Release Notes: Junos® OS Release 19.3R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

22 April 2021

Contents

Introduction

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

These release notes accompany Junos OS Release 19.3R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

New Features in 19.3R2

Feature	Release Note Section
MX2K-MPC11E (MX2010 and MX2020 routers)	“What's New” on page 84
MX-SPC3 Services Card (MX240, MX480, and MX960)	“What's New” on page 84
Hierarchical Class of Service (CoS) parity with MPC11 (MX2000, MPC11E) and MX2000-SF83	“What's New” on page 84
L2 CoS Classifiers and Rewrite feature on MX Routers (MX240, MX480, MX960) with the MPC10 line card	“What's New” on page 84
Forwarding CoS parity (L2 classifiers, rewrite) support for MX Series routers with MPC11 line cards	“What's New” on page 84
Seamless MPLS CoS support for pseudo-wires from access node (AN) and Multi-Services-Edge (MSE) node (MX480, MX960, MX2010, MX2020)	“What's New” on page 84
Class of Service (CoS) parity support for forwarding-class (FC) counters (MPC10 and MPC11)	“What's New” on page 84
Support for EVPN functionality (MX Series with MPC10 line card)	“What's New” on page 84
Support for EVPN functionality on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for load balancing on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84

Feature	Release Note Section
Support for BFD on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
MX2000-SFB3 Switch Control Board (MX2010 and MX2020 routers)	“What's New” on page 84
Chassis and power management for MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for Rate Selectability on MX2K-MPC11E line cards (MX2010 and MX2020 routers)	“What's New” on page 84
MPC Protocol and Application Support for MX2K-MPC11E line cards	“What's New” on page 84
MPC11E supports Junos node slicing (MX2010, MX2020)	“What's New” on page 84
gNMI-based streaming telemetry support for Packet Forwarding Engine sensors on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Packet Forwarding Engine support for JTI on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
CPU/NPU sensors support using JTI on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Layer 2 services on PWHT (MX2K-MPC11E line card)	“What's New” on page 84
Supported features on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for basic Layer 2 features on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for an increase in MAC table size on the MPC10E-15C-MRATE line cards (MX Series)	“What's New” on page 84
Support for new MPC line card MPC11E (MX Series)	“What's New” on page 84
Support for Seamless MPLS Layer 3 features on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Error recovery, fault handling, and resiliency support for MX2K-MPC11E (MX2020 and MX2010)	“What's New” on page 84
Support for MPLS features on MX2K-MPC11E line card	“What's New” on page 84

Feature	Release Note Section
Support for multicast forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for CCC and Layer 3 firewall forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for firewall forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Support for port mirroring (MX2K-MPC11E line card on MX2010 and MX2020 routers)	“What's New” on page 84
Support for tunnel interfaces (MX2K-MPC11E line card on MX2010 and MX2020 routers)	“What's New” on page 84
Support for multicast forwarding on MX2K-MPC11E line cards (MX2010 and MX2020 routers)	“What's New” on page 84
Fabric support on MX2K-MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Configure next-hop-based dynamic tunnels on MX2K-MPC11E line card (MX2010 and MX2020 routers)	“What's New” on page 84
Inline active flow monitoring on MX2K-MPC11E line card (MX2010 and MX2020 routers)	“What's New” on page 84
Support for Two-Way Active Measurement Protocol (TWAMP) and Hardware timestamping of RPM probe messages on MX2010 and MX2020 routers	“What's New” on page 84
L2TPv2 silent failover on peer interface for L2TPv2 subscriber services on MX2K-MPC11E line card (MX2010 and MX2020 routers)	“What's New” on page 84
Port mirroring support on MX2K-MPC11E (MX2010 and MX2020 routers)	“What's New” on page 84
FlowTapLite suppport on MX2K-MPC11E line cards (MX2010 and MX2020 routers)	“What's New” on page 84
Distributed denial-of-service on MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84
Subscriber services uplink support on MPC11E line cards (MX2010 and MX2020)	“What's New” on page 84

Feature	Release Note Section
CoA messages support Session-Timeout attribute (MX Series)	“What's New” on page 84

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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.

What's New in Release 19.3R2

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 19.3R2.

What's New in 19.3R1-S1

Hardware

- **New ACX5448-M Universal Metro Routers**—In Junos OS Release 19.3R1-S1, we introduce the ACX5448-M, a top-of-rack router with support for advanced security capabilities such as Media Access Control Security (MACsec). A compact 1 U model, the ACX5448-M provides a system throughput of up to 800 Gbps through the following port configuration:
  - Forty-four 10-Gigabit Ethernet SFP+ or 1-Gigabit Ethernet SFP ports (**0** through **43**). The ACX5448-M supports MACsec only on these ports.
  - Six 100-Gigabit Ethernet QSFP28 or 40-Gigabit Ethernet QSFP+ ports (**44** through **49**). You can channelize each QSFP28 port into four 25-Gbps interfaces and each QSFP+ port into 10-Gbps interfaces using breakout cables (and the **channelization** configuration).

  The ACX5448-M routers have redundant fan modules and redundant AC or DC power supply modules.
- **New ACX5448-D Universal Metro Routers**—In Junos OS Release 19.3R1-S1, we introduce the ACX5448-D, a top-of-rack router for aggregation environments. Designed for packet-optical convergence, this compact 1 U router provides wire-speed packet performance, very low latency, and a rich set of Layer 2 and Layer 3 features.

  The ACX5448-D provides a system throughput of up to 800 Gbps through the following port configuration:
  - Thirty-six 10-Gigabit Ethernet SFP+ or 1-Gigabit Ethernet SFP ports (**0** through **35**).
  - Two 100-Gigabit Ethernet QSFP28 or 40-Gigabit Ethernet QSFP+ ports (**36** and **37**). You can channelize each QSFP28 port into four 25-Gbps interfaces and each QSFP+ port into four 10-Gbps interfaces using breakout cables (and the **channelization** configuration).
  - Two 200-Gigabit Ethernet CFP2-DCO ports (**38** and **39**).

  The ACX5448-D routers have redundant fan modules and redundant AC or DC power supply modules.

What's New in Release 19.3R1

Class of Service

- **Class of Service (CoS) parity support for ACX5000 routers**—Starting in Junos OS 19.3R1, the Class-of-Service feature set is supported on ACX5000 devices to enable users to configure classification, rewrite, shaping, queueing, and scheduling parameters for traffic flow.

  For more information regarding CoS, see Understanding Class of Service.

- **Support for Class-of-Service (CoS) for ACX5448 devices**—Starting in Junos OS 19.3R1, support is provided for Class-of-Service (COS) on ACX5448 devices to include firewall filter families (ANY, VPLS, ethernet-switching, CCC, IPv6, IPv4, Lo0-IPv6, Lo0-IPv4, and MPLS), and CoS (classification, policing, forwarding policy, forwarding class to queue map, WRED and Tail drop profiles, fabric queue and scheduling configuration, scheduler, deep buffers, and remarking).

  For more information regarding CoS, see Understanding Class of Service.

High Availability (HA) and Resiliency

- **VRRP support (ACX5448)**—Starting in Junos OS Release 19.3R1, the ACX5448 router supports the Virtual Router Redundancy Protocol (VRRP) over aggregated Ethernet and integrated routing and bridging (IRB) interfaces. The VRRP queue size is limited, so it doesn't disturb other protocols such as Bidirectional Forwarding Detection (BFD) and connectivity fault management (CFM). The ACX5448 supports 16 VRRP groups.

  [See Understanding VRRP.]

- **Software Support (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, ACX5448-D and ACX5448-M routers support:

  - Chassis management software—Manages the onboard FRUs
  - Upgradable common BIOS software—Initializes all the devices on the hardware
  - FPC and PIC management

Interfaces and Chassis

- **Hardware resiliency support (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, ACX5448-D and ACX5448-M routers support the resiliency feature, which includes handling of hardware failure and faults. Resiliency on an ACX5448-D enhances its debugging capability in the case of hardware failure of its components such as Routing Engine, solid-state drive (SSD), and PCI Express. For example, the resiliency feature enables the router to recover from inter-integrated circuit (I2C) failure, and improves its voltage monitoring, temperature monitoring, and PCI Express error handling and reporting. The resiliency feature also provides DRAM single-bit and multibit error checking and correction (ECC) capabilities.

  [See show chassis fpc errors.]

- **Interface speed, channelization, and MACsec support (ACX5448-M)**—In Junos OS Release 19.3R1, we introduce the ACX5448-M Universal Metro Router with support for advanced security capabilities such as Media Access Control Security (MACsec). The ACX5448-M has the following port types:
  - Forty-four 10-Gigabit Ethernet or 1-Gigabit Ethernet SFP+ ports (**0** through **43**). Based on the optics plugged in, the ports come up either as 1-Gbps or 10-Gbps.
  - Six 100-Gigabit Ethernet QSFP28 ports (**44** through **49**). These ports support 100-Gbps (the default) and 40-Gbps speeds. You can channelize these ports into four 25-Gbps or four 10-Gbps interfaces.

  The 44 SFP+ ports on the ACX5448-M support MACsec; however, the six QSFP28 ports do not support MACsec.

  [See Channelize Interfaces on ACX5448-D and ACX5448-M Routers.]

Layer 2 Features

- **Support for Layer 2 Features (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, Junos OS supports Layer 2 bridging, Q-in-Q tunneling, no-local switching, Layer 2 protocol tunneling, Spanning Tree Protocols (RSTP,MSTP), Bridge Protocol Data Unit (guard, root and loop protect), Ethernet OAM, VPLS, BGP, LDP, pseudowire ping, and Bidirectional Forwarding Detection (BFD) support for virtual circuit connectivity verification (VCCV) on ACX5448-D and ACX5448-M routers.
- **Support for Layer 2 switching cross-connects (ACX5440)**—Starting in Junos OS Release 19.3R1, you can leverage the hardware support available for cross-connects on the ACX5448 device with the Layer 2 local switching functionality using certain models. With this support, you can provide the EVP and Ethernet Virtual Private Line (EVPL) services.

  [See Configuring Layer 2 Switching Cross-Connects Using CCC.]

Layer 3 Features

- **Support for Layer 3 features (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, Junos OS supports Layer 3 protocols, multicast, and MPLS as the transport mechanism on ACX5448-D and ACX5448-M routers.

Management

- **OpenConfig AAA data model support (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 19.3R1 supports the configuration leaves specified in the OpenConfig AAA data model. Mapping the OpenConfig AAA configuration to the Junos AAA configuration using the following YANG files in the data model makes this support possible:
  - openconfig-aaa.yang
  - openconfig-aaa-types.yang
  - openconfig-aaa-tacacs.yang
  - openconfig-aaa-radius.yang

The configuration model supporting the OpenConfig data model includes:

- A translation script **(.py / .slax)** that maps each configuration leaf in the OpenConfig schema to one or more configuration leafs in the JUNOS OS schema.
- A deviation file **(.yang)** that specifies how much the implementation deviates from the vendor-neutral model.

[See Mapping OpenConfig AAA Commands to Junos Configuration.]

Network Management and Monitoring

- **Support for adding custom YANG data models to the Junos OS schema (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, ACX5448-D and ACX5448-M routers support loading custom YANG data models on the device, which enables you to add RPCs or configuration hierarchies that are customized for your operations. The ability to add data models to a device is beneficial when you want to create device-agnostic RPCs and configuration models that can be used on different devices from one or more vendors.

  [See Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS.]

Port Security

- **Media Access Control Security (MACsec) support (ACX5448)**—Starting with Junos OS Release 19.3R1, ACX5448 routers support MACsec on 1-Gigabit Ethernet SFP and 10-Gigabit Ethernet SFP+ ports. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is standardized in IEEE 802.1AE.

  [See Understanding Media Access Control Security (MACsec).]

Routing Policy and Firewall Filters

- **Match condition support for IPv6 firewall filters (ACX6360)**—Starting in Junos OS Release 19.3R1, the ACX6360 router supports the following firewall filter match conditions for IPv6 traffic: **address**, **destination-address**, **destination-port**, **destination-port-except**, **destination-prefix-list**, **port**, **port-except**, **icmp-code**, **icmp-code-except**, **icmp-type**, **icmp-type-except**, **next-header**, **next-header-except**, **prefix-list**, **source-address**, **source-port**, **source-port-except**, and **source-prefix-list**.

  [See Firewall Filter Match Conditions for IPv6 Traffic on ACX Series Routers and ACX6360 Documentation.]

Routing Protocols

- **Clocking and Synchronous Ethernet support (ACX5448)**—Starting in Junos OS Release 19.3R1, ACX5448 routers support frequency synchronization using the Synchronous Ethernet and Ethernet Synchronization Message Channel (ESMC) protocols. The routers also support phase and time synchronization through Precision Time Protocol (PTP).

  [See Synchronous Ethernet Overview.]

- **Transparent clock over IPv6 support (ACX5448)**—Starting with Junos OS Release 19.3R1, ACX5448 routers support transparent clock functionality for PTP over IPv6. To configure the transparent clock functionality, you must include the **e2e-transparent** statement at the **[edit protocol ptp]** hierarchy level. Use the **show ptp global-information** command to check the status of the transparent clock functionality configured on the router.

  [See Understanding Transparent Clocks in Precision Time Protocol.]

- **Support for RIPv2 (ACX5448)**—Starting in Junos OS Release 19.3R1, Junos OS supports RIP version 2 (RIPv2) for both IPv4 and IPv6 packets.

Services Applications

- **Support for Two-Way Active Measurement Protocol or TWAMP (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, you can configure TWAMP on your ACX5448-D and ACX5448-M routers. TWAMP enables you measure the IP performance between two devices in a network. The ACX5448-D and ACX5448-M routers support only the reflector side of TWAMP.

  [See Two-Way Active Measurement Protocol on ACX Series.]

- **Support for virtualization (ACX5448-D and ACX5448-M)**—Starting in Junos OS Release 19.3R1, the Routing Engines on the ACX5448-D routers and ACX5448-M routers support virtualization.

  On Routing Engines of ACX5448-D routers and ACX5448-M routers, one instance of Junos OS, which runs as a guest operating system, is launched by default. The user needs to log in to this instance for operations and management.

  With virtualization of the Routing Engine, Junos OS supports new **request** and **show** commands associated with host and hypervisor processes. The commands are related to:

  - Reboot, halt, and power management for the host
  - Software upgrade for the host
  - Disk snapshot for the host

  [See What Are VM Hosts?.]

- **Port mirroring support for the IPv6 address family (ACX6360)**—Starting in Release 19.3R1, you can configure port mirroring on the ACX6360 router for the inet6 family. Port mirroring copies packets entering or exiting a port and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, and correlating events.

  [See Configuring Port Mirroring.]

Software Installation and Upgrade

- **Migration of Linux kernel version**—Starting in Junos OS Release 19.3R1, the following devices support the Wind River Linux 9 (WRL9) kernel version:

Platforms	Routing Engine Supported
ACX5448-D	RE-ACX-5448
MX240, MX480, and MX960	RE-S-X6-64G
MX2020 and MX2010	REMX2K-X8-64G
MX204	RE-S-1600x8
MX10003	RE-S-1600x8
MX2008	RE-MX2008-X8-64G
MX10016	RE X10
MX10008	RE X10
PTX5000	RE-PTX-X8-64G
PTX3000	RCBPTX
PTX10016	RE-PTX-2X00x4/RE X10
PTX10008	RE-PTX-2X00x4/RE X10
PTX1000	RE-PTX1000
PTX10002-XX	RE-PTX10002-60C
EX9208	RE-S-EX9200-2X00x6
EX9251	EX9251-RE
EX9253	EX9253-RE
EX9204	RE-S-EX9200-2X00x6
EX9214	RE-S-EX9200-2X00x6
QFX10002	RE-QFX10002-60C

Platforms	Routing Engine Supported
QFX10008	RE-QFX10008
QFX10016	RE-QFX10016

Starting in Junos OS Release 19.3R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following routers:

- MX Series—MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008
- PTX Series—PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

If you perform a software upgrade on a router with i40e NVM version earlier than 6.01, the upgrade fails and the following error message is displayed:

ERROR: i40e NVM firmware is not compatible ,please upgrade i40e NVM before installing this package

ERROR: Aborting the installation

ERROR: Upgrade failed

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

System Management

- **Transparent clock functionality support on (ACX5448)**—Starting in Junos OS Release 19.3R1, transparent clock functionality and a global configuration for enabling it are supported on the ACX5448 router. Transparent clock functionality works for PTP over both IPv4 and Ethernet packets. To check the status of transparent clock, use the **show ptp global-information** command.

  [See Understanding Transparent Clocks in Precision Time Protocol.]

- **Synchronous Ethernet and PTP support (ACX 5448)**—Starting in Junos OS Release 19.3R1, the ACX5448 router supports the following features:

  - Frequency synchronization using Synchronous Ethernet
  - Ethernet Synchronization Message Channel (ESMC)
  - Phase and time synchronization using Precision Timing Protocol (PTP)

  [See Synchronous Ethernet Overview.]

SEE ALSO

What's Changed | 21

Known Limitations | 23

Open Issues | 26

What's Changed

IN THIS SECTION

See what changed in the Junos OS main and maintenance release for ACX Series routers.

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

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

  - If a successful **commit** operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the **<source-daemon>** element as a child of the **<error-info>** element instead of the **<rpc-error>** element.

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

See Configuring RFC-Compliant NETCONF Sessions.

What's Changed in Release 19.3R2

There are no changes in behavior and syntax for ACX Series in Junos OS Release 19.3R2.

What's Changed in Release 19.3R1

General Routing

- **Support for gigether-options statement (ACX5048, ACX5096)**—Junos OS supports the gigether-options statement at the edit interfaces interface-name hierarchy on the ACX5048 and ACX5096 routers. Previously, support for the gigether-statement was deprecated.

  [See gigether-options.]

Junos OS XML API and Scripting

- **Range defined for confirm-timeout value in NETCONF and Junos XML protocol sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, the value for the **<confirm-timeout>** element in the Junos XML protocol **<commit-configuration>** operation must be in the range 1 through 65,535 minutes, and the value for the **<confirm-timeout>** element in the NETCONF **<commit>** operation must be in the range 1 through 4,294,967,295 seconds. In earlier releases, the range is determined by the minimum and maximum value of its unsigned integer data type.

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Release 19.3R1 and later, ACX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

  In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

- **Monitoring information available only in trace log (ACX Series)**—In Junos OS Release 19.3R1 and later, the Ethernet link fault management daemon (lfmd process) in the peer router stops monitoring the locally occurred errors until unified ISSU completes. You can view the monitoring-related details only through the trace log file.

System Logging

- **Preventing system instability during core file generation (ACX Series)**—Starting with Release 19.3R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when a change in number of supported history records is detected (ACX Series and MX Series)**—In Junos OS Release 19.3R1, when Ethernet connectivity fault management starts, it detects the number of history records supported by the existing performance monitoring history database and if there is any change from the number of history records supported (that is, 12) in Release 19.3R1, then the existing performance monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

SEE ALSO

What's New | 13

Known Limitations | 23

Open Issues | 26

Resolved Issues | 28

Documentation Updates | 32

Migration, Upgrade, and Downgrade Instructions | 33

Known Limitations

IN THIS SECTION

- General Routing | 24

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

- All PTP packets goes to the best-effort queue instead of the Network Control queue. This is due to the limitation on Qumran where the DSCP values are not preserved. PR1361315
- For an et-interface, only PRE_FEC_SD defect is raised, and no OTN alarm is raised. PR1371997
- The clock class to QL level mapping must be changed through CLI in Arb GM with network option-1. PR1384968
- The clock class must be changed to quality level mapping for Arb GM mode with network option-2. PR1386916
- When the timing configuration and the corresponding interface configuration are flapped for multiple times in iteration, PTP is stuck in Initialize state where the ARP for the neighbor is not resolved. In issue state, BCM hardware block gets into inconsistency state, where the lookup is failing. PR1410746
- Hardware-based fragmentation or reassembly is not supported. Software-based fragmentation rates are going to be extremely slow depending on CPU load. PR1419371
- On the ACX Series platform, the input packets account for all the incoming frames, including the oversized frames, whereas the oversized frame accounts only for the oversized frames. PR1425748
- You might see error messages if the optics is unplugged during the EEPROM read. This is expected and does not impact any functionality. PR1429016
- Packet rates are not seen for aggregated Ethernet logical interface. PR1429590
- Multicast packets are flooded in a BD if snooping is not enabled. If interfaces x and y belongs to a BD, then all multicast packets will be flooded to both x and y interfaces. If packets are received from interface x, packets will be flooded to x and y at ingress but discarded in the egress path for interface x because packet is received from the same interface. But these packets are also counted in the VOQ and hence we are seeing more queue statistics. It is a known Hardware limitation. **monitor interface xe-0/0/30 Input packets: 177958 (64 pps) [0] Output packets: 357306 (128 pps) [0] monitor interface xe-0/0/12 Input packets: 361161 (128 pps) [642] Output packets: 179878 (63 pps) [320] user@router> show interfaces queue xe-0/0/30 Queue: 0, Forwarding classes: best-effort Queued: Packets : 544032 192 pps . => Sum of 64 + 128pps user@router> show interfaces queue xe-0/0/12 Queue: 0, Forwarding classes: best-effort Queued: Packets : 550929 192 pps . => Sum of 64 + 128pps.** PR1429628
- Any packet with size greater than the MTU size are accounted for as oversized packets. Packets exceeding MTU sizes are not considered for Jabber check. PR1429923
- The statistics are accessed through Broadcom API, which is the same for both tagged and untagged packets. This cannot be changed in accordance with MX since it is direct access from Broadcom without

any statistics changes specific to tagging from ACX5448 side. It will impact other statistics if the change is made. PR1430108

- The port LEDs glowing during system/vmhost halt state is the expected behaviour across all ACX Series platforms. Even the system LED glows during halt state. PR1430129
- The initial transient messages are displayed however, these messages does not have any functional impact. PR1430355
- Packets dropped due to MTU checks in the output interface are not accounted for as MTU errors. All packets with sizes greater than the MTU size are accounted for as Oversized-packets in the input interface. PR1430446
- The 1-Gigabit interfaces are displayed as xe. This issue is observed with respect to auto-negotiation parameters even though there is no functional impact. PR1430835
- Layer 2 cpd core is seen continuously after loading the 19.3 DCB image. PR1431866
- If Layer 2 VPN sessions have OAM control-channel option set to router-alert-label, the no-control-word option in Layer 2 VPN shouldn't be used for BFD sessions to come up.PR1432854
- The new rate 1.8 MB/s that takes 16 to 17 minutes to copy an ACX5448 image size of 1.9-Gigabit, the ACX5448 Series takes lesser time comapared to Misha. Since, the rate limiter is in bps and does not support pps based, which is a hardware limitation from DNX, the average size is 512. This rate is approximate to one-third of the Misha rate. For file copy cases, the normal packet size seen are 1500. PR1439960
- The hold timer expiry is common across all platforms. It is not specific to the ACX5448 platforms. PR1439980
- With an asymmetric network connection, EX: a 10-Gbps MACsec port connected to a 10-Gbps channelized port, high and asymmetric T1 and T4 time errors are seen. This situation introduces a high two-way time error and also different CF updates in forward and reverse paths. PR1440140
- By default, the management interface speed is always displayed as 1000 Mbps in Junos OS command output. PR1440675
- With the MACsec feature enabled and introduction of traffic, the peak-to-peak value varies with the percentage of traffic introduced. Find the maximum and mean values of the time errors with different traffic rates (for example, two router scenarios). They can have max value as high as 1054ns with 95% traffic, 640ns for 90% traffic and 137ns with no traffic. PR1441388
- Synchronous Ethernet jitter tolerance test fails for MACsec ports. For Synchronous Ethernet and MACsec, there seems to be additional framing header and footer that would get added by the MACsec protocol. The impact on the jitter test is not obvious and appears undefined in the standards and not qualifiable with a single DUT and Calnex. PR1447296

SEE ALSO

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

- The switchover time is observed to be more than 50 ms under certain soak test conditions with an increased scale along with the multi-protocol and multi-router topology. PR1387858
- L2 rewrite happens on a regular Bridge domain and Vlan interfaces. Although, there are some service dependencies, VPLS in this case, in which Egress interface map table is not updated properly with the L2 rewrite map id that causes rewrite to stop working. PR1414414
- Clock class value is wrong in the default data as it shows ptp clock when the slave interface is down in PTP-OC device. PR1416421
- With ACX5448 platform devices, the ztp process will proceed with an image upgrade. This will occurs even when there is a mismatch in the platform name of the software image stored on ftp or ztp server, and on actual platform where the ztp process is running. PR1418313
- On Ethernet bridge, the L2 filters might not work as expected when you try to match vlan based fields for untagged packets. PR1423214
- DHCP clients are not able to scale to 96000. PR1432849
- With scaled interface configuration, delete/add or deactivate/activate interface takes time. This occurs because we did not clear the mac-entries associated with the L2 logical interface and BCM api's, which consumes time. PR1433426

- The time consumed on 1-Gigabit performance is not equal to as compared to 10-Gigabit. Compensation is done to bring the mean value under class A but the peak to peak variations are high and can go beyond 100ns. It has a latency variation with peak to peak variations of around 125ns-250ns without any traffic. (for example, 5-10% of the mean latency introduced by the each phy which is of around 2.5us). PR1437175
- Errors can be seen if the CFP2 optics are not plugged in. PR1438039
- Memory leaks are expected. PR1438358
- When the interface is flapped between channelized configurations, 25-Gigabit to 100-Gigabit, the parent AE interface configuration does not get cleaned up properly. This leads to traffic failure in that interface. The issue happens with channelized interfaces AE and delete sequence. Only channelized interface deletion does not create the issue, both the interface and chassis deletion creates the issue. PR1441374
- Recovery of JUNOS volume is not possible from the OAM menu. PR1446512
- Drop profile max threshold might not reached its full limit when the packet size is other than 1000 bytes. This is due to the current design limitation. PR1448418
- If the client et-interface is up and transported state is in init state, the restart transported process updates the state to normal. This scenario is seen when the interfaces are deleted, recreated, and applied configurations. PR1449937
- Red Drops seen on the 25-Gigabit Channelized AE Interfaces after some events (For example, deactivate and activate.) on the PEER box. PR1450674
- You cannot form 125000 IGMP groups with ACX5448 receiving 125000 IGMP v2 reports per second. This is a product limitation from BCM and CPU host path queuing model. PR1454465
- When there is a failure of i2c daemon, the riofeb might crash on ACX5448. PR1455928
- The CFM REMOTE MEP does not come up after configuration or if the state remains in Start state. PR1460555
- On ACX5448 Series, when you enable the local Loopback on 10-Gigabit interface, the link does not come up. The link comes up properly on 1-Gigabit and 100-Gigabit after the local loopback is enabled. As a workaround, disable or enable the interface to keep the link up when the local loopback is enabled. PR1460715
- When you remove one member of AE and reboot the device with an AE configuration, an Arp issue occurs with AE. PR1461485

SEE ALSO

What's New | 13

What's Changed | 21

Known Limitations | 23

Resolved Issues

IN THIS SECTION

Learn which issues were resolved 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.

