an EVPN routing instance blocks the unexpected traffic. PR1472511

SEE ALSO

What's New | **87**

What's Changed | **113**

Open Issues | **119**

Resolved Issues | **128**

Documentation Updates | **145**

Migration, Upgrade, and Downgrade Instructions | **146**

Open Issues

IN THIS SECTION

Learn about open issues in this release for MX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Class of Service (CoS)

- CoS forwarding class default fields are missing in the schema file, resulting in the error log **LIBCOS_COS_ATTRIBUTE_RETRIEVE_FAILED**. PR1470252
- CoSEXP classifier and rewrite with protocol option **mpls-inet-both-non-vpn** is not working as expected. PR1479575
- If you have to take an interface out of aggregated Ethernet bundle and configure it to operate in a stand alone mode, then doing this in single commit might render the operation ineffective and could lead to connectivity issues. This is seen due to a race condition between Routing Engine daemons (COSD, DCD/Chassisd), Packet Forwarding Engine, and kernel. This issue is found when there is explicit CoS configuration is made on the interface. However, the problem can be seen without explicit CoS too as there is default CoS always present. In some cases, it is possible that a single shot commit sends out multiple operational messages down to kernel and might confuse the kernel to do unintended optimization that could lead to a message being consumed at kernel and not being sent to Packet Forwarding Engine. The result is the same even in this case. PR1504287

EVPN

- VXLAN OAM host-bound packets are not throttled with DDoS policers. PR1435228
- When IRB stanza is added with disable, IRB MACs are advertised to RPD. As a workaround, enable and disable the IRB interface if the router goes into this state. PR1510954
- On MX and all EVO platforms, when multicast snooping is enabled in EVPN for VLAN-based and vlan-bundle service scenario, the host under PE might not get the gateway MAC due to arp is broken. The service will be impacted. PR1515927
- In a VXLAN static VTEP tunnels scenario (including static VXLAN without EVPN), after Routing Engine switchover or restart of Layer 2 learning, if you create a new VTEP interface, the interface might not work. PR1520078

Forwarding and Sampling

- When an IPv4 prefix is added to a prefix list referenced by an IPv6 firewall filter, then the log message **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized** is not seen. PR1395923
- Syslog error **rpd[2191]: krt_flow_dfwd_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out** is seen. Restart the firewall process in both Routing Engines when firewall error logs are noticed along with **SSD hardware failure** logs. PR1397171
- Verify event CP down or up is long enough to trigger an EP timeout for CoS hierarchy model 2, failing as expected DHCP subscribers are not bound. PR1505409

General Routing

- On a vMX platform, the performance of an X710 NIC is lower compared to the performance of an 82599 NIC. A 10-Gbps line rate can be achieved at a 512-byte packet size for the X710 NIC compared to 256 bytes for the 82599 NIC. PR1281366
- The **chain-composite** statement does not bring in a lot of gain because TCNH is based on an ingress rewrite premise. PR1318984
- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections and hence the reactions to failure situations might not be handled in a graceful way. For example, TCP connection timeout because of jlock hog crossing the boundary value (5 seconds) can cause bad consequences in MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling the marker infrastructure in a MX Series Virtual Chassis setup. PR1332765
- In an MS-MPC or MS-MIC in ALG scenario, the **MAC_STUCK** message might be observed and traffic might be dropped. PR1335956

- Packet Forwarding Engine error message **localttp_offload_tx_errcheck: failed to send packet 11 times in last one second** might be seen on a node slicing deployment. There is no known impact. PR1359149
- craftd messages are generated on MX10003 and MX204 platforms. They do not have the craft interface. Hence, these errors are expected, and can safely be ignored. When the craftd daemon tries to open the device, it fails with a junk character in the fatal error message because the error number is not mapped to a string in the kernel code. The error messages are **MX craftd[xxxx]: craftd detected platform mx10002**, **MX craftd[xxxx]: LIBJSNMP_SA_IPC_REG_ROWS: ns_subagent_register_mibs: registering 1 rows**, and **MX craftd[xxxx]: fatal error, failed to open smb device: ‚JÎÈ""**. PR1359929
- On MX2010 and MX2020 routers equipped with SFB2, some error messages are seen in the logs. There is no operational impact. PR1363587
- A few xe- interfaces go down with the error message **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error**. PR1377840
- The virtio throughput remains the same for multiqueue and single-queue deployments. PR1389338
- CPU overuse on FPC might be observed if the load-balance adaptive feature is enabled for an aggregate Ethernet interface. PR1399369
- FPC core files are generated on multiple additions or deletions of hierarchical CoS from pseudowire devices. As a workaround, remove the pseudowire devices without changing the hierarchical CoS configuration. PR1414969
- The **show services hybrid-access sessions**, **show services hybrid-access statistics**, and **show services hybrid-access tunnels** commands display values of zero for hybrid access gateway traffic statistics even when traffic is active in the gateway sessions and the tunnels. PR1419529
- Dynamic tunnel summary displays an incorrect count of up and total tunnels after multiple iterations of activating and deactivating the dynamic tunnel configuration. It is just a display issue, and there is no problem with the functionality. PR1429949
- On MPC10E 3D MRATE-15xQSFPP, L2 over GRE is not supported. Although the configuration gets committed, the feature does not work. PR1435855
- The FPC might crash when Packet Forwarding Engine memory usage for a partition such as NH/DFW is high. Under low Packet Forwarding Engine memory condition, the **Safety Pool below 25% Contig Free Space** or **Safety Pool below 50% Contig Free Space** log might be observed. PR1439012
- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card. Because of vendor limitations, achieving subsecond hold timers is not possible. PR1444516
- On the MX Series platforms with NG-RE installed, the process vehostd might crash without generating core files and automatic restart of vehostd might fail. The crash of vehostd impacts the management of Junos VMs. PR1448413
- Physical interface policer is not supported on the MPC11 line card. PR1452963
- FPC crash on MX240 and MX2020 routers or Packet Forwarding Engine crash on MX104 routers might happen when the MIC-3D-8OC3-2OC12-ATM is installed and ATM interface is configured. PR1453893

- On the MPC11E line card, FIB download rates are lower than MPC10E by 30 percent. PR1456816
- With the scaled filter-based forwarding (FBF) configuration, two instances seem unable to forward the traffic to the respective routing instances. It appears that the FBF programming is incorrect for the two FBF instances. PR1459340
- With multiple different fixed-sized traffic streams configured at 1,000,000 fps (40 Gbps combined rate) on aggregated Ethernet0 along with another independent aggregated Ethernet interface (aggregated Ethernet1, 50 percent line rate 4 streams bidirectional => 118 Gbps combined traffic rate), both hosted on a single Packet Forwarding Engine instruction of an MPC11E line card, small varying packet drops occurs for every iteration on aggregated Ethernet1 on disabling aggregated Ethernet0. The drops might vary from 200 to certain 1000 frames. PR1464549
- For the MPC10E card line, the IS-IS and micro-BFD sessions do not come up during baseline. PR1474146
- Changing framing modes on a CHE1T1 MIC between E1 and T1 on an MPC3E NG HQoS line card causes the MIC to go offline. PR1474449
- If the LDP is scaled to 480,000 over four interface, the mirror subsystem times out and goes down. PR1474965
- Error messages **[Error] L2alm : l2alm_mac_process_hal_delete_msg:667 Ignoring MAC delete with ifl index 355, fwd_entry has 7888** are seen after performing configuration removal/restore with IP/MPLS configurations in the MX480 box. PR1475785
- mgd core files are generated during startup. 64-bit cMGD must be used if cMGD is running on a 64-bit Junos OS. PR1481335
- Invalid packets are dropped by DUT with TCC encapsulation configuration as intended but the statistics counters get incremental. PR1481698
- Error message **fpc0 user.crit aftd-trio: [Critical] Em: Possible out of order deletion of AftNode #012#012#012 AftNode details - AftIndirect token:230791 group:0 nodeMask:0xffffffffffffffff indirect:333988 hwInstall:1#012** is seen while executing the IP/MPLS script in baseline. PR1486158
- High-scale login and logout (around 1M bearers) can prevent some sessions from logging in again. PR1489665
- Support for upgrading PSMs firmware on the MX2000 platforms. PR1489939
- On MX Series with MS-MPC and MS-MIC, if there are self-generated packets such as TCP-tickle and UDP-logging, there might be data congestion on the data path due to no throttling functionality for such types of packets. When the data path is blocked, prolonged flow control might happen with the service interfaces being brought down and a PIC reboot. At the same time, the mspmand core file will be generated if **dump-on-flow-control** is enabled. PR1489942
- The **show system firmware** command shows the upgrade status on the PSMs when the firmwares are upgraded. PR1493045
- On an MX2020 router with SFB3/MPC11, the AER image is required for uncorrectable or correctable PCIe error. PR1493065

- On MX platforms with MPC10 and MPC11 line cards, if the shared bandwidth policer is referenced by an interface specific firewall filter, and the filter is bound to the interface of the affected line card, during the line card reboot, there is a chance that firewall filter information updates come out of order when being sent from Routing Engine to Packet Forwarding Engine. Due to handling the out of order update messages, the line card might crash and get into continuous reboot state. PR1493084
- The component sensor does not export data under CBO/1 in the expected time. PR1493579
- In DS-Lite scenario, some B4 devices might not be able to establish softwire with an AFTR device if more then one DS-Lite softwire concentrator address is configured. This issue happens when AFTR functionality is provided by MS-MPC line card. PR1496211
- When the NETCONF session is established over an outbound SSH connection, the high rate of pushing configuration to an ephemeral database might result in flapping of the outbound SSH connection or a memory leak issue. PR1497575
- After backup Routing Engine halt, CB1 goes offline and comes back online; this leads to backup Routing Engine boot up and shows reboot reason as **0x1:power cycle/failure**. There is no functional impact of this issue. PR1497592
- In a Layer 3 VPN scenario, there are 2 EBGP (1 multihop and 1 single-hop) and 1 IBGP between the CE device and the PE device, then EIBGP might not work correctly with Layer 3 VPN **composite-next-hop** statement is enabled. PR1501935
- On deactivating and activating a routing instance, packets from nonexisting sources on GRE and UDP designated tunnel are accepted where they are supposed to be dropped. PR1503421
- On an MX2020 device, the MPC11 line card is not supported. PR1503605
- After access interface flaps, VPLS, Layer 2 VPN, and Layer 2 circuit goes down. PR1505307
- On all Junos OS platforms with the Juniper Telemetry Interface configured, the rpd might crash when there is telemetry streaming is in progress and meanwhile there is a network churn. This is a timing issue, and the rpd recovers automatically. PR1505425
- This issue is related to CPCD service on spc3 card. Http-get request is not redirected even after service is attached to the subscriber. PR1505438
- In an EVPN scenario with VRRPv6 is used, the Ethernet source MAC address might be used for IPv6 mac-ip binding when the NA is sent from VRRPv6 master. AS this unexpected behavior is triggered on regular intervals, it causes the entries to keep refreshing in the EVPN database because NS from VRRPv6 master changes the mac-ip binding. This impacts the traffic. PR1505976
- In a scenario where QSFP is used as a single interface or child link of the aggregated Ethernet interface, if the interface is disabled and enabled frequently, the write errors might happen on inter-integrated circuit of QSPF. Then the laser of QSFP might not be enabled. PR1510994
- On MPC11 line card, DFW crash is seen after you perform remove and restore configurations on backup Routing Engine. PR1512770

- Tunable parameters wavelength set through the CLI configuration is not set on SFP+-10G-T-DWDM-ZR optics when the optics is placed on MPC7E 3D 40XGE line card. PR1513321
- On some boot, external 1PPS cTE might see up to 22ns, while the 2way TE stays within 20ns. PR1514066
- Traffic drop observes when multicast traffic on a group with 4000 egress aggregated Ethernet ports is sent. The drop is always on the egress port that are on same the Packet Forwarding Engine as ingress. PPE times out before the multicast packet is processed and causes the packet to drop. PR1514646

High Availability (HA) and Resiliency

- When an ZPL is done while traffic is running, some BGP sessions might flap and leads to traffic loss. The drop is transient and traffic recovers after the ZPL. PR1487144

Interfaces and Chassis

- The same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating the incorrect configuration. PR1221993
- With CFM configuration, if you perform an upgrade between releases that use different db formats, continuous cfmd crashes might be seen after the upgrade. This is because of the old version of **/var/db/cfm.db**. PR1281073
- After Routing Engine switchover without nonstop routing (NSR) on the broadband network gateway (BNG), some VRF routing instances might experience silent dropping of traffic destined to the hosts behind a static PPPoE subscriber's CPE device. The affected routing instances are configured without the **vrf-table-label** statement and should have a static route configured with the pp0.xxx interface as a next hop. PR1488302
- All logical interfaces are not programmed when any logical interface with VLAN bridge encapsulation and not under a bridge domain or EVPN/virtual-switch routing instance. PR1501414
- Input and output bytes count mismatch in IPv6 traffic statistics while issuing the **show interface extensive** command. PR1505100
- Loss of PPPoE subscribers during unified ISSU, out of 28,000 PPPoE subscribers only 375 where logged in again after unified ISSU. PR1514152
- Commit error is observed when you delete all the units under ps0 interface and keeps **flexible-vlan-tagging** configuration. PR1514319

Layer 2 Ethernet Services

- The DHCP DECLINE packets are not forwarded to the DHCP server when **forward-only** is set within **dhcp-reply**. PR1429456
- Subscriber recovery in relay using LQ/BLQ fails. PR1504266
- DHCPV6_LEASEQUERY counter might not be as expected in the output of the **show dhcpv6 server statistics** command. PR1506418

MPLS

- Aggressive switchovers due to MBB/CSPF computations causes traffic loss on all branches in the tree even if a single branch fails to come up due to remerge detection on the transit router. PR1487916
- GRES or NSR followed by restart routing on the master device does not honour remerge behavior. PR1489168
- In an MPLS-TE scenario with high scale LSPs (for example, 20,000), CSPF job might get stalled for new/existing LSP if some configuration changes (which impacts the rpd process) are done when constrained shortest path first (CSPF) job is suspended and pending. Traffic Engineering Database CSPF job goes in a state where it is not able to recover till the time rpd process is restarted. PR1502993
- When the **container-label-switched-path** statement is configured with LDP tunneling, LDP targeted adjacency might go down and stay down after a configuration not related to **container-label-switched-path** is modified. PR1509578
- The rpd might continuously crash after upgrading pre Junos OS Release 18.1 to Junos OS Release 18.1 and later while graceful-restart and RSVP/static LSP are configured. This is because there is a change in the data structure written to the restart database file from Junos OS Release 18.1 and later. So, when rpd comes up and tries to read the restart database file written by pre Junos OS Release 18.1 image, the rpd might crash. PR1517018

Network Management and Monitoring

- On MX Series platforms, SNMPv3 informs do not work after restart due to a problem in open source library. The issue can be detected using the **monitor traffic interface** command from an interface that can reach the SNMP collector. PR1497841

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. PR1054798
- For the bridge domains configured under an EVPN instance, ARP suppression is enabled by default. This enables the EVPN to proxy the ARP, and reduces the flooding of ARP in the EVPN networks. Because of that, storm control does not take effect on ARP packets on the ports under such bridge domains. PR1438326
- The CFM remote MEP is not coming up after configuration or remains in Start state. PR1460555
- When traffic is received from 1000 different VRF instances on a PE device from a CE device, then a few flows (4 to 5) are dropped at the PE device. As a workaround, disable the particular VRF instance or particular instance's PIM protocol and enable it back. PR1460471
- NPC core file is generated at **trinity_rt_iff_attach,pfe_ifl_family_attach,ifrt_ifl_family_adder,ifrt_ifl_family_add_vector,ifrt_command_handler**. PR1461892
- In NTP with the boot-server scenario, when the router or switch boots, the NTP daemon will send an ntpdate request to poll the configured NTP boot-server to determine the local date and time. If the ntpdate is not activated correctly while the device is booting, the ntpdate might not work successfully. Then, some cosmetic error messages of time synchronization might be seen, but there is no impact on the time update because the ntp daemon updates the time eventually. PR1463622
- A few OAM sessions are not established with scaled EVPN ETREE and CFM configurations. PR1478875
- If the interface is newly added as a CE interface, the existing broadcast, unicast, and multicast (BUM) traffic can be looped. The loop prevention feature is designed to start working whenever a new CE interface is added by configuration. But the existing BUM traffic can be distributed to a new CE interface earlier, before enabling the loop prevention feature. PR1493650
- Traffic loss is observed after unified ISSU, when you enable or disable and activate or deactivate the interface. PR1493723
- On MX series platforms, when a route's next hop is an IRB interface with lt- as the underlying Layer 2 interface, it does not get programmed on the Packet Forwarding Engine, resulting in packet drop. PR1494594

- On MX platforms with any MPC in enhanced network service mode, if VRRP is configured on aggregated Ethernet interface, after **set chassis fpc X pic X traffic-manager mode ingress-and-egress** is enabled, traffic sent to virtual IP/MAC might be dropped and the forward traffic might be affected. PR1501014
- On MPC7, MPC8, and MPC9, an interface configured with **vlan-tags outer** as 88a8 sends out 8100. In this scenario, the IP traffic arrives at the PE router gets destination route from IRB interface inside a VPLS instance, and then forwards to a CE device. PR1502867
- Kernel crash causes the router to reboot when making VRRP related changes. PR1511833
- Due to rare timing issue, the FPC might crash because of route table object fetch failure in an EVPN multihoming scenario. PR1513509
- The **show jnh qmon queues-sensor 0** command output has no content. The issue is seen only when the interface and sensor configurations are committed in a single commit, and that too only during the first time when the line card is up. This is a corner case scenario caused by the timing issue between the queue configuration and sensor configuration which will not be seen in the field. You can avoid this corner case by committing the interface and sensor configurations in a two step commit. PR1514881

Routing Protocols

- Even when the **protocols mpls traffic-engineering bgp-igp** statement is configured, the UDP tunnel routes are not added to inet.0. The UDP tunnel routes are added only to inet.3 table whether the statement is configured or not. PR1457426
- On an MX Series router that is running as a Layer 3 VXLAN gateway, if the **igmp-snooping** statement is enabled in partial but not for all bridge domains, multicast traffic loss could be observed in a non-igmp snooping bridge domains. PR1481987
- Upstream state is found with joint to source and no prune to RP while checking source and receiver in same VPNs attached to same PE customer-risk. PR1508401
- When Multicast Source Discovery Protocol (MSDP) is configured, if there is a huge number of source-active (SA) messages present in the network (for example, around 20,000 or more), the rpd process might crash. PR1517910

VPNs

- In an MVPN environment with the SPT-only option, if the source or receiver is connected directly to the c-rp PE device and the MVPN data packets arrive at the c-rp PE device before its transition to SPT, the MVPN data packets might be dropped. PR1223434
- The **c-multicast-source-pe** and **c-multicast-pim-source-pe** XML tags are not under **c-multicast-address** hierarchy in the XML output. So any scripts expecting or associating these values to the **c-multicast** does not map the information to the corresponding given **c-multicast-address** entry. There is no functional or other impact. In normal show output (that is without xml), it displays correctly. PR1509948

- HRS with min-rate works fine with selective ingress replication provider-tunnel. In case of RSVP-TE point-to-multipoint as provider-tunnel, MBB might happen during network triggers and it might leads to more than 50ms traffic loss. Here in this particular scenario where RSVP-TE P2MP I-PMSI + S-PMSI provider tunnel is used when we do access-link failure, MBB is happening and leading to more than 3s traffic loss. Even with S-PMSI provider-tunnel also we are seeing the same issue. PR1520568

SEE ALSO

What's New | 87

What's Changed | 113

Known Limitations | 116

Resolved Issues | 128

Documentation Updates | 145

Migration, Upgrade, and Downgrade Instructions | 146

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in Junos OS Release 20.2R1 for MX Series routers.

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

Application Layer Gateways (ALGs)

- SIP messages that need to be fragmented might be dropped by the SIP ALG. PR1475031
- FTPS traffic might be dropped on MX Series platforms if FTP ALG is used. PR1483834

Class of Service (CoS)

- The MX Series generated OAM/CFM LTR messages are sent with a different priority than the incoming OAM/CFM LTM messages. PR1466473
- The MX10008 and MX100016 routers might generate cosd core files after executing the **commit/commit check** command if the **policy-map** configuration is set. PR1475508
- Error message **GENCFG write failed (op, minor_type) = (delete, Scheduler map definition) for tbl id 2 ifl 0 TABLE Reason: No such file or directory** is observed. PR1476531
- MX Series platforms with MPC1-Q and MPC2-Q line cards might report memory errors. PR1500250

EVPN

- Remote MAC address present in EVPN database might be unreachable. PR1477140
- Deleting a Layer 2 logical interface generates an error if the interface is not deleted first from EVPN. PR1482774
- The ESI of IRB interface does not update after autonomous-system number change if the interface is down. PR1482790

- Dead next hops might flood in a rare scenario after remote PE devices are bounced PR1484296
- The ARP entry gets deleted from the kernel after adding and deleting the virtual-gateway-address. PR1485377
- The rpd core file might be generated when doing Routing Engine switchover after disabling BGP protocol globally. PR1490953
- VXLAN bridge domain might lose VTEP logical interface after restarting chassisd. PR1495098
- The VXLAN function might be broken due to a timing issue. PR1502357
- The MAC address of the LT interface might not be installed in the EVPN database. PR1503657

Forwarding and Sampling

- IP-IP de-encapsulation fails if de-encapsulation filter is applied on loopback interface. PR1469219
- Traffic might be forwarded into the default queue instead of the correct queue when the VPLS traffic has three or more VLAN tags with VLAN priority 5. PR1473093
- The filter might not be installed if the **policy-map xx** is present under the filter. PR1478964

General Routing

- Syslog error message **PFEIFD: Could not decode media address with length 0** might be generated by the Packet Forwarding Engine. PR1341610
- The nondefault routing instance is not supported correctly for NTP packets in a subscriber scenario. PR1363034
- Egress monitored traffic is not mirrored to destination for analyzers on MX Series routers. PR1411871
- **FPC x Voltage Tolerance Exceeded** alarm raised and cleared upon bootup of JNP10K-LC2101. PR1415671
- The pccd starts running from the system start. PR1417052
- Resetting the Playback Engine logs are seen on the MPC5E line cards. PR1420335
- PF core voltage is not set according to the required e-fuse value and remains as default value of 0.9V on the JNP10008-SF and JNP10016-SF Switch Interface Boards (SIBs). PR1420864
- FPC might crash after GRES when you commit the changes in firewall filter with the **next term** statement in the subscriber scenario. PR1421541
- PTP might not work on the MX104 platform if phy-timestamping is enabled. PR1421811
- When you run the **show route label X | display json** command, two **nh** keys are present in the output. PR1424930
- PTP and show warning are disabled when hyper mode is configured. PR1429527

- Interfaces on the MPC-3D-16XGE-SFPP might go down due to CB0 clock failure. PR1433948
- ZF interrupts for out-of-range destination Packet Forwarding Engine INTR for Gnt are observed when the MPC6 or MPC9 line card is brought up. PR1436148
- System reboot is required when GRES is enabled or disabled with the **mobile-edge** configuration. PR1444406
- On the MPC10E-15C-MRATE with 25-Gigabit Ethernet ports, FEC statistics are not getting reset after changing FEC mode. PR1449088
- RE-MX2008-X8-128G secure BIOS version mismatch alarms. PR1450424
- Need to add support for drop flows when the packet drops. PR1451921
- When MVLAN interface (OIF map) is changed, the existing multicast subscribers with membership reports in place experience loss of multicast traffic until traffic is forwarded to a new OIF map. PR1452644
- Interfaces shutdown by 'disable-pfe' action might not be up using MIC offline or online command. PR1453433
- When scale configurations are applied from approximately 10 minutes, chassisd CLI will either have a delay in response or will time out. PR1454638
- On 4-port 1-Gigabit Ethernet using QSFP28 optics, continuous logging in chassisd process occurs when speed 1-Gigabit Ethernet is configured with **pic_get_nports_inst** and **ch_fru_db_key**. PR1456253
- On the MPC11E line card, need to add the support of optics-options low light. PR1456894
- LSP statistics are not getting reset after restart routing. PR1458107
- Inline S-BFD packets are dropped on MPC6E MIC1/PIC1 ports: 0-11. PR1459529
- Occasional warning message such as **TCP Connect error** can be seen during FPC reboot. PR1460153
- Multiple leaf devices and prefixes are missing when LLDP neighbor is added after streaming is started at the global level. PR1460347
- Support of del_path for the LLDP neighbor change at various levels. PR1460621
- When you receive IPv6 over IPv4 IBGP session, the IPv6 prefix is hidden. PR1460786
- Explicit deletion notification (del_path) is not received when LLDP neighbor is lost as a result of disabling local interface on the DUT through CLI (gNMI). PR1461236
- On the MPC10E line cards, more output packets than expected are seen when ping function is performed. PR1461593
- The **show dynamic-tunnel database** CLI command output does not filter IP-IP tunnels based on destination. PR1461659
- The **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed** message appears when both DIP switches and power switch are turned off. PR1462065
- Inline BFD session might flap on renegotiation of timers from slow to aggressive interval. PR1462775