Resolved Issues: 19.3R2

General Routing

- The AE interface with LACP stays down after the router reboots if the link-speed is configured. PR1357012
- The optic comes with Tx enabled by default. Because the port is administratively disabled, the port is stopped however, because the port has not been started, it does not disable Tx. PR1411015
- The ACX5448:40G FEC on ACX5448, which is FEC enabled by default, must be aligned with the MX and QFX platforms, where FEC is NONE. PR1414649
- On the ACX5448-X:SKU and ACX5448-D, 96000 ARPs gets populated however, only 47000 NH entries are present. Around 50 percentage of packet drop is observed. PR1426734
- The deviation in traffic rate in the queue is around 8 to 10 percentage in some cases. PR1436297
- The interface on ACX1100 remains down while using SFP-1FE-FX (740-021487). PR1439384
- In an ACX Series, the auto exported route between VRFs might not respond for icmp echo requests. PR1446043
- The ACX5048 l2circuit with backup neighbor configuration stops passing traffic after the link flaps and move back from backup neighbor to primary. PR1449681

- On the ACX5448-X NPI platforms, after disabling 100-Gigabit and 40-Gigabit interface laser output power, the ACX5448-D interface diagnostics optics display some values. PR1452323
- The ACX5448 FPC crashes due to segmentation fault. PR1453766
- Wrong operating state is displayed in the snmp trap for removing the fan. PR1455577
- On PTX10001-20C and ACX6360, you must enable gigether option to configure the Ethernet FEC on the client ports. PR1456293
- The ACX5448-D and ACX5448-M devices does not display airflow information and temperature sensors as expected. PR1456593
- The ACX5448 L2VPN with encapsulation type Ethernet stops passing traffic after adding a random port with vlan configuration. PR1456624
- Route resolve resolution does not occur when the packet size is 10000. PR1458744
- The traffic might be blockhead during link recovery in an open Ethernet access ring with ERPS configured. PR1459446
- On the ACX5000, SNMP mib walk for jnxOperatingTemp does not return any value for FPC in the new versions. PR1460391
- On the ACX platform, the LLDP neighbor does not recover from a lag after the software is upgraded to 18.2R3-S. PR1461831

Platform and Infrastructure

- The REST API process becomes non-responsive when a number of requests come with a high rate. PR1449987

Routing Protocols

- The rpd might crash continuously due to memory corruption in the ISIS setup. PR1455432

Resolved Issues: 19.3R1

Class of Service

- When the forwarding-class is configured under firewall policer, the dfwd might crash. PR1436894

General Routing

- On QFX1000 Series platforms (QFX10002, QFX10008, and QFX10016), the 1-Gigabit copper module interface on 10G line card (QFX10000-60S-6Q) incorrectly displays with 'Link-mode: Half-duplex'. PR1286709
- Upon classifying the Layer3 packets, DSCP will not be preserved or lost at the egress due to the limitations of broadcom chipset. PR1322142

- START_BY_START_ERR interrupt handler was not available with the previous version of bcm sdk code. This lead to the status checking of this flag continously by bcmDPC process leading to high CPU utilization. This has been fixed in this release by adding a handler for this interrupt. PR1329656
- On ACX5000 platforms with Junos 16.2 onwards, if the ECC Errors occur, the FPC/fxpc process might use high CPU. This issue can be hit after the upgrade in some cases. PR1360452
- As part of the pic_periodic, before setting the port to master/slave mode, AN bit is checked if AN is complete and this would return if AN is still in progress. Since An was disabled, this port wasn't set to either mode and this was going on in a loop causing the CPU to go high. PR1360844
- On ACX1x00/ACX2x00/ACX4x00 running in 15-releases previous to 15.1R8, when configuring "mac-table-size" under bridge-domain, a wrong commit error appear not allowing the commit to pass. PR1364811
- ACX led on GE interface goes down when speed 10M is added. PR1385855
- Link Fault Signaling (LFS) feature is not supported on ACX5448 10/40/100GbE interfaces. PR1401718
- If user configures invalid speed config on TIC ports (PIC slot 1) on ACX6360-OR/OX, TIC interfaces are not created. regress@tron# show chassis fpc 0 { pic 1 { port 0 { speed 40g; } } }. PR1403546
- On ACX 1000/2000/4000/5048/5096 platforms, after a new child IFL with VLAN and filter is added on an AE IFD or changing the VLAN ID of a child IFL with filter, traffic over the AE IFD might get filtered with that filter on the child IFL. Example: ae-0/0/0 is an IFD and ae-0/0/0.100 is an IFL. PR1407855
- The **show services inline stateful-firewall flow** or **show services inline stateful-firewall flow extensive** command may cause the memory leak, which may cause the inline Nat issue. PR1408982
- When using PCEP (Path Computation Element Protocol), if a PCE (Path Computation Element) generates a PCUpdate or PCCreate message which contains a metric type other than type 2, the Junos device acting as PCC (Path Computation Client) may fail to process the message and reject the PCUpdate or PCCreate message from the PCE. When the issue occurs the LSPs' (Label-Switched Path) characteristics cannot be updated hence it may cause traffic impact. PR1412659
- On the ACX-5448, the BFD timer value are not as per the configured 900ms with multiplier 3, its showing 6.000 with multiplier 3 instead for most of the sessions. PR1418680
- On ACX5K platform, the fxpc process high CPU usage might be seen under rare condition if parity errors are detected in devices. It has no direct service/traffic impact. However since CPU utilization is high during this issue, there are some side-effects. For example, it could impact time-sensitive features such as BFD. PR1419761
- Copying images from WAN interface to RE of ACX5448 takes long time. PR1422544
- On ACX5448 box, traffic with VLAN tag which doesn't match any of the configured interfaces will be dropped. While after changing interface encapsulation from ethernet-bridge to vlan-bridge, the unmatched traffic can enter an erroneous bridge-domain which the bridge index (VSI) is the same as the vlan-id of the unmatched traffic. PR1423610

- On ACX5448 platforms, the JUNIPER_SOURCE LR4T2 optics may not work properly due to the fact that an internal defect causes it to not output power, as a result, the interface may not become up. PR1424814
- Due to BCM sdk design, EEDB hardware entry is not freed for unicast next-hop creation. This leads to resource leakage and is not allowing to higher scale. PR1426734
- In a rare condition, due to a timing issue, the FPC or fxpc might crash if the AE interface flaps. or example, deactivating or activating of the AE interface. PR1427362
- Multiple HW i2c failure observed because of intermittent I2C access failure on main board switches. PR1429047
- Chassisd can crash with unsupported hcos configuration when mx104 is used as fusion aggregation device. PR1430076
- The tx laser was enabled by default in CPLD. Therefore, the link is shown up on the peer as soon as the pfe starts. PR1430910
- L4 Hashing will work for both IPv4 & IPv6 packets, if any one of the two CLIs is enabled. To disable L4 hashing for any one of IPV4 or IPV6, both CLIs needs to be in disabled state. CLIs for reference, set forwarding-options hash-key family inet layer-4 set forwarding-options hash-key family inet6 layer-4. PR1431206
- On ACX5448, if egress link is 40G/100G, small size packets are encapsulated improperly and causing remote interface drops the packets as runts. PR1434900
- No-vrf-propagate-ttl may not work after activate or deactivate of COS configuration in acx. PR1435791
- Maximum theoritical that can go for shaping rates on a queue will be upto 10%. PR1436297
- 1PPS performance metrics (class A) of G.8273.2 are not met for 1G interfaces because of the variable latency added by the Vittesse PHY. PR1439231
- Transit DHCP packets are not punted to CPU and are transparently passthrough. PR1439518
- In an ACX5448 platforms, when the PFE failed to allocate packet buffer, portion of packet memories may not be freed. PR1442901
- On an ACX5448 box, link flaps or CoS configuration changes (specific to temporal value changes) might result in traffic drop on all interfaces and recorded as RED drops. PR1443466
- ACX5448/18.3R1-S4.1 not performing proper dot1p CoS rewrite on interfaces configured with l2circuit/local-switching/family ccc. PR1445979
- ACX5448 FPC crashed due to segmentation fault, due to timing issue. There is very low chance of this core occurring. PR1453766

Interfaces and Chassis

- On ACX series, in CFM ethernet OAM scenario, after the upgrade from 17.4 onwards, the cfmd coredump might be seen after committing configuration on CFM (connectivity-fault-management). PR1425804

Layer 2 Ethernet Services

- In 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 may need more time to get IP address assigned. The subscriber may remain in this state until it's lease expires if it had previously bound with the address in the option 50. PR1435039

MPLS

- Dynamically configured RSVP LSPs for LDP link protection may not come up after disabling/enabling protocol mpls. PR1432138

Routing Protocols

- On ACX platforms, the loopback address exported into other VRF instance might not work. PR1449410
- On ACX platforms, when there is MAC change for LDP neighbor and IP remains the same, ARP update is proper but MPLS LDP may still use the stale MAC of the neighbor. If there is any application/service such as MP-BGP using LDP as next-hop, all transit traffic pointing to the stale MAC will be dropped. PR1451217

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for the ACX Series.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Router. 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

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of 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 | 21
- Known Limitations | 23
- Open Issues | 26
- Resolved Issues | 28
- Documentation Updates | 32

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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 the Junos OS main and maintenance releases for EX Series switches.

NOTE: The following EX Series switches are supported in Release 19.3R2: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in 19.3R2

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.3R2.

What's New in 19.3R1

Hardware

- **Support for two 100-Gigabit Ethernet QSFP28 transceivers on the 2-port QSFP+/QSFP28 uplink module (EX4300-48MP and EX4300-48MP-S switches)**—Starting in Junos OS Release 19.3R1, you can install two 100-Gigabit Ethernet QSFP28 transceivers in the 2-port QSFP+/QSFP28 uplink module (model number: EX-UM-2QSFP-MR) for EX4300-48MP and EX4300-48MP-S switches. You can install two QSFP+ transceivers, two QSFP28 transceivers, or a combination of one QSFP+ transceiver and one QSFP28 transceiver in the uplink module.

  If you configure both the ports on the uplink module to operate at 100-Gbps speed, the four QSFP+ ports on the switch are disabled.

  [See EX4300 Switch Hardware Guide.]

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **802.1X trunk port and multidomain authentication (EX4300-48MP switches)**—Starting with Junos OS Release 19.3R1, 802.1X trunk port and multidomain authentication is supported on EX4300-MP switches. Authentication on the trunk port supports only single supplicant and single-secure supplicant modes.

  Multidomain authentication is an extension of 802.1X authentication for multiple supplicants, which authenticates multiple clients individually on one authenticator port. Multidomain authentication allows one VoIP client and multiple data clients to authenticate to different VLANs while on the same port. The VoIP client is authenticated to the voice VLAN while the data clients are authenticated to the data VLAN.

  [See Understanding 802.1X and VoIP on EX Series Switches.]

EVPN

- **Support for DHCP relay in an EVPN-MPLS network (EX9200 switches, MX Series, and vMX)**—Starting in Junos OS Release 19.3R1, EX9200 switches, MX Series routers, and vMX virtual routers support DHCPv4 and DHCPv6 relay in an EVPN-MPLS network. We support this feature in a data center architecture that includes a layer of spine devices that perform EVPN Layer 2 and Layer 3 functions. These devices are connected to a layer of leaf devices that perform EVPN Layer 2 functions. In this architecture, DHCP clients are connected to leaf devices, and DHCP servers are connected to spine devices. The DHCP relay functions are centralized at the spine layer. As a result, this architecture is known as the centrally routed bridging architecture.

  [See DHCP Relay Agent in EVPN-MPLS Network.]

- **IGMP snooping support for EVPN-VXLAN (EX9200 switches, MX Series, vMX)**—Starting in Junos OS Release 19.3R1, you can configure IGMP snooping on EX9200 switches, MX Series routers, and vMX virtual routers in an EVPN-VXLAN network. Enabling IGMP snooping helps to constrain multicast traffic to interested receivers in a broadcast domain.

  The listed devices support these IGMP snooping use cases in a centrally routed bridging overlay (an EVPN-VXLAN network with a two-layer IP fabric):

  - Forwarding multicast traffic within a VLAN (intra-VLAN)
  - Routing multicast traffic between VLANs (inter-VLAN) using one of the following methods:
    - IRB interfaces configured with Physical Interface Module (PIM) on an elected designated router (DR) device
    - A PIM gateway with Layer 2 or Layer 3 connectivity
    - An external multicast router

  The listed devices support these IGMP versions and membership report modes:

  - IGMPv2 with Any-Source Multicast (ASM) (*,G) mode only.
  - IGMPv3 in either of the following modes:
    - ASM (*,G)—the default behavior.
    - Single-Source Multicast (SSM) (S,G)—you must explicitly configure by including the **evpn-ssm-reports-only** configuration statement at the **[edit protocols igmp-snooping]** hierarchy level.

  [See Overview of IGMP Snooping in an EVPN-VXLAN Environment.]

Forwarding and Sampling

- **Customizing hashing parameters and shared-buffer alpha values for better load balancing (EX4650 and QFX5120 switches)**—These switches achieve load balancing through use of a hashing algorithm, which determines how to forward traffic over LAG bundles or to next-hop devices when ECMP is enabled. The hashing algorithm makes hashing decisions based on values in various packet fields. Starting with Junos OS Release 19.3R1, you can explicitly configure some hashing parameters to make hashing more efficient. The shared-buffer pool is a global memory space that all ports on the switch share dynamically as they need buffers. The switch uses the shared-buffer pool to absorb traffic bursts after the dedicated-buffer pool is exhausted. The shared-buffer pool threshold is dynamically calculated based on a factor called alpha. Also starting with Junos OS Release 19.3R1, you can specify the alpha, or dynamic threshold, value to determine the change threshold of shared buffer pools for both ingress and egress buffer partitions.

  To specify hashing parameters:

  ```
  user@switch# set forwarding-options enhanced-hash-key hash-parameters (ecmp | lag)
  ```

To specify a threshold value for a particular queue:

`user@switch#` **set class-of-service shared-buffer (ingress|egress) buffer-partition *buffer* dynamic-threshold *value***

[See hash-parameters and buffer-partition.]

Interfaces and Chassis

- **Power over Ethernet IEEE 802.3bt (EX4300-48MP switches)**—The IEEE 802.3bt standard for Power over Ethernet (PoE) is supported on EX4300-48MP switches. The IEEE 802.3bt standard enables delivery of up to 90 W over all four pairs of wire in a standard RJ-45 Ethernet cable.

  [See Understanding PoE on EX Series Switches.]

Junos Telemetry Interface

- **JTI support extended for Junos kernel GRES and RTSOCK (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, and PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel graceful Routing Engine switchover (GRES) and routing socket (RTSOCK) information using remote procedure call (gRPC) services. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel GRES and routing socket information:

  - /junos/chassis/gres/
  - /junos/kernel/rtsock/

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

- **JTI support extended for Junos kernel LAG, NSR, and TCP (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos Telemetry Interface (JTI) extends support for streaming Junos kernel Link Aggregation Group (LAG), nonstop Routing (NSR) Junos socket replication (JSR), and Transport Control Protocol (TCP) information using remote procedure call (gRPC) services. Device monitoring and network analytics applications can use Junos kernel sensors to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel LAG, NSR, and TCP information:

  - /junos/chassis/aggregated-devices/
  - /junos/routing-options/nonstop-routing/
  - /junos/kernel/tcpip/tcp/

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

- **JTI support extended for Junos kernel IPv4 and IPv6 (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, and PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel IPv4 and IPv6 information using remote procedure call (gRPC) services. Device monitoring and network analytics applications can use Junos kernel sensors to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel IPv4 and IPv6 information:

  - /junos/kernel/tcpip/arp/ — Address Resolution Protocol cache
  - /junos/kernel/tcpip/ndp/ — Neighbor Discovery Protocol cache
  - /junos/kernel/tcpip/netisr/ — NETISR network queues
  - /junos/kernel/tcpip/nhdix/ — Next-hop index space exhaustion
  - /junos/kernel/tcpip/rtb/ — Route tables

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

- **JTI support extended for Junos kernel IP multicast, tunnels, TNP, and VPLS (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel IP multicast, tunnels, Trivial Network Protocol (TNP), and virtual private LAN service (VPLS) information using remote procedure call (gRPC) services. Device monitoring and network analytics applications can use Junos kernel sensors to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel IP multicast, tunnels, TNP, and VPLS information:

  - /junos/kernel/multicast/
  - /junos/kernel/tunnel/
  - /junos/kernel/tnp/
  - /junos/kernel/vpls/

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

Layer 2 Features

- **Configuring Q-in-Q tagging behavior for the native VLAN (EX4300 and EX4300-MP switches and Virtual Chassis)**—Starting in Junos OS Release 19.3R1, when Q-in-Q tunneling is configured and an untagged packet is received on a C-VLAN interface, you can configure these switches to add either one or two tags before sending the packet out of the S-VLAN interface. To send two tags, set the configuration statement *input-native-vlan-push* to "enable" and ensure that the *input-vlan-map* configuration is set to "push".

  [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translations.]

Management

- **OpenConfig AAA data model support (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 19.3R1 supports the configuration leafs specified in the OpenConfig AAA data model. Mapping the OpenConfig AAA configuration to the Junos AAA configuration using the following YANG files in the data model makes this support possible:

  - openconfig-aaa.yang
  - openconfig-aaa-types.yang
  - openconfig-aaa-tacacs.yang
  - openconfig-aaa-radius.yang

  The configuration model supporting the OpenConfig data model includes:

  - A translation script (.py / .slax) that maps each config leaf in the OpenConfig Schema to one or more config leafs in the JUNOS Schema.
  - A deviation file (.yang) that specifies how much the implementation deviates from the vendor-neutral model.

[See Mapping OpenConfig AAA Commands to Junos Configuration.]

Multicast

- **MLDv1, MLDv2, and MLD snooping (EX4650 and QFX5120-48Y switches and Virtual Chassis)**—Starting in Junos OS Release 19.3R1, you can configure Multicast Listener Discovery (MLD) version 1 (MLDv1), MLD version 2 (MLDv2), and MLD snooping on EX4650 and QFX5120-48Y switches and Virtual Chassis. With MLD snooping enabled, these switches or Virtual Chassis replicate and forward IPv6 traffic for a multicast group only to the interfaces in a VLAN with listeners who joined the group, rather than flooding to all interfaces in the VLAN.

  [See Examples: Configuring MLD and Understanding MLD Snooping.]

Routing Policy and Firewall Filters

- **Support for IPv6 filter-based forwarding (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 19.3R1, you can use stateless firewall filters in conjunction with filters and routing instances to control how IPv6 traffic travels in a network on EX4650 and QFX5120 switches. This is called IPv6 filter-based forwarding. To set up this feature, you define a filtering term that matches incoming packets based on the source or destination address and then specify the routing instance to send packets to. You can use filter-based forwarding to route specific types of traffic through a firewall or security device before the traffic continues on its path. You can also use it to give certain types of traffic preferential treatment or to improve load balancing of switch traffic.

  [See Firewall Filter Match Conditions for IPv6 Traffic and Filter-Based Forwarding Overview.]

Routing Protocols

- **RIPng routing protocol supported (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 19.3R1, EX4650 and QFX5120 switches support the RIPng routing protocol.

  [See Basic RIPng Configuration.]

Services Applications

- **Support for real-time performance monitoring or RPM (EX4650)** —Starting in Junos OS Release 19.3R1, you can configure active probes to track and monitor traffic across the network and to investigate network problems on EX4650 switches.

  You can use RPM in the following ways:

  - Monitor time delays between devices.
  - Monitor time delays at the protocol level.
  - Set thresholds to trigger SNMP traps when values are exceeded.

    You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test.

- Determine automatically whether a path exists between a host router or switch and its configured BGP neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

[See Understanding Real-Time Performance Monitoring on Switches.]

Software Installation and Upgrade

- **Migration of Linux kernel version**—Starting in Junos OS Release 19.3R1, the following devices support the Wind River Linux 9 (WRL9) kernel version:

Platforms	Routing Engine Supported
ACX5448-D	RE-ACX-5448
MX240, MX480, and MX960	RE-S-X6-64G
MX2020 and MX2010	REMX2K-X8-64G
MX204	RE-S-1600x8
MX10003	RE-S-1600x8
MX2008	RE-MX2008-X8-64G
MX10016	RE X10
MX10008	RE X10
PTX5000	RE-PTX-X8-64G
PTX3000	RCBPTX
PTX10016	RE-PTX-2X00x4/RE X10
PTX10008	RE-PTX-2X00x4/RE X10
PTX1000	RE-PTX1000
PTX10002-XX	RE-PTX10002-60C
EX9208	RE-S-EX9200-2X00x6
EX9251	EX9251-RE

Platforms	Routing Engine Supported
EX9253	EX9253-RE
EX9204	RE-S-EX9200-2X00x6
EX9214	RE-S-EX9200-2X00x6
QFX10002	RE-QFX10002-60C
QFX10008	RE-QFX10008
QFX10016	RE-QFX10016

Starting in Junos OS Release 19.3R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following routers:

- MX Series—MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008
- PTX Series—PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

If you perform a software upgrade on a router with i40e NVM version earlier than 6.01, the upgrade fails and the following error message is displayed:

ERROR: i40e NVM firmware is not compatible ,please upgrade i40e NVM before installing this package

ERROR: Aborting the installation

ERROR: Upgrade failed

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

Virtual Chassis

- **Virtual Chassis support (EX4650 and QFX5120-48Y switches)**—Starting in Junos OS Release 19.3R1, you can interconnect two EX4650 or two QFX5120-48Y switches into a Virtual Chassis, which operates as one logical device managed as a single chassis.
  - Member switches must be two EX4650 or two QFX5120 switches (no mixed mode).
  - Both member switches take the Routing Engine role with one as master and one as backup.
  - You can use any of the 100-Gbps QSFP28 or 40-Gbps QSFP+ ports on the front panel (ports 48 through 55) as Virtual Chassis ports (VCPs) to connect the member switches.
  - You can run nonstop software upgrade (NSSU) to update the Junos OS release on both member switches with minimal traffic disruption during the upgrade.
  - EX4650 and QFX5120 Virtual Chassis support the same protocols and features as the standalone switches in Junos OS Release 19.3R1 except for the following:

- IEEE 802.1X authentication
- EVPN-VXLAN (QFX5120)
- Layer 2 port security features, DHCP, and DHCP snooping
- Junos telemetry interface (JTI)
- MPLS
- Multichassis link aggregation (MC-LAG)
- Redundant trunk groups (RTG)
- Priority-based flow control (PFC)

Configuration parameters and operation are the same as for other non-mixed EX Series and QFX Series Virtual Chassis.

[See Virtual Chassis Overview for Switches.]

SEE ALSO

What's Changed

IN THIS SECTION

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

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

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

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

  See Configuring RFC-Compliant NETCONF Sessions.

What's Changed in 19.3R2

There are no changes in behavior and syntax for EX Series in Junos OS Release 19.3R2.

What's Changed in 19.3R1

General Routing

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the show interfaces mc-ae extensive command. The output now displays the following two additional fields:

  - Local Partner System ID?LACP partner system ID as seen by the local node.
  - Peer Partner System ID?LACP partner system ID as seen by the MC-AE peer node.

  Previously, the **show interfaces mc-ae extensive** command did not display these additional fields.

See [show interfaces mc-ae.].

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Release 19.3R1 and later, EX Series switches support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

  In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

Junos OS XML API and Scripting

- **Range defined for confirm-timeout value in NETCONF and Junos XML protocol sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, the value for the **<confirm-timeout>** element in the Junos XML protocol **<commit-configuration>** operation must be in the range 1 through 65,535 minutes, and the value for the **<confirm-timeout>** element in the NETCONF **<commit>** operation must be in the range 1 through 4,294,967,295 seconds. In earlier releases, the range is determined by the minimum and maximum value of its unsigned integer data type.

Layer 2 Features

- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—From Junos OS Release 19.3R1, the configuration statement **input-native-vlan-push** at the **[edit interfaces *interface-name*]** hierarchy level is introduced. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement **input-vlan-map** with a **push** operation is configured.

  [See input-native-vlan-push.]

System Logging

- **Preventing system instability during core file generation (EX Series)**—Starting with Release 19.3R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

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 a VLAN uses an IRB interface as the routing interface, the **vlan-id** parameter must be set to **none** to ensure proper traffic routing. This issue is platform-independent. PR1287557

Platform and Infrastructure

- During software upgrade to a more recent Junos OS Release 19.3 image, the system hung right after the **request system software add /var/tmp/ *image.gz*** was issued. Power-cycling can recover the device. PR1405629
- When the device is loaded and unloaded with MACsec configuration multiple times with operations made continuously, Layer 3 connectivity is lost and hence the system is halted. You need to reboot the device to resume operations. PR1416499
- Filters are installed only during the addition of routes if there is enough space. If it fails due to nonavailability of TCAM space, those routes will not be processed for filter addition later when space becomes available. PR1419926

- On deactivating and activating PoE, the PoE interfaces draw more power (as per display) for quite some time. PR1431647
- The **set class-of-service shared-buffer ingress buffer-partition lossless-headroom percent 0** command is not supported in a Virtual Chassis, as the VCP ports should have some heardroom to support PFC. The configuration will be rejected at the hardware layer with a log message. PR1448377

SEE ALSO

What's New | 35

What's Changed | 44

Open Issues | 48

Resolved Issues | 50

Documentation Updates | 58

Migration, Upgrade, and Downgrade Instructions | 59

Open Issues

IN THIS SECTION

Learn about open issues 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.

Authentication and Access Control

- Before running the **load ssl-certificate path PATHNAME** command, we have to now configure the path using the **set protocols dot1x ssl-certificate-path PATHNAME** command, if the pathname is not the default path **/var/tmp/**. PR1431086

Infrastructure

- On EX3400 and EX2300 during ZTP with configuration and image upgrade with FTP as file transfer, image upgrade is successful but sometimes VM core files are observed. PR1377721
- On EX Series platforms, if configuring large-scale number of firewall filters on some interfaces, the FPC might crash with the generation of core files. PR1434927
- On EX2300 and EX2300-C platforms, if Junos OS is with FreeBSD kernel version 11 with the build date on or after 2019-02-12, the switch might stop forwarding traffic or responding to the console. To recover the switch, power-cycle it. PR1442376

Interfaces and Chassis

- On GRES, VSTP port cost on aggregated Ethernet interfaces might get changed, leading to a topology change. PR1174213

Platform and Infrastructure

- On a EX2300 switch, the output of the **show chassis routing-engine** command might display an incorrect value of **mac reset** for the **last reboot reason** field. PR1331264
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur. Even with enhanced convergence configured, there is no guarantee that subsecond convergence will be achieved. PR1371493
- On EX9208, 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
- You can see multiple reconnect logs, such as **JTASK_IO_CONNECT_FAILED**, during the device initialization. There is no functionality impact due to these messages. These messages can be ignored. PR1408995
- uRPF check in strict mode does not work properly. PR1417546
- On EX9208, traffic loss is observed if ingress and egress ports are in different FPCs. PR1429714
- Traffic drop is seen from SD to AD. Loss is intermediate and not seen regularly. This is because a few packets transmitted from the egress of AD1 are short of FCS (4 bytes) + 2 bytes of data due to which the drops occur. It is seen that the normal data packets are of size 128 bytes (4 bytes FCS + 14 bytes

Ethernet header + 20 bytes IP header + 90 bytes data), while the corrupted packet is 122 bytes (14 bytes Ethernet header + 20 bytes IP header + 88 bytes data). PR1450373

- On EX9208, 33% degradation with MAC learning rate in Junos OS Release 19.3R1 is seen while comparing with Junos OS Release 18.4R1. PR1450729
- On EX9208, l2ald and eventd are hogging 100% after issuing **clear ethernet-switching**, and the following syslog errors are continuously seen: **l2ald[18605]: L2ALD_IPC_MESSAGE_INVALID: Invalid message received (message type 0, subtype 0): null message**. PR1452738
- In case you use a specific DAC cable, EX3400 might not recognize it properly and Virtual Chassis function does not work as a result. PR1463635

Routing Protocols

- BGP IPv4/IPv6 convergence and RIB install or delete time is degraded in Junos OS Release 19.1R1, 19.2R1, 19.3R1, and 19.4R1. PR1414121

SEE ALSO

What's New | 35

What's Changed | 44

Known Limitations | 47

Resolved Issues | 50

Documentation Updates | 58

Migration, Upgrade, and Downgrade Instructions | 59

Resolved Issues

IN THIS SECTION

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

Resolved Issues: 19.3R2

EVPN

- ARP request or neighbor solicitation message might be sent back to the local segment by the designated forwarder (DF) router. PR1459830

Infrastructure

- Certain EX Series platforms might generate VM core files by panic and reboot. PR1456668
- Traffic is dropped on EX4300-48MP device acting as a leaf device in a Layer 2 IP fabric EVPN VXLAN environment. PR1463318

Interfaces and Chassis

- VRRP-V state is flapping with init and idle states after configuring **vlan-tagging**. PR1445370

J-Web

- Some error messages might be seen when using J-Web. PR1446081

Layer 2 Features

- ERPS nodes do not converge to idle state after failure recovery or reboot. PR1431262
- MAC or ARP learning might not work for copper base SFP-T transceiver on EX4600. PR1437577

Platform and Infrastructure

- LACP DDoS policer is incorrectly triggered by other protocols traffic on all EX92XX/T4000 and MX platforms. PR1409626
- The interface on a failed member FPC of EX2300 and EX3400 Virtual Chassis might stay up for 120 seconds. PR1422507
- LED turns on even after the Virtual Chassis members are powered off. PR1438252