- The MVPN traffic might be dropped after performing switchover. PR1463302
- The **native-vlan-id** functionality does not work and untagged traffic does not pass with the **native-vlan-id** configuration. PR1463544
- The jdhcpd process might consume high CPU use, and no further subscribers can be brought up if there are more than 4000 dhcp-relay clients in the MAC-MOVE scenario. PR1465277
- On the MPC10E and MPC11E line cards, the bandwidth-percent with shaping-rate might not work as expected on aggregated Ethernet interfaces after shaping-rate change. PR1465766
- The bbe-smgd process generates core files on the backup routing engine. PR1466118
- ICMP error messages are still unreceived after enabling the **enable-asymmetric-traffic-processing** configuration statement. PR1466135
- A few DHCP INFORM packets specific to a particular VLAN might be taking the incorrect resolve queue. PR1467182
- On the MPC11E line card, the DOM MIB alarm for the channelized 10-Gigabit Ethernet interface is not showing any alarm for LF/RF. PR1467446
- Daemons might not be started if **commit** is executed after **commit check**. PR1468119
- PPP IPv6 NCP fails to negotiate during the PPP login. PR1468414
- The rpd process might crash if BGP sharding is enabled. PR1468676
- The tcp-log connections fail to reconnect and get stuck in "Reconnect-In-Progress" state. PR1469575
- Unable to set up 26M sessions (NAPT44) at 900,000 pps. PR1470833
- In rare occasions the router might send out one extra URR quota value for a bearer. PR1470890
- Syslog message **fpcX user.notice logrotate: ALERT exited abnormally with [1]** pops at 04:02:01. PR1471006
- DHCP relay with forward-only might fail to send OFFER messages when DHCP client is terminated on logical tunnel interface. PR1471161
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card might cause major alarms on other FPCs in the system. PR1471372
- The clksyncd crash might be seen when PTP over aggregated Ethernet is configured on the MX104 platform. PR1471466
- On the MPC11E line card, locating a specific 100-Gigabit Ethernet, 40-Gigabit Ethernet, and 10-Gigabit Ethernet port in the card by blinking the corresponding port LED does not work. PR1471894
- Chassis alarm on BSYS might be observed: **RE0 to one or many FPCs is via em1: Backup RE**. PR1472313
- Performing back-to-back rpd restarts might cause rpd to crash. PR1472643
- Manually configured ERO on NS controller might lost when PCEP session bounced. PR1472825

- SDB goes down very frequently if the **reauthenticate lease-renewal** statement is enabled for DHCP. PR1473063
- Some routes might not be installed into the FPC after it gets restarted. PR1473079
- On the MPC11E line card, **show dynamic-tunnels database** command does not show traffic statistics. PR1473096
- On MPC11, oversubscription drops are not accounted in Routing Engine CLI under resource drops when Flow control is disabled. PR1473191
- Dynamic-profile for VPLS-PW pseudowire incorrectly reports Dynamic Static Subscriber Base Feature license alarm. PR1473412
- On the MPC11E line card, after doing Routing Engine switchover on BSYS , the AF interface on peer router shows status as down with the reason being that the Packet Forward Engine is down on the GNF. PR1473555
- When both MSTP and ERP are enabled on the same interface, then ERP does not come up properly. PR1473610
- Drops counter does not increment for the aggregated Ethernet even after the member link shows the drops. PR1473665
- Ingress multicast replication does not work with GRES configuration. PR1474094
- DHCP-server RADIUS-given mask is being reversed. PR1474097
- On MX150 platform, core files are not seen under **show system core-dumps**. PR1474118
- A newly added LAG member interface might forward traffic even though its micro BFD session is down. PR1474300
- Upon external X86 node slicing server reboot, the host SNMP configuration gets overwritten by the JDM SNMP configuration settings. PR1474349
- When traffic loss is observed on a 100-Gigabit Ethernet logical interface, the MACsec sessions are up and live. PR1474714
- On the MPC11E line card, basic circuit cross-connect traffic flow does not occur with the logical systems. PR1474983
- The clksyncd process generates core file after the GRES. PR1474987
- Memory leak leads to restart of the MPC10E line card. PR1475036
- Stateful firewall rule configuration deletion might lead to memory leak. PR1475220
- The full list should be returned. A leaf should be considered atomic, regardless of whether it is a single value or a list for on-change event. PR1475293
- The RADIUS accounting updates of the service session have incorrect statistic data. PR1475729
- When xSTP protocols are enabled on interface all, it might run on **vlan-tagging**/**flexible-vlan-tagging** Layer 3 interfaces and lead to blocking of SXE interface. PR1475854

- Traffic loss might be seen as backup Routing Engine takes around 20 seconds to acquire mastership. PR1475871
- Traffic drop might be observed while performing a unified ISSU on the MX2020, MX2010, and MX960 platforms. PR1476505
- The bbe-mibd might crash on MX Series platform in subscriber environment. PR1476596
- On the MPC10 or MPC11 line cards, Routing Engine might not be able to send packets with traffic-manager enhanced-priority-mode configuration enabled. PR1476683
- The host-generated packets which might get dropped at the other end. PR1476764
- Traffic loss might occur to the LNS subscribers in case the **routing-service** statement is enabled under the dynamic profile. PR1476786
- Traffic loss might be seen in SAEGW scenario after the daemon restarts or after the GRES operation. PR1477461
- In NAT-T scenario, IKE version 2 IPsec tunnel flaps if the tunnel initiator is not behind NAT. PR1477483
- The rpd process might crash when the JET RIB API is used to set the "bandwidth" attribute. PR1477745
- On the MX2010 platform, syslog message **spmb0 cmty_sfb_temp_check: sfb[0] is powered OFF" & "spmb0 cmty_sfb_voltage_check_one: sfb[0] is powered OFF** is flooding even though SFBs are online. PR1477924
- Error log message **chassisd[7836]: %DAEMON-3-CHASSISD_IOCTL_FAILURE: acb_get_fpga_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device)** is observed after every commit. PR1477941
- Packet Forwarding Engine might be disabled because of the major error on MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9. PR1478028
- The **show evpn statistics instance** command gets stuck on multihomed scenario. PR1478157
- At-scale logins of both default and dedicated bearers might require retries from the control plane. PR1478191
- The ukern-platformd process might crash on MX2000 platforms with MPC11 line card. PR1478243
- Output chain filter counters are not proper. PR1478358
- MX Series-based MPC line card might crash when there is bulk route update failure in a corner case. PR1478392
- The FPC with **vpn-localization vpn-core-facing-only** configuration might be stuck in ready state. PR1478523
- On MX240, MX480, MX960, MX2000, MX10003, MX10008, and MX10016 with the MPC7E, MPC8E, and MPC9E line cards, hardware sensor information is logged every 30 minutes. PR1478816
- The protocol MTU might not be changed on lt- interface from the default value. PR1478822

- The TCP-log sessions might be in Established state but no logs get sent out to the syslog server. PR1478972
- Mobile-edge sessions might be lost if GRES is being performed while sessions are logged in with URR enabled. PR1478985
- The SCBE3 fabric plane gets into check state in MX Series Virtual Chassis. PR1479363
- Interface states are not showing correctly between main and shards on one of the interfaces in cRPD. PR1479801
- After kmd restarts, IPsec SA comes up but the traffic fails for some time in certain scenarios. PR1480692
- 100-Gigabit interface might randomly fail to come up after maintenance operations. PR1481054
- Issue with binding non-default routing instance to existing soft-gre group. PR1481278
- After unified ISSU on the master and the backup Routing Engine, **ISSU enhanced-mode: Performing action get-state for error /fpc/5/pfe/0/cm/0/PCIe_Error/0/PCIE_CMERROR_UNCORRECTABLE (0x190001)** error message is generated. PR1481859
- The rpd might crash when you execute the **show route protocol l2-learned-host-routing** or **show route protocol rift** CLI command on a router. PR1481953
- Log in to some PPPoE subscribers through aggregate Ethernet interface might cause the device to reboot. PR1482431
- Fragmentation limit and reassembly timeout configuration under services option is missing for SPC3. PR1482968
- When checking the BFD fuctionality over L2VPN client, BFD session is not coming up. PR1483014
- Link errors might be seen after restarting the FPC or fabric plane. PR1483124
- Traffic impact might be seen when the **policy-multipath** is configured without LDP on SPRING-TE scenario. PR1483585
- The downstream IPv4 packet greater than BR MTU are getting dropped in MAP-E. PR1483984
- Traffic rate is not as expected on aggregated Ethernet interface when child links are from MPC11 and MPC9 line card after applying a policer. PR1484193
- ARP entry might not be created in the EVPN-MPLS environment. PR1484721
- The logical tunnel interface might not work on MPC10 line card. PR1484751
- Fix and enhancement for **request rift package activate** for the junos-rift package. PR1485098
- Attribute sending zero value should be compressed because it uses too much bandwidth in periodic streaming. PR1485257
- Interface input error counters are not increasing on the MX150 platforms. PR1485706
- The **krt-nexthop-ack-timeout** command might not automatically be picked up on restarting the rpd process. PR1485800

- MPC10E line card installed in the FPC slot 4 might drop host outbound traffic. PR1485942
- Command completion help text for LLDP-MED coordinate configuration statement contains spelling errors. PR1486327
- The aftd process might crash when MPC10 line card is installed. PR1487416
- Incorrect frame length of 132 bytes might be captured in packet header. PR1487876
- cMGD/cRPD: XML is not properly formated. PR1488036
- Add support for PSM firmware upgrade on the MX2000 platform. PR1488575
- During multiple login and logout of 250,000 sessions, there can be daemon restart due to mishandling of data. PR1489512
- NAT rule-sets processing order is not getting processed based on the order configured under **service-set**. It is getting processed based on the NAT rules defined under **[services nat source]** hierarchy level configuration. PR1489581
- With 4 member AMS used in the service-set, commit check fails when /30 subnet address is used as NAT pool IP. PR1489885
- Error syslog message **Failed to connect to the agentx master agent (/var/agentx/master): Unknown host (/var/agentx/master) (No such file or directory)** is continuously being generated with dns-sinkholing. PR1490487
- When NAT/SFW rule is configured with application-set with multiple applications having different TCP inactivity-timeout, sessions are not getting TCP inactivity-timeout as per the configured application order. PR1491036
- The DAC cable is not detected after reboot or plug out or plug in. PR1491116
- The unified ISSU is not supported on next-generation MPC cards. PR1491337
- Multiple deactivating and activating of security traceoptions along with clear single NAPT44 session could result in generation of flowd core file. PR1491540
- MS-MIC is down after loading some releases in MX Virtual Chassis scenario. PR1491628
- FPCs might stay down or restart when you swap the MPC7, MPC8, and MPC9 line cards with the MPC10 and MPC11 line cards or vice versa in the same slot. PR1491968
- User-configured MTU might be ignored after the unified ISSU upgrade uses **request vmhost software in-service-upgrade**. PR1491970
- Behavior change in clients with multiple gRPC channels to same target. PR1492088
- The delay of LT interfaces coming up is seen on MPC11E line card after you configure scaled PS interfaces anchoring to RLT. PR1492330
- On MX10008 platform, SNMP table entPhysicalTable does not match the PICs shown for the **show chassis hardware** command. PR1492996
- DHCP subscribers do not come up as expected after deactivating Virtual Chassis port. PR1493699

- The **ptp-clock-global-freq-tracable** leaf value becomes false and does not change to true when the internal lock is in the **Acquiring** state. PR1493743
- The LSP might not come up in LSP externally-provisioned scenario. PR1494210
- Error message **PFE_ERROR_FAIL_OPERATION: Unable to unbind cos scheduler from physical interface 147** is observed on the MPC9E line card after restarting the MPC11E line card. PR1494452
- Missing firmware image file in **usr/share/pfe/firmware**. PR1494557
- In node slicing setup after GRES, RADIUS interim updates might not carry actual statistics. PR1494637
- Group address is not programmed back after deactivating and activating the bridge domain. PR1495480
- Flood next-hop ID is not same in both master and backup Routing Engine. PR1495925
- Error message **PFEIFD: Could not decode media address with length 0** is generated by Packet Forwarding Engine when subscribers come up over a pseudowire interface. PR1496265
- Port numbers logged in ALG syslog are incorrect. PR1497713
- Subscribers might be disconnected after one of the aggregated Ethernet participating FPCs comes online in a Junos OS node slicing scenario. PR1498024
- SNMP polling does not show correct PSM jnxOperatingState when one of the PSM inputs failed. PR1498538
- The rpd might crash when multiple VRFs with 'IFLs link-protection' are deleted at a single time. PR1498992
- The commit check might fail when adding IFL into a routing instance with the **no-normalization** statement enabled under **[routing-instances]** hierarchy. PR1499265
- The heap memory leak might be seen on the MPC10 and MPC11 line cards. PR1499631
- The SPC3 card might crash if SIP ALG is enabled. PR1500355
- On MX2020 and MX2010, the **pem_tiny_power_remaining** message will be continuously logged in chassisd log. PR1501108
- Application ID does not display under NAT/SFW rule configured with application 'any' rule. PR1501109
- Support license start and end date in MIBs. PR1503790
- The **show bridge statistics** command does not display the statistics information for pseudowire subscriber interfaces. PR1504409
- The l2cpd crash might be seen if you add or delete ERP configuration and then restart l2cpd. PR1505710
- GnmiJuniperTelemetryHeader incompatibility introduced in Junos OS Release 19.3. PR1507999
- The host generated packets might get dropped if the **force-control-packets-on-transit-path** statement is configured. PR1509790
- The multicast traffic might be dropped if ALB is enabled on the aggregated Ethernet interface. PR1512157

High Availability (HA) and Resiliency

- Unified ISSU might fail on MX204 and MX10003 Virtual Chassis with an error message. PR1480561

Infrastructure

- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. PR1462986
- Add F-label veto code checks for per-pfe f-label pools. PR1466071

Interfaces and Chassis

- Syslog error **scchassisd[ ]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x** is observed after MX Virtual Chassis local or global switchover. PR1428254
- Decoupling of Layer 2 logical interfaces from bridge and EVPN configuration. PR1438172
- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. PR1459201
- On the MPC11E line card, the IPv6 local stats are counted against the IPv6 transit traffic statistics as well. PR1467236
- When you configure ESI on a physical interface, the traffic drops when you disable the logical interface under the physical interface. PR1467855
- Executing commit might hang because of stuck dcd process. PR1470622
- Traffic is not forwarded properly when traffic-control-profiles with logical interface queues are configured. PR1475350
- Commit error is not thrown when member link is added to multiple aggregation group with different interface specific options. PR1475634
- The interface on MIC3-100G-DWDM might go down after performing an interface flap. PR1475777
- When you delete and add a logical interface (both the logical interfaces with the same VLAN ID) in a single commit, the configuration check fails with the error **duplicate VLAN-ID**. PR1477060
- A stale IP address might be seen after a specific order of configuration changes in logical systems scenario. PR1477084
- Traffic is seen for 248 seconds when an aggregated Ethernet member link is brought down with minimum link configuration. PR1477821
- MC-AE interface might be shown as unknown status if you add the subinterface as part of the VLAN on the peer MC-AE node. PR1479012

- For ATM interfaces configuration, if any logical interface has the **allow-any-vci** configuration, then the commit operation might fail. PR1479153
- PPPoE subscribers are not up while verifying static IPv4 subscriber in passive mode. PR1483395
- CFM over BD along with negative events lead to restart and CFM DM two-way verification fails. PR1489196
- The **vrrp-inherit-from** change operation leads to packet loss when traffic is forwarded to the VIP gateway. PR1489425

Intrusion Detection and Prevention (IDP)

- The CLI now provides helpful remarks about IDP's tunable detector parameters. PR1490436
- When creating custom IDP signatures that match on raw bytes (hexadecimal), the commit check fails if the administrator has configured the depth parameter. PR1506706

J-Web

- Junos OS security vulnerability in J-Web and Web-based (HTTP/HTTPS) services. PR1499280

Junos Fusion for Enterprise

- SDPD core file is found at **vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry**. PR1454335
- Loop detection might not work on extended ports in Junos fusion scenarios. PR1460209

Junos Fusion Satellite Software

- Temperature sensor alarm is seen in Junos fusion scenario. PR1466324

Layer 2 Ethernet Services

- On MX2010 and MX2020 platforms, no alarm is generated when FPC is connected to master Routing Engine through backup Routing Engine/CB. PR1461387
- Member links state might be unsynchronized on a connection between a PE device and a CE device in an EVPN active/active scenario. PR1463791
- Telemetry data for relay/bindings/binding-state-v4relay-binding and relay/bindings/binding-state-v4relay-bound is not correct. PR1475248
- On MX204 platform, the Vendor-ID is set as MX10001 in factory-default configuration and DHCP client messages. PR1488771

- With ALQ and VRRP configurations, DHCP subscribers are not coming up. PR1490907
- Issues with DHCPv6 relay processing confirm and reply packets. PR1496220
- The MC-LAG might become down after disabling and then enabling the **force-up**. PR1500758

Layer 2 Features

- Connectivity is broken through LAG because of the members configured with **hold-time** and **force-up**. PR1481031

MPLS

- Traffic loss might be seen if p2mp with NSR enabled. PR1434522
- P2MP LSP might flap after VT interface in MVPN routing instance is reconfigured. PR1454987
- The RSVP interface bandwidth calculation rounds up. PR1458527
- The rpd might crash in PCEP for the RSVP-TE scenario. PR1467278
- The fast reroute detour next-hop down event might cause the primary LSP go in the **Down** state in a particular scenario. PR1469567
- The rpd process might crash during shutdown. PR1471191
- The LDP and BFD sessions are not coming up in a scaled setup. PR1474204
- The RSVP LSPs might not come up in a scaled network with a very high number of LSPs if NSR is used on the transit router. PR1476773
- PCC might flood with event logs to controller. PR1476822
- Kernel crash and device might restart. PR1478806
- The rpd process crashes on the backup Routing Engine when LDP tries to create LDP P2MP tunnel upon receiving corrupted data from the master Routing Engine. PR1479249
- On MX Series with MPC10E line card, rpd core files **in rsvp_copy_route (rt=< optimized out>, rtparms_p=< optimized out>) at ../../../../../../../../../src/junos/usr.sbin/rpd/mpls_te/proto/rsvp/proto/rsvp_route.c:3033** are seen after GRES. PR1485985
- The rpd might crash on restart of master Routing Engine or backup Routing Engine when chain-NH has inner and outer labels in SR-TE scenario. PR1486077
- High CPU utilization for rpd might be seen if RSVP is implemented. PR1490163
- The rpd might crash when BGP with FEC 129 VPWS enabled flaps. PR1490952
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. PR1493431

- The rpd might crash in a rare condition under SR-TE scenario. PR1493721
- The rpd core files are generated during unified ISSU. PR1493969
- The rpd process might crash when SNMP polling is done using oid jnxMplsTeP2mpTunnelDestTable. PR1497641
- The rpd process might crash with RSVP configured in a rare timing case. PR1505834

Platform and Infrastructure

- Core.vmxt.mpc0 is seen at **0x096327d5 in l2alm_sync_entry_in_pfes (context=0xd92e7b28, sync_info=0xd92e7a78) at ../../../../../src/pfe/common/applications/l2alm/l2alm_common_hw_api.c:1727**. PR1430440
- With chained composite next hop enabled, the MPLS CoS rewrite does not work for IPv6 PE device traffic. PR1436872
- Traffic loss might be seen in case of Ethernet frame padding with VLAN. PR1452261
- Modifying the REST configuration might cause the system to become unresponsive. PR1461021
- On MX204 platform, Packet Forwarding Engine errors might occur when incoming GRE tunnel fragments get sampled and undergo inline reassembly. PR1463718
- The CoS might not work on MPC10E and MPC11E line cards. PR1465870
- VXLAN packet might be discarded with flow caching enabled on MX150 and vMX. PR1466470
- All the subscriber services might be unavailable on vBNG running on MX150 and vMX running in payg mode. PR1467368
- The JNH memory leaks after CFM session flap for LSI and VT interfaces. PR1468663
- The switch might not be able to learn MAC address with **dot1x** and **interface-mac-limit** configured. PR1470424
- SSH login might hang and the TACACS+ server closes the connection without sending any authentication failure response. PR1478959
- Remote MEPs are not coming up as expected while verifying MIP functionality with bridge domains. PR1484303
- The **show system buffer** command displays all zeros in the MX104 chassis. PR1484689
- MAC learning under bridge domain stops after MC-LAG interface flaps. PR1488251
- MAC malformation might happen in a rare scenario under MX Series Virtual Chassis setup. PR1491091
- In node slicing setup MPLS TTL might be set to zero when the packet goes through af interface configured with CCC family. PR1492639
- A specific IPv4 packet might lead to FPC restart. PR1493176

- Python or Slax script might not be executed. PR1501746
- MPCs might crash when there is a change on routes learned on IRB interface configured in VPLS and EVPN instance. PR1503947
- Traffic convergence failed with ICL failure case. PR1505465

Routing Policy and Firewall Filters

- The router-id from martian address range cannot be committed even if the range is allowed by configuration. PR1480393

Routing Protocols

- The BGP session might be stuck with high BGP OutQ value after GRES on both sides. PR1323306
- PIM RPF selection for the specific multicast group might get incorrectly applied to other multicast groups. PR1443056
- TI-LFA might be unable to install backup path in the routing table in a specific case. PR1458791
- BGP NSR with more than 40,000 IPv6 peers is not qualified or supported. PR1461436
- IS-IS IPv6 routes might flap when there is an unrelated commit under protocol stanza. PR1463650
- The rpd might crash if IPv4 routes are programmed with IPv6 next hop through JET APIs. PR1465190
- BGP peers might flap if the parameter of hold-time is set small. PR1466709
- The configured BGP damping policy might not take effect after BGP is disabled and then enabled followed by **commit**. PR1466734
- The rpd might stop when both instance-import and instance-export policies contain the as-path-prepend action. PR1471968
- Removing cluster from BGP group might cause prolonged convergence time. PR1473351
- Adjacency SID might be missed and not be advertised to peer/controller/BMP monitor in BGP-LS NLRI. PR1473362
- SFTP does not connect properly and the following error is displayed: **Received message too long**. PR1475255
- BGP TCP MD5 authentication support is not available. PR1476669
- The rpd process might crash with BGP multipath and route withdraw occasionally. PR1481589
- The rpd process crashs due to specific BGP UPDATE packets. PR1481641
- The rpd process might crash when deactivating logical systems. PR1482112

- BGP multipath traffic might not fully load-balanced for a while after adding a new path for load sharing. PR1482209
- The rpd might be crashed after BGP peer flapping PR1482551
- RIPv2 packets stop transmitting when changing interface-type configuration from p2mp to broadcast. PR1483181
- The rpd process crashes if the same neighbor is set in different RIP groups. PR1485009
- On MX Series, MSDP memory leak is observed. PR1485206
- The BGP-LU routes do not have the label when BGP sharding is used. PR1485422
- Removal of the BGP and rib-sharding configuration might cause routing protocols to become unresponsive. PR1485720
- Layer 3 VPN RR with **family route-target** and **no-client-reflect** statement does not work as expected. PR1485977
- Traffic loss is seen on a scaled MPLS setup after unified ISSU in enhanced mode. PR1486657
- The rpd process crashes if the BGP LLGR with RIB sharding and traceoptions for graceful-restart are configured. PR1486703
- The rpd might crash when you perform GRES with MSDP configured. PR1487636
- High CPU utilization might be observed when the outgoing BGP updates are sent slowly. PR1487691
- The rpd process might generate core file after **always-compare-med** is configured for BGP path-selection. PR1487893
- BGP RIB sharding feature cannot be run on a system with a single CPU. PR1488357
- The rpd crashes when reset OSPF neighbors. PR1489637
- The BGP route target family might prevent route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- The rpd might crash because of rpd resolver problem of INH. PR1494005
- The static route in inet6.0 or inet6.3 RIB might fail to delete. PR1495477
- For SPRING support SRv6, continuous rpd core files are generated at **isis_set_rt_pfx_sid_tsi,isis_route_change_rt** after configuring **[set protocols isis topologies ipv6-unicast]**. PR1495994
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. PR1497721
- The rpd might crash if the import policy is changed to accept more routes that exceed the teardown function threshold. PR1499977
- The rpd process crashs when processing a specific BGP packet. PR1502327
- The **show bgp neighbors** command shows change in x-path output for **input-updates** value. PR1504399
- BGP might not advertise routes to peers after a peer flap. PR1507195

Services Applications

- **flow-tap** add function might not work after the dynamic flow capture services process restarted. PR1472109
- On an MX Series router, L2tp LTS fails to forward the **agentCircuitId** and **agentRemoteId** AVP toward the LNS. PR1472775
- The kmd might crash due to the incorrect IKE SA establishment after the remote peer's NAT mapping address has been changed. PR1477181
- NPC core files are found at **services_inline_handle_svc_set_add services_inline_gencfg_handler gencfg_specific_handler**. PR1502527

Subscriber Access Management

- The authd process might crash after the unified ISSU from Junos OS Release 18.3 and earlier to Junos OS Release 18.4 and later. PR1473159
- Syslog messages **pfe_tcp_listener_open_timeout: Peer info msg not received from addr: 0x6000080. Socket 0xfffff804ad23c2e0 closed** is observed. PR1474687
- The delete request of a specified service session through CoA could fail. PR1479486
- The CoA request might not be processed if it includes the **proxy-state** attribute. PR1479697
- The **mac-address** CLI option is hidden under the **access profile profile-name radius options calling-station-id-format** statement. PR1480119
- The authd logs events might not be sent to syslog host when **destination-override** is used. PR1489339

VPNs

- Traffic loss might be observed when the inter-AS next-generation MVPN VRF is disabled on one of the ASBRs. PR1460480
- The rpd might crash when "link-protection" is added or deleted from LSP for MVPN ingress replication selective provider tunnel. PR1469028
- On MVPN scenario, the LSP might stay down on removing all VT interfaces from a single hop egress. PR1474830
- The MPC10E-15C-MRATE next-generation MPVN ingress replication flushing out is not proper when in egress the ingress replication configuration is deactivated. PR1475834
- The l2circuit neighbor might be stuck in RD state at one end of MG-LAG peer. PR1498040

- The rpd core files are generated while disabling l2ckt with connection protection, backup neighbor configuration, and l2ckt trace logs enabled. PR1502003
- The rpd might crash when you delete l2circuit configuration in a specific sequence. PR1512834

SEE ALSO

What's New | 87

What's Changed | 113

Known Limitations | 116

Open Issues | 119

Documentation Updates | 145

Migration, Upgrade, and Downgrade Instructions | 146

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 20.2R1 documentation for MX Series.