- The rpd might generate core files during router bootup due to a file pointer issue as there are two code paths that can close the file. We are attempting to close the file without validating the file pointer. PR1438597
- EX4600 Virtual Chassis does not come up after you replace the fiber connection on the Virtual Chassis port with a DAC cable. PR1440062
- MAC addresses learned on RTG might not be aged out after a Virtual Chassis member is rebooted. PR1440574
- On Ex4300-MP, the following log is generated continuously: **rpd[6550]: task_connect: task AGENTD I/O.: Connection refused** PR1445618
- Traffic might be dropped when a firewall filter rule uses **then vlan** as the action in a Virtual Chassis scenario. PR1446844
- Unicast ARP requests are not replied to with **no-arp-trap** option. PR1448071
- Tunneling-encapsulated packets are dropped on the Layer 3 VPN MPLS PE-CE interface. PR1451032
- DHCP snooping static binding does not take effect after deleting and re-adding the entries. PR1451688
- Configuration change in VLAN all option might affect the per-VLAN configuration. PR1453505
- Version compare in PHC might fail, making the PHC to download the same image. PR1453535
- REST API process becomes nonresponsive when a number of requests come with a high rate. PR1449987
- Packet drop might be seen after removing and reinserting the SFP transceiver of the 40G Uplink Module ports. PR1456039
- The syslog **timeout connecting to peer database-replication** is generated when the **show version detail** command is issued. PR1457284
- Overtemperature SNMP trap messages appear after an update even though the temperature is within the system thresholds. PR1457456
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used. PR1458559
- Storage space limitation leads to image installation failure during Phone home on EX2300 and EX3400 platforms. PR1460087
- ERP might not revert to idle state after reload or reboot of multiple switches. PR1461434

Routing Protocols

- Junos BFD sessions with authentication flap after a certain time. PR1448649

User Interface and Configuration

- EX4600 and QFX5100 are unable to commit baseline configuration after zeroization. PR1426341
- Problem with access to J-Web after an update from Junos OS Release 18.2R2 to Release 18.2R3. PR1454150

Virtual Chassis

- Current MAC address might change when deleting one of the multiple Layer 3 interfaces. PR1449206

Resolved Issues: 19.3R1

EVPN

- The device may proxy the ARP probe packets in an EVPN environment. PR1427109
- Configuring ESI on a single-homed 25-Gbps port might not work. PR1438227

Forwarding and Sampling

- Enable interface with input/output vlan-maps to be added to a routing-instance configured with a VLAN ID or VLAN tags (instance type virtual-switch/vpls). PR1433542

Infrastructure

- The traffic to the NLB server may not be forwarded if the NLB cluster works in multicast mode. PR1411549
- The operations on the console might not work if the **system ports console log-out-on-disconnect** statement is configured. PR1433224

Interfaces and Chassis

- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. PR1402606
- The IFLs in EVPN routing instances might flap after committing configurations. PR1425339
- EX9214 unexpected **duplicate VLAN-ID** commit error. PR1430966

Layer 2 Ethernet Services

- DHCP request may get dropped in DHCP relay scenario. PR1435039
- On EX9200 switches, DHCP Relay is stripping the 'GIADDR' field in messages toward the DHCP clients. PR1443516

Network Management and Monitoring

- Overtemperature trap does not sent out even though there is the Temperature Hot alarm. PR1412161

Platform and Infrastructure

- [SIRT]Certain QFX and EX Series devices are vulnerable to 'Etherleak' memory disclosure in Ethernet padding data (CVE-2017-2304). PR1063645
- Transit OSPF traffic over Q-in-Q tunneling might be dropped if a firewall filter is applied to Lo0 interface. PR1355111
- The l2ald process might crash and generate a core file on EX2300 Virtual Chassis when a trunk port is converted to dot1x access port with tagged traffic flowing. PR1362587
- QFX5120 and EX4650 : Convergence delay between PE1 and P router link is more than expected delay value. PR1364244
- IPv6 router advertisement (RA) messages might increase internal kernel memory usage. PR1369638
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface. PR1403528
- MAC address movement might not happen in Flexible Ethernet Services mode when family **inet/inet6** and **vlan-bridge** are configured on the same IFD. PR1408230
- EX9251, EX9253, and EX9208: DDoS violation for LLDP, MVRP, provider MVRP and dot1x is incorrectly reported as LACP DDoS violation. PR1409626
- EX2300-24P, error message: **dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family**. PR1410717
- EX4300-48MP : Chassis Status LED glow yellow instead of amber. PR1413194
- The upgrade of the PoE firmware might fail on EX3400. PR1413802
- EX3400 : **show chassis environment** repeats "OK" and "Failed" at short intervals. PR1417839
- The EX3400 Virtual Chassis status might be unstable during the bootup of the Virtual Chassis or after the Virtual Chassis port flaps. PR1418490
- EX4300-48MP-18.3R1.9 //Over Temperature SNMP trap generated wrongky for LC (EX4300-48P) based on the master Routing Engine (EX4300-48MP) temperature threshold value. PR1419300
- EX4300: Runt counter never incremented. PR1419724
- The pfex process might crash and generates core files when you reinsert SFP. PR1421257
- Commit of configurations involving **interface-range** defined over wildcard range such as ge-*/*/* not supported. PR1421446

- Virtual Chassis may become unstable and FXPC core files may be generated when there are a lot of configured filter entries. PR1422132
- Traffic loss when one of the logical interfaces on the LAG is deactivated or deleted. PR1422920
- Ensure phone-home works in factory-default configuration. PR1423015
- Adding the second IRB interface to an aggregated Ethernet interface and rolling it back might cause the first IRB interface to stop working. PR1423106
- IPv6 multicast traffic received on one Virtual Chassis member might be dropped when egressing on other Virtual Chassis member if MLD snooping is enabled. PR1423310
- EX3400 : Auto negotiation status shows incomplete on ge-0/2/0 using SFP-SX. PR1423469
- Multicast traffic might be silently dropped on ingress port with **igmp-snooping** enabled. PR1423556
- MACsec connection on EX4600 platforms might not come back up after interface disconnect/reconnect. PR1423597
- On MX204 optics "SFP-1GE-FE-E-T" I2C read errors are seen when an SFP-T is inserted into a disabled state port. PR1423858
- The auditd crashed when Accounting RADIUS server was not reachable. PR1424030
- The native VLAN ID of packets might fail to be removed when leaving out. PR1424174
- MAC overlapping between different switches. PR1425123
- SNMP (ifHighSpeed) value is not getting displayed appear properly only for VCP interfaces, it is getting displayed as zero. PR1425167
- The jdhcpd might consume 100% CPU and crash if **dhcp-security** is configured. PR1425206
- Interface flapping scenario might lead to ECMP next-hop installation failure on EX4300s. PR1426760
- Virtual Chassis split after network topology changed. PR1427075
- The fxpc/Packet Forwarding Engine might crash on EX2300 or EX3400 platforms. PR1427391
- Rebooting or halting a Virtual Chassis member might cause 30 seconds down on RTG link. PR1427500
- IPv6 traffic might be dropped when static /64 IPv6 routes are configured. PR1427866
- VIP might not forward the traffic if VRRP is configured on an aggregated Ethernet interface. PR1428124
- EX2300-24P : l2ald core files observed after removal and re-addition of multiple supplicant mode with PVLAN on interface. PR1428469
- Data port LEDs are off even while interfaces are up. PR1428703
- CI-PR: Verification of ND inspection with a dynamically bound client, moved to a different VLAN on the same port is failing. PR1428769
- The delay in transmission of BPDUs after GRES might result in loss of traffic on EX2300/3400 Virtual Chassis. PR1428935

- When forward-only is set within **dhcp-reply**, dhcp declines are not forwarded to server. PR1429456
- EX4300 does not drop FCS frames with CRC error on xe- interfaces. PR1429865
- Unicast ARP requests are not replied to with the **no-arp-trap** option. PR1429964
- EX4300 without soft error recovery(parity check, correction and memscan) enable. PR1430079
- The jdhcpd_era log files constantly consume 121M of space out of 170M, resulting into file system full and traffic impact. PR1431201
- EX4300-48MP switch cannot learn MAC address through some access ports that are directly connected to a host when auto-negotiation is used. PR1430109
- Disabling DAC QSFP port may not work on MX204, MX10003, or EX9251. PR1430921
- Incorrect model Information while polling through SNMP from Virtual-Chassis. PR1431135
- The ERPS failover does not work as expected on an EX4300 device. PR1432397
- Native VLAN might not take into effect when it is enabled with flexible VLAN tagging on a Layer 3 subinterface. PR1434646
- The device might not be accessible after the upgrade. PR1435173
- The mc-ae interface may get stuck in waiting state in a dual mc-ae scenario. PR1435874
- i40e NVM upgrade support for EX9200 platform. PR1436223
- The FPC/pfex crash may be observed due to DMA buffer leaking. PR1436642
- The **/var/db/scripts** directory might be deleted after executing **request system zeroize**. PR1436773
- Commit check error for VSTP on EX9200s: **xSTP:Trying to configure too many interfaces for given protocol**. PR1438195
- The DHCP Snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it. PR1438351
- The dot1x might not work when **captive-port** is also configured on the interface on backup/non-master FPC. PR1439200
- DHCPv6 relay binding is not up while verifying the DHCP snooping along with DHCPv6 relay. PR1439844
- The ports of the EX device might stay in up state even if the EX46XX/QFX51XX series device is rebooted. PR1441035
- Clients in an isolated VLAN might not get IP addresses after completing authentication when both **dhcp-security** and **dot1x** are configured. PR1442078
- EX3400 fan alarm (Fan X not spinning) appears and disappears repeatedly after the fan tray is removed (absent). PR1442134
- DHCPv6 client might fail to get an IP address. PR1442867
- Non-designated port is not moving to backup port role. PR1443489
- **/var/host/motd does not exist** message is flooded every 5 seconds in chassisd logs. PR1444903

- On EX4300-MP, log generated continuously: **rpd[6550]: task_connect: task AGENTD I/O.128.0.0.1+9500 addr 128.0.0.1+9500: Connection refused**. PR1445618
- CI-PR: On EX3400 - dot1xd core files found @ **macsec_update_intf macsec_destroy_ca**. PR1445764
- Major alarm log messages for temperature conditions for EX4600 at 56 degrees Celsius. PR1446363
- The traffic might be dropped when a firewall filter rule uses 'then vlan' as the action in a Virtual Chassis scenario. PR1446844
- The PoE might not work after upgrading the PoE firmware on EX4300 platforms. PR1446915
- The firewall filters might not be created with error logs after reboot. PR1447012
- Phone home on EX3400 fails because sysctl cannot read the device serial number. PR1447291
- Added CLI configuration **on-disk-failure** on EX3400. PR1447853
- Unicast ARP requests are not replied with the **no-arp-trap** option. PR1448071
- On EX3400, IPv6 routes received through BGP do not show the correct age time. PR1449305
- Incoming Layer3-encapsulated packets are dropped on Layer 3VPN MPLS PE-CE interface. PR1451032

Routing Protocols

- Host-destined packets with filter log action might not reach the Routing Engine if log/syslog is enabled. PR1379718
- Sometimes, IGMP snooping may not work. PR1420921
- The multicast traffic might be dropped when proxy mode is used for **igmp-snooping**. PR1425621
- The error message **RPD_DYN_CFG_GET_PROF_NAME_FAILED: Get profile name for session XXX failed: -7**, may be seen in syslog after restarting routing daemon. PR1439514
- The bandwidth value of the DDoS-protection might cause the packets loss after the device reboot. PR1440847
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. PR1443507
- Loopback address exported into other VRF instance might not work on EX Series, QFX Series, or ACX Series platforms. PR1449410
- MPLS LDP may still use stale MAC of the neighbor even the LDP neighbor's MAC changes. PR1451217

Subscriber Access Management

- EX4300 **/var** showing full **/var/log/dfcd_enc file** grows in size PR1425000

User Interface and Configuration

- EX4600 and QFX5100 were unable to commit baseline configuration after zeroization. PR1426341

Virtual Chassis

- Current MAC address might change when deleting one of the multiple Layer 3 interfaces. PR1449206

VPNs

- MVPN using PIM dense mode does not prune the OIF when PIM prune is received. PR1425876

SEE ALSO

- What's New | 35
- What's Changed | 44
- Known Limitations | 47
- Open Issues | 48
- Documentation Updates | 58
- Migration, Upgrade, and Downgrade Instructions | 59

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for the EX Series switches.

SEE ALSO

- What's New | 35
- What's Changed | 44
- Known Limitations | 47

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for the EX Series. 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

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://support.juniper.net/support/eol/software/junos/.

SEE ALSO

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

NOTE: For a complete list of all hardware and software requirements for a Junos Fusion Enterprise, including which Juniper Networks devices can function as satellite devices, see Understanding Junos Fusion Enterprise Software and Hardware Requirements .

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 19.3R2 for Junos fusion for enterprise.

> **NOTE:** For more information about the Junos Fusion Enterprise features, see the Junos Fusion Enterprise User Guide.

SEE ALSO

- What's Changed | 61
- Known Limitations | 62
- Known Issues | 62
- Resolved Issues | 63
- Documentation Updates | 64
- Migration, Upgrade, and Downgrade Instructions | 65

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 19.3R2 for Junos fusion for enterprise.

SEE ALSO

- What's New | 61
- Known Limitations | 62
- Known Issues | 62
- Resolved Issues | 63
- Documentation Updates | 64
- Migration, Upgrade, and Downgrade Instructions | 65

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 19.3R2 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

IN THIS SECTION

Learn about open issues in this release 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.

Junos Fusion for Enterprise

- On EX4300 when 10G fiber port is using 1G Ethernet SFP optics, auto-negotiation is enabled by default. To bring up the satellite device, BCM recommends to disable the auto-negotiation for PHY84756 ports. PR1420343
- In a Junos fusion for enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, loop-detect feature does not work for ports 0-23, since the loop detect filter is not properly applied. PR1426757
- In a Junos fusion for enterprise environment, when traffic originates from a peer device connected to the aggregation device and the ICL is a LAG, there might be a reachability issue if the cascade port is disabled and traffic has to flow through the ICL LAG to reach the satellite device. As a workaround, use single interface as the ICL instead of a LAG. PR1447873
- Traffic drop seen from satellite device to the aggregation device on L3 unicast OSPF routes. PR1450373

SEE ALSO

What's New | 61

What's Changed | 61

Known Limitations | 62

Resolved Issues | 63

Documentation Updates | 64

Migration, Upgrade, and Downgrade Instructions | 65

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in Junos OS main and maintenance releases for Junos Fusion 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: 19.3R2

There are no resolved issues in Release 19.3R2.

Resolved Issues: 19.3R1

- Traffic might be dropped in Junos Fusion Enterprise scenario with dual aggregation devices. PR1417139
- The 1G SFP in 10G upstream port on EX3400 and EX4300 satellite devices might not come up. PR1420343
- The loop-detect feature does not work in Junos Fusion Enterprise. PR1426757

SEE ALSO

- What's New | 61
- What's Changed | 61
- Known Limitations | 62
- Known Issues | 62
- Documentation Updates | 64
- Migration, Upgrade, and Downgrade Instructions | 65

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 for documentation for Junos fusion for enterprise.

SEE ALSO

- What's New | 61
- What's Changed | 61
- Known Limitations | 62
- Known Issues | 62
- Resolved Issues | 63
- Migration, Upgrade, and Downgrade Instructions | 65

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 Enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos Fusion 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:

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

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

`user@host>` **request system software add validate reboot *source*/*package*-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 Enterprise. See Configuring or Expanding a Junos Fusion Enterprise.

For satellite device hardware and software requirements, see Understanding Junos Fusion 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 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

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html

Downgrading from Junos OS

Junos Fusion 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 Enterprise from Junos OS Release 19.2R1, follow the procedure for upgrading, but replace the 19.2 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 for the Junos Fusion 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

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.3R2.

SEE ALSO

- What's Changed | 71
- Known Limitations | 72
- Open Issues | 72
- Resolved Issues | 73
- Documentation Updates | 74
- Migration, Upgrade, and Downgrade Instructions | 75

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 Junos OS Release 19.3R2 for Junos fusion for provider edge.

SEE ALSO

- What's New | 71
- Known Limitations | 72
- Open Issues | 72
- Resolved Issues | 73
- Documentation Updates | 74
- Migration, Upgrade, and Downgrade Instructions | 75

Known Limitations

IN THIS SECTION

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

Junos Fusion Provider Edge

- To configure TACACS+ authentication in the JUNOS device, you must define a local user using an authorization template. The TACACS+ server is configured with the local user so that the user can inherit the privileges after the user is authenticated. In the Fusion setup, if the defined local user used by TACACS+ contains a capital letter, the TACACS+ users might not be able to get access the satellite devices from the aggregate device. PR1431897

SEE ALSO

- What's New | 71
- What's Changed | 71
- Open Issues | 72
- Resolved Issues | 73
- Documentation Updates | 74
- Migration, Upgrade, and Downgrade Instructions | 75

Open Issues

IN THIS SECTION

This section lists the Open Issues in hardware and software in Junos OS Release 19.3R2 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.

Junos Fusion Provider Edge

- The IGMP membership is not getting learnt by the AD fully even when the IGMP queries are been sent out. PR1419265
- The traffic drops from SD to AD. The loss is intermediate and not seen regularly. This occurs because few packets that are transmitted from the egress of AD1 is short of FCS of data. It is observed that the normal data packets are of size 128 bytes while the corrupted data packet is of size 122 byte. PR1450373

SEE ALSO

What's New | 71

What's Changed | 71

Known Limitations | 72

Resolved Issues | 73

Documentation Updates | 74

Migration, Upgrade, and Downgrade Instructions | 75

Resolved Issues

IN THIS SECTION

- Resolved Issues: 19.3R2 | 74
- Resolved Issues: 19.3R1 | 74

Learn which issues were resolved in Junos OS main and maintenance releases for Junos Fusion Provider Edge.

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

Resolved Issues: 19.3R2

There are no fixed issues in Junos OS 19.3R2 for ACX Series.

Resolved Issues: 19.3R1

Junos Fusion for Provider Edge

- Auto-negotiation is not disabled in hardware after setting **no-auto-negotiation** option in the CLI. PR1411852
- Junos fusion: Incorrect power values for extended optical ports. PR1412781
- The sdpd process may continuously crash if there are more than 12 cascade ports configured to a satellite device. PR1437387
- The aggregated Ethernet interface might flap whenever a new logical interface is added to it. PR1441869
- Deprecate Junos fusion support on QFX10000. PR1448245

SEE ALSO

What's New | 71

What's Changed | 71

Known Limitations | 72

Open Issues | 72

Documentation Updates | 74

Migration, Upgrade, and Downgrade Instructions | 75

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for Junos fusion for provider edge.

SEE ALSO

What's New | 71

What's Changed | 71

Known Limitations | 72

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 19.3R2 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 Provider Edge.

  ```
  user@host> request system software add validate reboot
  source/jinstall64-19.3R2.SPIN-domestic-signed.tgz
  ```

- For 32-bit software:

  ```
  user@host> request system software add validate reboot
  source/jinstall-19.3R2.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 Provider Edge.

  ```
  user@host> request system software add validate reboot
  source/jinstall64-19.3R2.SPIN-export-signed.tgz
  ```

- For 32-bit software:

  ```
  user@host> request system software add validate reboot
  source/jinstall-19.3R2.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 19.3R2 **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 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 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 19.3R2, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html.

Downgrading from Junos OS Release 19.3

To downgrade from Release 19.3 to another supported release, follow the procedure for upgrading, but replace the 19.3 **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

- What's New | 71
- What's Changed | 71
- Known Limitations | 72
- Open Issues | 72
- Resolved Issues | 73
- Documentation Updates | 74

Junos OS Release Notes for MX Series 5G Universal Routing Platform

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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 main and maintenance releases for MX Series routers.

What's New in Release 19.3R2

Hardware

> **NOTE:** The MX2K-MPC11E line card is supported in Junos OS 19.3R2 and later 19.3 releases and in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

- **MX2K-MPC11E (MX2010 and MX2020 routers)**—Junos OS Release 19.3R2 introduces the MX2K-MPC11E Modular Port Concentrator (MPC). The MX2K-MPC11E is a fixed-configuration MPC which delivers bandwidth up to 4-Tbps per MPC slot for MX2020 and MX2010 routers. Optimized for speeds of 100-Gigabit Ethernet and higher, the MPC11E provides 400-Gigabit Ethernet-ready ports with multi-rate QSFP-DD interfaces for seamless 100-Gigabit Ethernet to 400-Gigabit Ethernet upgrades. The MX2K-MPC11E supports native MACsec and IPsec encryption that can originate and terminate thousands of IPsec sessions with no impact on performance.

  The MX2K-MPC11E port panel has 40 rate-selectable (multi-rate) QSFP ports. You can use any port as a 100-Gigabit Ethernet interface, 40-Gigabit Ethernet interface, or 10-Gigabit Ethernet interface (using breakout cables). The 40 QSFP-DD optical transceiver ports are grouped into eight sets of five ports. A port group compose a logical PIC in the software. You can configure different data rates for each port group as long as the total throughput for the port group does not exceed 500 Gbps.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

- **MX-SPC3 Services Card (MX240, MX480, and MX960)**—Starting with Junos OS Release 19.3R2, the MX-SPC3 Services Card is available on MX240, MX480, and MX960 routers. The MX-SPC3 card provides additional processing power to run Next Gen Services. The MX-SPC3 contains two Services Processing Units (SPUs) with 128 GB of memory per SPU. Line cards such as DPCs, MICs, and MPCs intelligently distribute all traffic traversing the router to the SPUs to have services processing applied to it.

  Next Gen Services provide the best of both routing and security features on MX Series routers MX240, MX480, and MX960. All Next Gen Services are provided by the MX-SPC3 services card Next Gen Services provide capabilities for manipulating traffic before it's delivered to its destination. Next Gen Services features run on the MX Series, and are based on a different software architecture than legacy MX Series services. You can run Next Gen Services on MX240, MX480 and MX960 routers. Some Next Gen Services features use different Junos CLI statements than the equivalent legacy service.

  > **NOTE:** The only services card that supports Next Gen Services is the MX-SPC3. Next Gen Services use their own software architecture, which is not compatible with legacy services.

Table 1 on page 86 summarizes the Next Gen Services supported in this release.

Table 1: Next Gen Services Summary

Next Gen Services Supported by MX-SPC3 Services Card	
Carrier Grade NAT	6rd Softwires
	Deterministic NAT
	Dynamic Address-Only Source NAT
	Global System Logging
	IPv4 Connectivity Across IPv6-Only Network Using 464XLAT
	Network Address Port Translation
	Port Forwarding
	Static Source NAT
	Stateful NAT64
	Static Destination NAT
	Stateless Source Network Prefix Translation for IPv6
	Twice NAPT
	Twice Static NAT
	Class of Service
Stateful Firewall Services	
Intrusion Detection Services	
Traffic Load Balancing	
DNS Request Filtering	
Aggregated Multiservices Interfaces	
Inter-chassis High Availability	NAT, Stateful Firewall, and IDS Flows

Table 1: Next Gen Services Summary *(continued)*

Next Gen Services Supported by MX-SPC3 Services Card
See Protocols and Applications Supported by MX-SPC3 Services Card for information about the protocols and applications that this SPC3 supports.

The MX-SPC3 services card is compatible end-to-end with the MX Series Switch Fabrics, Routing Engines and MS-MPC line cards. See Table 2 on page 87:

Table 2: MX-SPC3 Services Card Compatibility with MX Series Switch Fabrics, Routing Engines and MPC Line Cards

Switch Fabric	Route Engine	MPC Line Cards
SCBE	RE-S-1800X4-16G-UPG-BB RE-S-1800X4-32G-UB	MPC2E-3D MPC2-3D-NG MPC3E and MPC3E-3D-NG MPC4E-3D MPC-3D-16XGE
SCBE2	RE-S-1800X4-16G-UPG-BB RE-S-1800X4-32G-UB RE-S-X6-64G-UB	MPC2E-3D MPC2-3D-NG MPC3E and MPC3E-3D-NG MPC4E-3D MPC5E and MPC5EQ MPC7E, MPC7EQ, and MPC-3D-16XGE MPC-3D-16XGE

Refer to our TechLibrary for all MX router documentation. For Next Gen Services, refer to the following documentation: See

- Next Gen Services Interfaces Overview for Routing Devices
- Next Gen Services Interfaces User Guide for Routing Devices
- Broadband Subscriber Services Feature Guide
- Monitoring, Sampling, and Collection Services Interfaces Feature Guide
- MX240 Universal Routing Platform Hardware Guide

- MX480 Universal Routing Platform Hardware Guide
- MX960 Universal Routing Platform Hardware Guide
- MX Series 5G Universal Routing Platform Interface Module Reference

Class of Service

- **Hierarchical Class of Service (CoS) parity with MPC11 (MX2000, MPC11E) and MX2000-SF83**—Starting in Junos OS Release 19.3R2, hierarchical class of service (CoS) parity is supported on MPC11 (MX2000, MPC11E) and MX2000-SF83.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Class of Service.]

- **L2 CoS classifiers and rewrite feature on MX Series routers (MX240, MX480, MX960) with the MPC10 line card**—Starting in Junos OS Release 19.3R2, support is provided for L2 Class of Service (CoS) classifiers and rewrite on MX Series routers (MX240, MX480, MX960) with the MPC10 line card.

  [See Protocols and Applications Supported by the MPC10E.]

- **Forwarding COS parity (L2 classifiers, rewrite) support for MX Series routers with MPC11 line cards**—Starting in Junos OS Release 19.3R2, Junos OS supports forwarding CoS parity (L2 classifiers, rewrite) for MX Series routers with the MPC11 line card.

  [See Understanding Class of Service.]

- **Seamless MPLS CoS support for pseudo-wires from access node (AN) and Multi-Services-Edge (MSE) node (MX480, MX960, MX2010, MX2020)**—Starting with Junos OS Release 19.3R2 support is provided for pseudo-wires from access node (AN) and Multi-Services-Edge (MSE) node for MX Series (MX480, MX960, MX2010, MX2020) routers to include seamless MPLS class of service(CoS) (BA and MF classifier, rewrite, schedulers, drop-profiles, policers, HQoS support - interface-set, PS IFD level, S-VLAN level, logical unit/C-VLAN level, traffic-control profile).

  [See Understanding Class of Service.]

- **Class of Service (CoS) parity support for forwarding-class (FC) counters (MPC10 and MPC11)**—Starting in Junos OS Release 19.3R2, class of service (CoS) parity support is provided for forwarding-class (FC) counters on MX Series routers with MPC10 and MPC11 line cards. Feature was originally introduced in Junos OS Release 14.1.

  [See Understanding Class of Service.]

EVPN

- **Support for EVPN functionality (MX Series with MPC10 line card)**—Starting in Junos OS 19.3R2, you can configure MPC10 line cards on a MX Series router to support single-homed devices on an EVPN-MPLS network.

  [See EVPN Multihoming Overview.]

- **Support for EVPN functionality on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS 19.3R2, you can configure MX2K-MPC11E line cards on MX2010 and MX2020 routers to support single-homed devices on an EVPN-MPLS network.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See EVPN Multihoming Overview.]

Forwarding and Sampling

- **Support for load balancing on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, the following advanced Layer 2 features are supported on MX2010 and MX2020 routers with MX2K-MPC11E line cards and Enhanced Switch Fabric Boards (SFB3): enhanced hash key options, consistent flow hashing, symmetrical load balancing over 802.3ad LAGs, source IP only hashing, and destination IP only hashing.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Configuring Per-Flow Load Balancing Based on Hash Values.]

High Availability (HA) and Resiliency

- **Support for BFD on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, MX2010 and MX2020 routers with the MX2K-MPC11E line card support the following Bidirectional Forwarding Detection (BFD) features:
  - Centralized BFD
  - Distributed BFD
  - Inline BFD (single-hop only)
  - Single-hop BFD
  - Multihop BFD
  - Micro BFD

- BFD over integrated routing and bridging (IRB) interfaces
- BFD over pseudowire over logical tunnel and redundant logical tunnel interfaces
- Virtual circuit connectivity verification (VCCV) BFD for Layer 2 VPNs, Layer 2 circuits, and virtual private LAN service (VPLS)

Micro BFD at the Packet Forwarding Engine level behaves slightly differently on MX2K-MPC11E line cards. If micro BFD is enabled on an aggregated Ethernet (ae) interface, the micro BFD packets are not subjected to firewall filters for both tagged and untagged ae interfaces.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding BFD for Static Routes for Faster Network Failure Detection and Understanding Distributed BFD.]

Interfaces and Chassis

- **MX2000-SFB3 Switch Control Board (MX2010 and MX2020 routers)**—In Junos OS Release 19.3R2, we introduce a new MX2000-SFB3 Switch Fabric Board. The MX2000-SFB3 creates a highly scalable and resilient "all-active" centralized switch fabric that delivers up to 4 Tbps of full-duplex switching capacity per slot for MX2010 and MX2020 routers. The MX2000-SFB3 along with the MX2K-MPC11E line card provides industry-leading 100-Gigabit Ethernet, 200-Gigabit Ethernet, and 400-Gigabit Ethernet port density.