Advanced Subscriber Management Provider

- The Broadband Subscriber Services User Guide incorrectly stated that for Routing Engine-based, converged HTTP redirect services, a CPCD service rule can include both a redirect term and a rewrite term. It also incorrectly stated that you can include separate rewrite and redirect rules in the same service profile.

SEE ALSO

What's New | 87

What's Changed | 113

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the MX Series. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

Starting in Junos OS 17.4R1 release, FreeBSD 11.x is the underlying OS for all Junos OS platforms which were previously running on FreeBSD 10.x based Junos OS. FreeBSD 11.x does not introduce any new Junos OS related modifications or features but is the latest version of FreeBSD.

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.x-based Junos OS	FreeBSD 11.x-based Junos OS
MX5,MX10, MX40,MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 20.2R1

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the juniper.conf and ssh files might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the Installation and Upgrade Guide.

For more information about the installation process, see Installation and Upgrade Guide and Upgrading Junos OS with Upgraded FreeBSD.

Procedure to Upgrade to FreeBSD 11.x-based Junos OS

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install FreeBSD 11.x-based Junos OS:

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. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.

4. Select the Software tab.

5. In the Install Package section of the Software tab, 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 routing platform or to your internal software distribution site.

10. Install the new jinstall package on the routing platform.

> **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.

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-32-20.2R1.9-signed.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-20.2R1.9-signed.tgz
  ```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-32-20.2R1.x-limited.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-20.2R1.9-limited.tgz
  ```

Replace source with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname*/*pathname***

- **http://*hostname*/*pathname***
- **scp://*hostname*/*pathname***

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.x) to Junos OS (FreeBSD 11.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 11.x, and Junos OS (FreeBSD 6.x) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

Use the **reboot** command to reboot the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the junos-vmhost-install-x.tgz image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the Installation and Upgrade Guide.
- Starting in Junos OS Release 20.2R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following MX Series routers:
  - MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008

  [See https://kb.juniper.net/TSB17603.]

NOTE: After you install a Junos OS Release 20.2R1 jinstall package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the jinstall package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. See the VM Host Software Administrative Commands in the Installation and Upgrade Guide.

Procedure to Upgrade to FreeBSD 6.x-based Junos OS

Products impacted: MX5, MX10, MX40, MX80, MX104.

To download and install FreeBSD 6.x-based Junos OS:

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. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.

4. Select the Software tab.

5. In the Install Package section of the Software tab, 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 routing platform or to your internal software distribution site.

10. Install the new jinstall package on the routing platform.

> **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.

- All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

  `user@host>` **request system software add validate reboot *source*/jinstall-ppc-20.2R1.9-signed.tgz**

- Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/jinstall-ppc-20.2R1.9-limited-signed.tgz
```

Replace source with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname/pathname***
  - **http://*hostname/pathname***
  - **scp://*hostname/pathname***

The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Use the **reboot** command to reboot the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 20.2R1 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases.

Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform the following Junos OS installation on each Routing Engine separately to avoid disrupting network operation:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine, and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the Installation and Upgrade Guide.

Downgrading from Release 20.2R1

To downgrade from Release 20.2R1 to another supported release, follow the procedure for upgrading, but replace the 20.2R1 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

SEE ALSO

Junos OS Release Notes for NFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for NFX Series.

> **NOTE:** For information about NFX product compatibility, see NFX Product Compatibility.

Application Security

- **AppQoE multihoming with active-active deployment (NFX150, NFX250, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and vSRX)**—Starting In Junos OS Release 20.2R1, AppQoE is enhanced to support multihoming with active/active deployment. In previous releases, AppQoE supports multihoming with active/standby deployment.

  In active/active deployment, the spoke device connects to multiple hub devices. Application traffic can transit through any of the hub devices if the link to the hub device meets SLA requirements. Application traffic can switch seamlessly between the hub devices in case of SLA violation or if the active hub device is not responding.

  To support active/active mode, you must enable the BGP multipath to allow the device to select multiple equal-cost BGP paths to reach a given destination.

  [Application Quality of Experience (AppQoE).]

- **Packet capture for unknown application traffic (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.2R1, you can generate packet capture information for unknown application traffic on your security device. You can use this information to get more insight on unknown applications.

  After you configure packet capture for the application traffic on your device, the packet capture function captures the packet details and stores the information in a packet capture (**.pcap**) file. You can use the packet capture details of an unknown application to define a new custom application signature and create a security policy rule to manage the application traffic more efficiently.

You can submit the packet capture information to Juniper Networks to debug why an application is not detected, and if required, request to create an application signature.

[See Application Identification.]

High Availability

- **High availability on NFX250 NextGen devices**—Starting in Junos OS Release 20.2R1, NFX250 NextGen devices support the high availability feature. You can configure a cluster of two NFX250 NextGen devices to act as primary and secondary devices for protection against device failures. The high availability feature supports Layer 2 and Layer 3 features in dual CPE deployments.

  By default, the ge-0/0/0 interface functions as the control interface. You can configure one of the remaining front panel interfaces as the fabric interface. On the LAN, the active/backup mechanism is used. If the primary device fails, the secondary device takes over the operation. On the WAN, both active/active and active/backup mechanisms are supported.

  [How to Configure the NFX250 NextGen.]

Interfaces

- **ADSL and VDSL interfaces on NFX350 devices**—Starting in Junos OS Release 20.2R1, NFX350 devices support ADSL and VDSL interfaces.

  [How to Configure the NFX350.]

SEE ALSO

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for NFX Series devices.

What's Changed in Release 20.2R1

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 20.2R1 for NFX Series devices.

SEE ALSO

- What's New | 154
- Known Limitations | 156
- Open Issues | 157
- Resolved Issues | 159
- Documentation Updates | 161
- Migration, Upgrade, and Downgrade Instructions | 162

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for NFX Series devices. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

High Availability

- On an NFX250 chassis cluster, commit fails for LAG deployment on a reth interface. PR1487857

Platform and Infrastructure

- With an SRX1500 device used as a hub device and an NFX350 device as spoke device, IPsec replay errors are seen with HTTP traffic when the AppQoE passive probing is enabled. As a workaround, use SRX4200 as the hub device. PR1461068

SEE ALSO

What's New | **154**

What's Changed | **156**

Open Issues | **157**

Resolved Issues | **159**

Documentation Updates | **161**

Migration, Upgrade, and Downgrade Instructions | **162**

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.2R1 for NFX Series devices.

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

High Availability

- In an NFX250 chassis cluster, successive FPC0 manual restarts using the **request chassis fpc slot 0 restart** command must be 120 seconds apart. If restart is attempted within this interval, it is rejected with an error message, **Router is in transition, try again**.

  As a workaround, wait for 120 seconds between successive FPC0 restarts. PR1486155
- For an NFX250 chassis cluster, MAC learning should be disabled on fabric VLANs. We also recommend that you have only one L2 and L3 interface per node as part of the fabric VLAN. PR1495188

Interfaces

- When you issue a **show interface** command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. PR1306191
- On NFX150 and NFX250 NextGen devices, when you add, modify, or delete a VNF interface that is mapped to an L2 or L3 data plane, kernel traces might be observed on the NFX Series device console. PR1435361
- The heth-0-4 and heth-0-5 ports do not detect traffic when you try to activate the ports by plugging in or unplugging the cable. As a workaround, perform a link flap or enable or disable the interface using the CLI. PR1449278
- The link disable option puts the analyzer interface in an inconsistent state with link state as DOWN and admin state as UP. PR1442224

Platform and Infrastructure

- On NFX150 devices, MAP-E customer edge (CE) configurations do not perform validation to check whether the suffix part is nonzero. The configuration must ensure that the suffix part of configurations involving MAP-E prefixes are zeros. PR1457927
- On NFX350 devices, traffic drop is seen with fragmented traffic, and the log reports **FLOW_REASSEMBLE_FAIL**. PR1475023
- On NFX150 devices, srxpfe core file is observed while testing the ADSL interface. PR1485384
- Login access to JDM through TACACS failed after upgrade to Junos OS Release 18.4R3

  As a workaround, log in as a local user. PR1504915

Virtual Network Functions (VNFs)

- On NFX Series devices, while configuring **vmhost vlans** using **vlan-id-list**, the system allows duplicate VLAN IDs in the VLAN ID list. PR1438907.

SEE ALSO

What's New | **154**

What's Changed | **156**

Known Limitations | **156**

Resolved Issues | **159**

Documentation Updates | **161**

Migration, Upgrade, and Downgrade Instructions | **162**

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS Release 20.2R1 for NFX Series devices.

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

Application Security

- AppQoE is sending active prob packets for the deleted **active-probe-params**. PR1492208

High Availability

- On NFX250 chassis cluster, L3 interfaces are not getting created after secondary automatic reboot when control port recovery is enabled. PR1502449

Interfaces

- On NFX150 devices, no error is displayed when the commit fails after you configure **native-vlan-id** on an access VNF interface. PR1438854
- On NFX250 NextGen devices, the **monitor interface traffic** command might not display the pps output for SXE and physical interfaces. PR1464376
- On NFX350 devices, the **clear interface statistics all** command takes a longer time to execute. PR1475804
- On NFX350 devices, if you delete and add an SXE interface, the SXE interface moves to the Spanning Tree Protocol blocking (STP BLK) state, and the traffic drops on that interface. PR1475854

Mapping of Address and Port with Encapsulation (MAP-E)

- On NFX Series devices, IP identification (IP ID) is not changed after MAP-E NAT44 is performed on fragment packets when the packets reach the customer edge (CE) device.

  PR1478037

Platform and Infrastructure

- On NFX150 devices, MAC aging does not work. You must remove aged MAC entries from the CLI. PR1502700
- On NFX350 devices, if you execute the **show vmhost mode** command multiple times, JDM might crash and cause the **show vmhost mode** commands to stop working. PR1474220
- Core files on NFX250 while adding the second LAN subnet. PR1490077
- After initiation of zeroization, the NFX250 device is going into a reboot loop. PR1491479
- The **request vmhost power-off** command reboots the NFX250 NextGen device instead of powering off the device. PR1493062

Virtualized Network Functions (VNFs)

- On NFX150 and NFX250 NextGen devices, when two flowd interfaces are mapped to the same physical interface and if you delete the interface mapping to VF0, the traffic flow is disrupted. Even though the mapping is moved to VF0, the MAC address is not cleared in VF1, which disrupts the traffic. PR1448595
- On NFX350 devices, VNF instantiation is not working properly. PR1478456

SEE ALSO

- What's New | 154
- What's Changed | 156
- Known Limitations | 156
- Open Issues | 157
- Documentation Updates | 161
- Migration, Upgrade, and Downgrade Instructions | 162

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for NFX Series devices.

SEE ALSO

- What's New | 154
- What's Changed | 156
- Known Limitations | 156
- Open Issues | 157
- Resolved Issues | 159
- Migration, Upgrade, and Downgrade Instructions | 162

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the NFX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

NOTE: For information about NFX product compatibility, see NFX Product Compatibility.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases

End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

Basic Procedure for Upgrading to Release 20.2

When upgrading or downgrading Junos OS, use the **jinstall** package. For information about the contents of the **jinstall** package and details of the installation process, see the Installation and Upgrade Guide. Use other packages, such as the **jbundle** package, only when so instructed by a Juniper Networks support representative.

> **NOTE:** The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the device, such as configuration templates and shell scripts (the only exceptions are the juniper.conf and ssh files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the device. For more information, see the Software Installation and Upgrade Guide.

> **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.

To download and install Junos OS Release 20.2R1:

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. Select the name of the Junos OS platform for the software that you want to download.

3. Select the **Software** tab.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the Download Software page.

5. In the Install Package section of the Software tab, 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 Juniper Networks representatives.

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.

SEE ALSO

Junos OS Release Notes for PTX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in Junos OS Release 20.2R1 for PTX Series routers.

High Availability (HA) and Resiliency

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes mastership. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

  [See Understanding the Ephemeral Configuration Database.]

- **Unsupported hardware for unified ISSU (MX240, MX480, MX960, MX10003, and PTX3000)**—The following cards do not support unified ISSU upgrading to Junos OS Release 20.2R1:
  - MPC7E-MRATE
  - MPC8E with MRATE MIC
  - MPC9E with MRATE MIC
  - MPC10E-10C-MRATE
  - MPC10E-15C-MRATE
  - PTX5000 with 24-Port 10-Gigabit Ethernet, 40-Gigabit Ethernet PIC with QSFP+ or 15-Port 10-Gigabit, 40-Gigabit Ethernet, 100-Gigabit Ethernet PIC with QSFP28
  - MX10003 with QSFP28 Ethernet TIC

Interfaces and Chassis

- **Support for 1-Gbps speed on QFX-60S line card (PTX10008 and PTX10016)**—In Junos OS Release 20.2R1 and later, the QFX10000-60S-6Q line card supports 1-Gbps speed on its ports (0 to 59). The QFX10000-60S-6Q line card contains 60 SFP+ ports that support 10 Gbps, two dual-speed QSFP28 ports that support either 40 Gbps or 100 Gbps, and four QSFP+ ports that support 40 Gbps. You can individually configure ports 0 to 59 for 10-Gbps or 1-Gbps port speed. Use the **set chassis fpc *fpc-slot-number* pic *pic-number* port *port-number* speed 1G** command to change the mode of a port from 10 Gbps to 1 Gbps. The transceivers supported for 1 Gbps are QFX-SFP-1GE-LX, QFX-SFP-1GE-SX, and QFX-SFP-1GE-T.

  By default, QFX1000-60S-6Q line card (ports 0 to 59) operates at 10-Gbps speed.

  [See QFX10000 Line Cards for details on the combination of modes supported on the ports.]

Juniper Extension Toolkit (JET)

- **RIB service APIs support dynamic next-hop interface binding (MX Series, PTX Series, and vMX)**—Starting in Junos OS Release 20.2R1, programmed RIB routes react to Up, Down, Add, and Delete events for direct next-hop interfaces. When all direct next-hop interfaces are unusable, the route becomes inactive. This prevents traffic from being dropped and keeps inactive routes from being propagated through the network.

  This feature applies to all routes programmed using the rib_service JET API where an interface is configured as a direct next hop, including interfaces that are part of a flexible tunnel. It also applies to tunnels configured with the flexible_tunnel_service JET API.

  To disable this feature, use **edit routing-options programmable-rpd rib-service dynamic-next-hop-interface disable**.

[See rib-service (programmable-rpd), Juniper Extension Toolkit Developer Guide, and Juniper Engineering Network website.]

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

  [See language (Scripts), Develop Off-Device JET Applications, and Develop On-Device JET Applications.]

Junos Telemetry Interface

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

  [See Mapping OpenConfig Routing Policy Commands to Junos Configuration and Mapping OpenConfig Network Instance Commands to Junos Operation.]

- **ON-CHANGE BGP peer information statistics support for JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 provides BGP peer sensor support using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON_CHANGE statistics are sent to an outside collector.

  The following resource paths are supported:

  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/active** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/received** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/sent** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/rejected** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/admin-state** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/established-transitions** (ON_CHANGE)

- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/last-established** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/received/notification** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/messages/received/update** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/notification** (stream
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/update** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/session-state** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/supported-capabilities** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/local-address** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-address** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-port** (ON_CHANGE)

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

- **Telemetry support for LDP and MLDP traffic statistics (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, the following LDP and multipoint LDP native sensors are added for the Junos telemetry interface:
  - /junos/services/ldp/label-switched-path/ingress/usage/
  - /junos/services/ldp/label-switched-path/transit/usage/
  - /junos/services/ldp/p2mp/interface/receive/usage/
  - /junos/services/ldp/p2mp/interface/transmit/usage/
  - /junos/services/ldp/p2mp/label-switched-path/usage/

  You must enable telemetry streaming with the **sensor-based-stats** option at the **[edit protocols ldp traffic-statistics]** hierarchy level.

  The **show ldp traffic-statistics** command is enhanced to display upstream LDP traffic statistics and to display multipoint LDP traffic statistics per interface.

On PTX Series routers, this feature is not supported for the following variants:

- PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine
- PTX10003
- PTX10008 with the PTX10K-LC1201-36CD line card
- FPC2 line cards do not support ingress multipoint LDP statistics.

[See sensor (Junos Telemetry Interface).]

- **CPU statistics support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 supports streaming various CPU statistics and process parameters using remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) services and Junos telemetry interface (JTI). You can stream CPU usage per process (statistics are similar to output from the **show system process detail** operational mode command), as well as CPU usage per Routing Engine core.

  This feature supports the private data model openconfig-procmon.yang.

  To stream statistics to an outside collector, include the following resource paths in a gRPC or gNMI subscription:

  - Individual process level information (resource path **/system/processes/process**)
  - Individual Routing Engine core information (resource path **/components/component/cpu/**)

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

- **Packet Forwarding Engine sensor support with INITIAL_SYNC on JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000 line of routers, QFX5100, and QFX5200)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export Packet Forwarding Engine statistics from devices to an outside collector using gNMI submode INITIAL_SYNC. When an external collector sends a subscription request for a sensor with INITIAL_SYNC (gnmi-submode 2), the host sends all supported target leaves (fields) under that resource path at least once to the collector with the current value. This is valuable because:

  - The collector has a complete view of the current state of every field on the device for that sensor path.
  - Event-driven data (ON_CHANGE) is received by the collector at least once before the next event is seen. In this way, the collector is aware of the data state before the next event happens.
  - Packet Forwarding Engine sensors that contain zero counter values (zero-suppressed) that normally do not show up in streamed data are sent, ensuring that all fields from each line card (also referred to as source) are known to the collector.

  NOTE: ON_CHANGE data is not available for native (UDP) Packet Forwarding Engine Sensors.

INITIAL_SYNC submode requires that at least one copy to be sent to the collector; however, sending more than one is acceptable.

INITIAL_SYNC submode is supported for the following sensors:

- Sensor for CPU (ukernel) memory (resource path **/junos/system/linecard/cpu/memory/**)
- Sensor for firewall filter statistics (resource path **/junos/system/linecard/firewall/**)
- Sensor for physical interface traffic (resource path **/junos/system/linecard/interface/**)
- Sensor for logical interface traffic (resource path **/junos/system/linecard/interface/logical/usage/**)
- Sensor for physical interface queue traffic (resource path **/junos/system/linecard/interface/queue/**)
- Sensor for physical interface traffic except queue statistics (resource path **/junos/system/linecard/interface/traffic/**)
- Sensor for NPU memory (resource path **/junos/system/linecard/npu/memory/**)
- Sensor for NPU utilization (resource path **/junos/system/linecard/npu/utilization/**)
- Sensor for packet statistics (resource path **/junos/system/linecard/packet/usage/**)
- Sensor for software-polled queue-monitoring statistics (resource path **/junos/system/linecard/qmon-sw/**)

[See Understanding OpenConfig and gRPC and gNMI on Junos Telemetry Interface and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

MPLS

- **Support for MPLS ping and traceroute for segment routing (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.2R1, we extend the MPLS ping and traceroute support for all types segment routing--traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

  We also support the following features:

  - FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
  - ECMP traceroute support for all types of SR-TE paths.

  We do not support the following:

  - Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
  - OAM for IPv6 prefix.
  - BFD

[See traceroute mpls segment-routing spring-te and ping mpls segment routing spring-te.]

Network Management and Monitoring

- **SNMP support for multicast LDP MIB objects (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS SNMP extends support for the following multicast LDP MIB tables and objects:
  - mplsMldpInterfaceStatsTable
  - mplsMldpFecUpstreamSessPackets
  - mplsMldpFecUpstreamSessBytes
  - mplsMldpFecUpstreamSessDiscontinuityTime

  The multicast LDP standard MIB builds on the objects and tables that are defined in RFC3815, which only supports LDP point-to-point label-switched paths (LSPs). This multicast LDP MIB provides support for managing multicast LDP point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs.

  [See Standard SNMP MIBs Supported by Junos OS and SNMP MIB Explorer.]

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

  [See NETCONF Sessions over Outbound HTTPS.]

- **Enhanced on-box monitoring support on the control plane (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure traceoptions to track all events related to system-level and process-level memory monitoring. You can also view the history of the actions taken for system-level and process-level memory monitoring by using the **show system monitor memory actions** command.

Routing Policy and Firewall Filters

- **Support for additional route filter qualifiers in a policy statement (PTX1000 and PTX10000)**—Starting in Junos OS Release 20.2R1, the following list-level qualifiers are supported: **exact**, **longer**, **orlonger**, **prefix-length-range**, and **upto**.

  You can use route filter lists to group individual route filters created at the **[edit policy-options]** hierarchy level. Each item in a list consists of a complete route filter statement, including a destination prefix, a match type, and an optional action. Reuse the list in different policies, adding whatever qualifiers you need, instead of re-creating a different one for every use case.

  [See Understanding Route Filters for Use in Routing Policy Match Conditions.]

Routing Protocols

- **TI-LFA SRLG protection for IS-IS (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure Shared Risk Link Group (SRLG) protection in topology-independent loop free alternate (TI-LFA) networks. IS-IS computes the fast reroute path that is aligned with the post-convergence path and excludes the SRLG of the protected link. All local and remote links that share any SRLG with the protecting link are excluded. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface.

  To enable TI-LFA SRLG protection with segment routing for IS-IS, include the **srlg-protection** statement at the **[edit protocols isis interface *name* level *number* post-convergence-lfa]** hierarchy level.

  [See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS.]

- **Support for BGP-LU over SR-TE for color-based mapping of VPN Services (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, we are extending support to BGP labeled unicast service for color-based mapping of VPN services over Segment Routing-Traffic Engineering (SR-TE). This enables you to advertise BGP-LU IPv6 and IPv4 prefixes with an IPv6 next-hop address in IPv6-only networks where routers do not have any IPv4 addresses configured. With this feature, BGP-LU can now resolve IPv4 and IPv6 routes over SR-TE core. BGP-LU constructs a colored protocol next hop, which is resolved on a colored SR-TE tunnel in the **inetcolor.0** or **inet6color.0** table. Currently we support BGP IPv6 LU over SR-TE with IS-IS underlay.

  See [Understanding Static Segment Routing LSP in MPLS Networks.]

- **Support for BGP-SR-TE rearchitecture (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS provides support for controller-based BGP segment routing--traffic engineering (SR-TE) routes to be installed as source packet routing traffic-engineered (SPRING-TE) routes. BGP installs the SR-TE policy in the routing tables bgp.inetcolor.0 and bgp.inet6color.0, and these routes are subsequently installed in the routing tables inetcolor.0 or inet6color.0 by SPRING-TE.

  In releases before Junos OS Release 20.2R1, controller-based BGP SR-TE routes are installed as BGP routes in the routing table. To maintain consistency and for easy maintenance, all SR-TE based routes appear as SPRING-TE routes irrespective of the source.

You need to enable **source-packet-routing** at the **[edit protocols]** hierarchy level to see the routes installed in inetcolor.0 or inet6color.0. A new option **detail** is introduced under **traceoptions (Protocols Spring-TE)** to trace the detailed information.

[See Segment Routing Traffic Engineering at BGP Ingress Peer Overview.]

System Logging

- **Support to track the maximum number of routing and forwarding (RIB/FIB) routes and VRFs (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can track and display the high-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system (RPD) using the **show route summary** CLI command. High-water mark refers to the maximum number of routing and forwarding (RIB/FIB) table routes and VRFs that was present in the RPD system. The high-water mark data can also be viewed in the syslog at the **LOG_NOTICE** level.

  You can configure the interval of the high-water mark data using the **highwatermark-log-interval** CLI configuration statement at the [**edit routing-options**] hierarchy level. The minimum time gap at which the high-water mark data logged in the syslog is 30 seconds. You can configure the value for **highwatermark-log-interval** CLI configuration statement between 5 to 1200 seconds.

  [See routing-options and show route summary.]

SEE ALSO

What's Changed

IN THIS SECTION

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.2R1 for PTX Series routers.

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

  [See commit (System).]

Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

  [See Develop Off-Device JET Applications and Develop On-Device JET Applications.]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

  [See Juniper EngNet.]

Network Management and Monitoring

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

SEE ALSO

What's New | 165

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Release 20.2R1 for PTX Series routers.

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

General Routing

- On the PTX10008 or PTX10016 routers, the GRES takes more than 3 minutes to complete when shutdown is initiated by the internal **vmhost init 0** command. PR1312065
- The filter-based GRE encapsulation does not work in the egress direction when the filter attachment interface and the interface to reach the next hop are the same. PR1465837
- The **sflow record** command shows incorrect output interface for the egress sampling during the incoming MPLS|IPv4 and outgoing IPv4 with ECMP. PR1478012
- The PTX10000 routers includes the incoming MPLS label stack length also in the jvision counters when acting as the PE device egress counter. PR1482408
- On the PTX1000 routers, the following error message is observed when the sampling MPLS+IPv4/IPv6 traffic is forwarded over the IP-IP tunnel: **dlu.ucode.jflow_not_routable pechip**. PR1485770
- The following error messages are seen after configuring **set chassis maximum-ecmp 64**: **JPRDS_NH:jprds_nh_alloc(),990: JNH[3] failed to grab new region for EGRESS**. PR1490813
- The **show dynamic-tunnels database statistics <dest>** command must be structured so that the statistics are fetched deterministically for the IPv4 and IPv6 based tunnels. PR1488715

Routing Protocols

- Router receives and discards traffic for 3 and half minutes after bootup when IGP overload is configured. PR1495435

SEE ALSO

What's New | **165**

What's Changed | **173**

Open Issues | **176**

Resolved Issues | **178**

Documentation Updates | **181**

Migration, Upgrade, and Downgrade Instructions | **181**

Open Issues

IN THIS SECTION

Learn about open issues in the Junos OS Release 20.2R1 for PTX Series routers.

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

General Routing

- Alarm action does not work for minor errors after changing the threshold to 1. PR1345154
- CPU performance might become slow. PR1399369
- The firewall counter for lo0 interface might not increase. PR1420560

- Memory leakage is observed while running the longevity check. PR1438358
- The vehostd process might crash without generating a core file and automatic restart might fail. PR1448413
- Mirroring does not work in Junos OS Release 19.4R2. PR1491789
- On the PTX5000 line of routers, the **show filter index < number> counter** vty command displays values as zero at **28-02-HOSTBOUND_NDP_DISCARD_TERM**. The counter does not increase for the NDP packets. The issue is only with the **show filter index** command, which is a debug tool in vty. This issue has no impact on the NDP functionality for the user traffic. PR1420057
- The outbound SSH connection flap or memory leak issue might be observed during pushing the configuration to the ephemeral database with high rate. PR1497575
- Mirrored packets are corrupted when filter is applied with action port-mirror and discarded. PR1437546
- MPLS sensor does not receive Jvision data on the server. PR1514959
- Flow Session Count for In-line Jflow with IPv4 and IPv6 does not work as expected. PR1510150
- SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at collector. PR1484322

Interfaces and Chassis

- Identical IP address are configured on different logical interfaces from different physical interfaces in the same routing instance including the master routing-instance. PR1221993
- The cfmd process might continuously crash after the upgrade. PR1281073

MPLS

- The rpd process crashes at **rsvp_ing_rt_nh_remove_path**, **rsvp_delete_ing_path**, and **rsvp_delete_lp_headend_route** after routing restarts. PR1498457
- Ingress LSP setup rate is lower than 30 percent lower in Junos OS Release 18.2X75-D410 compared to Junos OS Release 18.2X75-D30.26. PR1457992

Routing Protocols

- The aggregated Ethernet interface and BFD session remain down after the interface is disabled or enabled. PR1354409
- The **show dynamic-tunnels database** command does not reflect the current value of traffic statistics. It shows the cached value of traffic statistics, which might not be equal to the current value. PR1445705
- The traffic loss is observed when the FRR is triggered from BFD. PR1516411

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS Release 20.2R1 for PTX Series routers.

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

General Routing

- PTX interface stays down after the maintenance. PR1412126
- With Junos OS Release 19.4R1 on PTX10008 device along with 4x1GE feature, continuous logging in the chassisd file is observed. PR1456253
- Upgrading fails due to communication failure between the Junos VM and host OS. PR1438219
- The local-loopback test fails with the gigether options. PR1458814
- The PTX1000 or PTX10002 router might discard traffic silently after the transient SIB or FPC voltage alarms. PR1460406
- On the PTX5000 for FPC3, optics-options syslog and link-down do not work as expected. PR1461404

- The sample, syslog, or log action in the output firewall filter with packet size less than 128 might cause ASIC wedge (all packet loss). PR1462634
- On modifying TNL DST NETWORK (more specific TNL DST NETWORK), the IP-IP tunnel gets flushed but fails to get created even though a less specific matching TNL DST NETWORK exists. PR1462805
- On the PTX10000 line of routers, FPC might restart during runtime. PR1464119
- The PTX5000 SIB3 might fail to come up in the slot 0 with or without slot 8 when the Routing Engine 1 is the master. PR1471178
- The input-vlan-map or output-vlan-map might not work properly in the Layer 2 circuit local-switching scenario. PR1474876
- Sampling process might crash when the MPLS or MPLS over the UDP traffic is sampled. PR1477445
- Multicast routes add or delete events might cause adjacency and LSPs to go down. PR1479789
- FPC might crash when dealing with the invalid next hops. PR1484255
- In the StrictPriority mode, the MedH and MedL should be of separate priorities, the StrcH and High become one priority. PR1490505
- The BFD sessions flap when the firewall filter in the loopback0 is changed. PR1491575
- Traffic impact might be seen when policy-multipath is configured without LDP on the Spring-TE scenario. PR1483585
- On a dual Routing Engine GRES or NSR enabled PTX10008 or PTX10016 router, a few TCP-based application sessions like BGP or LDP might flap upon Routing Engine mastership switch. PR1503169
- The router might become nonresponsive and bring traffic down when the disk space becomes full. PR1470217
- Unable to bring the ports up when plugging the optic QSFP-100G-LR4-T2(740-061409) to PTX3000 or PTX5000. PR1511492
- PHP device has NH mis-programming for members of ECMP for SR label route used for reaching the IPV6 destinations. PR1457230

Infrastructure

- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. PR1462986

Layer 2 Ethernet Services

- Member links state might be asychronized on a connection between the PE device and the CE devices in the EVPN A/A scenario. PR1463791

MPLS

- Kernel crash and device restart might occur. PR1478806
- The BGP session might keep flapping between two directly connected BGP peers because of the wrong usage of the TCP-MSS. PR1493431
- The rpd process might crash in a rare condition under the SR-TE scenario. PR1493721

Routing Protocols

- The BGP NSR must be able to synchronize 4000 or more IPv6 sessions. PR1461436
- On the PTX3000 or PTX5000 line of routers, the ppmd process generates a core file after configuring the sbfd responder on the RE-DUO-2600. PR1477525
- The rpd process might crash with the BGP multipath and route withdraw occasionally. PR1481589
- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- BGP multi-pathed traffic might not fully load-balance for a while after adding a new path for the load sharing. PR1482209
- LSP auto-bandwidth adjust-interval change does not get detected on commit in some cases. PR1484801

SEE ALSO

What's New | **165**

What's Changed | **173**

Known Limitations | **175**

Open Issues | **176**

Documentation Updates | **181**

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for PTX Series routers.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the PTX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

Basic Procedure for Upgrading to Release 20.2

When upgrading or downgrading Junos OS, use the **jinstall** package. For information about the contents of the **jinstall** package and details of the installation process, see the Installation and Upgrade Guide. Use

other packages, such as the **jbundle** package, only when so instructed by a Juniper Networks support representative.

> **NOTE:** Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:
>
> ```
> user@host>request system snapshot
> ```

> **NOTE:** The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the juniper.conf and ssh files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the Installation and Upgrade Guide.

> **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.

To download and install Junos OS Release 20.2R1:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage:

   https://support.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the Install Package section of the Software tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new jinstall package on the router.

> **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.

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-20.2R1.9.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-20.2R1.9-limited.tgz
```