  The MX2000-SFB3 interoperates with these legacy MPCs: MPC6E, MPC8E, MPC9E.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See MX Series 5G Universal Routing Platform Interface Module Reference.]

- **Chassis and power management for MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, the MX2010 and MX2020 routers with the MX2K-MPC11E line card support chassis management features, including field- replaceable unit (FRU) management, power budgeting and management, and environmental monitoring.

  The MX2K-MPC11E line card supports the following configuration:

  - The ambient temperature is less than 46°C.
  - The ports on the MX2K-MPC11E line cards operate at various modes or speeds (10-Gbps, 40-Gbps, or 100 Gbps). The pic-mode specifies the speed of the active ports. If pic-mode is not specified, then the default mode is 100 Gbps.

- Supports dynamic power management.
- Supports both hyper mode (the default mode) and normal mode.
- Supports both normal mode (the deafult mode) and enhanced priority mode for interface schedulers.
- Supports interface queueing modes, namely WAN port queueing mode (the default mode), limited queueing mode, and enhanced queueing mode.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding How Configuring Ambient Temperature Helps Optimize Power Utilization and Understanding How Dynamic Power Management Enables Better Utilization of Power.]

- **Support for Rate Selectability on MX2K-MPC11E line cards (MX2010 and MX2020 routers)**—In Junos OS Release 19.3R2, we introduce a new fixed-configuration, rate-selectable line card, MX2K-MPC11E, with built-in ports that can operate at 100 Gbps speed. The ports on the MX2K-MPC11E are called rate-selectable or multi-rate ports as they support multiple port speeds. You can choose to configure all ports in a PIC to operate at the same speed or configure all the ports at different supported speeds. The default port speed is 100 Gbps for all ports. Only Port 0, of every PIC, can operate at all three speeds (10 Gbps, 40 Gbps, and 100 Gbps). For the 10 Gbps speed, you will need a breakout cable.

  On the MX2K-MPC11E, you can choose to configure:

  - Port 0 of every PIC in 4x10-Gigabit Ethernet mode using QSFPP-4x10-Gigabit Ethernet optics and 4x10-Gigabit Ethernet breakout cables.
  - Port 0 of every PIC in 40-Gigabit Ethernet mode using QSFPP optics.
  - All other ports of every PIC in 100-Gigabit Ethernet mode using QSFP28 optics.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Introduction to Rate Selectability.]

- **MPC Protocol and Application Support for MX2K-MPC11E line cards**—Starting with Junos OS Release 19.3R2, MX2020 and MX2010 routers with MX2K-MPC11E line cards support many MPC protocols and applications. For a complete list, see *Protocols and Applications Supported by the MX2K-MPC11E*.

  - Standard Generic Routing Encapsulation (GRE)
  - Bidirectional Forwarding Detection protocol (BFD)
  - Internet Control Message Protocol (ICMP) and ICMPv6
  - Border Gateway Protocol (BGP)

- BGP/MPLS virtual private networks (VPNs)
- Logical system and Virtual routing and forwarding (VRF) routing instances
- Load Balancing
- Class of Service (CoS)—per port, virtual LAN (VLAN), Point-to-Point Protocol over Ethernet (PPPoE) or Dynamic Host Configuration Protocol (DHCP), Egress hierarchical class-of-service (CoS) shaping
- Layer 2 Features
- Firewall filters and policers

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See MX Series 5G Universal Routing Platform Interface Module Reference.]

- **MPC11E supports Junos node slicing (MX2010, MX2020)**—Starting in Junos OS Release 19.3R2, the MPC11E supports Junos node slicing and abstracted fabric (af) interfaces. Using Junos node slicing, you can create multiple partitions in a single physical MX Series router. Each partition, referred to as a guest network function (GNF), behaves as an independent router. An Abstracted Fabric interface is a pseudointerface that exhibits a first-class Ethernet interface behavior. The abstracted fabric interface facilitates routing control and management traffic between GNFs through the switch fabric. In a Junos node slicing deployment, the MPC11E interoperates with all MPCs that support the abstracted fabric interfaces.

  > **NOTE:**
  > - The MPC11E interoperates only with the Switch Fabric Board SFB3.
  > - The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [SeeUnderstanding Junos Node Slicing.]

Junos OS XML, API, Scripting

- **Support for 64-bit architecture added for use of management interface in a nondefault routing instance in op scripts and JET applications (MX Series)**—Junos OS Release 19.3R2 supports 64-bit architecture for Junos operating scripts and on-box JET applications being able to use the function set_routing_instance() to program the protocol software (TCP/UDP) to use a nondefault routing instance instead of the default management routing interface.

  [See set_routing_instance() Function (Python).]

Junos Telemetry Interface

- **gNMI-based streaming telemetry support for Packet Forwarding Engine sensors on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, gRPC Network Management Interface (gNMI) service support is available to export Packet Forwarding Engine statistics for telemetry monitoring and management using Junos telemetry interface (JTI). Using gNMI and JTI, data is exported from devices to outside collectors at configurable intervals. This feature includes support (SensorD daemon) to export telemetry data for the OpenConfig model called AFT platform.

  Use the following resource paths to export sensor data for interface information and traffic, logical interface traffic, firewall filter counters, and policer counters:

  - **/junos/system/linecard/interface/**
  - **/junos/system/linecard/interface/traffic/**
  - **/junos/system/linecard/interface/queue/**
  - **/junos/system/linecard/interface/logical/usage/**
  - **/junos/system/linecard/firewall/**
  - **/junos/system/linecard/services/inline-jflow/**

  To provision the sensor to export data through gNMI services, use the **Subscribe** RPC. The **Subscribe** RPC and subscription parameters are defined in the gnmi.proto file. Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

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

- **Packet Forwarding Engine support for JTI on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, Junos telemetry interface (JTI) supports streaming of Packet Forwarding Engine statistics for MX2010 and MX2020 routers using Remote Procedure Calls (gRPC). gRPC is a protocol for configuration and retrieval of state information.

To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

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

- **CPU and NPU sensors support using JTI on MX2K-MPC11E line cards (MX2010 and MX2020)**—Junos OS Release 19.3R2 supports Junos telemetry interface (JTI) CPU and network processing unit (NPU) sensors on MX Series routers with MX2K-MPC11E line cards. JTI enables the export of statistics from these sensors to outside collectors at configurable intervals using remote procedure call (gRPC) services.

  Unlike the previous Junos kernel implementation for these sensors in earlier Junos OS releases, this feature uses the OpenConfig AFT model. Because of this, there is a difference in the resource path and key-value (kv) pair output compared to the Junos kernel output.

  Use the following resource paths to export statistics:

  /junos/system/linecard/cpu/memory/

  /junos/system/linecard/npu/memory/

  /junos/system/linecard/npu/utilization/

  To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support JTI.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

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

Layer 2 Features

- **Layer 2 services on PWHT (MX2K-MPC11E line card)**—Starting in Junos OS Release 19.3R2, the following Layer 2 services are supported with pseudowire headend termination (PWHT):
  - Service provider style bridging
  - MAC learning and MAC table aging

- BUM traffic handling including split horizon
- MAC move
- Statistics
- Mesh groups
- Static MAC
- MAC learning and forwarding on aggregate Ethernet interfaces
- Bridging on untagged interfaces
- Simple dual tag bridging (with VLAN-translation and VLAN map operations)

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See Layer 2 VPNs and VPLS User Guide for Routing Devices.]

- **Supported features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, the following advanced Layer 2 features are supported on MX2K-MPC11E line cards:
  - **Forwarding CoS (Q-depth monitoring)**—You can configure a Junos telemetry interface sensor that exports queue depth statistics for ingress and egress queue traffic. Telemetry data is exported directly from the line card. You can also apply one or more regular expressions to filter data. Only UDP streaming of data is supported. gRPC streaming of queue depth statistics is not currently supported. [See sensor (Junos Telemetry Interface).]
  - **Layer 2 firewall forwarding support**. [See Layer 2 Port Mirroring Firewall Filters.]
  - **Layer 2 forwarding**—IRB, VLAN handling, and Q-in-Q tunneling. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.] Firewall filters for Layer 2 and MAC filters. [See Layer 2 Forwarding Tables.]
  - **Multicast features**—P2MP (RSVP-TE P2MP and multipoint LDP inband) and P2MP interface support for PIM, Rosen multicast VPNs, and multicast-only fast reroute (MoFRR). [See Multicast Overview.]

    > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

- **Support for basic Layer 2 features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, MPC11E line card on MX2020 and MX2010 routers supports the following basic Layer 2 features:

- Layer 2 bridging with trunk and access modes
- MAC learning and aging
- Handling BUM (broadcast, unknown unicast and multicast) traffic, including split horizon
- MAC move
- Layer 2 forwarding and flooding statics
- Mesh groups
- Static MAC addresses
- MAC learning and forwarding on AE interfaces
- Bridging on untagged interfaces
- Basic Q-n-Q tunneling (without VLAN-translation and VLAN map operations)

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding Layer 2 Bridge Domains, Understanding Layer 2 Learning and Forwarding.]

- **Support for an increase in MAC table size on the MPC10E-15C-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.3R2, MX Series routers with MPC10E-15C-MRATE line cards support MAC table size of upto 1 million entries per PFE for Layer 2 services.

  You can configure the MAC limit size at global level at the [**edit protocols l2-learning global-mac-limit**] hierachy level.

  You can also configure the MAC table size using bridge domains at the [**edit bridge-domains *bridge-domain-name* bridge-options mac-table-size**] hierarchy level.

  [See Understanding Layer 2 Bridge Domains, Understanding Layer 2 Learning and Forwarding.]

Layer 3 Features

- **Support for new MPC line card MPC11E (MX Series)**—Starting in Junos OS Release 19.3R2, a new MPC line card MPC11E is introduced.

  The following Layer 3 features are supported on MPC11E in 19.3R2:

  - BGP
  - IS-IS
  - L3VPN
  - OAM - LSP/VPN ping, Trace Route, Auto Bandwidth, and MPLS-FRR Link node protection
  - OSPF

- RIP
- Tunnel (Generic Routing Encapsulation (GRE), Logical Tunnel (LT), and Virtual Tunnel (VT))

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

- **Support for Seamless MPLS Layer 3 features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, the following MPLS Layer 3 features are supported on MX2020 and MX2010 routers with MPC11E line cards:
  - Redundant logical tunnel interfaces.
  - Pseudowire subscriber interfaces using either logical tunnel or redundant logical tunnel interfaces as anchor point.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See Redundant Logical Tunnels Overview and MPLS Pseudowire Subscriber Logical Interfaces.]

MPLS

- **Support for MPLS features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, some of the MPLS features are supported on the new MX2K-MPC11E line card.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

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

Multicast

- **Support for multicast forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, multicast forwarding is fully supported on MX2010 and MX2020 routers with MX2K-MPC11E line cards and Enhanced Switch Fabric Boards (SFB3).

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Multicast Overview.]

Routing Policy and Firewall Filter

- **Support for CCC and Layer 3 firewall forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting with Junos OS Release 19.3R2, circuit cross-connect (CCC) traffic and Layer 3 firewall forwarding features are supported on MX2K-MPC11E line cards.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See CCC Overview.]

- **Support for firewall forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, firewall forwarding is fully supported on MX2010 and MX2020 routers with MX2K-MPC11E line cards and Enhanced Switch Fabric Boards (SFB3).

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Filter-Based Forwarding Overview.]

Routing Protocols

- **VRRP support for MPC10E-15C-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.3R2, MPC10E-15C-MRATE line cards support VRRP.

  [See Understanding VRRP.]

- **Support for disable-4byte-as and minimum-hold-time configurations (MX Series)**—Starting in Junos OS Release 19.3R2, you can use the **minimum-hold-time** and **disable-4byte-as** configurations. By configuring **minimum-hold-time**, you can now prevent BGP session establishment toward BGP peers trying to negotiate a BGP session hold-time lower than the configured **minimum-hold-time**, which helps reducing the load on a router by avoiding sending constant keepalive messages at a high frequency. You can use **disable-4byte-as** configuration to enable a BGP peer that uses a 4-Byte to interact with another BGP peer old speaker that uses 2-Byte.

  NOTE:
  - We recommend using Bidirectional Forwarding Detection (BFD) rather than lowering BGP hold timers and also recommend configuring a meaningful **minimum-hold-time** value (for example, 20 seconds or higher) for all BGP peers (for example, at the BGP group level). If a BGP remote node does not support BFD, and therefore a reduced BGP hold-time is easier for quicker discovery of a BGP neighbor failure, you can configure a smaller minimum-hold-time value. Use it with caution and only for a limited number of BGP peers.
  - We recommend that you configure the **disable-4byte-as** configuration only if the BGP peer does not support or ignores the capability advertisement of 4byte-as, and brings up the session as a 2byte AS.

  Configure **minimum-hold-time** at any of the following hierarchy levels:

  set **protocols bgp minimum-hold-time** *time-in-seconds*

  set **protocols bgp group** *group-name* **minimum-hold-time** *time-in-seconds*

  set **protocols bgp group** *group-name* **neighbor** *ip-address* **minimum-hold-time** *time-in-seconds*

  Configure **disable-4byte-as** at any of the following hierarchy levels:

  set **protocols bgp disable-4byte-as**

  set **protocols bgp group** *group-name* **disable-4byte-as**

  set **protocols bgp group** *group-name* **neighbor** *ip-address* **disable-4byte-as**

Services and Applications

- **Support for port mirroring (MX2K-MPC11E line card on MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, you can configure port mirroring on the MX2K-MPC11E line card to mirror a copy of a packet to a configured destination, in addition to the normal processing and forwarding of the packet. The MX2K-MPC11E supports IPv4 (inet) and IPv6 (inet6) address families only.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Configuring Port Mirroring.]

- **Support for tunnel interfaces (MX2K-MPC11E line card on MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, Junos OS supports three tunnel interfaces: generic routing encapsulation (GRE) tunnel, logical tunnel (LT), and virtual tunnel (VT) on the MX2K-MPC11E line card.

  - The GRE tunnel interface supports the **tunnel** statement with these options: **destination**, **key**, **source**, **traffic-class** and **ttl**. The **copy-tos-to-outer-ip-header** statement is also supported.
  - The LT interface supports **family inet**, **family inet6**, and **family iso** options. The **encapsulation** statement supports the Ethernet and VLAN physical interface options only.
  - The VT interface supports the **family inet** option only.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Tunnel Services Overview.]

- **Support for multicast forwarding on MX2K-MPC11E line cards (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, multicast forwarding is fully supported on MX2010 and MX2020 routers with MPC11E line cards and Enhanced Switch Fabric Boards (SFB3).

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Multicast Overview.]

- **Fabric support on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, the MX2K-MPC11E line card is introduced. It is composed of 8 Packet Forwarding Engines per FPC. Each Packet Forwarding Engine on the MX2K-MPC11E line card has 3 fabric planes per SFB, which is a total of 24 fabric planes. All Packet Forwarding Engines have fabric connectivity with the SFB3. The

fabric links are monitored for cyclic redundancy check (CRC) errors. Each Packet Forwarding Engine supports 500G fabric throughput when all 24 fabric planes are operational.

> **NOTE:**
> - Fabric redundancy is not supported on MX2K-MPC11E line card. The MX2K-MPC11E line card interoperates only with SFB3.
> - The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

- **Configure next-hop-based dynamic tunnels on MX2K-MPC11E line card (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, on MX2010 and MX2020 routers with an MX2K-MPC11E line card, you can configure next-hop-based dynamic tunnels for the following configurations:
  - MPLS-over-UDP—You can configure a dynamic MPLS-over-UDP tunnel that includes a tunnel composite next hop.

    In a dynamic tunnel configuration, where the Routing Engine forwards a lot of routes to the Packet Forwarding Engine, the FIB convergence may take more time resulting in traffic loss. Also, when you restart an FPC restart in a dynamic tunnel configuration, traffic flow may not resume.
  - MPLS-over-GRE—You can configure MPLS LSPs to use generic routing encapsulation (GRE) tunnels to cross routing areas, autonomous systems, and ISPs.

    > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See dynamic-tunnels.]
- **Inline active flow monitoring on MX2K-MPC11E line card (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, Junos OS supports inline active flow monitoring. Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Version 9 template is supported for IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPv6. IPFIX template is supported for IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPv6 flows. Both IPFIX and version 9 templates use UDP as the transport protocol.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Inline Active Flow Monitoring.]

- **Support for Two-Way Active Measurement Protocol (TWAMP) and Hardware timestamping of RPM probe messages on MX2010 and MX2020 routers**—Starting in Junos OS Release 19.3R2, MX2K-MPC11E line card supports:
  - TWAMP— Enable measurement of IP performance between two devices in a network.
  - Hardware timestamping of RPM probe messages— Enable timestamps on RPM probes messages in the Packet Forwarding Engine host processor.

    NOTE: The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Two-Way Active Measurement Protocol on RoutersUnderstanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers.]

- **L2TPv2 silent failover on peer interface for L2TPv2 subscriber services on MX2K-MPC11E line card (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, you can configure L2TPv2 silent failover and peer interface support for L2TPv2 subscriber services on MX2K-MPC11E line card.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Peer Resynchronization After an L2TP Failover.]

- **Port mirroring support on MX2K-MPC11E (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, you can configure port mirroring on the MX2K-MPC11E line card. You can configure port mirroring for the CCC, bridge, and family any only.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Port Mirroring.]

- **FlowTapLite suppport on MX2K-MPC11E line card (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.3R2, you can configure FlowTapLite on an MX2K-MPC11E line card. FlowTapLite enables interception of IPv6 packets on MX Series, M120, and M320 routers.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

[See Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs.]

Subscriber Management and Services

- **Distributed denial of service protection on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, the MPC11E-MRATE line cards support DDoS protection.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MPC11E.]

- **Subscriber services uplink support on MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R2, you can use the MX2K-MPC11E line cards for uplink connections to the core network. This support requires you to enable enhanced subscriber management.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MPC11E.]

- **CoA messages support Session-Timeout attribute (MX Series)**—Starting in Junos OS Release 19.3R2, you can apply a session timeout for subscriber sessions with a RADIUS CoA message that includes the Session-Timeout attribute (27). This capability is useful, for example, when subscribers purchase Internet access for a specific period of time and must log out when the session expires. In earlier releases, the router does not recognize the attribute if it is included in a CoA message.

  [See Understanding Session Options for Subscriber Access.]

What's New in Release 19.3R1

Authentication and Access Control

- **Support for DYCE randomizer (MX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports the Dyce randomizer on MX960 routers. Dyce is a cryptographically secure pseudo-random number generator (CSPRNG) in the Junos OS kernel. To enable the Dyce randomizer, use the **dyce** statement at the **[edit system rng]** hierarchy level and reboot the Routing Engine when prompted.

  set system rng dyce

  > **NOTE:** This configuration is available only at the administrator privilege level.

[See rng].

- **Support for remote authorization on tacplus for locally authenticated users (MX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports remote authorization on TACPLUS server for the locally authenticated users by using their locally configured parameters. You can use this feature through **password-options**, a new configuration statement, to configure options for local authentication; and its option **tacplus-authorization**, to choose the remote authorization for locally authenticated users.

  > **NOTE:** You must configure the password under **authentication-order** when **password-options** is configured.
  >
  > The feature does not work in a local fallback scenario because password is not configured under **authentication-order** for a local fallback scenario.

  To configure remote authorization, include the **tacplus-authorization** option under the **password-options** configuration statement at the **[edit system]** hierarchy level.

  [See password-options].

Class of Service (CoS)

- **Support for default interface set traffic control profiles (MX Series)**—Starting with Junos 19.3R1, you can set a single traffic control profile (TCP) to be the system-wide default input TCP for a dynamic interface-set and either the same or a different TCP to be the system-wide default output TCP for a dynamic interface set. When a dynamic interface-set is not explicitly assigned an output TCP, a static TCP configured as the default output TCP for a dynamic interface-set is automatically assigned to this interface-set. Likewise, when a dynamic interface-set is not explicitly assigned an input TCP, a static TCP configured as the default input TCP for a dynamic interface-set is automatically assigned to this interface-set. To set the default input TCP and output TCP, add **interface-set-input *profile-name*** and **interface-set-output *profile-name*** options at the **[edit class-of-service system-defaults traffic-control-profiles]** hierarchy level.

  [See system-defaults.]

- **Support for input traffic control profile assignment to dynamic logical interface sets (MX Series)**—To support assigning both static and dynamic logical interfaces to the same interface set (either static or dynamic), starting with Junos 19.3R1, you can apply an input traffic control profile (TCP) to a dynamic logical interface set in 4-level hierarchical scheduling or to two dynamic logical interface sets in 5-level hierarchical scheduling. Thus, Junos CoS enables you to dynamically assign a static input TCP with shaping-rate to a dynamic interface-set to enforce a customer's SLA. If no such SLA enforcement is needed, you can configure a static TCP that is designated as the default input TCP assigned to any dynamic interface-set that does not already have an explicitly assigned input TCP.

  [See Understanding Hierarchical CoS for Subscriber Interfaces.]

- **Class of Service (CoS) parity support for forwarding-class (FC) counters on MX-series routers (MPC10/MPC11)**—Starting in Junos OS 19.3R1, class of service (CoS) parity support is provided for forwarding-class (FC) counters on MX-series routers with MPC10 and MPC11 line cards. Feature was originally introduced in Junos OS 14.1.

  [For more information regarding CoS, see Understanding Class of Service.]

EVPN

- **Support for DHCP relay in an EVPN-MPLS network (EX9200 switches, MX Series, vMX)**—Starting in Junos OS Release 19.3R1, EX9200 switches, MX Series routers, and vMX virtual routers support DHCPv4 and DHCPv6 relay in an EVPN-MPLS network. We support this feature in a data center architecture that includes a layer of spine devices that perform EVPN Layer 2 and Layer 3 functions. These devices are connected to a layer of leaf devices that perform EVPN Layer 2 functions. In this architecture, DHCP clients are connected to leaf devices, and DHCP servers are connected to spine devices. The DHCP relay functions are centralized at the spine layer. As a result, this architecture in known as the centrally routed bridging architecture.

  [See DHCP Relay Agent in EVPN-MPLS Network.]

- **IGMP snooping support for EVPN-VXLAN (EX9200 switches, MX Series, and vMX)**—Starting in Junos OS Release 19.3R1, you can configure IGMP snooping on EX9200 switches, MX Series routers, and vMX virtual routers in an EVPN-VXLAN network. Enabling IGMP snooping helps to constrain multicast traffic to interested receivers in a broadcast domain.

  The listed devices support these IGMP snooping use cases in a centrally-routed bridging overlay (an EVPN-VXLAN network with a two-layer IP fabric):

  - Forwarding multicast traffic within a VLAN (intra-VLAN)
  - Routing multicast traffic between VLANs (inter-VLAN) using one of the following methods:
    - IRB interfaces configured with Physical Interface Module (PIM) on an elected designated router (DR) device
    - A PIM gateway with Layer 2 or Layer 3 connectivity
    - An external multicast router

  The listed devices support these IGMP versions and membership report modes:

  - IGMPv2 with Any-Source Multicast (ASM) (*,G) mode only.
  - IGMPv3 in either of the following modes:
    - ASM (*,G)—the default behavior.
    - Single-Source Multicast (SSM) (S,G)—you must explicitly configure by including the **evpn-ssm-reports-only** configuration statement at the **[edit protocols igmp-snooping]** hierarchy level.

  [See Overview of IGMP Snooping in an EVPN-VXLAN Environment .]

- **Multiple routing instance support for ping overlay and traceroute overlay on VXLAN (MX Series routers and vMX virtual routers)**—Starting in Junos OS Release 19.3R1, Junos OS supports using **ping overlay** and **traceroute overlay** commands on a static VXLAN tunnel with multiple routing instances. The OAM packets created for the **ping overlay** and **traceroute overlay** commands follow the same underlay network path as the data packets. This allows you to verify the connectivity and to detect fault in the underlay network for an overlay segment between two VTEPs.

  [See Understanding Overlay ping and traceroute Packet Support.]

General Routing

- **Seamless BFD inline mode support for static segment routing LSPs (MX Series)**—Starting in Junos OS Release 19.3R1, MX Series routers support inline mode for seamless BFD on static segment routing LSPs. Inline mode helps increase the number of supported sessions.

  [See Routing Engine-Based S-BFD for Segment Routing.]

- **BFD support (MX240, MX480, MX960)**—Starting in Junos OS Release 19.3R1, MX240, MX480, and MX960 routers with the MPC10E-15C-MRATE line card support the following BFD features:

  - Inline BFD with the **ps-over-lt** interface
  - Inline BFD with the **ps-over-rlt** interface
  - Inline BFD over GR tunnels
  - Inline BFD sessions over IRB interface
  - BGP Prefix-Independent Convergence (PIC) Edge

  [See Understanding BFD for Static Routes for Faster Network Failure Detection.]

- **Distributed BFD support (MX240, MX480, MX960)**—Starting in Junos OS Release 19.3R1, MX240, MX480, and MX960 routers with the MPC10E-15C-MRATE line card support the following distributed BFD features:

  - Micro-BFD
  - VCCV BFD
  - BFD sessions that are supported in inline mode can also run in distributed mode.

  Micro-BFD at the Packet Forwarding Engine level behaves slightly differently on MPC10E-15C-MRATE line cards. If micro-BFD is enabled an aggregated Ethernet (ae-) interface, the micro BFD packets are not subjected to firewall filters, for both tagged and untagged ae interfaces.

  [See Understanding BFD for Static Routes for Faster Network Failure Detection.]

Interfaces and Chassis

- **10-Gigabit Ethernet WAN PHY mode implementation (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Release 19.3R1, Junos OS follows the RFC 3635 specification while implementing WAN PHY (physical layer) mode of framing speed reporting. In earlier releases, the 10-Gigabit Ethernet WAN PHY mode implementation does not consider the payload data rate over the WAN interface sublayer. As specified in RFC 3635, for Ethernet-like interfaces on WAN interfaces operating at speeds greater than 1000 Mbps, **IfHighSpeed** now represents the current operational speed of the interface—9.294 Gbps— for WAN implementations.

  [See 10-Gigabit Ethernet Framing Overview.]

- **New bidirectional SFP transceivers for MX240, MX80, MX104, MX480, and MX960**—Starting in Release 19.3R1, Junos OS supports the new bidirectional SFP transceivers (part numbers: **740-088382: 1GE-BX80-T15R14-D** and **740-088384: 1GE-BX80-T14R15-U**), which provide a speed of 1 Gbps and a reach of up to 80 km. You can use the existing show commands such as **show chassis pic** and **show chassis hardware** to view the inventory details of the transceivers.

  [See show chassis pic and show chassis hardware.]

- **Support for 25-Gbps port speed on MPC10E (MX240, MX480, and MX960)**—In Junos OS Release 19.3R1, you can configure port speed of 25 Gbps on MPC10E-10C-MRATE and MPC10E-15C-MRATE on MX240, MX480, and MX960 routers. Use QSFPP-4x25GE breakout cables to configure 25-Gbps port speed on:

  - MPC10E-10C-MRATE: Any of the ten ports.
  - MPC10E-15C-MRATE: Any of the fifteen ports.

  > **NOTE:** The MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards support the following 4x25 Gigabit Ethernet transceivers:
  >
  > - QSFP-100GBASE-SR4 (with optical breakout)
  > - QSFP-100GBASE-PSM4 (with optical breakout)

  [See MPC10E-15C-MRATE Rate Selectability Overview, MPC10E-10C-MRATE Rate Selectability Overview].

- **High-voltage second-generation universal PSM for MX960**—Starting in Junos OS Release 19.3R1, MX960 routers support the new high-voltage second-generation universal power supply module (PSM). This single feed PSM provides a maximum output power of 5100 W, and supports either AC or DC input. The PSM supports a 1+1 redundancy.

  [See MX960 Power System Overview.]

Junos OS XML API and Scripting

- **Program management interface in a nondefault routing instance in op scripts and JET applications (MX Series)**—Junos OS Release 19.3R1 supports, on 32-bit architecture, Junos operating scripts and on-box JET applications can use the function set_routing_instance() to program the protocol software (TCP/UDP) to use a nondefault routing instance instead of the default management routing interface.

  [See set_routing_instance() Function (Python).]

- **IPv6 support in Python automation scripts (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.3R1, devices running Junos OS with upgraded FreeBSD support using IPv6 in:

  - Python automation scripts, including commit, event, op, and SNMP scripts
  - Juniper Extension Toolkit (JET) scripts
  - YANG action scripts

  IPv6 support enables Python scripts to establish connections and perform operations using IPv6 addresses.

  [See IPv6 Support in Python Automation Scripts.]