Replace the source with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname*/*pathname***
  - **http://*hostname*/*pathname***
  - **scp://*hostname*/*pathname***

The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the junos-vmhost-install-x.tgz image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the Installation and Upgrade Guide.

NOTE: After you install a Junos OS Release 20.2 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you

can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the Installation and Upgrade Guide.

SEE ALSO

What's New | **165**

What's Changed | **173**

Known Limitations | **175**

Open Issues | **176**

Resolved Issues | **178**

Documentation Updates | 181

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for QFX Series switches.

NOTE: The following QFX Series platforms are supported in Release 20.2R1: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.

Junos on White Box runs on Accton Edgecore AS7816-64X switches in this release. The software is based on Junos OS running on QFX5210 switches, so release-note items that apply to QFX5210 switches also apply to Junos on White Box.

What's New in Release 20.2R1-S1

Flow-Based and Packet-Based Processing

- **Support for user-defined flex hashing for MPLS traffic flows (QFX5210; Accton AS7816 running Junos OS on White Box)**—Starting in Junos OS Release 20.2R1-S1, you can configure user-defined flex hashing to load balance MPLS traffic based on TCP or UDP source/destination port information. User-defined flex hashing, which supports protocol versions IPv4 and IPv6, enables you to set byte offsets in packet headers to influence hashing computation. You specify two offsets, each 2 bytes in length, from the first 128 bytes of a packet. Configure the selected bytes to be directly used for hashing or to be used only when the data pattern in these bytes matches with specific values (conditional match). To provide load balancing in spine layers, configure flex hashing and encapsulate the traffic in VXLAN, thus enabling entropy at UDP source ports. At de-encapsulation, configure the **no-inner-payload** statement to load balance based on the outer UDP header.

  To configure user-defined flex hashing:

  ```
  set forwarding-options enhanced-hash-key flex-hashing name ethtype mpls num_labels source-port hash-offset
    offset1 base_offset1 offset1_value offset1_mask offset2 base_offset2 offset2_value offset2_mask
  ```

  To configure a conditional match (repeat the command below with values for offsets and match data 2-4):

  ```
  set forwarding-options enhanced-hash-key conditional-match name offset1 base_offset1 offset1_value
    matchdata1 matchdata1_mask
  ```

  To enable load balancing on VXLAN transit traffic based on the outer UDP header:

  ```
  set forwarding-options enhanced-hash-key vxlan no-inner-payload
  ```

  To troubleshoot, use **show forwarding-options enhanced-hash-key**.

Limitations:

- Use a maximum of two MPLS labels.
- Use only even values for **offset1** and **offset2**.
- If you are using conditional matches, configure the conditions before you attach them to the flex-hashing entry.
- An aggregated Ethernet (AE), or LAG, interface is not supported as an input interface. You *can* configure input interfaces on LAGs by configuring the same user-defined flex-hashing data and the same conditional-match data on all *member* interfaces of a LAG interface. Use unique flex-data profile names and unique conditional-data profile names for each member interface—for example:
  - ...**enhanced-hash-key conditional-match COND_L1_V6_UDP_SRC_PORT_1**...
  - ...**enhanced-hash-key conditional-match COND_L1_V6_UDP_SRC_PORT_2**...

Software Installation and Upgrade

- **Zero touch provisioning (ZTP) with IPv6 support (EX3400, EX4300, QFX5100 and QFX5200 switches, MX-Series routers)**—Starting in Junos OS Release 20.2R1-S1, ZTP supports the DHCPv6 client. 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 will continue to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device will check for DHCPv6 bindings and follow the same process as for DHCPv4 until the device can be provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

  The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

  > **NOTE:** Only HTTP and HTTPS transport protocols are supported EX3400, EX4300, QFX5100, and QFX5200 devices.

  [See Zero Touch Provisioning.]

What's New in Release 20.2R1

Hardware

- **New QFX5120-48T Ethernet Switch (QFX series)**—Starting with Junos OS Release 20.2R1, the QFX5120-48T is a 10GbE/100GbE data center switch offering 48 10GbE RJ-45 ports and six 40GbE/100GbE QSFP28/QFSP+ ports. The 48 cooper ports support 1-Gbps and 10-Gbps speed and the last 6 ports (port 48 to 53) support 40-Gbps and 100-Gbps speed. By default, the first 48 ports operate at 10-Gbps speed and the last six ports 100-Gbps speed.

  QFX5120-48T switches supports both manual and auto-channelization, but manual CLI channelization always takes precedence. [See Port Settings.]

  To install the QFX5120-48T switch hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the QFX5120 Switch Hardware Guide. See Feature Explorer for the complete list of features for any platform.

  Table 2 on page 190 summarizes the software features supported in this release.

  Table 2: Features Supported by QFX5120-48T Switches

Feature	Description
Authentication and Acccess Control	• IEEE 802.1X authentication support. [See User Access and Authentication User Guide.] • IP source guard. [See Configuring IP Source Guard (ELS).] • Local password authentication support for password change policy. • Storm control support (broadcast, unicast, and multicast). [See Understanding Storm Control.] • Radius and TACACS+ authentication. [See Authentication Order for RADIUS, TACACS+, and Local Password.] • Role-based access control (RBAC), and role-based CLI management.
BGP	• Support for BGP Monitoring Protocol (BMP) Version 3 and IPv6 BGP standards. [See Understanding the BGP Monitoring Protocol and Supported IPv6 Standards.] • BGP advertising aggregate bandwidth across external BGP links for load balancing. [See Load Balancing for a BGP Session.] • Support for BGP large communities, link-state distribution, multipath at global level, and support for 4-byte autonomous system numbers. [See Routing Policies for BGP Communities.] • EBGP route support, multiprotocol BGP (MBGP) extensions, and frequent BGP keepalive messages with a short BGP hold time. [See BGP Overview.] • Routing protocol process (rpd) recursive resolution over multipath. [See BGP Route Resolution Overview.] • BGP labeled-unicast. [See labeled-unicast (Protocols BGP.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Class of Service	• Standard class of service (CoS) feature support including configuring classification, rewrite, queuing, shaping, buffering, and scheduling parameters for traffic management. [See CoS Support on QFX Series Switches.] • IEEE 802.1p rewrite and classification. • Class-based queuing with prioritization. [See Understanding CoS Output Queue Schedulers.] • Single-rate two-color marking, single-rate three-color marking, and two-rate three-color marking. [See Overview of Policers.] • Separate unicast and multi-destination classifiers, forwarding classes, and output queues. [See Understanding Junos CoS Components.] • Direct port scheduling. [See Understanding CoS Port Schedulers on QFX Switches.] • Queue shaping using the shaping-rate statement. [See Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth).] • Priority-based flow control (PFC) with 802.3x Ethernet PAUSE and explicit congestion notification (ECN). [See Understanding CoS Flow Control (Ethernet PAUSE and PFC) and Understanding CoS Explicit Congestion Notification.] • CoS support for link aggregation groups (LAGs). • Weighted random early detection (WRED) packet drop profiles and tail drop. [See Understanding CoS Congestion Management and Understanding CoS WRED Drop Profiles.] • Rewrite rule (marking) of bridged packets. [See Understanding Junos CoS Components.] • Policing or rate limiting of traffic to apply limits to traffic flow. [See Overview of Policers.]
DHCP	• Client link-layer address option 79 for DHCPv6. [See mac-address (DHCP Relay Agent).] • DHCP server, DHCP smart relay configuration, DHCP relay with DHCP server, and DHCP client in separate routing instances. [See DHCP Message Exchange Between DHCP Clients and DHCP Server in Different Virtual Routing Instances.] • DHCP relay with option 82 for Layer 2 VLANs and Layer 3 interface. [See DHCP Relay Agent Information Option (Option 82).] • DHCP and DHCPv6 snooping. [See DHCP Snooping.] • DHCP static addresses. [See Configuring Static DHCP IP Addresses.] • Extended DHCP (also referred to as virtual router (VR) aware DHCP. [See Legacy DHCP and Extended DHCP.] • Textual interface description using DHCP relay agent option 82 (circuit ID). [See DHCP Relay Agent Information Option (Option 82).]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
EVPN and VXLAN	• EVPN proxy ARP and ARP suppression. [See EVPN Proxy ARP and ARP Suppression Proxy.] • EVPN control plane and VXLAN data plane support. [See Understanding EVPN with VXLAN Data Plane Encapsulation.] • EVPN pure type-5 route support. [See EVPN Type-5 Route with VXLAN encapsulation for EVPN-VXLAN.] • LACP in EVPN active-active multihoming. [See Example: Configuring LACP for EVPN VXLAN Active-Active Multihoming.] • Automatically generated Ethernet segment identifiers in EVPN-VXLAN and EVPN-MPLS networks. [See Understanding Automatically Generated and Assigned ESIs in EVPN Networks.] • EVPN-VXLAN support of Virtual Chassis and Virtual Chassis Fabric. [See Integrating a Virtual Chassis Fabric into an EVPN-VXLAN Environment.] • Support for VMTO for ingress traffic. [See Configuring EVPN Routing Instances.] • MAC filtering, storm control, and port mirroring support in EVPN-VXLAN overlay networks. See MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.] • Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface. See [Understanding Flexible Ethernet Services Support With EVPN-VXLAN.] • Support for multihomed proxy advertisement. [See EVPN Multihoming Overview.] • Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network. [See Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network.] • Support for graceful restart and graceful restart protocol extension support for unicast and type 5 messages on EVPN-VXLAN. [See Graceful Restart in EVPN.] • Standard class-of-service (CoS) features—classifiers, rewrite rules, and schedulers are supported on VXLAN interfaces. [See Understanding CoS on OVSDB-Managed VXLAN Interfaces.] • Firewall filtering and policing on EVPN-VXLAN traffic. [See Understanding VXLANs and Overview of Firewall Filters.] • Configurable VXLAN UDP port. • Support for IGMP snooping for EVPN-VXLAN in a multihomed environment. [See Overview of Multicast Forwarding with IGMP Snooping in an EVPN-VXLAN Environment.] • Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks. [See Supported Protocols on an IRB Interface in EVPN-VXLAN .] • VXLAN Layer 2 gateway (static, OVSDB, EVPN), Q-in-Q tag manipulation, dynamic load balance, and hashing options. [See OVSDB-VXLAN User Guide for QFX Series Switches.] • BPDU protection in EVPN-VXLAN. [See Supported Protocols on an IRB Interface in EVPN-VXLAN.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Firewall Filters and Policers	• Support for firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces. [See Overview of Firewall Filters.] • Single-rate two-color marking, single-rate three-color marking, and two-rate three-color marking. [See Overview of Policers.] • Dynamic allocation of firewall filters. • Enhanced filter classification of CPU-generated packets. • Firewall filter actions. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches.] • Firewall filter flexible match conditions and firewall filters on loopback and management interface. [See Firewall Filter Flexible Match Conditions.] • Port firewall filters (egress and ingress) and routed firewall filters (egress and ingress). [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • VLAN firewall filters (egress and ingress). [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • TCP/UDP port ranges in classification. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • Filter-based GRE de-encapsulation. [See Configuring a Firewall Filter to De-Encapsulate GRE Traffic.] • Loopback firewall filter scale optimization. [See Planning the Number of Firewall Filters to Create.]
High Availability (HA) and Resiliency	• Automatic recovery for port error disable condition. [See disable-timeout (Port Error Disable).] • Operating system resiliency to recover the Junos software using device recovery mode. [See Rescue Configuration.] • Partial resiliency for errors, machine-check exception (MCE), and advanced error reporting (AER). • Ethernet ring protection switching (ERPS). [See Ethernet Ring Protection Switching Overview.] • Graceful protocol restart for BGP and OSPF. [See Understanding Graceful Restart for BGP, graceful-restart (Protocols BGP) and Configuring Graceful Restart for OSPF.] • Nonstop software upgrade (NSSU), Nonstop bridging, and Nonstop active routing (NSR) for IPv6 and OSPFv2. • Virtual Chassis support. [See Understanding QFX Series Virtual Chassis.] • Virtual Chassis with NSSU support. You can interconnect two QFX5120-48T switches into a Virtual Chassis that operates as one logical device managed as a single chassis. [See Virtual Chassis Overview for Switches.] • Network Device Collaborative Protection Profile (NDcPP) certification.

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Interfaces and Chassis	• Dynamic ARP inspection (DAI) and static ARP support. [See Understanding and Using Dynamic ARP Inspection (DAI).] • Support for dynamic load balancing. [See Understanding Load Balancing for Aggregated Ethernet Interfaces.] • Proxy ARP per VLAN and unrestricted proxy ARP. [See Restricted and Unrestricted Proxy ARP Overview.] • Link protection support on aggregated Ethernet interfaces and updated behavior in static link protection mode. • Automatic detection of MDI and MDIX port connections. Auto MDI/MDIX is enabled by default. [See no-auto-mdix.] • Digital optical monitoring (DOM). [See show interfaces diagnostics optics.] • Support for fibre channel over Ethernet (FCoE), FCoE initialization protocol (FIP), FIP snooping, and up to 2500 total FIP snooping sessions supported on an interface. [See Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch.] • Filter-based GRE decapsulation. • IPv4 generic routing encapsulation (GRE) support. [See Configuring Generic Routing Encapsulation Tunneling.] • Auto-negotiation and port speed. [See auto-negotiation.] • Configure speed of Gigabit Ethernet copper SFP interfaces. [See Gigabit Ethernet Interface.] • IEEE 802.3ah link fault management (LFM). [See OAM Link Fault Management.] • Interface ranges. [See Interface Ranges.] • Jumbo frames (up to 9216 bytes) and jumbo frames on routed VLAN interfaces (RVIs). [See Configuring Routed VLAN Interfaces on Switches (CLI Procedure).] • Layer 3 logical interfaces. [See Layer 3 Logical Interfaces.] • Support for Network-to-network interface (NNI) and user network inteface (UNI) on the same physical interface. [See Configuring Q-in-Q Tunneling.] • Channelizing Ethernet interfaces. [See Channelizing Interfaces Overview.] • Dynamic port swap from 40G to 100G without restarting the packet forwarding engine. • PVLAN and Q-in-Q on the same interface. [See Configuring Q-in-Q Tunneling on QFX Series Switches.] • Link aggregation static and dynamic with LACP (fast and slow LACP), LLDP, and MC-LAG with configuration sync. • Uplink failure detection debounce interval. [See Uplink Failure Detection.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
IPv6	• BGP support for advertising multiple paths to IPv6 addresses. [See Example: Advertising Multiple Paths in BGP.] • Configure per-interface neighbor discovery protocol (NDP) cache protection. [See Neighbor Discovery Cache Protection Overview.] • IPv6 specific SSH and Telnet. • Support for IPv6 filter-based forwarding. [See Understanding Filter-Based Forwarding.] • Firewall filter support for IPv6 traffic: IPv6 fields for ingress port and VLAN firewall filters and policer action for MPLS firewall filters. [See Firewall Filter Match Conditions for IPv6 Traffic.] • Support for IPv6 L3 forwarding, IPv6 Layer 3 VPNs, IPv6 traceroute, IPv6 tunneling, and IPv6 attributes in RADIUS message and stateless auto configuration. • Support for IPv6 OSPFv3, IPv6 ping, secure IPv6 neighbor discovery protocol (NDP), and IPv6 source guard. [See OSPF Version 3 for IPv6 and IPv6 Neighbor Discovery User Guide.] • IPv6 access security (IPv6 neighbor discovery inspection, IPv6 stateless address auto-configuration (SLAAC) snooping, and understanding IPv6 router advertisement guard). [See IPv6 Neighbor Discovery Inspection,IPv6 Stateless Address Auto-configuration (SLAAC) Snooping and Understanding IPv6 Router Advertisement Guard.] • Support for IPv6 over MPLS (6PE), IPv6 over MPLS LSPs, IPv6 static routing, IS-IS for IPv6, path MTU discovery; SNMP, NTP, and DNS. [See Configuring Junos OS for IPv6 Path MTU Discovery.] • Virtual Router Redundancy Protocol (VRRP) and support for VRRP on IPv6 networks. [See VRRP and VRRP for IPv6 Overview.]
Junos OS XML API and Scripting	• Scripts: Python, SLAX, and XSLT commit, event, op, SNMP, and open-source Python modules supported in automation enhancement. • Support for REST API interfaces. • JET for Junos: modern programmatic interface for developers of third-party applications. [See Understanding JET Interaction with Junos OS.] • Configuration management: JSON format for configuration data. [See Defining the Format of Configuration Data to Upload in a Junos XML Protocol Session.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Junos Telemetry Interface (JTI)	• Support for the Junos Telemetry Interface. [See. Understanding OpenConfig and gRPC.] • Sensor level statistics support on Junos Telemetry Interface (JTI). [Guidelines for gRPC and gNMI Sensors.] • gNMI support for routing engine statistics for JTI. [See Guidelines for gRPC and gNMI Sensors.] • Enhancements to the sensor for BGP peer information. • Sensor for network discovery protocol (NDP) and Address Resolution Protocol table state information for IPv6 routes. • Sensor for memory utilization for routing protocol tasks. [See Guidelines for gRPC and gNMI Sensors.] • Sensor for LSP events and properties, LSP statistics, and gRPC streaming for LSP statistics. [See Guidelines for gRPC and gNMI Sensors.] • Packet Forwarding Engine statistics export using gNMI and JTI. • Aggregated Ethernet interfaces configured with the link aggregation control protocol (LACP), Ethernet interfaces configured with the link layer discovery protocol (LLDP), BGP peers, and RSVP interface events. [See Understanding OpenConfig and gRPC on Junos Telemetry Interface.] • OpenConfig LLDP model (v0.1.0). [See OpenConfig Data Model Version.] • OpenConfig to support operational models for VLANs. • OpenConfig Junos OS, OpenConfig, and Network Agent packages are delivered in a single TAR file. [See Installing the OpenConfig Package.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Layer 2 Features	• Data center bridging (DCB) application protocol TLV exchange. • Data Center Bridging Capability Exchange Protocol (DCBX) version support for IEEE DCBX version 1.01. [See Understanding DCBX.] • MAC address filtering, MAC table aging, and static MAC address assignment for interface. [See MAC Addresses and MAC Table Aging.] • Disable MAC learning, persistent MAC learning, MAC address limit per port, MAC limiting, MAC move limiting, MAC notification, and per VLAN (VLAN membership MAC limit). [See Understanding MAC Limiting and MAC Move Limiting for Port Security.] • Enhanced Layer 2 Software (ELS). [See Layer 2 Networking.] • IP directed broadcast traffic forwarding. • VLAN support, Link layer discovery protocol (LLDP), and Q-in-Q tunneling support. [See Configuring Q-in-Q Tunneling.] • Static LAG link protection. [See link-protection (Static LSPs).] • Redundant trunk groups (link redundancy). [See Understanding Redundant Trunk Links (Legacy RTG Configuration).] • L2PT, UDLD, 802.1AE/802.1x, Ethernet Local Management Interface (E-LMI), and Multiple MAC Registration Protocol (MMRP). [See layer2-protocol-tunneling.]
Layer 3 Features	• Configuring the GTP-TEID field for GTP traffic. [See Traffic Sampling, Forwarding, and Monitoring User Guide.] • Equal-cost multipath (ECMP) flow-based forwarding: 64 ECMP paths. [See Traffic Sampling, Forwarding, and Monitoring User Guide.] • Support to control traceroute over Layer 3 VPN. • Virtual routing and forwarding (VRF) support in IRB interfaces in a Layer 3 VPN. • Support for VRF-lite, BGP, IGMP, IS-IS, OSPF, PIM, and RIP.

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
MPLS	• MPLS support for label edge routers (LER) and label switch routers (LSR). [See MPLS Overview for Switches.] • Support for MPLS signaling protocols LDP and RSVP. [See LDP Overview and RSVP Overview.] • Fast reroute (FRR) support (a component of MPLS local protection for both one-to-one and many-to-one local protection). • Static LSPs. [See LSP Overview.] • MPLS node protection, link protection, and statistics for static LSPs. • MPLS OAM (LSP ping). • MPLS statistics. [See statistics (Protocols MPLS).] • MPLS automatic bandwidth allocation and dynamic count sizing. • MPLS with RSVP-based LSPs. • Support for IRB interfaces over an MPLS core network. [See Example: Configuring IRB Interfaces on QFX5100 Switches over an MPLS Core Network.] • MPLS stitching for virtual machine connections. [See Using MPLS Stitching with BGP to Connect Virtual Machines.] • MPLS over Layer 3 subinterfaces. [See MPLS Limitations on QFX Series and EX4600 Switches.] • Resource reservation protocol-traffic engineering (RSVP-TE), traffic engineering extensions (OSPF-TE, IS-IS-TE), Path Computation Element Protocol (PCEP), and PCE-initiated LSPs for the PCEP implementation. [See MPLS Applications User Guide.] • Equal-cost multipath (ECMP) operation on MPLS using firewall filters.
Multichassis Link Aggregation	• Resilient hashing support for link aggregation groups (LAGs) routes. [See Resilient Hashing on LAGs and ECMP groups.] • Keep a link up on a multichassis link aggregation group (MC-LAG) when LACP is not configured on one of the MC-LAG peers. [See Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up.] • Layer 3 unicast and multicast support for MC-LAG. [See Advanced MC-LAG Concepts.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Network Management	• IEEE 802.1ag OAM connectivity fault management. [See Understanding Ethernet OAM Connectivity Fault Management for Switches.] • Port mirroring (local and remote) and remote port mirroring to IP address (GRE). [See Understanding Port Mirroring and Analyzers.] • sFlow technology support. [See Understanding How to Use sFlow Technology for Network Monitoring on a Switch.] • Chef for Junos OS support. [See Chef for Junos OS Getting Started Guide.] • Puppet for Junos OS support. [See Puppet for Junos OS Administration Guide.] • Adding non-native YANG modules to the Junos OS schema. [See Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS.] • Enforcing RFC-compliant behavior in NETCONF sessions. [See Configuring RFC-Compliant NETCONF Sessions.] • Configuring the ephemeral database using the NETCONF and Junos XML protocols. [See Committing an Instance of the Ephemeral Configuration Database Using the NETCONF or Junos XML Protocol.] • Simple network management protocol (SNMP) remote monitoring (RMON) events, alarms, and history. [See SNMP MIB Explorer.] • Real-time performance monitoring (RPM). [See Understanding Real-Time Performance Monitoring on Switches.]
Open vSwitch Database (OVSDB)	• Automatic configuration of OVSDB-managed VXLANs with trunk interfaces. [See Understanding Dynamically Configured VXLANs in an OVSDB Environment.] • BFD in a VMware NSX for vSphere environment with OVSDB and VXLAN. [See Understanding BFD in a VMware NSX Environment with OVSDB and VXLAN.] • CoS on OVSDB-managed VXLAN interfaces. [See Configuring CoS on OVSDB-Managed VXLAN Interfaces.] • Firewall filters on OVSDB-managed interfaces. [See Understanding Firewall Filters on OVSDB-Managed Interfaces.] • MAC limiting on OVSDB managed interfaces. [See Features Supported on OVSDB-Managed Interfaces.] • OVSDB commit failures, schema updates, and support with Contrail. • OVSDB software in Junos OS software package. • OVSDB support with VMware NSX for vSphere. See [Understanding the Junos OS Implementation of OVSDB and VXLAN in a VMware NSX for vSphere Environment.] • Policers and storm control on OVSDB-managed interfaces. [See Understanding Firewall Filters on OVSDB-Managed Interfaces.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
Routing Protocols	• Bidirectional forwarding detection (BFD) support for BGP, IS-IS, and PIM. [See Example: Configuring BFD for BGP and Example: Configuring BFD for IS-IS.] • Static routing. [See Protocol-Independent Routing Properties User Guide.] • Unified Forwarding Table (UFT). [See Understanding the Unified Forwarding Table.] • IPv4 over GRE tunnels—encapsulation and de-encapsulation support. • IGMP version (v1/v2/v3), IGMP filter, IGMP snooping, proxy (relay), and querier. [See Understanding IGMP, IGMP Snooping Overview and igmp-querier.] • Remote support for LDP in IS-IS, static adjacency segment identifier for IS-IS, and alternate loop-free routes and topology-independent loop-free alternate for IS-IS. [See Understanding Remote LFA over LDP Tunnels in IS-IS Networks.] • Multicast Listener Discovery version 1 and 2. [See Configuring MLD.] • Multicast Source Discovery Protocol (MSDP) and multicast-only fast reroute (MoFRR). [See source (Protocols MSDP) .] • IPv6 protocol independent multicast (PIM), PIM Static RP and PIM dense mode (PIM DM), PIM source-specific multicast (PIM SSM), and PIM sparse mode (PIM SM). [See PIM Overview.] • Support for static multicast route leaking for VRF and virtual-router instances. [See Understanding Multicast Route Leaking for VRF and Virtual-Router Instances.] • Virtual routing instances for multicast and unicast protocols. [See Configuring Virtual Router Routing Instances.] • Remote LFA support for LDP tunnels in OSPF and alternate loop-free routes for OSPF and protocol independent multicast (PIM). [See Configuring Loop-Free Alternate Routes for OSPF.]
Spanning Tree Protocols	• Support for IEEE 802.1s Multiple spanning tree protocol (MSTP), IEEE 802.1w rapid spanning tree protocol (RSTP), IEEE 802.1D spanning tree protocol (STP), and IEEE 802.1ak multiple VLAN registration protocol (MVRP). [See Spanning-Tree Protocols User Guide.] • VSTP and RSTP and concurrent configuration. [See Configuring VSTP Protocol.] • Bridge protocol data unit (BPDU) protection, loop protection, and root protection. [See BPDU Protection for Spanning-Tree Protocols, Loop Protection for Spanning-Tree Protocols and Understanding Root Protection for STP, RSTP, VSTP, and MSTP.]
System Logging	• Support for forwarding structured system log messages to a remote system log server. [See Directing System Log Messages to a Remote Machine or the Other Routing Engine.] • System logging (syslog) over IPv4 and IPv6.

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
System Management	• Automatic software download, fast reboot, configuration and image rollback, commit process split into two steps, and rescue configuration. [See Software Installation and Upgrade Guide.] • Support for Precision Time Protocol (PTP) transparent clock. [See Configuring Transparent Clock Mode for Precision Time Protocol.] • Online insertion and removal (OIR). [See Removing an Expansion Module from a QFX5100 Device.] • Device recovery mode introduced with upgraded FreeBSD. [See How to Recover Junos OS with Upgraded FreeBSD.] • IPv4 support for Telnet. [See Configuring Telnet Service for Remote Access to a Switch.] • Secure boot with system security enhancement: secure bBoot. [See Software Installation and Upgrade Guide.] • Common BIOS support. • Licensing enhancements. [See Licenses for QFX Series.] • Zero touch provisioning (ZTP). [See Understanding Zero Touch Provisioning.]
Time Management	• Network Time Protocol (NTP). [See Understanding NTP Time Servers.] • Enhancement to NTP authentication method. [See Configuring NTP Authentication Keys.]
VLANs	• Configure tagged VLANs using the 802.1Q standard. [See Configuring Tagged VLANs.] • Default VLAN and multiple VLAN range support, dual VLAN tag translation, routed VLAN interfaces, and jumbo frames. • Support for 4096 VLAN IDs. [See 802.1Q VLAN IDs.] • Support to exclude RVIs from state calculations. [See Excluding a Routed VLAN Interface from State Calculations.] • Support for IRB interfaces on Q-in-Q VLANs. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.] • Static MAC address assignment for physical interface. • Support for Private VLANs and Q-in-Q on the same interface. [See Understanding Private VLANs.] • VLAN support for configuration and operational state models in Openconfig. [See OpenConfig Overview.]

Table 2: Features Supported by QFX5120-48T Switches *(continued)*

Feature	Description
To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported across all platforms, see the Hardware Compatibility Tool.	

Authentication, Authorization, and Accounting

- **802.1X authentication on Layer 3 interfaces (QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, and QFX5220)**—Starting in Junos OS Release 20.2R1, 802.1X authentication is supported on Layer 3 interfaces. The 802.1X IEEE standard for port-based network access control authenticates users attached to a LAN port. It blocks all traffic to and from a supplicant (client) at the interface until the supplicant's credentials are presented and matched on the RADIUS authentication server.

  [See 802.1X Authentication.]

EVPN

- **EVPN-VXLAN multicast support (QFX10002-60C)—**Starting in Junos OS Release 20.2R1, the QFX10002-60C switch supports the following multicast features:

  - Internet Group Management Protocol version 2 (IGMPv2) and IGMP snooping [See Overview of Multicast Forwarding with IGMP Snooping in an EVPN-VXLAN Environment.]
  - Selective multicast forwarding [See Overview of Selective Multicast Forwarding.]
  - Assisted replication [See Assisted Replication Multicast Optimization in EVPN Networks.]

  With the support of these multicast features, the QFX10002-60C switch can now perform the following:

  - Layer 2 intra-VLAN multicast forwarding
  - Layer 3 inter-VLAN multicast routing with:
    - An IRB interface running Protocol Independent Multicast (PIM)
    - A PIM gateway connected through a Layer 2 multicast VLAN (MVLAN) or a Layer 3 interface

- An external multicast router

High Availability (HA) and Resiliency

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes mastership. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

  [See Understanding the Ephemeral Configuration Database.]

Interfaces and Chassis

- **Support for 100-Gbps and 40-Gbps ports to operate at 10-Gbps or 1-Gbps speed (QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 20.2R1, you can use the Mellanox pluggable adapter (model number: MAM1Q00A-QSA) to convert quad-lane based ports to a single-lane based port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ cable connector. Use the QSA adapter to convert a 40GbE or a 100GbE port to a 10GbE or a 1GbE port. You can then plug-in an SFP+ transceiver or an SFP transceiver into the QSA adapter which is inserted into the QSFP+ or QSFP ports of the switch. You can use the commands **show chassis hardware** and **show chassis pic fpc-slot *slot-number* pic-slot *slot-number*** to view the optics inventory information for the QSFP ports.

  With this adapter, the QSFP Ports on QFX10002, QFX10008 and QFX10016 switches support the following transceiver types— 100-Mbps, 1-Gbps, 10-Gbps SFP+: SR, LR, ER, ZR, CWDM, DAC and T-SFP+.

  > **NOTE:** For this adaptor to work on the QSFP+ ports on the QFX10000-36Q line card in the QFX10008, you need to channelize the ports using the CLI command **set fpc fpc-slot pic pic-number port port-number port speed 10G**.

  [See show chassis hardware and show chassis pic.]

- **Support for multiple speeds and autonegotation (QFX5120-48Y, QFX5110-48S, and QFX5100-48S with JNP-SFPP-10GE-T transceiver)**—Starting in Junos OS Release 20.2R1, you can configure your switch to operate at multiple speeds when the JNP-SFPP-10GE-T transceiver is installed.

  On the QFX5110-48S and QFX5100-48S switches, you can configure 100-Mbps, 1-Gbps, and 10-Gbps speeds on the mge-0/0/*z* port by using the **set interfaces mge-0/0/*z* speed (100m|1g|10g)** command.

The switch ports operate at the configured speed and they can also switch to a supported lower speed (automatically) with the same transceiver installed, based on peer capability.

The QFX5120 operates at only two speeds–10 Gbps and 1 Gbps–when this transceiver is installed. By default, the switch comes up with 10-Gbps speed. To operate at 1-Gbps speed, use the **set chassis fpc 0 pic 0 port *port-number* speed 1G** command. Due to hardware limitations, you can configure the ***port-number*** value only in multiples of four, starting from port 0. You must also configure sets of four consecutive ports (for example, 0-3, 4-7, and so on) to operate at the common speed. After setting 1-Gbps speed, to revert to 10-Gbps speed, simply delete the **1G** speed configuration.

> **NOTE:** Only QFX5110-48S and QFX5100-48S switches support the multi-rate Gigabit Ethernet (mge) interface.

[See speed (Ethernet).]

Juniper Extension Toolkit (JET)

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

  [See language (Scripts), Develop Off-Device JET Applications, and Develop On-Device JET Applications.]

Junos Telemetry Interface

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

  [See Mapping OpenConfig Routing Policy Commands to Junos Configuration and Mapping OpenConfig Network Instance Commands to Junos Operation.]

- **ON-CHANGE BGP peer information statistics support for JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 provides BGP peer sensor support using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON_CHANGE statistics are sent to an outside collector.

  The following resource paths are supported:

  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/active** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes** (ON_CHANGE)

- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/received** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/sent** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/rejected** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/admin-state** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/established-transitions** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/last-established** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/received/notification** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/messages/received/update** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/notification** (stream
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/update** (stream)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/session-state** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/supported-capabilities** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/local-address** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-address** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-port** (ON_CHANGE)

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

- **EVPN statistics export using JTI (QFX5100, QFX5110, QFX5120, QFX5200, QFX10002-60C, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and using remote procedure call (gRPC) services to export EVPN statistics from devices to an outside collector.

  Use the following sensors to export EVPN statistics:

- Sensor for instance level statistics (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/**)
- Sensor for route statistics per peer (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/peer/**)
- Sensor for Ethernet segment information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/ethernet-segment/**). This includes EVPN designated forwarder ON_CHANGE leafs **esi** and **designated-forwarder**.
- Sensor for local interface information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/interfaces/**)
- Sensor for local IRB interface information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/irb-interfaces/**)
- Sensor for global resource counters and current usage (resource path **/junos/evpn/evpn-smet-forwarding/**)
- Sensor for EVPN IP prefix (resource path **/junos/evpn/l3-context/**)
- Sensor for EVPN IGMP snooping database (type 6) (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/sg-db/**)
- Sensor for EVPN IGMP join sync (type 7) ad leave sync (type 8) (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/sg-db/sgdb-esi**)
- Sensor to relate selected replicator on AR leaf on QFX5100, QFX5110, QFX5120, and QFX5200 switches (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/assisted-replication/**)
- Sensor for EVPN ON_CHANGE notifications (resource path **/network-instances/network-instance[instance-name='name']//protocols/protocol/evpn/ethernet-segment**)
- Sensor for overlay VX-LAN tunnel information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/vxlan-tunnel-end-point/**). This includes VTEP information ON_CHANGE leafs **source_ip_address, remote_ip_address, status, mode, nexthop-index, event-type** and **source-interface**.
- EVPN MAC table information (resource path **/network-instances/network-instance[instance-name='name']/mac_db/entries/entry/**)
- Sensor for MAC-IP or ARP-ND table (resource path **/network-instances/network-instance[instance-name='name']/macip_db/entries/entry/**)
- Sensor for MAC-IP ON_CHANGE table information (resource path **/network-instances/network-instance[name='name']/macip-table-info/**). Statistics include leafs **learning, aging-time, table-size, proxy-macip,** and **num-local-entries**.

- Sensor for MAC-IP ON_CHANGE entry information (resource path **/network-instances/network-instance[name='name']/macip-table/entries/entry/**). Statistics include leafs **ip-address, mac-address, vlan-id** and **vni**.
- Sensor for bridge domain or VLAN information (resource path **/network-instances/network-instance[instance-name='name']/bd/**)

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

- **CPU statistics support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 supports streaming various CPU statistics and process parameters using remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) services and Junos telemetry interface (JTI). You can stream CPU usage per process (statistics are similar to output from the **show system process detail** operational mode command), as well as CPU usage per Routing Engine core.

  This feature supports the private data model openconfig-procmon.yang.

  To stream statistics to an outside collector, include the following resource paths in a gRPC or gNMI subscription:

  - Individual process level information (resource path **/system/processes/process**)
  - Individual Routing Engine core information (resource path **/components/component/cpu/**)

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

- **Packet Forwarding Engine sensor support with INITIAL_SYNC on JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000 line of routers, QFX5100, and QFX5200)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export Packet Forwarding Engine statistics from devices to an outside collector using gNMI submode INITIAL_SYNC. When an external collector sends a subscription request for a sensor with INITIAL_SYNC (gnmi-submode 2), the host sends all supported target leaves (fields) under that resource path at least once to the collector with the current value. This is valuable because:

  - The collector has a complete view of the current state of every field on the device for that sensor path.
  - Event-driven data (ON_CHANGE) is received by the collector at least once before the next event is seen. In this way, the collector is aware of the data state before the next event happens.
  - Packet Forwarding Engine sensors that contain zero counter values (zero-suppressed) that normally do not show up in streamed data are sent, ensuring that all fields from each line card (also referred to as source) are known to the collector.

  > **NOTE:** ON_CHANGE data is not available for native (UDP) Packet Forwarding Engine Sensors.

INITIAL_SYNC submode requires that at least one copy to be sent to the collector; however, sending more than one is acceptable.

INITIAL_SYNC submode is supported for the following sensors:

- Sensor for CPU (ukernel) memory (resource path **/junos/system/linecard/cpu/memory/**)
- Sensor for firewall filter statistics (resource path **/junos/system/linecard/firewall/**)
- Sensor for physical interface traffic (resource path **/junos/system/linecard/interface/**)
- Sensor for logical interface traffic (resource path **/junos/system/linecard/interface/logical/usage/**)
- Sensor for physical interface queue traffic (resource path **/junos/system/linecard/interface/queue/**)
- Sensor for physical interface traffic except queue statistics (resource path **/junos/system/linecard/interface/traffic/**)
- Sensor for NPU memory (resource path **/junos/system/linecard/npu/memory/**)
- Sensor for NPU utilization (resource path **/junos/system/linecard/npu/utilization/**)
- Sensor for packet statistics (resource path **/junos/system/linecard/packet/usage/**)
- Sensor for software-polled queue-monitoring statistics (resource path **/junos/system/linecard/qmon-sw/**)

[See Understanding OpenConfig and gRPC and gNMI on Junos Telemetry Interface and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

Layer 2 Features

- **L2PT support (EX4650 and QFX5120-48Y switches, and QFX5100 and QFX5110 switches and Virtual Chassis)**—Starting in Junos OS Release 20.2R1, you can configure Layer 2 protocol tunneling (L2PT) to tunnel any of the following Layer 2 protocols: CDP, E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, STP (including RSTP and MSTP), UDLD, VSTP, and VTP.