Junos Telemetry Interface

- **Juniper AAA Model streaming telemetry support for subscriber services for JTI (MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX100016)**—Junos OS Release 19.3R1 supports streaming statistics for subscribers for the diameter application protocols Network Access Server Application (NASREQ), policy and charging rules function (PCRF), and Online Charging System (OCS). There are also new diameter peer sensors that provide response time measurements for messages exchanged between an MX router and the peer for each of the diameter applications. Statistics are exported using Junos telemetry interface (JTI) and the Juniper AAA Model, which covers telemetry export using remote procedure calls (gRPC), gRPC Management Interface (gNMI), or Juniper proprietary RPC or UDP.

  To stream diameter application statistics, include the resource paths in a subscription or using the **sensor** configuration statement:

  - For NASREQ statistics, **/junos/system/subscriber-management/aaa/diameter/clients/nasreq**
  - For PCRF statistics, **/junos/system/subscriber-management/aaa/diameter/clients/gx**
  - For OCS statistics, **/junos/system/subscriber-management/aaa/diameter/clients/gy**

  To stream response time measurements for the diameter applications, include the resource paths in a subscription or using the **sensor** configuration statement:

  - For NASREQ measurements, **/junos/system/subscriber-management/aaa/diameter/peers/peer[peer_address='*peer-address*']/nasreq/response-time**
  - For PCRF measurements, **/junos/system/subscriber-management/aaa/diameter/peers/peer[peer_address='*peer-address*']/gx/response-time**

- For OCS measurements, **/junos/system/subscriber-management/aaa/diameter/peers/peer[peer_address='*peer-address*']/gy/response-time**

To enable these statistics for an MX Series router for native (UDP) export, include the **sensors** statement at the [**edit services analytics**] hierarchy level.

To provision the sensor to export data through gNMI, use the Subscribe RPC defined in the gnmi.proto to specify request parameters.

To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module.

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

- **JTI support extended for Junos kernel GRES and RTSOCK (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos Telemetry Interface (JTI) extends support for streaming Junos kernel Graceful Routing Engine Switchover (GRES) and Routing Socket (RTSOCK) information using remote procedure call (gRPC) services. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel GRES and RTSOCK information:

  - /junos/chassis/gres/
  - /junos/kernel/rtsock/

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

- **Physical interface operational status sensor (int-exp) support on JTI (MX960, MX2010, and MX2020)**—Starting with Junos OS Release 19.3R1, Junos telemetry interface (JTI) uses the sensor int-exp (interface express) to export interface operational **UP** and **DOWN** status at a user-configurable rate. This sensor leverages statistics out of the physical interface sensor, providing faster and more frequent operational status statistics. Only the physical interfaces' operational status from the Flexible PIC Concentrator (FPC) is collected and reported. Statistics from the Routing Engine interface are not reported.

  You can apply the intf-exp sensor using the following path for either native (UDP) export using the **sensor** configuration statement or through OpenConfig using remote procedure calls (gRPC) services:

  - **/junos/system/linecard/intf-exp/**

  To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module.

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

- **gNMI support for Routing Engine statistics for JTI (MX960, MX2010, MX2020, PTX5000, PTX1000, and PTX10000)**—Junos OS Release 19.3R1 supports the Junos telemetry interface (JTI) export of Routing Engine sensors using gRPC Management Interface (gNMI). gNMI is a protocol for configuration and retrieval of state information. Both streaming and ON_CHANGE export is supported using gNMI.

  JTI exports the following statistics using gNMI:

  - Network discovery, ARP table state (resource path **/arp-information/**)
  - Network discovery, NDP table state (resource paths **/nd6-information/** and **/ipv6-ra/**)

  To provision the sensor to export data through gNMI, use the Subscribe RPC defined in the gnmi.proto to specify request parameters. Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **CPU/NPU sensors support using JTI (MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards)**—Junos OS Release 19.3R1 supports Junos telemetry interface (JTI) CPU and network processing unit (NPU) sensors on MX Series routers with on MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards. JTI enables the export of statistics from these sensors to outside collectors at configurable intervals using remote procedure call (gRPC) services.

  Unlike the previous Junos kernel implementation for these sensors in earlier Junos OS releases, this feature uses the OpenConfig AFT model. Because of this, there is a difference in the resource path and key-value (kv) pair output compared to the Junos kernel output.

  Use the following resource path to export statistics:

  /junos/system/linecard/cpu/memory/

  /junos/system/linecard/npu/memory/

  /junos/system/linecard/npu/utilization/

  To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support JTI.

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

- **gNMI-based streaming telemetry support for Packet Forwarding Engine sensors (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.3R1, gRPC Network Management Interface (gNMI) service support is available to export Packet Forwarding Engine statistics for telemetry monitoring and management using Junos telemetry interface (JTI). Using gNMI and JTI, data is exported from devices to outside collectors at configurable intervals. This feature includes support (SensorD daemon) to export telemetry data for integration with AFTTelemetry and LibTelemetry libraries in the OpenConfig model called AFT platform.

Use the following resource paths to export sensor data for interface information and traffic, logical interface traffic, firewall filter counters, and policer counters:

- **/junos/system/linecard/interface/**
- **/junos/system/linecard/interface/traffic/**
- **/junos/system/linecard/interface/queue/**
- **/junos/system/linecard/interface/logical/usage/**
- **/junos/system/linecard/firewall/**
- **/junos/system/linecard/services/inline-jflow/**

To provision the sensor to export data through gNMI services, use the **Subscribe** RPC. The **Subscribe** RPC and subscription parameters are defined in the gnmi.proto file. 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) and Understanding OpenConfig and gRPC on Junos Telemetry Interface.]

- **JTI support extended for Junos kernel LAG, NSR, and TCP (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, and PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel Link Aggregation Group (LAG), nonstop routing (NSR) Junos socket replication (JSR), and Transport Control Protocol (TCP) information using remote procedure call (gRPC) services. Device monitoring and network analytics applications can use Junos kernel sensors to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel LAG, NSR, and TCP information:

  - /junos/chassis/aggregated-devices/
  - /junos/routing-options/nonstop-routing/
  - /junos/kernel/tcpip/tcp/

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

- **JTI support extended for Junos kernel IPv4 and IPv6 (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, and PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel IPv4 and IPv6 information using remote procedure call (gRPC) services. Device monitoring and network analytics applications can use Junos kernel sensors provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel IPv4 and IPv6 information:

  - /junos/kernel/tcpip/arp/ — Address Resolution Protocol cache
  - /junos/kernel/tcpip/ndp/ — Neighbor Discovery Protocol cache

- /junos/kernel/tcpip/netisr/ — NETISR network queues
- /junos/kernel/tcpip/nhdix/ — Next-hop index space exhaustion
- /junos/kernel/tcpip/rtb/ — Route tables

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

- **JTI support extended for Junos kernel IP multicast, tunnels, TNP, and VPLS (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel IP multicast, tunnels, Trivial Network Protocol (TNP), and virtual private LAN service (VPLS) information using remote procedure call (gRPC) services. Device monitoring and network analytics applications can use Junos kernel sensors provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel IP multicast, tunnels, TNP, and VPLS information:

  - /junos/kernel/multicast/
  - /junos/kernel/tunnel/
  - /junos/kernel/tnp/
  - /junos/kernel/vpls/

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

Layer 2 Features

- **Supported features on MPC10E-15C-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.3R1, the MPC10E-15C-MRATE line card supports the following advanced Layer 2 features::
  - **Forwarding CoS (Q-depth monitoring)**—You can configure a Junos telemetry interface sensor that exports queue depth statistics for ingress and egress queue traffic. Telemetry data is exported directly from the line card. You can also apply one or more regular expressions to filter data. Only UDP streaming of data is supported. gRPC streaming of queue depth statistics is not currently supported. [See sensor (Junos Telemetry Interface).]
  - **Layer 2 firewall forwarding support**. [See Layer 2 Port Mirroring Firewall Filters.]
  - **Layer 2 forwarding**—IRB, VLAN handling, and Q-in-Q tunneling. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.] Firewall filters for Layer 2 and MAC filters. [See Layer 2 Forwarding Tables.]

- **Load balancing**—Enhanced hash key options, consistent flow hashing, symmetrical load balancing over 802.3ad LAGs, source IP only hashing, and destination IP only hashing. [See Configuring Per-Flow Load Balancing Based on Hash Values.]
- **Multicast features**—P2MP (RSVP-TE P2MP & mLDP inband) and P2MP interface support for PIM, Rosen multicast VPNs, and multicast-only fast reroute (MoFRR)(mLPD inband and PIM signaling). [See Multicast Overview.]

Management

- **OpenConfig AAA data model support (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 19.3R1 supports the configuration leafs specified in the OpenConfig AAA data model. Mapping the OpenConfig AAA configuration to the Junos AAA configuration using the following YANG files in the data model makes this support possible:
  - openconfig-aaa.yang
  - openconfig-aaa-types.yang
  - openconfig-aaa-tacacs.yang
  - openconfig-aaa-radius.yang

  The configuration model supporting the OpenConfig data model includes:
  - A translation script (**.py** / **.slax**) that maps each configuration leaf in the OpenConfig schema to one or more configuration leafs in the Junos schema.
  - A deviation file (**.yang**) that specifies how much the implementation deviates from the vendor-neutral model.

  [See Mapping OpenConfig AAA Commands to Junos Configuration.]

MPLS

- **Support for Entropy Label and FAT PW Label in Advanced Forwarding Toolkit (MX Series)**—Starting in Junos OS Release 19.3R1, the Junos OS supports entropy labels and Flow Aware Transport for Psuedowires (FAT) labels. Entropy label and FAT label when configured on the label-switching routers (LSRs) and label edge routers (LEs) perform load-balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAG) without the need for deep packet inspection of the payload.

  [See Configuring the Entropy Label for LSPs and FAT Flow Labels Overview].
- **Support for Seamless MPLS Layer 3 features (MX Series with MPC10E)**—Starting in Junos OS Release 19.3R1, the following MPLS Layer 3 features are supported on MX series routers with MPC10E line cards:
  - Redundant logical tunnel interfaces.

- Pseudowire subscriber interfaces using either logical tunnel or redundant logical tunnel interfaces as anchor point.

[See Redundant Logical Tunnels Overview, and MPLS Pseudowire Subscriber Logical Interfaces.]

Operation, Administration, and Maintenance (OAM)

- **Support for LFM (MPC10E)**—Starting in Junos OS Release 19.3R1, you can configure IEEE 802.3ah link fault management (LFM) for MPC10E-10C-MRATE and MPC10E-15C-MRATE on MX240, MX480, and MX960 routers. You can also configure the following supported LFM features:
  - Discovery and link monitoring
  - Distributed LFM
  - Remote fault detection and remote loopback

  NOTE: You cannot configure inline LFM for MPC10E on MX240, MX480, and MX960 routers.

  [See IEEE 802.3ah OAM Link-Fault Management Overview.]

Port Security

- **MACsec on logical interfaces (MX Series)**—Starting with Junos OS Release 19.3R1, you can configure Media Access Control Security (MACsec) at the logical interface level on the MIC-MACSEC-20GE in an MX Series router. This feature allows multiple MACSec Key Agreement (MKA) sessions on a single physical port.

  [See Media Access Control Security (MACsec) over WAN.]

- **PSK hitless rollover and fail-open mode for MACsec (MX Series)**—Starting with Junos OS Release 19.3R1, MX Series routers with the MPC10E-15C or MPC10E-10C line card support preshared key (PSK) hitless rollover and fail-open mode for Media Access Control Security (MACsec).

  [See Configuring Media Access Control Security (MACsec) on Routers.]

- **Support for configuring unicast MAC DA for MACsec (MX Series)**—Starting with Junos OS Release 19.3R1, you can configure the destination EAPoL address for MACsec as a unicast address. When MACsec Key Agreement PDUs (MKA PDUs) are exchanged over a provider network, they might be dropped or consumed if the default multicast address is used. You can configure the unicast MAC address to ensure that the MKA PDUs reach their destination.

  [See Media Access Control Security (MACsec) over WAN.]

Routing Protocols

- **Support for nondefault routing instance for outbound SSH (MX Series and SRX Series)**—Starting in Junos OS Release 19.3R1, you can specify the name of the routing instance on which the outbound SSH connectivity needs to be established using the **routing-instance** statement at the **[edit system services outbound-ssh]** hierarchy level. If you do not specify a routing instance, your device will establish the outbound SSH connection using the default routing table.

  [See outbound-ssh, Configuring Outbound SSH Service.]

- **Support for color mode in segment routing traffic engineering using BGP (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports color-only mode corresponding to color bits 01 and supports the steering fallback mechanism (in a limited manner) when color bits as set to 01 as described in IETF DRAFT-SPRING-SRTE. Use the **extended-nexthop-color** CLI configuration option to set color bits to 01 to enable color-only mode. Fall back to color-only SRTE policies is also supported and can be configured independently by configuring an import policy at the headend.

  [See Understanding Ingress Peer Traffic Engineering for BGP SPRING.]

- **Support for segment routing (SR) and segment routing traffic engineering (SRTE) statistics in Advanced Forwarding Toolkit (MX960, MX2010, and MX2020)**—Starting in Junos OS Release 19.3R1, the traffic statistics in a segment routing (SR) network can be recorded in an OpenConfig compliant format for Layer 3 interfaces. The statistics is recorded for the Source Packet Routing in Networking (SPRING) traffic only, excluding RSVP and LDP-signaled traffic, and the family MPLS statistics per interface is accounted for separately. The SR statistics also includes SPRING traffic statistics per link aggregation group (LAG) member, and per segment identifier (SID).

  Junos OS also supports SRTE telemetry statistics and BINDING-SID routes for uncolored SRTE policies. Uncolored SRTE LSP is characterized by the absence of color statement in its configuration. Junos OS now allows collection of traffic statistics for both ingress IP traffic and transit MPLS traffic that take uncolored SRTE paths. Also, you can install BINDING-SID labels even if the first hop of the segment list is a label. By default, traffic sensors and statistic collection are disabled for static SRTE routes. To enable provisioning of JVISION traffic sensors in Junos OS data plane to stream out traffic statistics on SR policies and their Binding-SID routes, use the existing statistics statement at the [**edit source-packet-routing telemetry**] hierarchy level.

  [See Understanding Source Packet Routing in Networking (SPRING).]

- **Support for OSPF TI-LFA back paths for Segment Routing (MX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports creation of OSPF topology-independent TI-LFA backup paths where the prefix SID is learned from a segment routing mapping server advertisement when the PLR and mapping server are both in the same OSPF area.

  [See Configuring Topology-Independent Loop-Free Alternate with Segment Routing for OSPF.]

Services Applications

- **Regulate and add frame and byte count for carrier-grade NAT syslog messages**—Starting in Junos OS Release 19.3R1, you can enable or disable the display of the carrier-grade NAT syslog message in certain

deployment scenarios. When you configure the **disable-session-open-syslog** statement at the **[edit services service-set *service-set-name* service-set-options]** hierarchy level:

- The JSERVICES_SESSION_OPEN logs are disabled and are no longer generated for the incoming sessions.
- The JSERVICES_SESSION_CLOSE logs, which are generated after the incoming sessions are closed, are printed with byte and packet count for the packets received and sent as part of the same session.

[See service-set (Services).]

- **Configure next-hop-based dynamic tunnels on MPC10E (MX240, MX480, and MX960 routers)**—Starting in Junos OS Release 19.3R1, on MX240, MX480, and MX960 routers with an MPC10E line card, you can configure next-hop-based dynamic tunnels for the following configurations:
  - MPLS-over-UDP—You can configure a dynamic MPLS-over-UDP tunnel that includes a tunnel composite next hop.
  - MPLS-over-GRE—You can configure MPLS LSPs to use generic routing encapsulation (GRE) tunnels to cross routing areas, autonomous systems, and ISPs.

  [See dynamic-tunnels.]

- **Inline active flow monitoring on MPC10E (MX240, MX480, and MX960 routers)**—Starting in Junos OS Release 19.3R1, you can configure inline active flow monitoring on the MPC10E line card to support:
  - MPLS, MPLS-IPv4, and MPLS-IPv6 traffic
  - IPv4 or IPv6 traffic on next-hop-based GRE tunnels and PS interfaces

  The flow records support both IPFIX and version 9 formats.

  [See Understanding Inline Active Flow Monitoring.]

- **L2TP silent failover on peer interface for L2TP subscriber services on MPC10E (MX Series)**—Starting in Junos OS Release 19.3R1, you can configure L2TP silent failover on peer interface for L2TP subscriber services on MPC10E line card.

  L2TP failover enables a failed L2TP endpoint to resynchronize with its nonfailed peer during recovery and restart of the L2TP protocol on the failed endpoint. L2TP failover is enabled by default.

  [See Peer Resynchronization After an L2TP Failover.]

- **FlowTapLite suppport on MPC10E (MX240, MX480, and MX960 routers)**—Starting in Junos OS Release 19.3R1, you can configure FlowTapLite on an MPC10E line card. FlowTapLite enables interception of IPv6 packets on MX Series, M120, and M320 routers.

  [See Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs.]

- **Support for TWAMP (ACX5448-M)**—Starting in Junos OS Release 19.3R1, you can configure Two-Way Active Measurement Protocol (TWAMP) on your ACX5448 routers. TWAMP enables you to measure the IP performance between two devices in a network. The ACX5448-M router supports only the reflector side of TWAMP.

[See Two-Way Active Measurement Protocol on ACX Series.]

- **Support for multiple features on tunnel interfaces on the MPC10E line card (MX Series)**—Starting in Release 19.3R1, Junos OS supports the multiple features on tunnel interfaces.

  On logical tunnels, you can now configure the CCC/MPLS family configuration.

  On GRE tunnel interfaces, you can:

  - Copy the ToS bits of the inner IP header to the outer IP packet header for traffic transiting the router.
  - Configure an IPv6 GRE tunnel interface.
  - Configure the flow-label field in the IPv6 header of the delivery protocol for an IPv6 GRE tunnel interface.
  - Enable or disable fragmentation of GRE-encapsulated packets for an IPv4 GRE tunnel interface.
  - Enable a GRE tunnel for keepalive messages for an IPv4 GRE tunnel interface.
  - Enable or disable path MTU discovery for an IPv4 GRE tunnel interface.

  [See GRE Keepalive Time Overview.]

  [SeeEnabling Fragmentation on GRE Tunnels.]

- **Support for ECMP routing with MAP-E (MX Series)**—Starting in Junos OS Release 19.3R1, you can configure ECMP routing with MAP-E. By configuring the **disable-auto-route** statement at the **[edit services softwire softwire-concentrator map-e <domain-name>]** hierarchy level, you can disable auto routes and configure static routes that point to an ECMP load balancer. As a result, the packets are distributed among different inline service interfaces.

  [See Understanding Mapping of Address and Port with Encapsulation (MAP-E).]

- **Enhance Juniper SkyATP URL filtering logging through sampling (MX240, MX480, and MX960 routers)**—Starting in Junos OS Release 19.3R1, the URL filtering daemon (url-filterd) supports inline sampling as a threat action. You can configure the following threat-level actions in the web filtering profile at the **[edit services web-filter- profile p1 security-intelligence-policy threat-level *threat-level* threat-action]** hierarchy level.

  - **drop-and-sample**
  - **drop-log-and-sample**
  - **log-and-sample**
  - **sample**

  The inline J-Flow samples the packets and sends it to a collector in IPFIX format. You can derive the threat level for the sampled packets received at the external collector by matching the received IP from the sampled packets with the corresponding IP entry in /var/db/url-filterd/urlf_si_cc_db.txt.

  [See web-filter.]

Software-Defined Networking (SDN)

- **Upgrade JDM to support WRL 9-based VM host for in-chassis Junos node slicing (MX2010, MX2020, MX2008, MX960, and MX480)**—Starting in Junos OS Release 19.3R1, you can upgrade the Juniper Device Manager (JDM) to support the Wind River Linux 9 (WRL 9)-based VM host for in-chassis Junos node slicing. For the upgrade, you don't need to change the configuration of the existing guest network functions (GNFs) in JDM, although you must stop the GNFs and JDM before you upgrade JDM. The JDM software version 19.3R1 supports both WRL 6-based and WRL 9-based VM host software versions.

  [See Junos Node Slicing Upgrade.]

Software Installation and Upgrade

- **Migration of Linux kernel version**—Starting in Junos OS Release 19.3R1, the following devices support the Wind River Linux 9 (WRL9) kernel version:

Platforms	Routing Engine Supported
ACX5448-D	RE-ACX-5448
MX240, MX480, and MX960	RE-S-X6-64G
MX2020 and MX2010	REMX2K-X8-64G
MX204	RE-S-1600x8
MX10003	RE-S-1600x8
MX2008	RE-MX2008-X8-64G
MX10016	RE X10
MX10008	RE X10
PTX5000	RE-PTX-X8-64G
PTX3000	RCBPTX
PTX10016	RE-PTX-2X00x4/RE X10
PTX10008	RE-PTX-2X00x4/RE X10
PTX1000	RE-PTX1000
PTX10002-XX	RE-PTX10002-60C
EX9208	RE-S-EX9200-2X00x6

Platforms	Routing Engine Supported
EX9251	EX9251-RE
EX9253	EX9253-RE
EX9204	RE-S-EX9200-2X00x6
EX9214	RE-S-EX9200-2X00x6
QFX10002	RE-QFX10002-60C
QFX10008	RE-QFX10008
QFX10016	RE-QFX10016

Starting in Junos OS Release 19.3R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following routers:

- MX Series—MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008
- PTX Series—PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

If you perform a software upgrade on a router with i40e NVM version earlier than 6.01, the upgrade fails and the following error message is displayed:

ERROR: i40e NVM firmware is not compatible ,please upgrade i40e NVM before installing this package

ERROR: Aborting the installation

ERROR: Upgrade failed

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

Subscriber Management and Services

- **Subscriber services uplink support on MPC10E line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.3R1, you can use the MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards for uplink connections to the core network. This support requires you to enable enhanced subscriber management.

  [See Protocols and Applications Supported by the MPC10E.]

- **Extended support for access models with heterogeneous subscriber types (MX Series)**—Starting in Junos OS Release 19.3R1, we support both four-level and five-level scheduler hierarchies for CuTTB, FTTH, and FTTB networks with both business and residential PPPoE subscribers. We support a network that includes all of the following subscriber types: PPPoE terminated residential, PPPoE business using ESSM op-scripts, PPPoE tunneled LAC, and L2BSA. The networks can include both conventional DSL and hierarchical access (CuTTB, FTTH, and FTTB) at the same time.

- **Dynamic interface sets for business subscribers in heterogeneous networks (MX Series)**—Starting in Junos OS Release 19.3R1, PPP can dynamically create business subscriber interface sets based on the physical interface name and the outer VLAN tag, matching the $junos-svlan-interface-set-name predefined variable. The AAA process sends the set name to the RADIUS server in the Juniper Networks QoS-Set-Name VSA (26-130) in the Access-Request message. The server returns the VSA in the Access-Accept only for business subscribers. RADIUS does not return the VSA for residential subscribers. The returned VSA is used to create a dynamic interface set for the business subscriber. Otherwise a default dynamic interface is created. You enable dynamic business subscriber interface sets with the **source-interface-set-at-login svlan** statement at the **[edit protocols ppp-service]** hierarchy level.

  [See Automatic Creation of Business Subscriber Interface Sets.]

- **Support for new PON and DSL G.fast DSL Forum VSA (MX Series)**—Starting in Junos OS Release 19.3R1, we have added support for the new G.fast and PON TLVs in ANCP Port Status Messages and PPPoE-IA tags per IETF draft extension to RFC 6320, *Protocol for Access Node Control Mechanism in Broadband Networks*. The CLI for traffic shaping based on DSL-Type (0x91) TLV or PON-Access-Type (0x92) TLV is now located in the **dsl** and **pon** stanzas at the **[edit system access-line]** hierarchy levels. The previously supported options are deprecated but redirect for a limited time to the new CLI. A new **type** option enables support for any future DSL or PON types by specifying the access line type value from the DSL-Type (0x91) TLV or PON-Access-Type (0x92) TLV.

  [See DSL Forum Vendor-Specific Attributes.]

- **New ANCP TLVs, PPPoE-IA tags, L2TP access line AVPs, and Juniper Networks VSAs (MX Series)**—Starting in Junos OS Release 19.3R1, new attributes have been added to support the new DSL G.fast and PON TLVs described in the IETF draft extension to RFC 6320, *Protocol for Access Node Control Mechanism in Broadband Networks*. The new TLVs are reported to RADIUS as VSAs and to the LNS in AVPs in alignment with RFC 5515, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*.

  [See ANCP DSL Attributes Mapped to Juniper Networks DSL Vendor-Specific Attributes.]

- **Support for OLT migration to PON (MX Series)**—Starting in Junos OS Release 19.3R1, use the **preference** statement at the **[edit system access-line attributes]** hierarchy level to configure how the router behaves when it receives both DSL TLVs and PON TLVs in ANCP Port Status messages or in PPPoE-IA tags. This situation occurs when the OLT redundantly reports the PON access line attributes both in PON TLVs and by overloading DSL TLVs. The DSL-Type TLV (0x91) is set to OTHER and PON rates for the subscriber access line are presented in the Actual-Net-Data-Rate-Upstream TLV 0x81 and Actual-Net-Data-Rate-Downstream TLV 0x82. The BNG saves and processes TLVs of the specified type and discards the other type.

  [See attributes (Access-Line Rate Adjustment).]

- **CoS adjusted shaping rate enhancements (MX Series)**—Starting in Junos OS Release 19.3R1, you can do the following:

  - Enable the router to automatically apply an adjusted shaping rate based on the DSL-Type TLV (0x91) and the PON-Access-Type TLV (0x92) to the following:

- PPPoE logical interface for a residential subscriber or to the dynamic business interface-set for a business subscriber.
- Backhaul node or shared media node that represents the PON tree (FTTH or FTTB) or a bonded Copper access-line (CuTTB).

- Define priority-level shaping rates as a percentage of the overall shaping rate to apply to an L2 CoS node dynamic interface set, where the interface set represents a backhaul node or a shared media node for multiple subscribers.

[See ANCP Agent Traffic Shaping and CoS.]

SEE ALSO

What's Changed

IN THIS SECTION

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

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

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

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

  See Configuring RFC-Compliant NETCONF Sessions.

What's Changed in Release 19.3R2-S5

Network Management and Monitoring

- **Change in startup notification after GRES (MX Series routers)**— The master Routing Engine sends a coldStart notification when a device comes up. The master Routing Engine also sends warmStart notifications for subsequent restarts of the SNMP daemon. After graceful routing engine switchover (GRES) the new master Routing Engine sends a single warmStart notification and the backup Routing Engine does not send any notification. In earlier releases, after GRES, the new master RE would sometimes send two notifications or a single notification. Of these, the first notification was always a coldStart notification and the second was either a coldStart notification or a warmStart notification.

  [See Standard SNMP Traps Supported by Junos OS.]

What's Changed in Release 19.3R2

Interfaces and Chassis

- **GNFs inherit hyper mode configuration from the BSYS (MX2010, MX2020, MX480, MX960, and MX2008)**—Starting in Junos OS Release 19.3R2, running the **show forwarding-options hyper-mode** command on a guest network function (GNF) displays the hyper mode status of the chassis (BSYS), not of the GNF. This is because hyper mode is a chassis-specific feature. This means that you cannot have different hyper mode configurations on individual GNFs as they inherit the hyper mode configuration from the chassis.

  [See show forwarding-options hyper-mode.]

Subscriber Management and Services

- **Enhancement to commands to display reason for Routing Engine disconnect (MX Series)**—Starting in Junos OS Release 19.3R2, several commands display the reason when the master and standby Routing Engines disconnect because of a memory mismatch error. On a chassis with two Routing Engines, a DRAM size mismatch error can result when both of the following are true:

  - The Routing Engines have different amounts of DRAM.
  - A 64-bit Junos OS image is loaded on the chassis.

  You can avoid this problem by doing either of the following:

  - Ensure that both Routing Engines have the same amount of DRAM.
  - Load a 32-bit image.

  The **show database-replication summary** and **show system subscriber-management summary** commands display the DRAM mismatch as the reason in the Disconnection field. The **request chassis routing-engine master switch check** command displays an error message if the DRAM size is different for the two Routing Engines.

- **Prevent queue-based throttling from stopping subscriber login (MX Series)**—Starting in Junos OS Release 19.3R2, you can specify a value of 0 with the **high-cos-queue-threshold** statement. This value prevents any subscriber from being throttled by queue-based throttling.