  [See Layer 2 Protocol Tunneling.]

Multicast

- **Static multicast route leaking for VRF and virtual router instances (EX4650 and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can configure the switch to statically share (leak) IPv4 multicast routes for IGMPv3 (S,G) traffic among different virtual router or virtual routing and forwarding (VRF) instances. You can only leak static multicast routes per group, not per source and group. The destination prefix length must be 32.

  To configure multicast route leaking to the VRF or virtual router instance *routing-instance-name*, configure the **next-table *routing-instance-name*.inet.0** statement at the **[edit routing-instances *routing-instance-name* routing-options static route destination-prefix/32]** hierarchy level.

  [See Understanding Multicast Route Leaking for VRF and Virtual Router Instances.]

- **Multicast-only fast reroute (MoFRR) (EX4650 and QFX5120-48Y)**—Starting in Junos OS Release 20.2R1, you can configure MoFRR to minimize multicast packet loss in PIM domains when link failures occur. With MoFRR enabled, the switch maintains primary and backup traffic paths, forwarding traffic from the primary path and dropping traffic from the backup path. If the primary path fails, the switch can quickly start forwarding the backup path stream (which becomes the primary path). The switch creates a new backup path if it detects available alternative paths. MoFRR applies to all multicast (S,G) streams by default, or you can configure a policy for the (S,G) entries where you want MoFRR to apply.

  [See Understanding Multicast-Only Fast Reroute.]

Network Management and Monitoring

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

  [See NETCONF Sessions over Outbound HTTPS.]

Routing Policy and Firewall Filters

- **Support for MPLS firewall filter on loopback interface (EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can apply an MPLS firewall filter to a loopback interface on a Label switching router (LSR). For example, you can configure an MPLS packet with **ttl=1** along with MPLS qualifiers such as **label**, **exp**, and Layer 4 **tcp/udp** port numbers. Supported actions include **accept**, **discard**, and **count**.

  You configure this feature at the **[edit firewall family mpls]** hierarchy level. You can only apply a loopback filters on **family mpls** in the ingress direction.

  [See Overview of MPLS Firewall Filters on Loopback Interface.]

Virtual Chassis

- **Virtual Chassis with NSSU support (QFX5120-48T)**—Starting in Junos OS Release 20.2R1, you can interconnect two QFX5120-48T switches into a Virtual Chassis that operates as one logical device managed as a single chassis. The Virtual Chassis:

- Has both switches in Routing Engine role (one master and one backup)
- Supports 100GbE QSFP28 or 40GbE QSFP+ ports (48 through 53) as Virtual Chassis ports (VCPs)
- Supports NSSU

A QFX5120-48T Virtual Chassis supports the same protocols and features as a standalone switch in Junos OS Release 20.2R1 except for the following:

- EVPN-VXLAN
- Junos telemetry interface (JTI)
- Multichassis link aggregation (MC-LAG)
- Priority-based flow control (PFC)

Configuration parameters and operation are the same as for other non mixed QFX Series Virtual Chassis.

[See Virtual Chassis Overview for Switches.]

- **802.1X authentication, Layer 2 port security, and MPLS support in a Virtual Chassis (QFX5120-48Y Virtual Chassis)**—Starting in Junos OS Release 20.2R1, the following protocol features are supported on a QFX5120-48Y Virtual Chassis:
  - IEEE 802.1X authentication
  - Layer 2 port security features, including IP source guard, IPv6 router advertisement (RA) guard, DHCP, and DHCP snooping
  - MPLS

  Configuration and operation are the same on the Virtual Chassis as on the standalone switch.

  [See 802.1X Authentication, MPLS Overview, DHCP Snooping, Understanding DHCP Snooping (ELS), Understanding IP Source Guard for Port Security on Switches, and Understanding IPv6 Router Advertisement Guard.]

SEE ALSO

What's Changed | 211

Known Limitations | 213

Open Issues | 215

Resolved Issues | 221

Documentation Updates | 226

Migration, Upgrade, and Downgrade Instructions | 227

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for QFX Series Switches.

Class of Service

- We've corrected the output of the **show class-of-service interface | display xml** command. Output of the following sort: **<container><leaf 1> data <leaf 2> data <leaf 3> data <leaf 1> data <leaf 2> data <leaf 3>** data now appears correctly as: **<container> <leaf 1> data <leaf 2> data <leaf 3> data <container> <leaf 1> data <leaf 2> data <leaf 3> data**.

General Routing

- Priority-based flow control (PFC) support (QFX5120-32C)—Starting with JunosOS 19.2R3, QFX5120-32C switches support priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic.
- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

  [See commit (System).]

Interfaces and Chassis

- **Autonegotiation status displayed correctly (QFX5120-48Y)**—In Junos OS Release 20.2R1, the **show interfaces *interface-name* <media> <extensive>** command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed.

  In the earlier Junos OS releases, incorrect autonegotiation status was displayed even when autonegotiation was disabled.

Junos Extension Toolkit

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

  [See Develop Off-Device JET Applications and Develop On-Device JET Applications.]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

  [See Juniper EngNet.]

- **IGMP snooping in EVPN-VXLAN multihoming environments (QFX5110)**—In an EVPN-VXLAN multihoming environment on QFX5110 switches, you can now selectively enable IGMP snooping only on those VLANs that might have interested listeners. In earlier releases, you must enable IGMP snooping on all VLANs associated with any configured VXLANs because all the VXLANs share VXLAN tunnel endpoints (VTEPs) between the same multihoming peers and require the same settings. This is no longer a configuration limitation.

Network Management and Monitoring

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

SEE ALSO

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Release 20.2R1 for QFX Series Switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Class of Service (CoS)

- On the QFX5100 line of switches, due to major third-party SDK upgrade in Junos OS Release 20.1R1 (from SDK 6.3.7 to 6.5.16), unified ISSU is not supported from any earlier releases to Junos OS Release 20.1 (image : jinstall-qfx-5-*). PR1479439

General Routing

- On the QFX5120 line of switches, one of the VCP ports of the throughput test result for most of the frame sizes is not close to 100 percent. PR1453709
- Convergence delay for link-protected MPLS LSP is more than 50ms. PR1478584
- During software validation Junos OS mounts the new image and validates the configuration against the new image. Since the TVP-based QFX5000 and QFX10000 line of switches are mounting the maximum 4 disks during normal execution, it cannot mount the extra disk for this purpose. Thus, QFX Series line of switches currently does not support configuration validation during upgrade on QFX5000 resulting in the syntax error when the image installation is triggered with "validation". PR1479753
- No option to upgrade firmware for the backup Routing Engine. PR1479925
- On a standalone device, the output of the **show snmp mib walk jnxFruName** command has an extra entry for Routing Engine. PR1483384
- After multiple/frequent GRES events on 2 member QFX5120 VC, **show chassis alarms** statement output shows PEM status incorrect for VC members. PR1486736
- In QFX100002, traffic drop during FRR might not be guaranteed to 50ms all the time. PR1486853
- Observing 100 percent L2 MAC scaling traffic loss in QFX10002-60C platform after loading EVPN-VXLAN collapsed profile configurations. PR1489753
- Abrupt power cycles is a disruptive action for storage device. There can be input and output events happening at any point of time and software will be unaware with a sudden power cycle and that could lead to file corruption. PR1507750

Layer 2 Ethernet Services

- If configuration or image file name has nonallowed special characters (like #%@) in it, ZTP over HTTP/HTTPS might not work. When HTTP/HTTPS URL is formed to download the file, the URL contains file name in it. HTTP/HTTPS does not expect any special characters in the URL. If special characters are present, the HTTP/HTTPS protocol returns "Bad request". In order to avoid the issue, do not use any nonallowed special characters in the filename. PR1503588

SEE ALSO

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.2R1 for QFX Series Switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Class of Service (CoS)

- The priority-based flow control (PFC) feature is not supported on 2-member Virtual Chassis currently because of the hardware limitation. PR1431895

EVPN

- EVPN-VXLAN : L2ald generates a core file at **l2ald_mem_free** while changing configurations in the DUT. PR1511165
- In all platforms with VXLAN static VTEP tunnels scenario (including static VXLAN without EVPN), creating a new VTEP interface might not work after the Routing Engine switchover or after restarting the l2-learning. PR1520078

General Routing

- QFX10000 platform drops the Aruba wireless access point (AP) heartbeat packets. As a result, the Aruba wireless AP cannot work. PR1352805
- USB upgrade of NOS image is not supported. PR1373900
- The **show chassis fpc** command displays an incorrect amount of available memory on a QFX10000 FPCs. PR1394978
- On QFX10000 Series platform, CPU overuse on PFC might be observed if the adaptive feature is enabled to load-balance for an aggregated Ethernet interface. PR1399369
- The **show chassis fpc** reports high CPU utilization in the **Steady** state. PR1492731
- With WRL7 on QFX5000 devices there is a possibility in reboot scenario the system going to DB prompt. This is due to a known issue in the QEMU version in WRL7. PR1411826
- On QFX5110 and QFX5120 platforms, unicast RPF check in strict mode might not work properly. PR1417546
- IPv6 neighbor solicitation packets for link-local address might be dropped when passing through QFX10002-60C via IRB interface. As a result, hosts inside VLANs could not communicate with each other using link-local addresses. PR1424244
- The issue occurs because of a PECHIP limitation when underlay is tagged. After de-encapsulation when the inner packet is recirculated it still retains the VLAN tag property from outer header because the outer header was tagged. Thus 4 bytes of inner tag got overwritten in the inner packet and the packet got corrupted, which will result in EGP chksum trap seen in PECHIP. Fixing PECHIP limitation in software has high risk. As a workaround, enable **encapsulate-inner-vlan** configuration. PR1435864
- The unified ISSU is not supported on QFX5200 switches and fails from Junos OS Release 17.2X75-D43.2 to some target versions. Also, dcpfe crash might be seen. PR1438690

- On QFX10000 platforms, in an EVPN-VXLAN (spine-leaf) scenario, the QFX10000 spine switches are configured with VXLAN Layer 3 gateway (utilizing the virtual gateway) on an IRB interface. If you enable and then subsequently remove the VXLAN Layer 3 gateway on this IRB interface on one or some of these spine switches, traffic drop might be observed. As a workaround, configure all virtual gateways with unique IPv4 or IPv6 MAC address. PR1446291
- On the Junos OS platforms with next-generation Routing Engine installed, the process vehostd might crash without generating a core file and automatic restart of vehostd might fail. The vehostd is a daemon for managing the life cycle of system-critical Junos OS VMs in the system. If the process vehostd gets in crash state, it will impact the management of Junos OS VMs. PR1448413
- On the QFX5000 line of switches, misleading ISSU logs are printed during the NSSU process even when the box does not perform ISSU. PR1451375
- Whenever any member in a remote Switch Port Analyzer (RSPAN) VLAN is removed from that VLAN, you must reconfigure the analyzer session for that RSPAN VLAN. PR1452459
- In overall commit time, the evaluation of mustd constraints is taking 2 seconds more than usual. This is because the persist-group-inheritance feature has been made a default feature in the latest Junos OS releases. Eventually, this feature helps improve the subsequent commit times for scaled configurations significantly. The persist-group-inheritance feature is useful in customer scenarios where groups and nested groups are used extensively. In those scenarios, the group inheritance paths are not built every time, thus subsequent commits are faster. PR1457939
- VXLAN VNI (multicast learning) scaling on QFX5110 traffic issue is seen from VXLAN tunnel to L2 interface. PR1462548
- "entPhysicalTable" MIB is not fetching expected data on QFX10002-72Q and QFX10002-36Q platforms. PR1462582
- On QFX5100 device, interface output counter is double counted for self-generated traffic. PR1462748
- The output of the **show chassis environment** command can be seen from backup members as well. The issue is common to all QFX Series platforms. PR1474520
- Dynamic IP-IP tunnels and filter-based IP-IP de-encapsulation filter on loopback interface cannot coexist together. If dynamic IP-IP tunnels were configured earlier, then FPC needs a reboot before it can be used for loopback IP-IP de-encapsulation filter. Also, the loopback interface might contain implicit filters. If these implicit filters get hit, the de-encapsulation filter might not get hit. PR1479613
- The **app-engin** configuration command does not show information for the backup member. PR1479900
- Instead of the FAN status, FPC status is checked and updated in JTI. PR1480259
- Redirects are used when a router determines that a packet is being routed suboptimally and seeks to inform the sending host that it should forward subsequent packets to that same destination through a different gateway. For QFX5110 and QFX5120, ICMP redirect message won't be generated in such cases. PR1481020

- The dcpfe process does not come up in some instances when the QFX5120 is abruptly powered off and then powered on. As a workaround, doing a power-cycle of the device or host reboot might recover the device. PR1481176
- On QFX Series platforms running Junos VM instance (excluding QFX10000 Series platforms), the laser signal might be transmitted on the disabled interfaces with QSFP and QSFP28 optics after device reboot. PR1487554
- On QFX10002 switches with MC-LAG configurations, traffic drops when you deactivate or activate physical interface trigger. PR1488166
- Commit fails on backup device of QFX5120-48T Virtual Chassis while removing the storm control with high availability (HA) configured. A warning is seen because the patch removes a statement that is not empty. PR1488847
- When the NETCONF session is established over outbound SSH, the high rate of pushing the configuration to the ephemeral DB might result in outbound SSH connection flap or a memory leak issue. PR1497575
- On QFX10008 platforms, if the BFD is configured over an aggregated Ethernet interface (member link across multiple FPCs), deactivating/activating the aggregated Ethernet interface or executing GRES might cause the BFD sessions to flap. PR1500798
- On QFX5100, ERPS might not work correctly on branch which as 1473610 fix, due to stp instance programming failure in hardware. PR1500825
- LLDP packets are not acquired when **native-vlan** configured is same as tagged vlan-id. PR1504354
- After deleting and adding logial switches on NSX-V setup repeatedly along with ovsdb configured, ping between VM to baremetal server fails intermittently. PR1506097
- On QFX5100, fxpc crash might be seen sometimes while installing image through ZTP. PR1508611
- Junos telemetry interface sensor **/junos/system/linecard/optics/** giving incorrect values for lane_laser_receiver_power_dbm, lane_laser_output_power_dbm while testing the diagnostic optics output. PR1509771
- Disruptive switchover (no GRES or NSR configured) can lead to stale PPM entries programmed on the new master Routing Engine and BFD sessions to remain down. PR1518106

High Availability (HA) and Resiliency

- The QFX5200-32C reboot time is degraded. A flush cache issue is seen because of the reliable SSD disk input/output change made for this platform. PR1511607

Infrastructure

- Device goes to db prompt with "panic: ffs_valloc: dup alloc" when you power on the device. It is recommended to run "fsck" since this is caused because of the FS mount failure. PR1480185

Interfaces and Chassis

- The same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface was assigned with the identical address after commit. There is no warning during the commit, only syslog messages indicating incorrect configuration. PR1221993
- Multicast traffic can be flooded for 15 to 20 seconds to both MC-LAG peers, after the following sequence of steps: 1. Disable or enable ICL. 2. Reboot one of MC-LAG peers. 3. Disable or enable a member link of ICL. This results in no traffic loss, and one of the MC-LAG nodes processes duplicate packets during this time period. PR1422473

Layer 2 Ethernet Services

- If forward-only is set within dhcp-reply in a Juniper Networks device as a DHCP relay agent, the DHCP DECLINE packets that are broadcasted from the DHCP client are dropped and not forwarded to the DHCP server. PR1429456

Layer 2 Features

- On QFX5120, during new tenant addition, there might be few transient packet drops (2 - 15 packets) for a couple of random intra-VNI traffic streams in an EVPN-VXLAN topology for the existing tenants. The drop is almost negligible and is automatically recovered. PR1455654
- On QFX5110 and QFX5120 platforms, changing lo0 IP address might sometimes result either in stale entry of IP in mpls_entry table or missing IP entry, which results in traffic drop for VXLAN traffic. PR1472333

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. PR1054798
- If interface is newly added as CE interface, existing bum traffic can be looped. Loop prevention features is designed to start working whenever new CE interface is added by configuration. But existing bum traffic can be distributed to new CE interface earlier than enabling of loop prevention feature. PR1493650

Routing Protocols

- If DDoS protection is disabled on QFX5100 Virtual Chassis and multicast traffic is being sent, the Virtual Chassis might become unstable, with high CPU usage and it might crash eventually, creating FXPC core files. Disabling DDoS protection will disable rate limiting for all host-bound traffic. We do not recommend disabling DDoS protection on the device, because, a high amount of control traffic can overwhelm the system, causing system instability. PR1238875
- On QFX5100 Virtual Chassis or Virtual Chassis fabric, when the **mini-PDT-base** configuration is issued, the following error message is seen in the hardware: **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:l3 nh 6594 unintsall failed**. There is no functionality impact because of this error message. PR1407175
- BGP route addition and deletion time and BGP, OSPF, and IS-IS link flap convergence time are increased in Junos OS Release 19.4 (forwarding plane). PR1464572
- With the **egress-to-ingress** configuration statement, the customer will not be able to configure 2000 scale and the scale is reduced to 1000. PR1514570
- Deleting the physical interface under an aggregated Ethernet interface might flap the BFD sessions formed on the remaining physical interfaces if same IP address is configured on multiple units in different routing-instances. PR1516556

Virtual Chassis

- On QFX5000 Virtual Chassis, DDoS violations that happen on the backup are not reported to Routing Engine. PR1490552

SEE ALSO

What's New | **186**

What's Changed | **211**

Known Limitations | **213**

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in Junos OS main and maintenance releases for QFX Series Switches.

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

Resolved Issues: 20.2R1

EVPN

- The ESI of IRB interfaces does not update after autonomous-system number change if the interface is down. PR1482790
- QFX10002-60C EVPN/VXLAN multicast: The **show** command issued for the VTEP interface did not show mesh-group id. PR1498052
- The VXLAN function might be broken due to a timing issue. PR1502357

Class of Service (CoS)

- Traffic might be forwarded to an incorrect queue when fixed classifier is used. PR1510365

General Routing

- The following error message is generated while booting: **CMQFX: Error requesting SET BOOLEAN, illegal setting 66**. PR1385954
- The configuration statement **show chassis errors active detail** is not supported for QFK5000 platforms. PR1386255
- The 10G fiber interfaces might flap frequently when they are connected to other vendor's switch. PR1409448

- The statement **show interface** indicates Media type: Fiber on QFX5100-48T running '-qfx-5e-' Junos OS image. PR1419732
- A vmcore is seen on QFX Series Virtual Chassis. PR1421250
- SFP-LX10 stay down until autonegotiate is disabled. PR1423201
- The default logical interfaces on channelized physical interfaces might not be created after ISSU/ISSR. PR1439358
- CRC error might be seen on the VCPs of the QFX5100 Virtual Chassis. PR1449406
- On QFX5000 no warning or error is shown when dual VLAN tag feature is configured on physical interface. PR1450455
- Members might stay disconnected from a QFX5120-32C and QFX5120-48T Virtual Chassis after a full-stack reboot. PR1453399
- Changing the VLAN name associated with access ports might prevent MAC addresses from being learned in an EVPN-VXLAN scenario. PR1454095
- The cosd crash might be observed if forwarding-class-set is directly applied on the child interface of an aggregated Ethernet interface. PR1455357
- Telemetry traffic might not be sent out when the telemetry server is reachable through a different routing instance. PR1456282
- Link up delay and traffic drop might be seen on mixed SP L2/L3 and EP L2 type configurations. PR1456336
- QFX5110 QSFP-100GBASE-SR4 made by the third party cannot link up. PR1457266
- An FPC might restart during runtime on the QFX10000 line of devices. PR1464119
- EPR iCRC errors in QFX10000 platforms might cause protocols to go down. PR1466810
- A few of DHCP INFORM packets specific to a particular VLAN might be taking the wrong resolve queue. PR1467182
- Traffic loss might be seen with framing errors or runts if MACsec is configured on EX4600/QFX5100 platforms. PR1469663
- The speed 10m might not be configured on the GE interface. PR1471216
- The traffic loss might occur when VTEP source interface is configured in multiple routing instances. PR1471465
- Egress ACL filter entries will be only 512 in Junos OS Release 19.4R1 on QFX5000. PR1472206
- The shaping of CoS does not work after reboot. PR1472223
- DSCP marking might not work as expected if the fixed classifiers are applied to interfaces on QFX5000/EX4600 platforms. PR1472771
- The detached interface in LAG might process the xSTP BPDUs. PR1473313

- On QFX5000, the **global-mac-table-aging-time** statement behavior with multi-homed EVPN-VXLAN ESI. PR1473464
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. PR1473610
- The RIPv2 packets forwarded across a L2 circuit connection might be dropped. PR1473685
- Continuous error log messages might be raised on QFX5000 platforms in EVPN/VXLAN scenario. PR1474545
- L2 circuit might fail to communicate through VLAN 2 on QFX5000 platforms. PR1474935
- On QFX Series platforms the system might stop new MAC learning and have impact on Layer 2 traffic forwarding. PR1475005
- DAC cables are not being properly detected in Packet Forwarding Engine in QFX5200. PR1475249
- There might be a traffic drop on QFX5110 and QFX5120 switches acting as leaf switches in a multicast environment with VXLAN. PR1475430
- FPC major error is seen after system boot up or FPC restart. PR1475851
- QFX Series platforms are exhibiting invalid Packet Forwarding Engine PG counter pairs to copy, src 0xfffff80, dst 0. PR1476829
- Continuous error logs on the device: **prds_ptc_wait_adoption_status: PECHIP[1] PTC[1]: timeout on getting adoption valid bit[8] asserted**. PR1477192
- The default Virtual Chassis MAC persistence timer is incorrectly set to 20 seconds instead of 20 minutes. PR1478905
- The remaining interface might be still in down state even though the number of channelized interfaces is no more than 5. PR1480480
- ARP request packets for unknown host might get dropped in remote PE device in EVPN-VXLAN scenario. PR1480776
- On QFX10000 and QFX5000, in SP style configuration, BUM traffic incorrectly gets blocked, while disabling or enabling a different logical interface. PR1482202
- On QFX5110, whenever the autonegotation is toggled on the interface, explicitly set the link-mode as well as the speed for the configuration to take effect. PR1484715
- The dcpfe core file might be seen with non-oversubscribed mode. PR1485854
- The 10GbE VCP ports will not be active in a QFX5100 Virtual Chassis scenario. PR1486002
- Virtual Chassis ports might go down in a mixed Virtual Chassis setup of QFX5100-24Q-2P/EX4300 and EX4600/EX4300. PR1489985
- After ISSU/ISSR, a port using SR4/LR4 optics might not come up. PR1490799
- BFD sessions start to flap when the firewall filter in the loopback0 is changed. PR1491575
- Traffic loss could be observed in a mixed Virtual Chassis setup of QFX5100 and EX4300. PR1493258

- Traffic loss could be seen in a MC-LAG scenario on QFX5120/EX4650. PR1494507
- SNMP polling for CPU utilization and CPU state of backup Routing Engine does not show in a two-member Virtual Chassis. PR1495384
- ARP do not get refreshed after timeout on QFX10002-60C. PR1497209
- Extra carrier transitions are seen on the peer when negative triggers are performed on QFX5100 and QFX5110. PR1497380
- An lcmd core file might be generated on QFX52100-64C. PR1497947
- Traffic might get dropped if aggregated Ethernet member interface is deleted and then added or a SFP of the aggregated Ethernet member interface is unplugged/plugged. PR1497993
- On QFX5210, unexpected behavior is seen for Port LED after upgrade. PR1498175
- Inter-VNI/VRF and intra-VNI/VRF traffic is dropped between the CE devices when the interfaces connected between TOR and multihomed PE devices are disabled. PR1498863
- The l2cpd crash might be seen while adding or deleting ERP configuration and then restarting l2cpd. PR1505710
- ARP replies might be flooded through the EVPN-VxLAN network as unknown unicast ARP reply. PR1510329

High Availability (HA) and Resiliency

- Unified ISSU will not be supported for QFX5000 for some versions. PR1472183

Interfaces and Chassis

- The MC-LAG configuration-consistency ICL-config might fail after committing some changes. PR1459201
- Executing commit might hang up because dcd process gets stuck. PR1470622
- Commit error is not thrown when member link is added to multiple aggregation group with different interface specific options. PR1475634
- MC-LAG consistency check fails if multiple IRB units are configured with the same VRRP group. PR1488681
- Error message is not getting generated while verifying GRE limitation. PR1495543

Junos Fusion for Enterprise

- Loop detection might not work on extended ports in Junos fusion scenarios. PR1460209

Layer 2 Ethernet Services

- EVPN-VXLAN ERB - dhcp relay-source lo0.1 is not used when enabled with anycast legacy IRB. PR1455076
- Member links state might be asychronized on a connection between PE and CE devices in an EVPN A/A scenario. PR1463791

- Issues with DHCPv6 relay processing confirm and reply packets. PR1496220

Layer 2 Features

- MAC learning might not work correctly on QFX5120. PR1441186
- The LLDP function might fail when a Juniper Networks device connects to a non-Juniper one. PR1462171
- A few MAC addresses might be missing from the MAC table in software on QFX5000 platform. PR1467466
- On QFX5120 switches QinQ, the third VLAN tag is not pushed onto the stack and SWAP is being done instead. PR1469149
- Traffic might be affected if composite next hop is enabled. PR1474142
- On QFX5200, MAC learning rate is degraded by 88 percent. PR1494072

MPLS

- Traffic might silently get dropped or discarded on the PE device when the CE device sends traffic to the PE device and the destination is resolved with two LSPs through one upstream interface. PR1475395
- The traffic might be lost over QFX5100 switch acting as a transit PHP node in the MPLS network. PR1477301
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. PR1493431

Platform and Infrastructure

- The SLAX script might be lost after upgrading software. PR1479803
- Traceroute monitor with mtr version v.69 shows a false 10 percent loss. PR1493824

Routing Protocols

- OSPF VRF sessions take a long time to come up when the host table is full and host routes are in LPM table. PR1358289
- BGP IPv4 or IPv6 convergence and RIB install/delete time degraded in Junos OS Release 19.1R1 and later mainline releases. PR1414121
- PIM (S,G) joins can cause MSDP to incorrectly announce source-active messages in some cases. PR1443713
- CRC errors might be seen on QFX5100 Virtual Chassis. PR1444845
- The core files might occur during adding or removing EVPN Type 5 routing instance. PR1455547
- [pfe_loadbalance] [pfeloadtag] flows not falling back to single link when inactivity-interval is set higher than IFG. PR1471729
- Traffic might not be forwarded over ECMP link in EVPN-VXLAN scenario. PR1475819
- ARP packets are always sent to CPU regardless of whether the storm-control is activated. PR1476708

- GRE transit traffic is not forwarded in VRRP scenario. PR1477073
- MUX State in LACP interface does not go to "collecting and distributing" and remains attached after enabling the ae interface. PR1484523
- FPC might go to "NotPrsnt" state after upgrading with non-QFX5100-24Q image in a Virtual Chassis/Virtual Chassis fabric setup. PR1485612
- CPU port queue gets full due to excessive pause frames being received on interfaces. This causes control packets from the CPU to all ports to be dropped. PR1487707
- The BGP route-target family might prevent RR from reflecting L2 VPN and L3 VPN routes. PR1492743
- The rpd might crash on QFX10000 due to rpd resolver problem of INH. PR1494005
- Firewall filter might not work in certain conditions under Virtual Chassis setup. PR1497133
- Traffic drop might be observed after modifying FBF firewall filter. PR1499918
- Change in x-path output for value "input-updates" in **show bgp neighbors**. PR1504399

SEE ALSO

- What's New | **186**
- What's Changed | **211**
- Known Limitations | **213**
- Open Issues | **215**
- Documentation Updates | **226**
- Migration, Upgrade, and Downgrade Instructions | **227**

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for the QFX Series Switches.

SEE ALSO

- What's New | **186**
- What's Changed | **211**
- Known Limitations | **213**
- Open Issues | **215**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

Upgrading Software on QFX Series Switches

When upgrading or downgrading Junos OS, always use the jinstall package. Use other packages (such as the jbundle package) only when so instructed by a Juniper Networks support representative. For information about the contents of the jinstall package and details of the installation process, see the Installation and Upgrade Guide and Junos OS Basics in the QFX Series documentation.

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to https://www.juniper.net/support/downloads/junos.html.

   The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.

3. Select **20.2** in the Release pull-down list to the right of the Software tab on the Download Software page.

4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.2 release.

   An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

   A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Download the software to a local host.

8. Copy the software to the device or to your internal software distribution site.

9. Install the new jinstall 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.

   Customers in the United States and Canada use the following command:

   `user@host>` **request system software add *source*/jinstall-host-qfx-5-x86-64-20.2-R1.n-secure-signed.tgz reboot**

   Replace ***source*** with one of the following values:

   - ***/pathname***—For a software package that is installed from a local directory on the switch.
   - For software packages that are downloaded and installed from a remote location:
     - **ftp://*hostname*/*pathname***
     - **http://*hostname*/*pathname***
     - **scp://*hostname*/*pathname*** (available only for Canada and U.S. version)

   Adding the **reboot** command reboots the switch after the upgrade is installed. When the reboot is complete, the switch displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 20.2 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

> **NOTE:** The QFX10002-60C switch supports only the 64-bit version of Junos OS.

> **NOTE:** If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add *<pathname><source>*** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-20.2R1.9.tgz
```

If the Install Package resides remotely from the switch, execute the **request vmhost software add *<pathname><source>*** command.

For example:

```
user@switch> request vmhost software add
ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-20.2R1.9.tgz
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R1.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add *<pathname><source>* reboot** command.

For example:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.2R1.n-secure-signed.tgz reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* reboot** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-20.2R1.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add *<pathname><source>*** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

   ```
   user@switch> configure
   ```

3. Disable Routing Engine redundancy:

   ```
   user@switch# delete chassis redundancy
   ```

4. Disable nonstop-bridging:

   ```
   user@switch# delete protocols layer2-control nonstop-bridging
   ```

5. Save the configuration change on both Routing Engines:

   ```
   user@switch# commit synchronize
   ```

6. Exit the CLI configuration mode:

   ```
   user@switch# exit
   ```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

   ```
   user@switch> request system software add validate
   /var/tmp/jinstall-host-qfx-10-f-x86-64-20.2R1.n-secure-signed.tgz
   ```

   For more information about the **request system software add** command, see the CLI Explorer.

9. Reboot the switch to start the new software using the **request system reboot** command:

   ```
   user@switch> request system reboot
   ```

   > **NOTE:** You must reboot the switch to load the new installation of Junos OS on the switch.
   >
   > To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete *<package-name>*** command. This is your last chance to stop the installation.

   All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

   While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

    ```
    user@switch> show version
    ```

    Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

    For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

    ```
    user@switch> request chassis routing-engine master switch
    ```

    For more information about the **request chassis routing-engine master** command, see the CLI Explorer.

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

    ```
    user@switch> show chassis routing-engine
    Routing Engine status:
      Slot 0:
        Current state                  Backup
        Election priority              Master (default)
    Routing Engine status:
      Slot 1:
        Current state                  Master
        Election priority              Backup (default)
    ```

14. Install the new software package using the **request system software add** command:

    ```
    user@switch> request system software add validate
    /var/tmp/jinstall-host-qfx-10-f-x86-64-20.2R1.n-secure-signed.tgz
    ```

    For more information about the **request system software add** command, see the CLI Explorer.

15. Reboot the Routing Engine using the **request system reboot** command:

```
user@switch> request system reboot
```

> **NOTE:** You must reboot to load the new installation of Junos OS on the switch.
>
> To abort the installation, do not reboot your system. Instead, finish the installation and then issue the **request system software delete jinstall *<package-name>*** command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the **show version** command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the CLI Explorer.

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
outing Engine status:
  Slot 1:
    Current state                  Backup
    Election priority              Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- Preparing the Switch for Software Installation on page 238
- Upgrading the Software Using Unified ISSU on page 238

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

  To verify that nonstop active routing is enabled:

  NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

  ```
  user@switch> show task replication
          Stateful Replication: Enabled
          RE mode: Master
  ```

  If nonstop active routing is not enabled (**Stateful Replication** is **Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the **request system snapshot** command.

Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the **/var/tmp** directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
   - On the switch, enter:

     ```
     user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
     ```

     where ***package-name*.tgz** is, for example, **jinstall-host-qfx-10-f-x86-64-20.1R1.n-secure-signed.tgz**.

     NOTE: During the upgrade, you cannot access the Junos OS CLI.

     The switch displays status messages similar to the following messages as the upgrade executes:

     ```
     warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
     lost!
     ISSU: Validating Image
     ISSU: Preparing Backup RE
     Prepare for ISSU
     ISSU: Backup RE Prepare Done
     Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
     Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz completed
     Spawning the backup RE
     Spawn backup RE, index 0 successful
     GRES in progress
     GRES done in 0 seconds
     Waiting for backup RE switchover ready
     GRES operational
     Copying home directories
     Copying home directories successful
     Initiating Chassis In-Service-Upgrade
     Chassis ISSU Started
     ISSU: Preparing Daemons
     ISSU: Daemons Ready for ISSU
     ```

```
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff
```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at **/var/log/vjunos-log.tgz**.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

SEE ALSO

What's New | **186**

What's Changed | **211**

Known Limitations | **213**

Open Issues | **215**

Resolved Issues | **221**

Documentation Updates | **226**

Junos OS Release Notes for SRX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.2R1 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for SRX Series devices.

Application Security

- **AppQoE multihoming with active/active deployment (NFX150, NFX250, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and vSRX)**—Starting In Junos OS Release 20.2R1, AppQoE is enhanced to support multihoming with active/active deployment. Previously, AppQoE supported multihoming with active/standby deployment.

  In active/active deployment, the spoke device connects to multiple hub devices. Application traffic can transit through any of the hub devices if the link to the hub device meets SLA requirements. Application traffic can switch seamlessly between the hub devices in case of service-level agreement (SLA) violation or the active hub device is not responding.

  To support active/active mode, you must enable the BGP multipath to allow the device to select multiple equal-cost BGP paths to reach a given destination.

  [See Application Quality of Experience (AppQoE).]

- **Packet capture of unknown application traffic (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.2R1, we've added new capability to your security device that allows you to capture unknown application traffic.

  Once you have configured the packet capture options on your security device, the unknown application traffic information is gathered and stored on the device in a packet capture file (**.pcap**). You can use the packet capture of an unknown application to define a new custom application signature. You can use this custom application signature in a security policy to manage the application traffic more efficiently.

  You can also send the **.pcap** file to Juniper Networks in cases where the traffic is incorrectly classified, or to request for the creation of an application signature.

  [See Application Identification.]

- **Application Quality of Experience (SRX4600)**—Starting in Junos OS Release 20.2R1, the SRX4600 supports AppQoE functionality. AppQoE enhances the user experience at the application level by monitoring the performance of business-critical applications. Based on the score, AppQoE selects the best possible link for that application traffic to meet performance requirements specified in the service-level agreement (SLA).

  The SRX4600 supports AppQoE in both the hub-and-spoke and the full mesh topologies.

  AppQoE support is already available on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and vSRX.

  [See Application Quality of Experience.]

Authentication and Access Control

- **Support to view user identify information in JIMS Active Directory (SRX Series)**— Starting in Junos OS Release 20.2R1, you can search and view user identity information such as logged users, connected devices and group list from Juniper Identity Management Service (JIMS) and Active Directory (AD) domain. The SRX Series device relies on JIMS to obtain user identity information.

  You can search the user identity information and validate the authentication source to provide access to the device. You can request JIMS to retrieve the group list for the Active Directory domain for identity information of an individual user.

  [See Configure Juniper Identity Management Service to Obtain User Identity Information.]

Flow-Based and Packet-Based Processing

- **NG-IOC cache increased (SRX4600, SRX5000 line of devices)**—Starting in Junos OS Release 20.2R1, we have increased the number of hash table entries for IOC3 from 2 million to 20 million wings, for IOC4 from 2 million to 10 million wings on SRX5000 line of devices and for IOC on SRX4600 from 2 million to 5 million wings.

  [See Express Path.]

General Packet Radio Switching (GPRS)

- **Support for Must-IE check and IE removal for GTPv1 and GTPv2 (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Release 20.2R1, Junos OS supports the following information element (IE) enforcement functions for GTPv1 and GTPv2:
  - **Must-IE check**: Use this function to check for the presence of IEs in GTPv1-C and GTPv2-C messages that helps to verify message integrity. The device check for the presence of Must-IEs of specific GTP messages and forwards the messages only if Must-IEs are present.

- **IE removal**: Use this function to remove IEs from GTPv1-C and GTPv2-C. This function helps to retain interoperability between Second-Generation Partnership Project (2GPP) and Third-Generation Partnership Project (3GPP) networks.

[See Example: Configure Must-IE check for GTPv1 and GTPv2, and Example: Configure IE removal for GTPV1 and GTPv2.]

Intrusion Detection and Prevention (IDP)

- **Policy-based threat profile for IDP (SRX Series)**—Starting from Junos OS Release 20.2R1, you can configure IDP rules with threat profiles to define attacker IP and target IP feeds.

  When traffic matches the feed data, IDP provides feed update to add the IP information in the Security Intelligence (SecIntel) module.

  This feature allows the SRX Series device to identify threats, and propagate intelligence for real-time enforcement and provides the ability to perform endpoint classification.

  [See IDP Policy Rules and IDP Rule Bases, security-intelligence, and Encrypted Traffic Analysis Overview.]

- **Signature Language Constructs (SRX Series)**—Starting in Junos OS 20.2R1, the following signature language constructs are supported in the IDP engine code to write more efficient signatures that help reduce false attacks:

  - Byte extract
  - Byte test
  - Byte jump
  - Byte math
  - Is-data-at
  - Detection filter

  [See IDP Signature Language Enhancements.]

Junos Telemetry Interface

- **Packet Forwarding Engine and Routing Engine sensor support on JTI (SRX5400, SRX5600, and SRX5800)**—Junos OS Release 20.2R1 provides streaming support for revenue interface statistics through Packet Forwarding Engine (PFE) sensors and pseudo interface statistics through Routing Engine sensors. Sensors are supported through Junos telemetry interface (JTI) and remote procedure calls (gRPC) or gRPC Network Management Interface (gNMI) services. gNMI service is also enabled for other supported Routing Engine sensors.

  Using JTI and gRPC or gNMI services, you can stream telemetry statistics to an outside collector.

  These interface sensors are supported:

- Physical interfaces (IFD) (resource path **/interfaces/interface/**).
- Logical interfaces (IFL) (resource path **/interfaces/interface/subinterfaces/**).

These Routing Engine sensors are supported using gNMI services (previously, only gRPC services were supported):

- System events (resource path **/junos/events**).
- BGP peer information (resource path **/network-instances/network-instance/protocols/protocol/bgp/**).
- Memory utilization for routing protocol task (resource path **/junos/task-memory-information/**).
- Operational state of Routing Engines, power supply modules, Switch Fabric Boards, Control Boards, Switch Interface Boards, Modular Interface Cards, and Physical Interface Cards (resource path **/components/**).
- Link Layer Discovery Protocol (LLDP) (resource path **/lldp/**).
- Address Resolution Protocol (ARP) statistics for IPv4 routes (resource path **/arp-information/**).
- Network Discovery Protocol (NDP) table state information for IPv6 routes (resource path **/nd6-information/**).
- NDP router-advertisement statistics (resource path **/ipv6-ra/**).
- IS-IS routing protocol statistics (resource path **/network-instances/network-instance/protocols/protocol/isis/levels/level/** and **network-instances/network-instance/protocols/protocol/isis/interfaces/interface/levels/level/**).

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

Juniper Extension Toolkit (JET)

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

  [See language (Scripts), Develop Off-Device JET Applications, and Develop On-Device JET Applications.]

J-Web

- **Improved VPN usability (SRX Series)**—Starting in Junos OS Release 20.2R1, we've refreshed the IPsec VPN page. You can see a new improved site-to-site VPN workflow configuration.

  [See About the IPsec VPN Page.]

- **Pass-through tunnel inspection is supported in TAP mode (SRX 300 line of devices, SRX550M, SRX1500, SRX4100, and SRX4200)**—Starting in Junos OS Release 20.2R1, the J-Web Setup Wizard TAP mode supports pass-through tunnel inspection. This allows the SRX Series device to inspect pass-through traffic over an IP-IP tunnel or GRE tunnel.

  [See Start J-Web.]

- **HTTP X-Forwarded for header support in IDP (SRX Series)**—Starting in Junos OS Release 20.2R1, IDP supports the HTTP X-Forwarded option. When you enable this option, during traffic flow, IDP saves the source IP addresses (IPv4 or IPv6) from the HTTP and SMTP traffic contexts and displays them in the attack logs.

  [See About the Sensor Page.]

- **Enhancements to custom application signatures (SRX Series)**—Starting in Junos OS Release 20.2R1, we've enhanced custom applications signatures with the following:

  - By default, the priority for the custom application is set to Low. This allows a predefined application to take precedence. If you want to override a predefined application, you must set the priority to High.
  - Depth option is supported. Use this byte limit for Application Identification (App ID) to identify custom application patterns for applications running over TCP or UDP or Layer 7 applications.
  - Custom Application Byte Limit is supported in Global Settings. This byte limit helps in understanding when to stop the identification of custom applications.

  [See Add Application Signatures and Global Settings.]

Juniper Sky ATP

- **Support for adaptive threat profiling**—Starting in Junos OS Release 20.2R1, you can configure adaptive threat profiling in Juniper Sky ATP. Adaptive Threat Profiling allows SRX Series devices to generate, propagate, and consume threat feeds based on their own advanced detection and policy-match events. You can generate adaptive threat profiling feeds with traditional policies, unified policies with application identification (AppID) or URL-based match criteria, and IDP. Navigate to **Configure > Adaptive Threat Profiling** in Juniper Sky ATP UI to configure adaptive threat profiling.

  [See Adaptive Threat Profiling Overview and Add Threat Feed for Adaptive Threat Profiling.]

- **Support for encrypted traffic analysis**—Starting in Junos OS Release 20.2R1, encrypted traffic analysis is supported in Juniper Networks Sky ATP. Encrypted traffic analysis helps you to detect malicious threats that are hidden in encrypted traffic without intercepting and decrypting the traffic. Navigate to **Monitor > Encrypted Traffic** in Juniper Sky ATP UI to view detailed information about encrypted traffic analysis-based detections. To configure encrypted traffic analysis, use the **security-metadata-streaming** command at **[edit services]** hierarchy level. Use the **show services security-metadata-streaming statistics** command to view the statistics of the sessions.

  [See Encrypted Traffic Analysis Overview and Encrypted Traffic Analysis Details.]

Logical Systems and Tenant Systems

- **Support for user firewall UAC authentication entries in shared mode for logical systems and tenant systems (SRX Series)**—Starting in Junos OS Release 20.2R1, logical systems and tenant systems support user firewall authentication with Unified Access Control (UAC).

  [See Understanding Integrated User Firewall Support in a Tenant System.]

- **User authentication support for tenant systems (SRX Series)**—Starting in Release 20.2R1, Junos OS introduces the following authentication support for tenant systems:

  - **address-assignment pools**: Creates centralized IPv4 and IPv6 address pools independent of the client applications that use the pools.
  - **access profiles**: Runs authentication and accounting requests.
  - **clear network-access aaa subscribers**: Clears AAA subscriber statistics and logs out subscribers. You can log out subscribers based on the username or on the subscriber session identifier.

  [See Firewall Authentication for Tenant Systems.]

Multicast

- **Strict packet order for multicast traffic (SRX345 and SRX1500)**—Starting in Junos OS Release 20.2R1, we have introduced a new mechanism to maintain multicast traffic order and resolve packet drop issue. Use the **strict-packet-order** command at the **[edit security flow]** hierarchy level to maintain the packet order.

  As part of this enhancement, you can configure the multicast route next-hop resolve attempts. When a multicast route next-hop resolve is unsuccessful, the SRX Series device attempts to resolve the next-hop route based on the specified retry counts. Use the **multicast-nh-resolve-retry** command at the **[edit security flow]** hierarchy level to specify the number of retry counts.

  [See flow.]

Network Address Translation (NAT)

- **Increased port block allocation size (SRX5000 line of devices with SPC2 and SPC3 cards)**—we've increased the port block allocation size so you can store more log files in the log server.
  - When you disable **interim log**, you can increase the size of port block allocation from *64* to *8* .
  - When you enable **interim log**, you can increase the size of port block allocation from *128* to *8*.

  If you configure the port block allocation size less than 8, the system displays the warning message **warning: To save system memory, the block size is recommended to be no less than 8**.

  [See Guidelines for Configuring Secured Port Block Allocation and Configure Port Block Allocation Size.]

Network Management and Monitoring

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

  [See NETCONF Sessions over Outbound HTTPS.]

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **Traffic log enhancement (SRX Series)**—Starting in Junos OS Release 20.2R1, we've enhanced the traffic log by supporting:
  - Escape in stream log forwarding and on-box reporting to avoid parsing errors. Stream mode supports escape in **sd-syslog** and **binary** format. Event mode supports escape only in **binary** format.
  - Different security log transport options for different streams.
  - Stream-event mode.
  - Increased maximum length of the stream mode **sd-syslog** format syslog message to 4*1472 bytes.
  - Different source addresses for different streams.
  - Year and millisecond in timestamps.

  [See log (Security) and mode (Security Log).]

- **CPU usage monitoring (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.2R1, you can use the following operational commands to monitor the average CPU usage information for the last minute, hour, or day of an SPC3 card:
  - **show security monitoring performance spu summary fpc *fpc-slot-number* pic *pic-slot-number***
  - **show security monitoring performance spu summary fpc *fpc-slot-number* pic *pic-slot-number* thread *thread-number***

  You can monitor the CPU usage information only when the PIC is online.

  We've introduced the new SNMP MIBs **jnxJsSPUMonitoringSPUThreadsNumber**, **jnxJsSPUMonitoringSPUThreadIndex**, **jnxJsSPUMonitoringSPUThreadLastMinUsage**, **jnxJsSPUMonitoringSPUThreadLastHourUsage**, and **jnxJsSPUMonitoringSPUThreadLastDayUsage** to monitor the CPU usage information of an SPC3 card.

  [See show snmp mib and show security monitoring performance spu.]

Platform and Infrastructure

- **Support for Application Quality of Experience (AppQoE) (SRX4600)**—Starting in Junos OS Release 20.2R1, AppQoE is supported on SRX4600 devices along with SRX300, SRX320, SRX340, SRX345, SRX550M, SRX4100, and SRX4200 devices.

  [See Security Policy for Controlling Traffic for VRF Routing-Instance, Flow Management in SRX Series Devices Using VRF Routing-Instance, Understanding ALG Support for VRF Routing-Instance, and Network Address Translation for VRF Routing-Instance.]

Port Security

- **Media Access Control Security (MACsec) (SRX380)**—Starting in Junos OS Release 20.2R1, MACsec is supported on high availability (HA) control and fabric ports of SRX380 devices in chassis cluster mode. MACsec provides secure communication for almost all types of Layer 2 traffic on Ethernet links. MACsec is capable of identifying and preventing most security threats at Layer 2 and can be used in combination with other security protocols to provide end-to-end network security. MACsec is standardized in IEEE 802.1AE.

  [See Media Access Control Security (MACsec) on Chassis Cluster.]

Security

- **Support for security feeds in security policies (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, you can add source and destination addresses to the security intelligence (SecIntel) profiles to generate security feeds in a security policy. You can accomplish this by configuring the **security-intelligence** configuration statements. After the feeds are generated, you can configure other security policies to use the feeds as a **dynamic-address˝˙þ** to match designated traffic and perform policy actions.