What's Changed in Release 19.3R1

EVPN

- **ARP suppression disabled by default (MX series)**—Starting in Junos OS Release 19.3R1, ARP suppression is disabled by default when you configure EVPN VLAN Bundle Services or when you include the **encapsulate-inner-vlan** option or the **decapsulate-accept-inner-vlan** option in the VLAN configuration.

General Routing

- **User confirmation prompt for configuring the suboptions of request vmhost commands (MX Series and PTX series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm with a [**yes** or a **no**] whether you want to configure the suboptions also.

  - **request vmhost reboot**
  - **request vmhost poweroff**
  - **request vmhost halt**

  In earlier Junos OS releases, the confirmation prompt is available only for the main options.

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Release 19.3R1 and later, MX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

  In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

- **Monitoring information available only in trace log (MX Series)**—In Junos OS Release 19.3R1 and later, the Ethernet link fault management daemon (lfmd) in the peer router stops monitoring the locally occurred errors until unified ISSU completes. You can view the monitoring-related details only through the trace log file.

Junos OS XML API and Scripting

- **Range defined for confirm-timeout value in NETCONF and Junos XML protocol sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, the value for the **<confirm-timeout>** element in the Junos XML protocol **<commit-configuration>** operation must be in the range 1 through 65,535 minutes, and the value for the **<confirm-timeout>** element in the NETCONF **<commit>** operation must be in the range 1 through 4,294,967,295 seconds. In earlier releases, the range is determined by the minimum and maximum value of its unsigned integer data type.

MPLS

- **Deprecated statement (MX Series)**—Starting in Junos OS Release 19.3R1, the **preference** statement is deprecated at the **[edit protocols source-packet-routing source-routing-path *name*]** hierarchy level. This is because you can have two different sequences of the same route, wherein the active route entry that is selected can be different.
- **IPv4 explicit-null label retained from the merged protocol MPLS label stack**—The IPv4 explicit-null label is retained from the merged protocol MPLS label stack, if the IPv4 explicit-null is at the bottom of the MPLS label stack.

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when a change in number of supported history records is detected (ACX Series and MX Series)**—In Junos OS Release 19.3R1, when Ethernet connectivity fault management starts, it detects the number of history records supported by the existing performance monitoring history database. If there is any change from the number of history records supported (that is, 12) in Release 19.3R1, then the existing performance monitoring history database is cleared and all performance-monitoring sessions are restarted with mi-index 1.

Routing Protocols

- **Change in the default behavior of advertise-from-main-vpn-tables configuration statement**—BGP now advertises EVPN routes from the main bgp.evpn .0 table. You can no longer configure BGP to advertise the EVPN routes from the routing instance table. In earlier Junos OS Releases, BGP advertised EVPN routes from the routing instance table by default.

  [See advertise-from-main-vpn-tables].

Services Applications

- **Hide HA information when the service set does not have HA configured**—When you run the **show services ha detail** command on a configuration with a service set, which does not have HA configured, the HA information is hidden for the service set in the output.
- **New syslog message displayed during NAT port allocation error (MX Series routers with MS-MPC)**—With address pooling paired (APP) enabled, an internal host is mapped to a particular NAT pool address. If all the ports under a NAT pool address are exhausted, further port allocation requests from the internal

host results in a port allocation failure. The following new syslog message is displayed during such conditions:

JSERVICES_NAT_OUTOF_PORTS_APP

This syslog message is generated only once for each NAT pool address.

- **Change in NAT port block syslog message display (MX Series routers)**—When you configure a softwire prefix other than 128, all the JSERVICES_NAT_PORT_BLOCK logs now display the prefixed B4 address. The following JSERVICES_NAT_PORT_BLOCK are modified:
  - JSERVICES_NAT_PORT_BLOCK_ALLOC
  - JSERVICES_NAT_PORT_BLOCK_RELEASE
  - JSERVICES_NAT_PORT_BLOCK_ACTIVE

  In earlier releases of Junos OS, when a softwire prefix is configured, some of the B4 addresses displayed in the JSERVICES_NAT_PORT_BLOCK log are /128 addresses (irrespective of the configured prefix). This change is not observed when the softwire prefix is not configured.

Software Defined Networking (SDN)

- **Increase in the maximum value of delegation-cleanup-timeout (MX Series)**—You can now configure a maximum of *2147483647* seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

  With the increase in maximum value of **delegation-cleanup-timeout** from *600* to *2147483647* seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

  [See delegation-cleanup-timeout.]

Subscriber Management and Services

- **Support for pseudowire physical interface for ANCP Autoconfiguration (MX Series)**—Starting in Junos OS Release 19.3R1, you can associate an ANCP neighbor with a subscriber-facing interface pseudowire physical interface for ANCP autoconfiguration of VLANs. When configured, ANCP Port Up and Port Down messages received on the interface trigger notifications to the auto configuration daemon (autoconfd) to initiate VLAN creation (Port Up) or removal (Port Down).

  [See Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs.]

System Logging

- **Preventing system instability during core file generation (MX Series)**—Starting with Release 19.3R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. Core files are generated only if there is sufficient available space. Otherwise, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

What's New | **84**

Known Limitations | **128**

Open Issues | **132**

Resolved Issues | **152**

Documentation Updates | **180**

Migration, Upgrade, and Downgrade Instructions | **181**

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 Routing Engine boots from the secondary disk when you:

  1. Press the reset button, on the RCB front panel, while the Routing Engine is booting up but before Junos OS is up.

  2. Upgrade software, by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network.

  3. Upgrade BIOS and the upgrade fails.

  4. Reboot and the system hangs before Junos OS is up. PR1344342

- First packet pertaining to J-Fflow Packet Forwarding Engine sensor in UDP mode is missing after line card reboot on PORTER-R platform. PR1344755
- During unified ISSU that warrants host upgrade, if the router is configured with 8 million IPv4 or IPv6 routes or more, the unified ISSU might fail, resulting in FPC restart. PR1348825
- If the MTU is configured to a value higher than 9500 which is the maximum permissible value, the configuration succeeds, but the actual value is set back to 1518B without any error. You can check the DCD log to verify the occurrence. PR1372690
- The MIC-MACSEC-20G supports 10-Gbps speed through the **set chassis fpc x pic y pic-mode 10G** configuration applied to both the PICs in that MIC. Any other PIC mode configuration should be removed and then the 10-Gbps PIC mode configuration is to be applied. PR1374680

- The dfe tuning failing at times is a known issue on MX10003. The only recovery option in this situation is to restart the FPC. PR1413233
- The MX Series Packet Forwarding Engine, does not account for the labels pushed onto the packet on the egress Packet Forwarding Engine, while the PTX Series Packet Forwarding Engine does. This results in a slight difference in the byte count for the same traffic stream across these two platforms. The packet count is still be the same across the platforms. Currently, this issue is noticed for uncolored SR-TE policies. PR1416738
- Since the loopback was created at the MACsec port (remote end) in this specific situation, the link itself is down at the EA port. Therefore, PRBS test fails with incrementing error counts. PR1421432
- HQoS configuration on a ps- interface anchored to a logical tunnel fails to commit with the following error: **[edit class-of-service interfaces ps0 unit 10]'output-traffic-control-profile cannot configure traffic control profile (pic has no CoS queuing) error: configuration check-out failed** . PR1429927
- Operational mode file checksum option is not available to specify the checksum for commit script. PR1431064
- [subscriber_services/cos-bbe] CLI **show class-of-service scheduler-hierarchy out-of-resource fpc slot** did not give OOR table in full scaling test. PR1433687
- [GRE] [MPC10] 100% traffic drop if I enable fragmentation on core interface side [by setting egress MTU < GRE MTU]. PR1433783
- CHASSISD core found @ fpc_sfxpc_la_ng_show_hw ui_sfxpc_show_hardware ms_parse_substring. PR1434188
- SPC3 cards are not supported with RE-2000, even if the RE-2000 is the backup Routing Engine. PR1435790
- Routing Engine interprets any input from the console port as interrupts. Depending on the frequency, console noise will impact the Routing Engine interrupt handling to different extents even with the current throttling mechanism. When the interrupt frequency is too high for the Routing Engine to handle, the impact can vary from line card reboot (partial impact) to Routing Engine reboot (chassis-wide impact). PR1436386
- In a large-scale setup (such as large number of routing instances or interfaces), if there are frequent changes in configuration and interface flapping when the rpd is restarted by deactivating and then activating the logical-system or restart routing, the rpd might crash. PR1438049
- The **number-of-sub-ports** configuration is not supported on the MPC11E interfaces. PR1442439
- vmx-zt: JDI-RCT:vRCT -- *,g route entries are not updated. PR1443515
- In a scaled scenario where Routing Engine pushes a lot of routes to Packet Forwarding Engine in presence dynamic tunnel configuration, FIB convergence might take more time leading to traffic drops. PR1454817
- Member lt- interfaces of an RLT must have same bandwidth configured. Bandwidth mismatch can lead to unexpected behavior. Changes to lt-/RLT interfaces should not be done if a PS interface is anchored over these tunnel interfaces. PR1458951

- The lt-interface scheduler remain in invalid state under egress physical interface list after changing the lt- tunnel to a different Packet Forwarding Engine, which has no functional impact. PR1458955
- DCD errors are seen with scaled interface configuration when interfaces are deleted and added. This is because of the fact that the Packet Forwarding Engine is slower compared to Junos OS kernel. So, we run out of hardware resources for a brief time. Once the Packet Forwarding Engine is able to program the entries, Junos kernel pushes the configuration once again and the state is restored in the Packet Forwarding Engine. PR1445778
- The core file is removed after upgrading to a new image. PR1447659
- Occasional warning messages such as **TCP connect error** can be seen during FPC reboot. These are generally inconsequential and have no impact on the FPC or the line card software functionality. PR1460153
- ifJnxOutQDrops counter value shows decreasing counter values as it seems to be rolling over. PR1460199
- When untagged frames are received on tagged interfaces, even though the traffic receives proper treatment, the **show interfaces extensive** CLI command output might not accurately reflect the input packet reject counter values under filter statistics section of the output. PR1463199
- Native VLAN cannot be configured for trunk or access ports in ZT line cards running Junos OS Release 19.3R2. PR1463544
- On MPC11E line card, the backup Routing Engine reflects as normal mode in cli **show forwarding-options hyper-mode**. PR1466719
- On MPC11E, in the DOM MIB, for the channelized 10-Gigabit Ethernet interface LF/RF alarms do not show up. PR1467446
- FPC online might take additional time during movement of MPC11 FPC from one GNF to another GNF. PR1469729
- When multiple PFE sensors are configured, sensor data export is not happening at the configured reporting interval. PR1469970
- When multiple Packet Forwarding Engine sensors are configured, sensor data export is not happening at the configured reporting interval. PR1469970
- Transit interface statistics for input and output bytes might be inaccurate when Layer 2 accounting overhead is configured at the PIC level or at the logical interface level on MPC10E and MPC11E. In this scenario, the total interface statistics will account for the configured Layer 2 accounting overhead, but the local interface statistics will not account for the configured Layer 2 accounting overhead. The transit interface statistics might be inaccurate because the Layer 2 accounting overhead for locally generated packets is not considered. PR1470150

Infrastructure

- On Juniper Networks Routing Engines with the Hagiwara CompactFlash card installed, after upgrade to Junos OS Release 15.1 and later, the failure message **smartd[xxxx]: Device: /dev/ada1, failed to read SMART Attribute Data** might appear on the messages log. PR1333855
- AUTO-CORE-PR : JDI-RCT vRCT : CRON core file is found @ **cron_popen child_process do_command**. PR1434152

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after the upgrade. This is because of the old version of the presence of the **/var/db/cfm.db**. PR1281073
- In a large-scale subscriber environment, changing aggregated Ethernet member link configuration might generate core files for the two Routing Engines. PR1375638
- When disabling the physical interface with JNP-100G-AOC-xM AOC cables, the port LED could turn red or go off depending on vendor. JNP-100G-AOC-xM cables sourced by Finisar cause the port LED to turn red when the physical interface is disabled. Cables sourced by Innolight causes the port LED to turn off. Transceiver vendor information can be obtained from the **show chassis pic fpc-slot <fpc-slot> pic-slot <pic-slot>** CLI command. The transceiver vendor field displays 'JUNIPER-FINISAR' for Finisar and 'JUNIPER-INNO' for Innolight. PR1415958

MPLS

- BUD node replicating duplicate packets toward egress PE device occurs when we have S-RSVP-TE P2MP with vt interfaces. PR1452864

Platform and Infrastructure

- On all platforms running Junos OS, execution of Python scripts through enhanced automation does not work on veriexec images. PR1334425
- Saltstack: Applying Salt state file with function junos.cli with invalid command reports status True and Succeeded: 1. PR1429675
- After reboot, the initial several packets might get lost if the traffic passes through an FT interface. PR1431983
- Saltstack: Salt execution module junos.install_os fails to remove the tmp image file **__salt.tmp.ecc1ME** on target minion after image installation. PR1432123

Routing Protocols

- [BGP] When scaling RIB to 80 million after FPC restart, cannot scale on backup Routing Engine. PR1444073

SEE ALSO

What's New | 84

What's Changed | 121

Open Issues | 132

Resolved Issues | 152

Documentation Updates | 180

Migration, Upgrade, and Downgrade Instructions | 181

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.

EVPN

- With Junos OS Release 19.3R1, VXLAN OAM host-bound packets are not throttled with DDoS policers. PR1435228
- When DHCP is used with EVPN, L2 learning daemon adds a destination route to the kernel with the **permanent remote** flag while dhcp process adds a destination route with the **permanent remote** flag. There could be a race condition where the l2 learning destination route is overwritten by DHCP route, causing the remote flag to get deleted. This subsequently leads to the ARP route to age out in the kernel. To ensure that DHCP routes are not added to the kernel, you must configure the **forward-only** statement under **forwarding-options dhcp-relay**. PR1439568
- The rpd process might crash with EVPN related configuration changes in static VXLAN to MPLS stitching scenario. PR1467309

Forwarding and Sampling

- The **skip-service** configuration does not work with IPv6 NDP negotiation or ping. PR1074853
- When IPv4 prefix is added on a prefix-list referred by 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
- Observed error message: **Process: dfwd, path: [edit interfaces ae2 unit 0 family inet], statement: filter, Index for referenced filter input_ipv4_ngn_filter is not defined**. PR1433146
- Traffic error is observed after locally switching for VLAN 100 and 101, while verifying selective local-switching functionality with 4000 VLANs. PR1436343
- The layer 2 learning CPU utilization might get high and remain stuck forever after switching configuration files several times between EVPN and non-EVPN (for example, VRRP) by loading the corresponding configuration file. Because of that, some of the data in the device is not successfully cleaned up when EVPN-configuration (virtual-switch) is removed and the Ethernet Segment Identifier (ESI) interface is configured in a non-EVPN routing instance. PR1446568
- For an aggregated Ethernet bundle of at least two members hosted at two different FPCs, if the aggregated Ethernet interface is with CoS, output-traffic-control-profile of shaping-rate, and the output filter of policer with logical-bandwidth-policer and bandwidth-percent, the aggregated Ethernet interface might have incorrect effective output policing rate. PR1466698

General Routing

- On MX Series routers with MPC5 or MPC6 installed, added support for periodic polling gigabit Ethernet switch errors and update chassis alarms if threshold is crossed: **alarmd[xxxx]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 2 Major Errors - GESW Error code: 0x18000d craftd[yyyy]: Major alarm set, FPC 2 Major Errors - GESW Error code: 0x18000d chassisd[zzzz]: FPC 2 sent action 0x44 that is not configured 0x2 Cmerror Op Sub Set: GE Switch: MAC TX Errors (47) Exceeded Raise Threshold (3) for Port (13:TSEC-1) Cmerror Op Sub Set: GE Switch: MAC Late Collisions (47) Exceeded Raise Threshold (3) for Port (13:TSEC-1)**. PR1004434
- The fxp0 is marked as "Dest-route-down" because of the specific operations like disabling and enabling operations. PR1052725
- If a Layer 3 interface is receiving a GRE encapsulated packet and the interface has the following two filters attached in ingres—**family any** with action as mirror and **family inet** with action as **decapsulate gre**—then the expected behavior is that the mirrored copy must have the GRE headers also. However, the configuration is not working as expected, because of the presence of the **family-inet** filter. If you want to mirror entire packets that reach the interface (that includes GRE header), the workaround is to deactivate or disable the de-encapsulate GRE action of that filter. PR1090854
- On the MX104 platform, when using snmpbulkget or snmpbulkwalk (for example, used by the SNMP server) on a chassisd-related component (for example, jnxOperatingEntry), chassis process (chassisd) high CPU usage and slow response might be seen because of a hardware limitation, which might also lead to a query timeout on the SNMP client. In addition, the issue might not be seen while using an SNMP query for interface statistics. As a workaround, to avoid the issue, use either of the following approaches:
  - Use snmpget or snmpwalk instead of snmpbulkget or snmpbulkwalk and include the -t 30 option when doing the SNMP query. For example, snmpget -v2c -c XX -t 30.
  - Use the -t 30 option with snmpbulkget or snmpbulkwalk. For example, snmpbulkget -v2c -c XX -t 30. PR1103870
- SIP session fails when the IPv4 SIP client in a public network initiates a SIP call with the IPv6 SIP client in a private network. PR1139008
- On dual Routing Engines with graceful Routing Engine switchover (GRES) enabled, after performing GRES, if the configuration synchronization on the backup Routing Engine fails when it becomes the new master Routing Engine, then in rare conditions, some interfaces cannot be deleted or configuration changes cannot be committed. PR1179324
- When same UID objects are used in both inet and inet6 services of the same subscriber session, deactivating the first session cause conditions which avoid releasing UID entry after deactivating second service session. This leads to having stale UID entry and can cause subscriber's connection problem in the future when UID pool might be completely exhausted. The probability of hitting the issue increases if amount subscriber to amount of unique services ratio is approaching 1 (that is, when almost every subscriber has a service with unique service objects). PR1188434

- **SPD_CONN_OPEN_FAILURE** and **SPC_CONN_FAILURE** log messages are seen in the log for SI interfaces when running SNMP walk on Service PIC NAT OIDs. The SNMP NAT OIDs are not used for the inline NAT solution. They are used for the NAT solution that runs on the service cards such as the MS-DPC, MS-PC, or MS-MIC. PR1310081
- During some operations, **krt_junos_sanity_check_ctrl_resp: rtsock request finally succeeded after error 16** syslog message causes the kernel to get busy. PR1310678
- On MX Series routers, the routing information base (RIB, that is routing table) and forwarding information base (FIB, that is forwarding table) might get out of sync in a very large scale network due to the kernel routing table (KRT) queue being stuck. The KRT queue is used by the rpd process to send forwarding information messages to Packet Forwarding Engines. With the stuck state, the queue can get into a state where no more messages can be sent to the Packet Forwarding Engines. This issue is applicable only to Junos OS Release 16.1 through Junos OS Release 17.3. PR1315212
- The **chain-composite** statement does not bring in a lot of gain because TCNH is based on 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. The reactions to failure situations might not be handled in a graceful way. TCP connection timeout occurs because the jlock hog crossing boundary value (5 seconds) causes bad consequences in the MX Series Virtual Chassis. Currently, there is no other easy solution that might reduce this jlock hog besides enabling marker infrastructure in the MX Series Virtual Chassis setup. PR1332765
- The log messages are generated on all Junos OS systems with transportd daemon support for optics monitoring. PR1337068
- First packet pertaining to J-Flow Packet Forwarding Engine sensor in UDP mode is missing after the line card reboots on MX150 platform. PR1344755
- With GRES enabled in a subscriber environment, if subscribers are logging in or logging out quickly, the service sessions in the session database (SDB) of the backup Routing Engine might leak. If the problem is not detected for a long time, the backup Routing Engine might not be able to synchronize with the master Routing Engine and will not be ready for GRES. PR1346300
- Backup Routing Engine might crash after more than 10 continuous GRES switchovers. PR1348806
- On next-generation Routing Engine, a failure of the Hardware Random Number Generator (HWRNG) will leave the system in a state where not enough entropy is available to operate. PR1349373
- In some cases, online insertion and removal (OIR) of a MIC on an FPC might silently drop the traffic destined to the FPC. The only way to recover from this is to restart the FPC. The issue is seen when corresponding CLI commands are used to go offline and then online the MIC. PR1350103
- The Packet Forwarding Engine selector stuck in a rerouted state on the unilist next hop after the primary aggregated Ethernet link is deactivated and then deactivated. PR1354786
- For configurations of bridging routing instances with aggregated Ethernet logical interfaces (6400 logical interfaces) and IRB instances, all from a single FPC, the CPU utilization of the FPC stays at 100 percent

for 4 minutes. The behavior from PFEMAN of the FPC has the processing time spiked on IF IPCs and this seems to be the case of MPC7E from Junos OS Release 16.1R1 (or even earlier). After 4 minutes, the CPU utilization comes down and the FPC is normal. Therefore, this scaled configurations on MPC7E takes a settling time of more than 4 minutes. PR1359286

- In rare circumstances, a faulty SFP transceiver installed in an MX104 might cause the AFEB to go offline. The backup Routing Engine and fan tray will also show an alarm. PR1360426
- When the GRES is enabled, if the Link Aggregation Control Protocol (LACP) is configured on the aggregated Ethernet interface, the aggregated Ethernet interface might flap when the Routing Engine switchover happens. For example, the unified ISSU is performed or the master Routing Engine is restarted. This might extend traffic loss (more than 5 seconds, depending on how fast the aggregated Ethernet interface recovers). See also https://kb.juniper.net/TSB17578 PR1365316
- When an FPC is booting up (either during unified ISSU, router reboot, or FPC restart), I2C timeout errors for SFP transceiver can be noticed. These errors are seen the I2C action is not completed because the device was busy. When the FPC is up, all the I2C transactions to the device are all right, so no periodic failure is observed. There is no functional impact, and these errors can be ignored. PR1369382
- If any of the log messages continue to appear in the MPC console, it indicates the presence of a faulty SFP/SFP+ transceiver, which is causing I2C transaction failure from main board CPU. There is no software recovery available to recover from this situation. These logs also indicate potential I2C transaction failure with any of the 10 ports available with GMIC2 in PIC 0 resulting in unexpected behavior. For example, a link does not come up or the MIC itself does not boot on restart. PR1375674
- A few xe interfaces are going down with **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error**. PR1377840
- Rsmon is blocking subscribers from logging in because the number of CoS scheduler queues is at 100 percent. PR1390424
- In a BGP Prefix-Independent Convergence (PIC) case, If a route R1 resolves on top of the multipath route R2, where R2 has primary and backup indirect next hops, it will be better if the backup leg is not used for the resolution of R1. There is no impact on any existing CLI commands. The backup path is never used when the primary path is available. PR1401322
- The MX104 router has the following limitations in error management. The **show chassis fpc error** CLI command is not available for MX104 in Junos OS Releases 13.3R7, 15.1R2,14.1R5,14.2R4, 13.3R8, and later. Junos OS does not initiate restart of the system on encountering a fatal error. Although you can configure the action to disable Packet Forwarding Engine for major errors, Junos OS does not disable the Packet Forwarding Engine on encountering a major error. PR1413314
- FPC core files are generated on multiple additions or deletions of hierarchical CoS from pseudowire devices. As a workaround, remove the pseudowire device without changing the hierarchical CoS configuration. PR1414969
- cRPD does not restrict the number of simultaneous JET API sessions. PR1415802

- Deleting **active-lease-query** configuration might sometimes lead to the generation of core files or the TCP connection might remain active. PR1415990
- On MX Series routers with MPCs, unicast traffic might get dropped when the destination is reachable over an integrated routing and bridging (IRB) interface and a label-switched interface (LSI) with two next hops. PR1420626
- Certain JNP10008-SF and JNP10016-SF SIBs manufactured between July 2018 and March 2019 might have an incorrect core voltage setting. As a workaround, reprogram the core voltage and update the setting in NVRAM memory. PR1420864
- If HTTP header enrichment function is used, the traffic throughput decreases when the traffic passes through header enrichment. PR1420894
- On an MX204 platform, the allocation of MAC address for the second PIC in the FPC might fall out of the MAC address pool, which might further cause a MAC conflict in the network. PR1422679
- In Junos OS Release 19.1 onward, Junos OS supports RFC 8231/8281 compliance by default. However, if the controller is not compliant with RFC 8231/8281, a backward compatibility statement can be configured to fall back to the pre-RFC 8231/8281 behavior. PR1423894
- The **show lldp neighbors interface** CLI command shows LLDP neighbor information only for particular interface. To know the information of all the neighbors, execute the **show** command for all the interfaces one by one. PR1426793
- On MX Series platforms, if you are configuring the sensors to stream data over UDP in static database, decoding of telemetry data at collector might not be correct. PR1426871
- NPC core file is generated at trinity_rtt_hw_bulk_helper,trinity_rt_delete,rt_entry_delete_msg_proc (rt_params=0x48803bd8) at **../../../../../../../../src/pfe/common/applications/route/hal/rt_entry.c:5210**. PR1427825
- On MX Series routers, the syslog **error: [Err] dfw_abstract_issu_stats_counters_restore:2222 Failed to find Index = 4613734? during ISSU with 19.3I-20190409_dev_common.0.2212** is observed. PR1429879
- On MX204 and MX10003 platforms, when a direct attach copper (DAC) QSFP port is disabled, the interface on the peer device might still remain up. It might have traffic impact in the network. PR1430921
- With 64-bit rpd running and traceoptions configured (for example, for BGP or MPLS statistics), the trace files are not rotating or rolling over as per the configured file size limit and the logs continue to be written to a single file continuously. PR1431033
- On MPC10, the error message **failed, Return code: 500** is seen with baseline. PR1431552
- Deleting the **chassis license bandwidth** configuration statement will not cause the device to default to maximum bandwidth available or entitled. PR1433157
- JSD crashes and generates a core file when subscribing and unsubscribing both gRPC and gNMI subscriptions aggressively from multiple sessions. PR1433744
- In gRIBI, programmed routes reference a next-hop group ID, which in turn points to one or more next-hop IDs. Each next- hop ID contains details of the actual next hop. Next-hop group ID and next-hop ID are

mapped to an IPv6 prefix (for example, FC01:: <GRP ID or NHOP ID>). In the case of an IPv4 indirect next hop, gRIBI needs to resolve IPv6 through IPv4 next hop over three levels of indirection. Junos OS does not support IPv6 over IPv4 multilevel next-hop resolution. Therefore, gRIBI cannot resolve next-hop GRPID <FC01::grpid> nexthop ID <FC01::nhopid> through an actual indirect IPv4 gateway address. PR1434050

- Some non-Juniper 40-Gigabit Ethernet SFPs might utilize 100-Gigabit Ethernet QSFP28 marking in their EEPROM, indicating CDR bypass mode, which enables the use of 100-Gigabit Ethernet optics at 40-Gigabit Ethernet speeds. On some 40-Gigabit Ethernet line cards, Junos OS detects an incorrect pluggable QSFP28 of type 0x11 (17 decimal) inserted into a QSFP+ of type 0x0d in the device and reports this error to syslog. PR1434183
- Packet Forwarding Engine major error is seen: **MQSS_CMERROR_DRD_RORD_ENG_INT_REG_CMD_FSM_STATE_ERR on MPC8/9E card**. PR1434278
- Traffic is not getting balanced by the weight specified in next hop group. PR1435286
- MPC10E 3D MRATE-15xQSFPP: L2overGRE is not supported in Junos OS Release 19.3R1. Though the configuration gets committed, the feature will not work. PR1435855
- ZF interrupts out-of-range dest PFE INTR for Gnt seen during MPC6 and MPC9 line card bringup, which has no functional impact. PR1436148
- Routing Engine interprets any input from the console port as interrupts. Depending on the frequency, console noise will impact the Routing Engine interrupt handling to different extents even with the current throttling mechanism. When the interrupt frequency is too high for the Routing Engine to handle, the impact can vary from line-card reboot (partial impact) to Routing Engine reboot (chassis-wide impact). PR1436386
- Mutiple interfaces on specific FPC are going down on MX480 after baseline profile configuration is verified. PR1437221
- On all Junos OS platforms, if hash-key is enabled under chassis, packets might be dropped because chassisd crashes the packets on other FPCs for which the hash-key is disabled. PR1437855
- In VSTP scenario, if **flexible vlan tagging** is configured on the interface and multiple logical interfaces are configured for the interface, if **vlan all interface all** is configured under VSTP, not all interfaces are enabled for this protocol. PR1439583
- If a route is programmed with Rib.RouteAdd/Update() with a VXLAN FlexibleTunnelProfile having destination MAC set to all zeros, then reading the route back via Rib.RouteGet() will result in destination MAC not being set in the RouteGetReply. PR1439940
- Syslog error message **UI_SCHEMA_MISMATCH_SEQUENCE: Schema header sequence numbers** is expected in Junos OS Release 19.3 but it has no functional impact. PR1440141
- In subscriber management, smihelper, subinfo, and jpppd crash due to the corruption in the secondary SDB. PR1440277
- When laser receiver power gets -inf, telemetry value corresponding to -infinity should be equivalent IEEE 754; that is, single precision float, 32bit value should be 0xff800000. PR1441015