  You can configure the **security-intelligence** configuration statements as permit, deny, or reject match conditions in a security policy at the following hierarchy levels:

  ```
  [edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
      application-services]
  [edit security policies from-zone zone-name to-zone zone-name policy policy-name then deny application-services]
  [edit security policies from-zone zone-name to-zone zone-name policy policy-name then reject application-services]
  ```

  [See security-intelligence and Encrypted Traffic Analysis Overview.]

- **Enhancements to configuring security policies (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, we have added advanced connection tracking options to security policies.

  You can configure the **advanced-connection-tracking** command at the**[edit security zones security-zone *zone name*]** hierarchy levels to generate a connection track table using source IP, destination IP (optional),

and destination port (optional) during session creation stage when traffic enters a given zone. This connection track mapping table also appears on the backup node in high availability (HA) pair.

You can configure the **advanced-connection-tracking** option under **[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]** to mandate that traffic matching given policy do a lookup in the *to-zone*'s connection track mapping table using the new session's key information. If there is no match, a new connection is not created.

[See advanced-connection-tracking.]

Software Installation and Upgrade

- **Zero-touch provisioning (ZTP) enhancements to support both DHCP options and phone-home client (SRX300, SRX320, SRX340, SRX345, SRX550 HM, and SRX1500)**—Starting in Junos OS Release 20.2R1, you can use zero-touch provisioning with DHCP options or the phone-home client to provision your device. As part of the factory default configuration, both ZTP and the phone-home client are included and are running at the same time when the device boots up in factory-default mode. ZTP with DHCP options is the first priority for provisioning. The device checks for DHCP bindings, and if there are DHCP bindings, but the DHCP bindings are not given the necessary ZTP-related options, (such as file server, and at least one image file or configuration file) the phone-home client will take over the provisioning process.

  [See Zero Touch Provisioning.]

Unified Threat Management (UTM)

- **UTM CLI test commands for Web Filtering and antispam feature (SRX Series)**— Starting in Release 20.2R1, Junos OS introduces the following test commands that help you to configure the Enhanced Web Filtering:
  - **test security utm enhanced-web-filtering url-check <*test-url*>**: Checks the category of a test string.
  - **test security utm web-filtering profile <*profile-name*><*test-url*>**: Checks the reputation of a test string.

  Junos OS introduces the following test command for the antispam feature:

  - **test security utm anti-spam ip-check <*test-IP*>**: Checks whether the IP address is a spam source.

  [See Unified Threat Management User Guide.]
- **CDF mode and inline-tap mode for AV**—Starting in Release 20.2R1, Junos OS introduces continuous delivery function (CDF) and inline-tap mode at the existing **[edit security utm default-configuration anti-virus]** hierarchy level. Continuous delivery function holds the last packet and sends out the other packets. This reduces system memory usage and speeds up the traffic. Inline-tap mode permits the traffic even if it is infected. Use inline-tap mode to check the antivirus feature without blocking or modifying the traffic.

[See Unified Threat Management User Guide.]

- **Safe search enhancement for Web filtering (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, we've introduced safe search UTM Web filtering on well-known search engines. This safe search enhancement enforces the safest Web browsing mode available, by default. You can disable the safe search option at the Web filtering-level and profile-level configurations. You can also block search engine cache on the well-known search engines. By blocking the search engine cache, you can hide your Web-browsing activities from other users if you are a part of an organization that has multiple Web users in educational, financial, health-care, banking, and corporate segments.

  [See Safe Search Enhancement for Web Filtering, feature-profile, websense-redirect, and juniper-local.]

SEE ALSO

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for SRX Series.

Application Security

- Junos OS Release 20.2R1 introduces a new CLI configuration statement **depth** under **set services application-identification application** ***application-name*** **over** ***application*** **signature** ***signature-name*** **member** ***number*** hierarchy. You can use this configuration statement to specify the byte limit for application identification (AppID) to identify the custom application pattern for the applications running over TCP or UDP or Layer 7 applications.

  Starting in Junos OS Release 20.2R1, you can display the configured **depth** value in J-Web using the **show services application-identification application detail** command.

```
user@host>  show services application-identification application detail application-1
```

```
Application Name: test
Application type: application-1
Description: N/A
Application ID: 16777221
Priority: high
Order: 65500
Disabled: No
Cacheable: No
Activation Date: N/A
Last Modified: N/A
Underlying consolidated Protocols/ports application is dependent on:
    Protocols:
        Protocol: junos:HTTP  / 67
        Protocol: junos:UDP   / 216
        Protocol: junos:TCP   / 205
        Protocol: junos:NET-PROXY / 2629
        Protocol: junos:SPDY  / 1469
        Protocol: junos:SSL   / 199
        Protocol: junos:LIBJINGLE-PSEUDOTCP / 3237
        Protocol: junos:STUN  / 201
        Protocol: junos:HTTPS / 68
        Protocol: junos:HTTP  / 67
        Protocol: junos:HTTP2 / 2553
        Protocol: junos:HTTP-TUNNEL / 750
        Protocol: junos:HTTP-PROXY / 2956
        Protocol: junos:HAPROXY / 3331
        Protocol: junos:COTP  / 22
    TCP Ports:
        Port: 80
        Port: 3128
```

```
        Port: 8000
        Port: 8080
Layer-7 Immediate Protocol(s):
    Protocol: HTTP        / 67
    Signature: fgnm
    Port range: N/A
    Member(s): 1
        Member m01
            Depth: 4
            Context: http-get-url-parsed-param-parsed
            Pattern: ads
            Direction: CTS
```

In the above sample, you can see the configured value of the depth is displayed as 4.

[See Application Identification].

- Starting in Junos OS Release 20.2R1, the syntax of the commands used for displaying the SLA profile details is changed as following:

Syntax in Junos OS Release Prior to 20.2R1	Syntax in Junos OS Release 20.2R1 or Later
show security advance-policy-based-routing sla profile *sla-profile-name* application application-name destination-group-name *destination-group-name* status	**show security advance-policy-based-routing sla profile *profile-name* application *application-name* next-hop *next-hop-id* status**
show security advance-policy-based-routing sla profile *sla-profile-name* application application-name destination-group-name *destination-group-name*	**show security advance-policy-based-routing sla profile *profile-name* application *application-name* next-hop *next-hop-id***

[See *show security advance-policy-based-routing sla profile (Application Name)*, *show security advance-policy-based-routing sla profile (Next-Hop)*, and *show security advance-policy-based-routing sla profile (Status)*.]

Flow-Based and Packet-Based Processing

- **ECMP load balancing in chassis cluster (SRX Series)**—Starting in Junos OS Release 20.2R1, in a chassis cluster setup, to avoid reroute flapping between primary and secondary sessions, add a logic to skip the reroute for backup sessions. But reroute can change the chassis interface of a flow session, so the session can be changed from backup session to primary session after reroute. You cannot skip reroute for such a session.

When you change the logic, the session reroute skips only the packets received from the chassis interface. So we can make sure the session continues as the backup session even after you reroute and change the out-going interface. Otherwise, reroute cannot be skipped for backup sessions.

- **Simplified HA (SRX Series)**—Starting in Junos OS Release 20.2R1, on SRX Series devices in a simplified HA setup, when you clear the session using the **clear security flow session** command, some warm sessions exist for an extended duration. To clear these warm sessions, a new CLI command **clear security flow session session-state warm** is introduced.

  clear security flow session all

Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

  [See Develop Off-Device JET Applications and Develop On-Device JET Applications.]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH_USAGE_INVALID**.

  [See Juniper EngNet.]

Juniper Sky ATP

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 20.2R1, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Sky Advanced Threat Prevention (ATP).

Network Management and Monitoring

- **Request support information for IPsec VPN (SRX Series)**—Starting in Junos OS Release 20.2R1, we've introduced the CLI **ipsec-vpn** option to the **request support information security-components** command. This new option displays all the configuration, states, and statistics information necessary for debugging IPsec VPN related issues.

  [See request support information.]

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

VPNs

- **New vendor ID for Internet Key Exchange (SRX Series)**—In Junos OS Release 20.2R1, we've introduced a new vendor ID **Juniper Networks** for Internet IKEv1 and IKEv2 which is advertised to the peer.

  [See Understanding IKE and IPsec Packet Processing.]

- **Change in CLI options help text description (SRX Series)**—Starting in Junos OS Release 20.2R1, we've changed the help text description as **NOT RECOMMENDED** for the following CLI options under **[edit security ike proposal *proposal-name*]**, [edit security ike policy *policy-name*], [edit security ipsec proposal *proposal-name*], and [edit security ipsec policy *policy-name*] hierarchies.

Hierarchy	CLI Options	Help Text Description
[edit security ike proposal *proposal-name* authentication-algorithm]	**md5**	**NOT RECOMMENDED**
	sha1	**NOT RECOMMENDED**
[edit security ike proposal *proposal-name* encryption-algorithm]	**3des-cbc**	**NOT RECOMMENDED**
	des-cbc	**NOT RECOMMENDED**
[set security ike proposal *proposal-name* dh-group]	**group1**	**NOT RECOMMENDED**
	group14	**NOT RECOMMENDED**
	group2	**NOT RECOMMENDED**
	group5	**NOT RECOMMENDED**
[edit security ike proposal *proposal-name* authentication-method]	**dsa-signatures**	**NOT RECOMMENDED**
[edit security ike policy *policy-name* proposal-set]	**basic**	**NOT RECOMMENDED**
	compatible	**NOT RECOMMENDED**
	standard	**NOT RECOMMENDED**

Hierarchy	CLI Options	Help Text Description
[edit security ipsec policy *policy-name* proposal-set]	**basic**	**NOT RECOMMENDED**
	compatible	**NOT RECOMMENDED**
	standard	**NOT RECOMMENDED**
[edit security ipsec proposal *proposal-name* encryption-algorithm]	**3des-cbc**	**NOT RECOMMENDED**
	des-cbc	**NOT RECOMMENDED**
[edit security ipsec proposal *proposal-name* authentication-algorithm]	**hmac-md5-96**	**NOT RECOMMENDED**
	hmac-sha1-96	**NOT RECOMMENDED**
[edit security ipsec policy *policy-name* perfect-forward-secrecy keys]	**group1**	**NOT RECOMMENDED**
	group2	**NOT RECOMMENDED**
	group5	**NOT RECOMMENDED**
	group14	**NOT RECOMMENDED**

[See authentication-algorithm (Security IPsec) and encryption-algorithm (Security IKE).]

- **Change in thread ID configuration (SRX Series)**—Starting in Junos OS Release 20.2R1, when you add, change, or delete the thread ID from distribution profile at **[edit security distribution-profile *profile-name* fpc *slot-number* pic *slot-number* thread-id]**, all tunnels part of modified distribution profile anchored on modified SPU member of distribution profile are teared down and re-negotiated.

  [See distribution-profile.]

SEE ALSO

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for SRX Series devices.

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

Authentication and Access Control

- When you use the **request services user-identification authentication-source jims groups domain <domain-name> (force-fetch|status)** command, the SRX Series device retrieve the complete group list, excluding the user list or device list.

  The secondary JIMS server is online and the related secondary JIMS validator is offline when the primary JIMS server is offline. Therefore, the connection to the JIMS validator reports an error message for the **group-query** or **validate-query** command.

  [See Querying JIMS for User Identity Information.]

Flow-Based and Packet-Based Processing

- Committing a large number of custom applications with a single member, a single context, and a varying pattern might result in significant time taken for completion of commit. Commit status can be checked using **show services application-identification commit-status**. PR1493127

J-Web

- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is autorefreshed. As a workaround, create a dynamic application as the last action while modifying the policy rule and click the Save button to avoid loss of configuration changes made to the policy rule. PR1460214

Routing Policy and Firewall Filters

- SecProfiling deployment starts from the root logical system and evolves to the user-defined logical system; currently the use-case under tenant is not mandated. PR1490071
- The J-Web IPsec VPN workflow only supports route-based VPNs. Policy-based VPNs are not supported. PR1498169

VPNs

- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with 60,000 tunnels up, when RG0 failover happens while an IPsec and/or IKE rekey is in progress, those rekeying tunnels might go down and traffic loss might be seen until the tunnel is reestablished. PR1471499
- On SRX Series device, the accounting stop message is not being sent after deactivating the access profile under the security IKE gateway. PR1485732

SEE ALSO

Open Issues

IN THIS SECTION

Learn about open issues in this release for SRX Series devices.

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

Flow-Based and Packet-Based Processing

- Use 512 antireplay window size for IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). Hence there are no out-of-order packets with 512 antireplay window size. PR1470637

J-Web

- On the SRX5000 line of devices, J-Web might not be responsive sometimes when you commit configuration changes after adding a new dynamic application while creating a new firewall rule. J-Web displays a warning while validating the configuration due to dynamic application or any other configuration changes. As a workaround, refresh the J-Web page. PR1460001
- For a spoke device of hub and Spoke topology, J-Web shows the VPN topology as **Site to Site**. PR1495973
- Configuration of global settings options of IPsec VPN such as TCP encap profile, IPsec power mode and IKE package installation are not supported from J-Web. PR1496439
- SSL proxy exempted URL categories list blank in HA setup. PR1516590
- Charts are appearing blank in generated **Threat Assessment Report** when J-Web is opened from Firefox browser v77.0.1. PR1517343

Routing Policy and Firewall Filters

- IP address that can't be divided exactly by three in **show security match-policies** can lead to matching failure. PR1483251

VPNs

- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, a new behavior has been introduced that differs from the behavior on the older SPC2 card. The SRX Series device with AutoVPN configuration can now accept multiple IPsec tunnels from a peer device (with the same source IP address and port number) using different IKE-IDs. PR1407356
- On the SRX5000 line of devices with an SPC3 card, sometimes IKE SA is not seen on the device when st0 binding on VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). PR1441411
- With NCP remote access solution, in a PathFinder case (for example, where IPsec traffic has to be encapsulated as TCP packets), TCP encapsulation for transit traffic is failing. PR1442145
- During 10,000 tunnel ramp-up, sometimes, IKED generates a core file. PR1479548

- On SRX Series devices with SPC3, when overlapping taffic-selectors are configured, multiple IPsec SAs get negotiated with peer device. PR1482446
- The SRX5000 line of devices with SPC3 was not supporting simultaneous IKE negotiation in Junos OS Release 19.2, 19.3, 19.4 or 20.1. PR1497297

SEE ALSO

What's New | 242

What's Changed | 253

Known Limitations | 259

Resolved Issues | 263

Documentation Updates | 268

Migration, Upgrade, and Downgrade Instructions | 268

Resolved Issues

Learn which issues were resolved in the Junos OS main and maintenance releases for SRX Series devices.

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

Application Layer Gateways (ALGs)

- RTSP data sessions are cleared unexpectedly during cold sync. PR1468001
- The flowd or srxpfe process might stop when an ALG creates a gate with an incorrect protocol value. PR1474942
- SIP messages that need to be fragmented might be dropped by SIP ALG. PR1475031
- FTPS traffic might get dropped on SRX Series or MX Series devices if FTP ALG is used. PR1483834

Authentication and Access Control

- SRX Series: Unified Access Control (UAC) bypass vulnerability (CVE-2020-1637). PR1475435

Flow-Based and Packet-Based Processing

- Command **show security pki local-certificate logical-system all** is not showing any output. PR1414628
- The trusted-ca and root-ca names or IDs should not be the same within an SSL proxy configuration. PR1420859
- Introduction of default inspection limits for application identification to optimize CPU usage and improve resistance to evasive applications. PR1454180
- TCP session might not time out properly upon receiving TCP RESET packet. PR1467654
- RPM test probe fails to show that round-trip time has been exceeded. PR1471606
- Support LLDP protocol on reth interface. PR1473456
- Certificate error when configuration is validated during Junos OS upgrade. PR1474225
- An unhealthy node might become primary in SRX4600 devices with chassis cluster scenario. PR1474233
- Packet drop might be observed on the SRX300 line of devices when adding or removing an interface from MACsec. PR1474674
- Stateful firewall rule configuration deletion might lead to memory leak. PR1475220
- The flowd or srxpfe process might stop when deleting user firewall local authentication table entry. PR1477627
- MPCs might stop when there is bulk route update failure in a corner case. PR1478392
- The nsd process pause might be seen during device reboots if dynamic application groups are configured in policy. PR1478608
- The flowd process core files might be seen when there is mixed NAT-T traffic or non-NAT-T traffic with PMI enabled. PR1478812
- When SRX5K-SPC3s or MX-SPC3s are installed in slots 0 or 1 in SRX5800 or MX960 devices, EMI radiated emissions are observed to be higher than regulatory compliance requirements. PR1479001
- The show mape rule statistics command might display negative values. PR1479165
- The wl-interface stays in ready status after you execute request chassis fpc restart command in Layer 2 mode. PR1479396
- Recent changes to JDPI's classification mechanism caused a considerable performance regression (more than 30 percent). PR1479684
- The flowd or srxpfe process might stop when advanced anti-malware service is used. PR1480005
- On Web proxy, memory leak in association hash table and DNS hash table. PR1480760

- The jsqlsyncd process synchronizes its databases every second even there is no change. PR1482428
- The firewall Web authentication graphics have been updated. PR1482433
- IMAP curl sessions get stuck in the active state if AAMW IMAP block mode is configured. PR1484692
- The show chassis temperature-thresholds command displays extensive FPC 0 output. PR1485224
- The configuration **set chassis psu redundancy n-plus-n** needs support on in high availability (HA) mode. PR1486746
- Commit does not work after the installation through boot loader. PR1487831
- If a cluster ID of 16 or multiples of 16 is used, the chassis cluster might not come up. PR1487951
- CPU board inlet increases after OS upgrade from Junos OS Release 15.1X49 to Junos OS Release 18.x. PR1488203
- All interfaces remain in the down status after the SRX300 line of devices power up or reboot. PR1488348
- There is a risk of service interruption on all SRX Series devices with a dual stacked CA server. PR1489249
- GRE or IPSec tunnel might not come up when **set security flow no-local-favor-ecmp** command is configured. PR1489276
- Sometimes multiple flowd core files are generated on both nodes of chassis cluster at the same time when changing media MTU. PR1489494
- Continuous drops seen in control traffic, with high data queues in one SPC2 PIC. PR1490216
- Phone client stop seen while doing SRX345 device ZTP with CSO. PR1496650
- Unexpected flow logging traffic beyond the packet filter. PR1497939
- Traffic interruption happens due to MAC address duplication between two devices running Junos OS. PR1497956
- Don't use capital characters for source-identity when using **show security match-policies** command. PR1499090
- J-Flow version 9 does not display correct outgoing interface for APBR traffic. PR1502432
- AppQoE support for dynamic-application. PR1503400
- The cfmd core observed when LTM is triggered for the session configured on ethernet-switching interface without bridge domain configuration. PR1503696

Intrusion Detection and Prevention (IDP)

- Configuring anomaly occurs in CLI. PR1490437

J-Web

- You cannot configure redundant PSU and power budget statistics on the SRX380 device that is in high availability (HA) mode through J-Web. PR1493713
- The J-Web users might not be able to configure PPPoE using PPPoE wizard. PR1502657

Layer 2 Ethernet Services

- Member links state might be asychronized on a connection between PE and CE devices in an EVPN active/active mode. PR1463791

Multiprotocol Label Switching (MPLS)

- BGP session might keep flapping between two directly connected BGP peers because of the wrong TCP-MSS in use. PR1493431

Network Address Translation (NAT)

- Issuing the **show security nat source paired-address** command might return an error. PR1479824

Network Management and Monitoring

- The flowd or srxpfe process might stop immediately after committing the J-Flow version 9 configuration or after upgrading to affected releases. PR1471524
- SNMP trap coldStart agent-address becomes 0.0.0.0. PR1473288

Platform and Infrastructure

- Modifying the REST configuration might cause the system to become unresponsive. PR1461021
- On SRX1500 and the SRX4000 line of devices, physically disconnecting the cable from fxp0 interface causes hardware monitor failure and redundancy group failover, when the device is the primary node in a chassis cluster. PR1467376

- The RGx might fail over after RG0 failover in a rare case. PR1479255
- The **/usr/libexec/ui/yang-pkg** and **/usr/libexec/ui/pyang** files not found in SRX Series devices during YANG installation. PR1496577

Routing Policy and Firewall Filters

- If a huge number of policies are configured on SRX Series devices and some policies are changed, the traffic that matches the changed policies might be dropped. PR1454907
- Support for dynamic tunnels on SRX Series devices was mistakenly removed. PR1476530
- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. PR1492436
- Traffic fails to hit the policies with matching source-end-user-profiles. PR1505002

Routing Protocols

- The rpd might stop when both instance-import and instance-export policies contain as-path-prepend action. PR1471968

Unified Threat Management (UTM)

- The utmd process might pause after deactivating UTM configuration with predefined category upgrading used. PR1478825

VPNs

- IKE SA does not get cleared and is showing very long lifetime. PR1439338
- IKED is treating all re-transmission of first IKE_INIT request packets as new connections when acting as responder. PR1460907
- The iked might crash when the IKE SA expires and the IPsec tunnel of expired IKE SAs still exists. PR1463501
- The newly configured IPsec tunnels might be stuck in VPNM verify-path state in a tunnel scaled scenario. PR1464353
- IPsec tunnels might flap when one secondary node is coming online after reboot in SRX Series high availability environment. PR1471243
- The kmd process might crash continually after the chassis cluster failover in the IPsec ADVPN scenario. PR1479738
- On SRX4200 device, 35 percent of drop is seen in all TPS cases. PR1481625

- Some options under IKE and IPsec policy and proposal help text description should change to **NOT RECOMMENDED**. PR1487515
- Use different XML tags for local and remote IKE ID to avoid confusion. PR1493368
- Issue with XML rpc **show security ipsec tunnel-distribution summary** output. PR1494274

SEE ALSO

- What's New | 242
- What's Changed | 253
- Known Limitations | 259
- Open Issues | 261
- Documentation Updates | 268
- Migration, Upgrade, and Downgrade Instructions | 268

Documentation Updates

There are no errata or changes in Junos OS Release 20.2R1 documentation for the SRX Series.

SEE ALSO

- What's New | 242
- What's Changed | 253
- Known Limitations | 259
- Open Issues | 261
- Resolved Issues | 263
- Migration, Upgrade, and Downgrade Instructions | 268

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

For information about software installation and upgrade, see the Installation and Upgrade Guide.

SEE ALSO

What's New | 242

What's Changed | 253

Known Limitations | 259

Open Issues | 261

Resolved Issues | 263

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the High Availability User Guide.

For additional information about using ISSU on security devices, see the Chassis Cluster User Guide for SRX Series Devices.

For information about ISSU support across platforms and Junos OS releases, see the In-Service Software Upgrade (ISSU) Web application.

Licensing

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

The major components of the framework are:

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

For information about the list of supported products, see Juniper Flex Program.

Compliance Advisor

For regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6 for Juniper Networks products, see the Juniper Networks Compliance Advisor.

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. https://apps.juniper.net/feature-explorer/
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home

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

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

Documentation Feedback

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

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

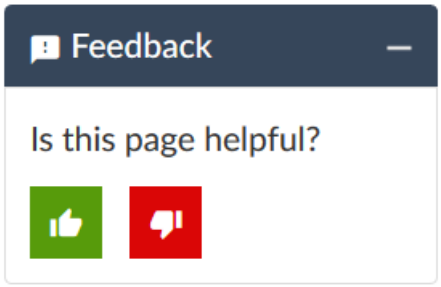

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

Revision History

20 July 2023—Revision 22, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 April 2023—Revision 21, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 November 2022—Revision 20, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 July 2022—Revision 19, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

30 June 2022—Revision 18, Junos OS Release 20.2R1– QFX Series.

23 December 2021—Revision 17, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 October 2021—Revision 16, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 September 2021—Revision 15, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 July 2021—Revision 14, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 April 2021—Revision 13, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 March 2021—Revision 12, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 10, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 December 2020—Revision 9, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 October 2020—Revision 8, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 October 2020—Revision 7, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 September 2020—Revision 6, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 5, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 August 2020—Revision 1, Junos OS Release 20.2R1-S1– EX Series, MX Series, and QFX Series.

30 July 2020—Revision 4, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 July 2020—Revision 3, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 July 2020—Revision 2, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

30 June 2020—Revision 1, Junos OS Release 20.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.