- Starting in Junos OS Release 19.3R1 and beyond the underlying OS is upgraded to WRL9. In general VMHost Architecture supports Junos Only Upgrade/Downgrade. We recommend to use VMHost Upgrade/Downgrade. Given the extent of changes when base OS has changed, doing Junos Only Upgrade/Downgrade to Junos OS Release 19.2 or earlier from WRL9 image is not possible. PR1441048
- Egress stream flush failure and silent dropping of traffic could occur in a rare occasion for a repeatedly flapping link on MPC7E, MPC8E, and MPC9E cards. PR1441816
- On routers running Junos OS and serving as EVPN gateways, FPC core files at heap_block_log due to NULL entries are also seen in the IFBD level list, which are typically added for flush list. So this seems to be the side effect of the relink logic failure flush logic for MACs when there is ifbd/bd delete. PR1441824
- In Junos OS PTP deployment where a configured child logical interface is in the PTP configuration and AE is in the interface configuration, during Packet Forwarding Engine initialization, Packet Forwarding Engine microcode is not able to find the correct outgoing interface (OIF) to send the packet to and takes the host route path, leading to congestion and interfaces being brought down. PR1442665
- The BGP session establishing over the GRE tunnel will fail when the router receives the BGP packets encapsulated as GRE and uses the firewall filter action to de-encapsulate the GRE header. PR1443238
- Interface hold-down timers cannot be achieved for less than 15 seconds on MPC11E at FRS. Because of vendor limitations, achieving subsecond hold timers is not possible. PR1444516
- When there is **no destination-port** statement enabled, by default the MX Series router rewrites the HTTPS request with destination-port 80. This behaviour is observed only when using converged CPCD. When using Static CPCD, the MX Series router is not modifying the destination-port. PR1446085
- On MX Series with MPC line cards, when using ps interface over rlt (Redundant Logical Tunnel) in l2circuit, the PW (Pseudowire) taffic gets blackholed if "no-control-word" is enabled. PR1447917
- In the vMX scenario, when the VRRPv3 is configured, the IPv6 packets that go through the VRRPv3 gateway (vMX) will be forwarded at lower bandwidth. If the IPv6 traffic throughput is beyond this bandwidth, the IPv6 packets might be dropped. PR1449014
- On the MPC10E-15C-MRATE with 25-Gigabit Ethernet ports, the FEC statistics are not getting reset after changing FEC mode. The old FEC statistics from previous FEC settings are carried over to the new FEC. PR1449088
- EX2300. EX3400, EX4300, EX4650, EX9250, MX150, MX204, the MX10000 line of routers, PTX1000, and the PTX10000 line of routers equipped with next-generation Routing Engines, if the helpers bootp command is enabled at the [forwarding-option] hierarchy, the Dynamic Host Configuration Protocol (DHCP) relay feature might not work as expected. Due to this issue, all the DHCP discovery packets could not be relayed. PR1449201
- In a vMX scenario, because the host CPU Core0 is not dedicated for virtualization, it might also be used by the other application. However, the vMX Virtual Routers' vCPU resource allocation is started from Core0, so when the usage of Core0 goes high, vMX might get less resources, which might lead to packet drop. PR1449289

- Currently IS-IS is sending system host-name instead of system-id in OC paths in lsdb or adjacency xpaths in periodic streaming and on-change notification. PR1449837
- The **show ddos-protection protocols arp statistics |display xml** command does not show APR violation packets and also is not incremented. PR1449968
- Traffic drop is seen from SD to AD. Loss is intermediate and not seen regularly. This issue occurs because of the few packets transmitted from the egress of AD1 are short of FCS (4 bytes) + 2 bytes of data, due to which the drops occur. Normal data packets are of size 128 bytes (4 bytes FCS + 14 bytes Ethernet header + 20 bytes IP header + 90 bytes data), while the corrupted packet is 122 bytes (14 bytes Ethernet header + 20 byte IP HEADER + 88 bytes data). PR1450373
- MPC10E dcd is unable to clean stale mt- logical interfaces while reloading Rosen configuration on the DUT. PR1450953
- Error messages are seen with certain configurations and these have no functionally impact. PR1451259
- When using the **replace pattern** CLI command to replace the name in the apply-group, the mgd will crash. PR1452136
- Physical interface policer are not supported in Junos OS Release 19.3 for MPC11. PR1452963
- With FPC restart done in presence of dynamic tunnels configuration, traffic might not resume sometimes. PR1454325
- Timestamp would be missed when command is edited and then run from CLI command history. PR1454387
- On MPC3E-NG cards with 100-Gigabit Ethernet interface in use, if the interface detects Loss of Lock (LOL) on the link without Loss of Signal (LOS), the interface goes down and might not come up again after the link is recovered. PR1454595
- Incorrect protocol number is displayed on issuing **show firewall log** command. PR1454792
- In a scaled scenario where the Routing Engine pushes a lot of routes to the Packet Forwarding Engine in a dynamic tunnel configuration, FIB convergence might take more time, leading to traffic drops. PR1454817
- Interface damping is not supported on MPC11E. PR1455152
- IPv6 accounting stop attributes are not correct for MLPPP subscribers. PR1455175
- The AFT **show nh detail index < id>** command is not showing similar output to non-AFT command **show nh id < id> extensive**. PR1455303
- With the latest provided Schema, **groups** and **services** configuration statements were missing in Junos Space GUI for both the MX10008 and MX10016. PR1455383
- After software upgrade, SNMP MIB walk does not poll or fetch any information. PR1455667
- Packet drops might be encountered with MPLSoUDP configurations in certain releases of Junos OS including Junos OS Release 18.3R3. PR1455753

- Multiple daemons might experience problems after a configuration change related to groups (core files are also possible). The following log is observed: **Aug 4 12:14:34.755 xxxxxxx dfwd[279]: %DAEMON-0: ../../../../../../src/ui/lib/access/model.c:239: insist 'model > 0 && model <= MODEL_MAX' failed Aug 4 12:14:34.935 xxxxxxx lrmuxd[294]: %DAEMON-0: ../../../../../../src/ui/lib/access/model.c:239: insist 'model > 0 && model <= MODEL_MAX' failed Aug 4 12:14:34.961 xxxxxxx rpd[277]: %DAEMON-0: ../../../../../../src/ui/lib/access/model.c:239: insist 'model > 0 && model <= MODEL_MAX' failed.** Issue prevents commit, but service might not be impacted. PR1455960
- With logical-system configuration, filter-based GRE encapsulation is not working. PR1456762
- On MPC11E, the FIB download rates are lower than MPC10E by 30 percent. PR1456816
- Syslog/link-down action on low-light using **set interfaces optics-options** is not supported on MPC11E. PR1456894
- During multiple configuration additions and deletions for the **flex-flow-sizing** statement, the Packet Forwarding Engine status might go to "Reconfiguring" state. PR1457282
- JSU package, when installed for LCMD daemon, might not restart the daemon with a new daemon package. PR1457304
- When an MX2K-MPC11E boots, the following harmless errors are seen and do not have functional impact. **Aug 27 22:29:15 sherekhan kernel: i2c i2c-100: (11/1:0x41) i2c transaction error (0x00000002) Aug 28 02:46:06 sherekhan kernel: i2c i2c-64: (7/1:0x41) i2c transaction error (0x00000002)**. PR1457655
- After more than 2 million multicast subscribers are activated without performing GRES or bbe-smgd restart, further multicast subscribers might not be able to log in. PR1458419
- The log **fpc4 Prolonged flow-control asserted by MAC on ms-4/0, bringing interface down** indicates that there could be a bug in the internal packet queuing of MS-MPC or MS-MIC interfaces. PR1459306
- With the scale filter-based forwarding configuration, two FBF filters seem to be unable to forward the traffic to respective routing instance. It appears that the FBF programming is incorrect for these two FBF filters. PR1459340
- Following are the issues observed:
  - PIC mode configuration support for the MX2K-MPC11E line card. The **set chassis fpc <> pic <> pic-mode <>** configuration is supported for only 100-Gigabit Ethernet on MPC11. Configuring other PIC modes (10-Gigabit Ethernet and 40-Gigabit Ethernet) will result in an invalid port profile chassis alarm.
  - The **set chassis fpc <> pic <> number-of-ports** configuration will not be supported on MPC11. Configuring this command will be a no-op/will be ignored.
  - set chassis fpc <> pic <> port <> number-of-sub-ports <> config will not be supported on MPC11. Configuring this knob will be no-op/will be ignored. PR1459433
  - The following syslog messages **TALUS(number) PCIe(number) DMA RX interrupt received. Queue stuck status 0xeeeeee0** are triggered because of the queue-back pressure or FPGA drops. PR1465888

- When VRRP is configured on MC-LAG groups, traffic destined to VRRP virtual MAC address might get dropped because the virtual MAC is not correctly programmed in the Packet Forwarding Engine. PR1459692
- When an MPC11E interface is administratively disabled and then enabled or when an optical module is plugged in, during the bring up of the interface, the interface might flap more than once. PR1459942
- In a subscriber management environment, subscriber statistics reported by CLI commands and RADIUS can be broken if a unified ISSU is performed from earlier to Junos OS Release 18.4, Junos OS Release 18.4, and later releases. PR1459961
- The **show ancp subscriber detail** command displays incorrect values for line attributes. PR1460812
- Some thread CPU information might not get exported for CPU memory sensor. PR1461155
- Error messages **MIB2D_RTSLIB_READ_FAILURE: check_rtsock_rc: failed in reading mac_db(xe-5/0/1:0): 0 (Invalid argument)** are seen. These messages have no functional impact. PR1461289
- On the MX Series platform, the rpd might crash if **dynamic-tunnel-anchor-pfe** is configured. PR1461980
- Traffic drops over aggregated Ethernet interfaces configured with Virtual Router Redundancy Protocol (VRRP). PR1462310
- A ukern hog might happen on MPC2ENG when there is a high rate of interface flaps (~30/40 flaps per second). This will eventually trigger a card reload with a next-generation MPC core file being generated. core. The issue might be caused by the external environment. PR1463859
- If a netconf session is initiated over in-band connection, the CPU utilization on mgd daemon might be stuck at 100 percent after the netconf session which is executing an RPC call for some commands gets interrupted by flapping interface. There is no impact observed to control-plane or forwarding-plane, the subsequent netconf session will continue to function. PR1464439
- The following syslog error messages are harmless and expected during unified ISSU or GRES or FPC offline or online scenarios. **[Oct 3 08:48:35.836 LOG: Err] ifl ps240.1 (1712): child ifl lt-1/0/0.32767 (7709) already there [Oct 3 08:48:35.836 LOG: Err] IFRT: 'Aggregate interface ifl add req' (opcode 87) failed [Oct 3 08:48:35.836 LOG: Err] ifl 1712, child ifl 7709; agg add failed**. PR1464524
- With multiple different fixed sized traffic streams configured at 1000000 fps (40 gbps combined rate) on ae0 along with another independent aggregated Ethernet interface (ae1, 50 percent line rate 4 streams bidirectional => 118 gbps combined traffic rate) both hosted on a single Packet Forwarding Engine inst of MPC11E, causes small varying packet drops every iteration on ae1 on disabling ae0. The drops might vary from 200 to certain 1000 frames. PR1464549
- On MPC11, failover time with LACP link-protection is in the order of 2sec+ when a PIC offline is performed. PR1464652
- A BFD session might flap when it is moving to aggressive interval after coming up with slow or non-aggressive interval. This issue is mainly seen in a scaled setup. PR1465285
- OGE0 dequeue watermark hit, when a Layer 2 traffic is sent over bridge domain that has no functional impact. PR1466027

- The following syslog error messages are harmless and expected during FPC offline/restart scenarios with PS-**RLT(with/without link protection) configuration. Nov 12 15:02:00 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x0 delete failed for ifl lt-3/0/0.32767 with err=2 Nov 12 15:02:00 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x1 delete failed for ifl lt-3/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x1 delete failed for ifl lt-5/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x0 delete failed for ifl lt-5/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_lp_handle_event: LP event = 6, child lt-5/0/0 err = 22 The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline/online scenarios. Nov 12 15:08:37 cleansing fpc3 user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIflName:ps1.0 memberIflName:lt-3/0/0.32767 Nov 12 15:08:37 cleansing fpc3 user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIflName:ps1.0 memberIflName:lt-5/0/0.32767**. PR1466531
- In the PPPoE subscriber management environment, due to the PPPoE inline keepalives timeout events might get dropped by the Routing Engine, the PPPoE subscribers get stuck. This issue might not allow the PPPoE subscribers to reconnect. PR1467125
- When user checks for interface level statistics, issue is coming for IPv6 counters at the originating router IPv6 local statistics counters not updating due to that IPv6 local statistics counters not incrementing. IPv6 transit statistics are derived from total statistics and local statistics (Transit stats = Total - Local). As the local statistics are not updating, total statistics and transit statistics will be the same. And this issue specific to MPC10E and MPC11E platforms. PR1467236
- When a connection between aftman and aft-ulcd is dropped, aft-ulcd will exit. As a result, FPC reboots. PR1467246
- L2 wholesale is not forwarding all client requests with stacked VLAN. PR1467468
- Swapping MPC11E with legacy MPC9, MPC8, MPC6, or vice versa will result in training failures. It is recommended to power-off the chassis before swapping the cards. PR1467725
- Crypto library shim memory utilization performance is improved by using data shim instead of control shim. PR1467874
- When the Packet Forwarding Engine power-off feature is used on MPC11, PIC 6 and 7 status might display as either "Ready" or as "Hardware Error". This display does not actually indicate a real hardware error. PR1468644
- For VLAN tags greater than 1 and with inner-list configuration, logical interface channel table will not get configured properly due to missing of TLV parsing. As a result, the logical interface configuration of more than one tag along with inner-list is not supported and traffic will not flow through. PR1469396
- The **monitor interface traffic** command might show some delay in some external server GNFs. PR1469733
- SNMP interface-mib stops working for PPPoE clients. In this scenario, SNMP works fine for standard queries on the MX Series router, but for subscriber statistics always returns value of zero. PR1470664

- On the MPC11E line card, MTU errors count captured in the **show pfe statistics traffic** are not matching exactly to the actual count of frames dropped. PR1471554
- On the MPC11E line card, locating a specific 100g-Gigabit Ethernet/40-Gigabit Ethernet/10-Gigabit Ethernet port in the card by blinking the corresponding port LED through the **request chassis port-led start/stop** does not work. PR1471894
- Some inconsistencies are seen in HA SYNC sessions and TLB real server flap after deletion and adding of member interface of AMS. PR1472550
- Swapping MPC11E with legacy MPC9, MPC8, MPC6, or vice versa will result in training failures. It is recommended to power-off the chassis before swapping the cards. PR1472641
- On certain MX Series platforms, **show dynamic-tunnels database** command might not show traffic statistics even though traffic is passing fine. PR1473096
- On MPC11E, when traffic on WAN is oversubscribed, oversubscription drops are not accounted in Routing Engine CLI under resource drops when flow control is disabled. PR1473191
- An occasional aftd core file is generated on setups with GNF configuration after rebooting another line card such as MPC8 within the same GNF. PR1473403
- When the MPLS interface is configured in logical router configuration, an FPC core file might be generated. PR1473502
- Creating more than 500 PS interfaces on MPC11 takes long time. PR1473559
- MX2K-MPC11E line card graceful offline takes approximately 3 minutes in a scaled setup. PR1474093
- An occasional aftd core is observed when loading baseline configurations with TCL profile. PR1474160
- Ping does not work on ixia connecting WAN interfaces and AF interfaces when input inet filter is applied on the interfaces. PR1474302
- On MPC11E, if the user has an invalid port profile, and performs Routing Engine switchover, then some of the interfaces gets into invalid state. PR1470153
- **show services sessions** and **show services sessions extensive output** are not displaying member interface of the AMS where the session got landed. It is displaying only AMS interface name. PR1474313

Infrastructure

- On MX480 devices, the following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. PR1315605
- On all Junos OS platforms, if there is an older OS package existing on an upgraded device, the Routing Engine might go to amnesiac mode after a reboot. The OS package is essential when toggling back-and-forth between Junos OS Release 15.1 and a later release. PR1445151

Interfaces and Chassis

- After GRES, the 1-Gigabit Ethernet interface changes to 10-Gigabit Ethernet. PR1326316
- In L2TP scenario when MX Series router functions as L2TP tunnel switch (LTS), there is a memory leak in jpppd process running on the backup Routing Engine, which will eventually generate jpppd core files because of the out of memory condition. There is no functional impact, though, because this issue occurs on the backup Routing Engine. PR1350563
- In large-scale subscriber environment, the two Routing Engines generate core files when changing aggregated Ethernet member link configuration. PR1375638
- If an aggregated Ethernet interface has VRRP configuration in the following use cases, member logical interfaces will not be created after the member physical interface comes up, and the aggregated Ethernet interface will be in the down state. The FPC restarts **request chassis fpc restart slot <>** and chassis-control restart **restart chassis-control** reboot both Routing Engine **request system reboot both-routing-engines**. So, before performing these operations, you should remove the VRRP configuration from the aggregated Ethernet interface. PR1429045
- In a VRRP scenario with aggregated Ethernet interfaces having mixed link speeds, if the aggregated Ethernet interface is logically divided into multiple subinterfaces, after adding or removing the subinterfaces, followed by an FPC restart, the aggregated Ethernet interface might remain "down" even though the member interfaces of the aggregated Ethernet interface are up and the FPC is working fine. PR1437670
- When all the routing instances configured under a logical system are deleted, the logical interfaces associated to those routing instances are deleted from respective routing instance but are not getting added to default routing instance. This unexpected behavior is seen because of a bug in cleanup of routing instances. PR1444131
- When an aggregated Ethernet link flaps because of the LACP timeout or LACP state reinitialization and layer 2 VPN comes-up within CCM timeout, CFM sessions on Layer 2 VPN or Layer 2 circuit logical interfaces might get stuck in a failed state. PR1454187
- The voltage high alarm might not be cleared when voltage level comes back to normal for MIC on MPC5. PR1467712

Layer 2 Ethernet Services

- On MX2010 and MX2020 routers, no alarm was generated when an FPC is connected to the master Routing Engine through the backup Routing Engine or CB. PR1461387
- On MX10008 and MX10016 platforms, if the **dhcp-relay** statement is enabled under the [**forwarding-option**] hierarchy, either in default or non-default routing instance, the Dynamic Host Configuration Protocol (DHCP) relay feature might not work as expected. Due to this issue, all the DHCP discovery packets might not be relayed. PR1447323

- On MX Series platform with DHCP subscribers scenario, the subscriber might not come online after the router is rebooted. PR1458150
- On the MX Series platform with the DHCP subscriber scenario, if a subscriber logging is happening during the ISSU process, the unified ISSU might fail. PR1465964

Layer 2 Features

MPLS

- When **vpn-localization vpn-core-facing-only** is configured and the configuration is removed completely or restored with baseline configuration, the FPC might get stuck. This is because of failure in cleaning up VT interface during complete configuration removal. PR1359087
- Dynamic SPRING-TE tunnel creation to LDP (non-SR) speaking nodes is not supported even in the presence of mapping server configurations. SPRING-TE internally converts the tunnel hop IP addresses (prefix/adjacency) into corresponding labels through the auto-translate feature. This feature internally makes use of the traffic engineering database where currently the mapping server entries are not present. PR1432791
- In LDP to BGP-LU stitching scenario, when the BGP route goes down, and MPLS ping is done before that route is pulled out of the routing table, the rpd will crash. PR1436373
- MPLS OAM for p2mp LSP is not yet implemented for MPC10E and MPC11E. Thus p2mp trace route does not work in MPC10E/MPC11E. PR1440636
- Previously, changing LDP protocol preference was considered catastrophic and all LDP targeted sessions were reset. PR1459301
- The rpd might generate a core file during shutdown because of the LDP prefix table not being cleaned up. PR1471191
- MPC11 transit LSPs are not up after reboot on redirector. PR1473383

Network Management and Monitoring

- Junos OS is used to send a cold trap from the new master just after the first GRES. This is because the cold_start timestamp file is not present/updated after the reboot. So, for the first GRES, it used to send cold start trap. PR1461839

Next Gen Services MX-SPC3 Services Card

- NAPT66 pool split is not supported with AMS. As a result, commit fails with V6 pool in AMS. PR1396634
- IPv6 throughput is not at part with IPv4. PR1439917

- IPv6 throughput numbers for NAT with HTTPs traffic is not at par with IPv4. PR1449435
- NAT performance is impacted with remote syslog enabled.PR1460211
- Drop in session setup rate for sfw is seen with syslog enabled. PR1462049
- "TALUS(number) PCIe(number) DMA RX interrupt received. Queue stuck status 0xeeeeee0" are spurious messages which are triggered in system logs due to queue-back pressure or FPGA drops.PR1465888
- On MX960, MX240 and MX480 with USF mode enabled and USF based services configuration, subsequent Junos vmhost upgrades will fail with an error. Moving to the baseline configuration and then upgrading works without error. PR1472287
- Unable to setup 26M sessions (NAPT44) at 900Kpps per second. PR1470833
- **show services sessions** and **show services sessions extensive output** do not display member interface of the AMS where the session got landed. It displays AMS interface name only. PR1474313
- Adding 100,000 CPS IPv6 SFW traffic on top 12 million PPS/50 gbps IPv6 SFW traffic resulting in PPS traffic getting reduced to 10 million PPS/42 gbps due to the latency created by CPS traffic processing. PR1472314
- The Next-Gen Services MX-SPC3 services card does not come online automatically when the junos-vmhost image is installed on the Next-Generation Routing Engine (NG-RE): RE-S-X6-64G-UB. PR1482334

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. PR1054798
- Because of the transient hardware events, fabric stream might report CPQ1: **Queue underrun indication - Queue <q>** continuously. For such events, all fabric traffic is queued for the Packet Forwarding Engine reporting the error, resulting in a high amount of fabric drops. PR1265385
- An accuracy issue occurs with three-color policers of both type single rate and two rate in which the policer rate and burst-size combination of the policer accuracy vary. This issue is present starting in Junos OS Release 11.4 on all platforms that use MX Series ASIC. PR1307882
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur. Even with enhanced convergence configured, there is no guarantee that subsecond convergence will be achieved. PR1371493
- On MX Series routers with MS-MPC cards, when FPC restart or routing instance type is changed (for example, virtual-router to vrf), or RD is changed, traffic from a group virtual private network (GVPN) tunnel to MPLS over UDP tunnel might fail to get decrypted on the MS-MPC. This might cause complete service loss. PR1422242
- On MX480 traffic loss is observed if ingress and egress ports are in different FPCs. PR1429714

- On all Junos OS platforms, if the **retry-options** statement exists in the configuration file after you upgrade the device to a new version, the older versions of login-attempts and login-locks exist on an upgraded device. Under these circumstances, the device might not be accessible through **ssh/telnet/console**. Meanwhile, the sshd process crash might be seen. PR1435173
- The **show ddos-protection protocols < protocol> statistics** command does not show arrival rates for FPC on MX Series routers. PR1438066
- For the bridge domains configured under an EVPN instance, the 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, the storm control is not taking effect to the ARP packets on the ports under such bridge domains. PR1438326
- Arrival rates are not seen at system level when **global-disable fpc** is configured on QFX Series. PR1438367
- The error message **LUCHIP(0) PPE_0 Errors sync xtxn error** continuously floods after the unified ISSU on the MX104. PR1446800
- Cm errors on certain MPC line cards are classified as major which should be minor or non-fatal. If these errors are generated, it might get projected as a bad hardware condition and therefore trigger Packet Forwarding Engine disable action. PR1449427
- A traffic loss for multicast traffic stream is observed for few of the receives. PR1450235
- A dual Routing Engine Juniper node slicing GNF with no GRES configured and with **system internet-options no-tcp-reset drop-all-tcp** configured might enter dual backup Routing Engine state upon manual GNF Routing Engine mastership switchover attempt with **request chassis routing-engine master [acquire|release|switch]** CLI command from either GNF Routing Engine CLI. PR1456565
- Expected PIM joins are not learned after performing a GRES. PR1457166
- Layer 2 traffic sent from one member to another member is corrupted on MX Series Virtual Chassis. PR1467764

Routing Policy and Firewall Filters

- On all Junos OS platforms, route resolution might be inconsistent if any route is resolving over the multipath route; for example, BGP over BGP. After the metric value of any PNH change (refers to the second PNH and using it to perform the second instance of next-hop resolution). Meanwhile, if the hash-selection changes, it might result in routes resolution inconsistency. Traffic drops are observed if the packages are still forwarding to the old. Any recursive resolving multipath scenario might trigger this issue. PR1453439

Routing Protocols

- When any of the peers flap, the flapping peer that is in established state might not come up at times when there are multiple BGP peers configured under a group. Also, if there is any other peer trying to come up, that peer might also not come up. The probability of this happening is rare, and there are no known triggers for this besides having two or more peers (IPv4 or IPv6 or both) under a group. PR967788
- If Bidirectional Forwarding Detection (BFD) is configured in fast mode (the parameter minimum-interval is configured with microseconds), during the initial phase of the BFD session, because of a network issue or certain filter, the device might drop the BFD packet with the final bit set. As a result, the BFD session get stuck with slow timers (for example, 2 seconds). This issue might impact the convergence of the network protocol related to dropping more packets. PR1254063
- LDP and OSPF are in sync state because IGP interface is down with **ldp-synchronization** enabled for OSPF: **user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity**. As per the current analysis, **IGP interface down** message is seen as the reason because although LDP notified OSPF that LDP synchronization is achieved, OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up yet. PR1256434
- At scale, a GNF with ps over rlt and multiple MPCs might show BFD flap at recovery. PR1386574
- On all devices running Junos OS, when auto-export is configured in two VPN routing and forwarding (VRF) instances, the routes get exported from/to each other. In this case, if add-path is also configured in BGP (even in an unrelated peer group), the rpd process might be stuck at 100 percent CPU utilization because of the infinite loop of route flashing in VRFs. PR1402140
- During NSR initial state replication on scaled setup, there could be cases where while BGP state replication is still going on, BGP task replication might get marked as completed. This is because BGP replication is triggered and controlled by the backup Routing Engine. PR1404470
- In a scenario with IS-IS running single shortest path first (SPF) for IPv4 and IPv6, (that is, the multitopology is not enabled) when a new IS-IS link comes up, interface address (IFA) for IPv4 comes up quickly and the route is installed. The interface address for IPv6 will not come up quickly, because duplicate address

detection (DAD) is enabled by default. Therefore, after SPF calculation, the next hop list for IPv6 remains empty for about 11 seconds. IS-IS ends up deleting the route. PR1430581

- On all Junos OS platforms working as the source node (for example, node S) where per-prefix loop free alternate (PP-LFA) is configured for Open Shortest Path First (OSPF) routing protocol, if the destination prefix is learned from two originator nodes (for example, node E and node F) with different costs, and both originator nodes E and F are directly connected with the source node S. PP-LFA might not work as expected in such a scenario where the last hop needs to be protected on the penultimate hop. As a result, an improper backup next hop might be selected that could not handle the node failure case and micro-loop might be seen. PR1432615
- The **request system core-dump routing** CLI command is not supported in cRPD. For restarting the rpd, use the kill command from the Linux shell. PR1433349
- The issue is seen when a direct change of route distinguisher is done on a routing instance. As a workaround, deactivate the instance before changing route distinguisher and then reactivate. PR1433913
- With **no-ipv6-routing** enabled under IS-IS, inet6.3 table is not getting created, which affected 6PE and 6VPE scenarios. PR1435298
- Three BGP replication flaps are seen on a new master Routing Engine after GRES. A route synchronization issue is also seen between Routing Engines without GRES. PR1441925
- On all Junos OS platforms, SSH login from automation tools to the Junos OS device is not successful when using authentication method 'password' (not 'keyboard-interactive'), if the username is configured both as a local user and also on a remote RADIUS/TACACS server, using the Juniper-Local-User-Name attribute pointing to a different local username. This issue is observed in Junos releases 17.4R3, 18.2R3, 18.3R2, 18.4R2, 19.1R2 and 19.2R1 and later releases. PR1454177
- In the Multicast Source Discovery Protocol (MSDP) scenario where the router acts as both rendezvous point (RP) and first-hop router (FHR) currently there is a lack of the loop avoidance mechanism in MSDP peer-RPF to avoid this special MSDP peer topology. Therefore, it might cause a source-active message to loop continuously and eventually causes the rpd memory leak. PR1454244
- With IS-IS configured, in a very rare case, memory corruption might occur. This might cause the rpd to crash continuously. PR1455432
- The rpd crash might be observed due to modification of router-id in OSPF NSSA with area-range configured. PR1459080
- Enabling IS-IS multitopology globally and not disabling it on the unsupported interfaces might result in route deletion or addition. There might also be a traffic drop during unrelated configuration change/commit. PR1463650
- When Bidirectional Forwarding Detection (BFD) configuration is removed, a BFD packet with session state set to "AdminDown" and diagnostic code set to some appropriate value should be sent to the peer end. But RFC does not mandate what diagnostic code should be sent and what action should be taken if different diagnostic code is received. Currently, if a BFD packet with session state set to "AdminDown" is received by Juniper Networks device, Juniper device will check both the session state and the diagnostic

code in the packet. If the session state is "AdminDown" and the diagnostic code is 7 (which means "diag AdminDown"), the BFD session will be set to "Down" and the BFD client (that is, the service which is protected by BFD) will be notified with "AdminDown" and the BFD client session will not flap. But if the BFD packet with session state set to "AdminDown" along with diagnostic code other than 7 is received, the BFD client will be notified with "Down" and the BFD client session will flap. Juniper device sets the diagnostic code to 7 for "AdminDown" packet, so no issue occurs between Juniper device and Juniper device. But if Juniper device is interworking with other vendor device (such as third party device) which does not set the diagnostic code to 7, the BFD client session might flap on Juniper device side when the BFD configuration is removed from the peer end of the BFD session. PR1470603

Services Applications

- MS-PIC core files might be generated if MS-DPC is also installed on the device when making a configuration change to an EXP type rewrite-rule applied to a SONET interface in an MX Series FPC Type 2 or MX Series FPC Type 3. PR1137941

Subscriber Access Management

- In a Gx-Plus for provisioning subscribers scenario, when the Policy and Charging Rules Function (PCRF) server is unreachable or the diameter protocol is down, the subscriber login might fail to successfully establish a session or the subscribers might fail to bind a service policy by Gx-Plus after the PCRF server connectivity is restored. PR1449064

User Interface and Configuration

- Changing nested apply groups does not take effect. PR1427962
- In MX Series platforms, the J-Web page might not get redirected to the login page after the session is expired with idle timeout. PR1459888

SEE ALSO

What's New | **84**

What's Changed | **121**

Known Limitations | **128**

Resolved Issues | **152**

Documentation Updates | **180**

Migration, Upgrade, and Downgrade Instructions | **181**

Resolved Issues

IN THIS SECTION

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

Resolved Issues: 19.3R2

General Routing

- On MX Series Virtual Chassis, suboptimal aggregated Ethernet load balancing occurs when an aggregated Ethernet bundle is part of an ECMP path. PR1255542
- Interface with Tri Rate Copper SFP (P/N:740-01311) in MIC 3D 20x 1GE(LAN)-E,SFP will stop forwarding traffic after ISSU upgrade. PR1379398
- Resource monitoring is not configured, but rsmon is blocking subscribers from logging in because the number of CoS scheduler queues is at 100 percent. PR1390424
- FPC crash and slow convergence upon HMC fatal error condition is seen when inline J-Flow is used. PR1407506
- Log severity level change is seen on MX150. PR1411846
- Redirect IP is not supported for BGP FlowSpec filters. If such an action is programmed for BGP FlowSpec rules, then it will not be reflected in filter term action. PR1413371
- DHCP/DHCPv6 subscribers might fail to establish sessions on PowerPC based MX Series platforms. PR1414333
- cRPD does not restrict the number of simultaneous JET API sessions. PR1415802
- The JSU package installation might fail. PR1417345
- Multiple ANs are created during PSK configuring or unconfiguring and lasting longer duration. PR1418448
- ROUTING_LOOP_DETECTED subcode is not generated under PATHERR_RECV code when strict path loop is created for LSP event telemetry notifications. PR1420763

- jnxFruState shows value as 10 for Routing Engine instead of 6 in response to .1.3.6.1.4.1.2636.3.1.15.1.8.9.1.0.0. PR1420906
- After control plane event a few IPsec tunnels failed to send traffic through the tunnel. PR1421843
- The XML output might not be hierarchically structured if issuing **show security group-vpn member ipsec statistics** command. PR1422496
- The **show system subscriber-management summary** command should include a failure reason for standby disconnect when primary and backup Routing Engine memories mismtach. PR1422976
- Ports might get incorrectly channelized if they are 10-Gigabit Ethernet already and they are channelized to 10G again. PR1423496
- MPC reboot or Routing Engine mastership switchover might occur on MX204 and MX10003. PR1426120
- The ppp session does not work properly on MX Series platform. PR1428212
- The **global-mac-limit** and **global-mac-ip-limit** configuration statements might allow more entries than the configured values. PR1428572
- “Protect core" configured router might send IPfix sampling packets with incorrect next-hop information. PR1430244
- Modify certain MIB objects to support 64 bits instead of existing 32 bits (to support higher memory and scale supported in SPC3). PR1430979
- MicroBFD 3x100ms flap upon inserting a QSFP to other port. PR1435221
- The local route in the secondary routing table gets stuck in the KRT. PR1436080
- A “FRU insertion" SNMP trap is sent whenever the Routing Engine is rebooted or offline/online. PR1436212
- Subscriber interim statistics might be reset to 0 and idle-timeout does not work in MX Series Virtual Chassis setup. PR1436419
- The rpd might generate a core file during router boot-up due to a file pointer issue because there are two code paths that can close the file. PR1438597
- Incorrect values are seen in JUNIPER-TIMING-NOTFNS-MIB table. PR1439025
- The **show jdaf service cmd statistics / clients** command is not available on EVO. PR1439118
- Frequent interface flapping might cause CPU hang of particular 100G port on MPC7, MPC8, MPC9, MPC10E, and MPC11E or similar MPCs. PR1440526
- FPC might get stuck in 100 percent CPU utilization due to continuous work of bulk manager thread. PR1440676
- The Layer2 dynamic VLAN might be missed when an interface is added or removed for an aggregated Ethernet interface. PR1440872
- Added new OID that calculates the buffer utilization where inactive memory is not considered as free memory. PR1441680

- In enhanced-ip or enhanced-Ethernet mode with destination-class-usage (DCU) accounting enabled, MS-DPC might drop all traffic that should egress through aggregated Ethernet interface. PR1442527
- The bbe-smgd might crash on MX Series platforms. PR1443109
- GRE packets that are larger than MTU get dropped on MX204 platforms when sampling is enabled on the egress interface. PR1444186
- Access route might get stuck in bbe-smgd and the rpd is not cleared. PR1445155
- MX Series router is rewriting the HTTPS request with destination-port 80 when using converged CPCD. PR1446085
- The static route for NAT might never come up if you switchover the service interface that has NAT and GR configuration. PR1446267
- The MX Series MPC line card might crash and restart during unified ISSU with large-scale logical interfaces. PR1446993
- The J-Flow version 5 stops working after changing input rate value. PR1446996
- Sonet option is enabled for the xe- interface. PR1447487
- Traffic silently drops when using ps- interface over RLT in a Layer 2 circuit with **no-control-word** enabled. PR1447917
- The rpd process might crash if BGP is activated or deactivated multiple times. PR1448325
- **vehostd application failed** minor alarm is observed. PR1448413
- Interface attributes might cause high CPU usage of dcd. PR1448858
- GMIC2-MACSEC-IFL: With master password configuration, DOT1Xd core files are seen on loading the configuration backup. PR1448965
- Changing the hostname will trigger LSP on-change notification, not an adjacency on-change notification. Also, currently IS-IS is sending hostname instead of system-id in OC paths. PR1449837
- Need to provide more meaningful error message, when performing commit on JDM without exchanging the SSH keys PR1449871
- **No localhost ifl for rtt 65535** can be seen on MX Series routers running Junos OS enhanced subscriber management feature. PR1450057
- Interfaces might flap forever after deleting the interface disable configuration. PR1450263
- MoFRR: Issue with MLD + IGMP scale. PR1450803
- VLAN configuration change with l2ald restart might cause kernel synchronization issues and impact forwarding. PR1450832
- On MPC10E, the dcd is unable to clean stale mt- logical interfaces while reloading Rosen configuration on the DUT. PR1450953

- SNMP query for IPsec decrypted/encrypted packets does not fetch correct values and **KMD_SNMP_FATAL_ERROR** error is observed. PR1451324
- RMPC core file is found after configuration changes are done on the network for PTP/clock synchronization. PR1451950
- Firmware upgrade for PSU (JNP10K-AC2 and JNP10K-DC2) on MX10000 systems with Routing Engine redundancy configuration enabled might fail due to lcmd being disabled by the firmware upgrade command. PR1452324
- **PLL_CMERROR_MPC_LMK04906_WAN_LD** and **PLL_CMERROR_MPC_LMK04906_WAN_LOS** errors are seen. PR1452604
- The **drop-flow** statement under **set services service-set <sset-name> service-set-options tcp-session tcp-non-syn drop-flow** is not supported. PR1452902
- On MPC10E there is an inconsistancy between AFT and non-AFT line cards in displaying ldp p2mp traffic-statistics on bud node. PR1453130
- The values displayed in the output of **show snmp mib walk jnxTimingNotfnsMIB.3** are not correct. PR1453436
- PTP out of synchronization when HWDB is not accessible during initialization. PR1453531
- Alarm is not sent to syslog on MX10003 platform. PR1453533
- Delay in freeing processed defragments buffers leads to prolonged flow control and could lead to crash. PR1453811
- The ANCP interface-set QoS adjusts might not be processed. PR1453826
- The FPC might crash when the severity of error is modified. PR1453871
- RADIUS interim accounting statistics are not populated on the MX204. PR1454541
- The access request for L2BSA port up might not be retransmitted if the RADIUS server used is unreachable. PR1454975
- CRAFTD syslogs fatal errors along with junk characters are seen upon its startup and exits after four startup attempts. PR1454985
- Interface damping is not supported on MPC11E. PR1455152
- JET/JSD RPC tag handling bug. PR1455426
- SmiHelperd process is not initialized in Junos OS PPC releases. PR1455667
- Device chooses incorrect source address for locally originated IPv6 packets in routing instance when destination address is reachable through static route with **next-table** statement. PR1455893
- BgpRouteInitialize API exits with error code 2. PR1455967
- There is high temperature from **show chassis environment** output after MPC4E insert to slot 5. PR1456457

- CLI command with **invoke-on** and **display xml rpc** results in unexpected multiple RPC commands. PR1456578
- Multicast traffic loss is seen in rare cases in seamless MPLS with MVPN configuration. PR1456905
- The bbe-statsd process continuously crashes if any parameter is set to 0 in the file mx_large.xml. PR1457257
- Default value of 2^32 replay-window size results in framing errors at an average of one in 2^32 frames received. PR1457555
- The chassisd process and all FPCs might restart after Routing Engine switchover. PR1457657
- The **show subscriber extensive** command incorrectly displays Domain Name System (DNS) address provided to DHCP clients. PR1457949
- The subscriber routes are not cleared from backup Routing Engine when session is aborted. PR1458369
- Subscribers unable to log in after more than 2 million multicast subscribers are being activated. PR1458419
- Traffic might silently get dropped or MPC might crash on the MPC10E line card during firewall filter terms changing on it. PR1458499
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used. PR1458559
- GRPC updates on_change are not working when performing delete operations. PR1459038
- Prolonged flow-control is asserted and a core file is created by **dump-on-flow-control**. PR1459306
- The error messages with **create_pseudos: unable to create interface device for pip0 (File exists)** might be seen after restarting chassisd. PR1459373
- Incomplete output of **show ancp subscriber access-aggregation-circuit-id < access aggregation circuit ID>**. PR1459386
- Port profile behavioral changes on MPC11. PR1459433
- Telemetry streaming of mandatory TLV 'ttl' learned from LLDP neighbor is missing. PR1459441
- The traffic might silently be dropped during link recovery in an open Ethernet access ring with ERPS configured. PR1459446
- In MC-LAG scenario, traffic destined to VRRP virtual MAC gets dropped. PR1459692
- Traffic silently drops upon interface flap after DRD auto-recovery. PR1459698
- The PPTP ALG does not work with DNAT44. PR1460027
- If **vlan-offload** is configured on the VMX platform, **input-vlan-map** might not work. PR1460544
- A bbe-smgd core file is generated when all RADIUS servers are unreachable. PR1461340
- The rpd core files are generated during system boot up. PR1461796
- There is a BBE statistics collection and management process, (bbe-statsd) memory issue seen on the backup Routing Engine. PR1461821

- **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed** when both DIP switches and power switch are turned off. PR1462065
- Flow stuck core file and flowd watchdog core file are generated while trying to ping DNS server 8.8.8.8 on the Internet through DUT configured with NAPT44. PR1462277
- RADIUS interim accounting statistics are not populated on MX204. PR1462325
- Mac-learning is broken for vlan-id all scenario. PR1463078
- The subscribers might not pass traffic after doing some changes to the dynamic-profiles filter. PR1463420
- The PPP IPv6CP might fail if the statement **routing-services** is enabled. PR1464415
- The **show task memory detail** shows incorrect cookie information. PR1464659
- The PPPoE session gets in terminated state and the accounting stops for the session that is delayed. PR1464804
- MPC5E and MPC6E might crash because of the internal thread hogging the CPU. PR1464820
- Memory leak on l2cpd process might lead to l2cpd crash. PR1469635
- PE interfaces on MPC11E go down on remove/add back of MPC from GNF along with syslog error messages. PR1471014

EVPN

- The specific source-ports of UDP packet are dropped on EVPN/VXLAN setup. PR1441047
- The EVPN type 2 routes might not be advertised properly in logical systems. PR1443798
- The bridge mac-table age timer does not expire for rbeb interfaces. PR1453203
- Instance type is changed from VPLS to EVPN and this results in packet loss. PR1455973
- Preference-based DF election algorithm does not work on LT interface. PR1458056
- ARP request/NS might be sent back to the local segment by DF router. PR1459830
- In EVPN scenario, memory leak might be observed when **proxy-macip-advertisement** is configured. PR1461677
- Traffic received from VTEP gets dropped if the VNI value used for type-5 routes is greater than 65,535. PR1461860

Forwarding and Sampling

- On MX204 input/output counters of aggregated Ethernet bundle/member links configured on non-default logical systems are not updated. PR1446762
- ARP packets are getting dropped by Packet Forwarding Engine after **restart chassis-control** in MX Series routers. PR1450928
- The pfed might crash and not be able to come up on the TVP platforms. PR1452363

- Commit error and dfwd core file generated are observed when applying a firewall filter with action **then traffic-class** or **then dscp**. PR1452435
- MX Series routers show the following message: **L2ALD_MAC_IP_LIMIT_REACHED_IF: Limit on learned MAC+IP bindings reached for .local.1048605; current count is 1024**. PR1462642

Infrastructure

- The duplex status of management interface might not be updated in the output of **show** command. PR1427233

Interfaces and Chassis

- Mgd processes increase as the mgd processes are not closed properly. PR1439440
- Continuous VRRP state transition (VRRP master/backup flaps) will be seen when one device drops VRRP packets. PR1446390
- Dual VRRP mastership might be seen after Routing Engine switchover occurs ungracefully. PR1450652
- LACP daemon crashes continuously. PR1450978
- The severity level log might be flooded when the QSFP-100GE-DWDM2 is inserted. PR1453919
- CFM UP MEP sessions does not come up in scaled scenario over L2VPN circuits on LAG interfaces. PR1454187
- The VRRP traffic loss is longer than one second for some backup groups after performing a GRES. PR1454895
- Mismatched MTU value causes the RLT interface to flap. PR1457460
- Need two configuration statement for EOAM CFM interoperability between MX10003 and Ciena CPE. PR1465608
- vrrpv3mibs are not working on QFX Series platform to poll VRRPv6 related objects. PR1467649

J-Web

- Junos OS: Session fixation vulnerability in J-Web. PR1410401
- Some error messages might be seen when using J-Web. PR1446081

Layer 2 Ethernet Services

- The jdhcpd process might go into infinite loop and cause 100 percent CPU usage. PR1442222
- PPPoE holding DHCPv6 prefix causes DHCPv6 binding failure because of the duplicate prefix. PR1453464
- DHCP subscriber might not come online after the router is rebooted. PR1458150
- DHCP packet might not be processed correctly if DHCP option 82 is configured. PR1459925

Layer 2 Features

- LSI interface not created and remote MACs not learned with RPD_KRT_Q_RETRIES: ifl iff add: Device busy. PR1295664
- In VC scenario traffic drop might be seen when one VC member reboots and rejoins the VC. PR1453430

MPLS

- The flow label is not pushed when **chained-composite-next-hop ingress l2ckt/l2vpn** is enabled. PR1439453
- LINX:lsi intf/Layer 2 Virtual Chassis goes down on one router in VPLS domain though the MPLS path is still available in inet.3.Reason shows as MPLS label out-of-range. PR1442495
- The rpd memory leak might be seen when the inter-domain RSVP LSP is in down state. PR1445024
- Traffic drop might be seen after traceoption configuration committed in MPLS/RSVP P2MP scenario. PR1447480
- The rpd crashes and generates a core file at **ted_delete_abstract_hop (instance=0x75d33c0, hop_name=< optimized out>)** during abstract-hop testing. PR1448769
- The LDP route timer is reset when committing unrelated configuration changes. PR1451157
- All LDP adjacencies flap after changing LDP preference. PR1459301
- Previously configured credibility preference is not considered by CSPF despite that the configuration is deleted or changed to prefer another protocol in the traffic engineering database. PR1460283
- High CPU usage and rpd core file might be observed if **ldp track-igp-metric** is configured and IGP metric is changed. PR1460292
- P2MP traceroute fails with aggregated Ethernet bundle over AFT. PR1470815

Network Management and Monitoring

- The **Wrong Type** error might be seen for the hrProcessorFrwID object. PR1446675

Platform and Infrastructure

- LACP DDoS policer is incorrectly triggered by other protocols traffic on all EX92XX/T4000 and MX Series platforms. PR1409626
- Packet drops, replication failure or ksyncd crashes might be seen on the logical system of a Junos OS device after Routing Engine switchover. PR1427842
- Some hosts behind unnumbered interface are unreachable after the router/FPC restarts. PR1449615
- FPC might reboot with vmcore because of the memory leak. PR1449664
- The DF flag BGP packets are dropped over MPLS LSP path. PR1449929
- REST API process will get nonresponsive when a number of requests come at a high rate. PR1449987

- In EVPN VXLAN scenario, sometimes host-generated packets are getting dropped as hitting "reject route" in the Packet Forwarding Engine. PR1451559
- The Routing Engine originated IPv6 packets might be dropped when an interface-group rule is configured under IPv6 filter. PR1453649
- Port mirroring is not happening with VPLS. PR1458856
- MX Series routers traceroute initiated from PE device does not show the tunnel endpoint hop in the output. PR1461441
- Make CLI configuration flag "version-03" optional. PR1462186

Routing Policy and Firewall Filters

- The rib-group might not process the exported route correctly. PR1450123
- Routes resolution might be inconsistent if any route is resolved over the multipath route. PR1453439

Routing Protocols

- The rpd crashes in Junos OS Release 16.1 or later during BGP convergence. PR1351639
- TI-LFA cannot find backup path when IS-IS overload bit is set on computing node. PR1412923
- Unsupported configuration (EPE with dynamic-next-hop GRE tunnels) continuously crashes the rpd process and generates core files. PR1431536
- Removing MVPN configuration from a VRF instance might crash the rpd process and generate a core file because it is considered to be a catastrophic change for routing daemon. PR1434347
- The rpd might crash with SR-TE configuration change. PR1442952
- The rpd crash might be seen after configuring OSPF nssa area-range and summaries. PR1444728
- The rpd CPU utilization gets 100 percent due to incorrect path-selection. PR1446861
- The multicast traffic might be dropped in PIM with BGP PIC setup. PR1447187
- L3VPN PE-CE link protection exhibits unexpected behaviour on MX2000 platforms. PR1447601
- Junos OS BFD sessions with authentication flap after a certain time. PR1448649
- The BGP routes might fail to be installed in routing instance if the "from next-hop" policy match condition is used in the VRF import policy. PR1449458
- SPRING-LDP interoperability issues occur with colocated SRMS+SR-client+LDP-stitching. PR1452956
- The rpd scheduler slip for BGP GR might be up to 120 seconds after the peer goes down. PR1454198
- MPC10E: MoFRR with MLDP inband signaling is not working. PR1454199
- The rpd memory might leak in a certain MSDP scenario. PR1454244
- The rpd process crashes permanently and generates a core file with BGP configuration option **optimal-route-reflection** set. PR1454803

- The rpd might crash when multipath is in use. PR1454951
- The rpd might crash continuously due to memory corruption in IS-IS setup. PR1455432
- Prefix SID conflict might be observed in IS-IS. PR1455994
- The rpd process generates a core file at **rt_nhn_tree_stop,rt_table_tree_free_family, bgp_sync_free_tsp** after deactivating protocols. PR1457358
- The rpd might crash when OSPF router-id gets changed for NSSA with area-range configured. PR1459080
- The rpd memory leak might be observed on backup Routing Engine due to BGP flap. PR1459384
- The rpd scheduler slips might be seen on RPKI route validation enabled BGP peering router in a scaled setup. PR1461602
- The rpd might generate a core file when BMP configured and scaled BGP peer is flapping. PR1462441
- The rpd might crash if both BGP add-path and BGP multipath are enabled. PR1463673

Services Applications

- The l2tpd core file is generated after restarting this daemon. PR1461335

Subscriber Access Management

- Subscriber filtering for general authentication services traceoptions might report debug messages for other users. PR1431614
- Subscribers login fails when PCRF server is unreachable. PR1449064
- DHCPv6 subscribers might be stuck in a state after the authd process crash. PR1460578
- Problem with linked-pool-aggregation after attempting to delete a pool in middle of the chain. PR1465253

VPNs

- P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuraton. PR1432434
- The rpd process might crash because of the memory leak in MVPN RPF Src PE block. PR1460625

Resolved Issues: 19.3R1

Class of Service (CoS)

- FPC generates core files with the reason **Data TLB Miss** after committing **rate-limit-burst** globally. PR1425667
- The host-inbound packets might be dropped if configuring host-outbound FC. PR1428144
- The dfwd crashes for the **forwarding-class** configuration in policers. PR1436894

EVPN

- The RA packets may be sent out without using the configured virtual gateway address. PR1384574
- EVPN-VXLAN: VTEP tunnel doesn't get deleted when EVPN peer goes down. PR1390965
- The process rpd crash may be observed with EVPN type-3 route churn. PR1394803
- Replace multihome advertisement proxy bit from L2_info community to ARP/ND extended community. PR1408055
- Traffic drop might be seen when the core facing link comes up in an EVPN-VXLAN scenario. PR1408840
- The next-hop is not cleaned up properly when one of the multihomed CE-PE links goes down. PR1412051
- rpd crash on backup Routing Engine after **nonstop-routing** is enabled with EVPN. PR1425687
- The device may proxy the ARP probe packets in an EVPN environment. PR1427109
- The CE interface IP address is missed in **mac-ip-table** of the EVPN database. PR1428581
- Incorrect MAC count with **show evpn/bridge statistics**. PR1432293
- Stale MAC addresses are present in the bridge MAC table in EVPN-MPLS scenario. PR1432702
- Asynchronous between ARP table and Ethernet switching table happens if the EVPN ESI link flaps multiple times. PR1435306
- IRB logical interface is not up when local Layer 2 member is down and IM next hop is present. PR1436207
- Configuring ESI on a single-homed 25-Gigabit Ethernet port might not work. PR1438227
- The specific source ports of UDP packet are dropped in an EVPN/VXLAN setup. PR1441047
- Restarting l2-learning might cause some remote MAC addresses to move into forwarding 'dead' state. PR1441565
- Traffic dropped at MX/EVPN-L3GW when VRRP switchover is initiated at the host side; irb_arp_ndp next hop is programmed as discard during the problem state. PR1442319
- Core isolation feature does not work after you configure and then delete the no-core-isolation statement on an MX Series router. PR1442973
- Local host (EVPN routes and MAC/IP) is missing from the EVPN database or **mac-ip-table** when **vlan-id** is removed from evpn and re-added. PR1443933
- Instance type is changed from VPLS to EVPN and this resulted in packet loss. PR1455973

Forwarding and Sampling

- You might be unable to apply the firewall filter configuration change after ISSU to Junos OS Release 16.1R1 or later. PR1419438
- EVPN enhancement for MAC flush mechanism in Junos OS. PR1421018
- Junos OS Release 19.1: Firewall filter and policers are not working correctly. PR1424183

- **rt-delay-threshold** can be set at less than 1 second, but **rt-marker-interval** is limited to 1 second. PR1425544
- The device is in Amnesiac mode after unified ISSU , with the generation of the the **mgd: error: configuration check-out failed** message. PR1432664
- Enable interface with input/output vlan-maps to be added to a **routing-instance** configured with a VLAN ID or VLAN tags (instance type virtual-switch/vpls). PR1433542
- High CPU utilization of l2ald is seen after replacing an EVPN configuration. PR1446568

General Routing

- In a BGP/MPLS scenario, if the next-hop type of label route is indirect, disabling and enabling the **family mpls** configuration of the next-hop interface might cause the route to go into a dead state. PR1242589
- **mspmand[190]: msvcs_session_send: Plugin id 3 not present in the svc chain for session**. PR1258970
- An enhancement for better accuracy on the drop statistic of the **show class-of-service fabric statistics** command. PR1338647
- BGP IPv4 PIC: Packet Forwarding Engine selector is stuck in rerouted state on unilist next hop after primary aggregated Ethernet link is deactivated and then activated. PR1354786
- The rpd scheduler slip might be seen when frequently deleting, modifying, or adding groups that are applied at the top level. PR1361304
- Interface with Tri-rate Copper SFP (P/N:740-01311) in "MIC 3D 20x 1GE(LAN)-E,SFP" will stop forwarding traffic after unified ISSU. PR1379398
- The unicast traffic from an IRB interface toward an LSI might be dropped due to Packet Forwarding Engine mismatch at egress processing. PR1381580
- Incorrect log message for chip errors (extra dash "-"). PR1385066
- Migrate from syslog API to Errmsg API - VM host messages on Junos OS. PR1387099
- The BNG might not respond with PADO and create any demultiplexing interface when a PPPoE PADI packet is received. PR1390989
- Commit error might be observed after adding additional sites to an existing group and **routing-instance** configuration. PR1391668
- Layer 3 gateway did not update ARP entries if IP or MAC quickly move from one router to another router in an EVPN-VXLAN environment. PR1395685
- Error messages such as **VMHost RE 0 Secure BIOS Version Mismatch** and **VMHost RE 1 Secure Boot Disabled** alarms are seen. PR1397030
- The service PIC might crash while changing CGNAT mode. PR1397294
- The PPPoE subscribers are unable to reconnect after FPC reboot. PR1397628
- Confirmation message is missing when issuing **request vmhost reboot re***. PR1397912

- The **show system firmware** CLI command gets hidden on MX Series platforms. PR1398022
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. PR1399683
- The traffic might be always taking the backup path although the primary path is available in a BGP PIC scenario. PR1401322
- The rpd generates core files: **cmgr_if_route_exists_condition_init, ctx_handle_node, task_reconfigure_complete**. PR1401396
- The na-grpcd log file is not rotated and keeps growing until the Routing Engine is out of disk space. PR1401817
- GNMI : sync_response received earlier for interface sensor subscribed in on-change mode. PR1403672
- Incorrect display of assigned prefixes to a subscriber in the output of **show interface < dynamic demux interface>**. PR1404369
- The rpd might crash due to a race condition with the combination of community actions done at both BGP import policy and a forwarding-table policy. PR1406357
- Change the default parameters for resource-monitor rtt-parameters. PR1407021
- The rpd might crash when a commit check is executed on LDP traceoptions filtering. PR1407367
- **openconfig-network-instance:network-instances** support for IS-IS must be hidden, unless supported. PR1408151
- MX-MPC2-3D-EQ and MPC-3D-16XGE-SFPP will now show "Exhaust A" temperature, rather than the intake temperature. PR1409406
- The nonexistent subscribers might appear in the output of the **show system resource-monitor subscribers-limit chassis extensive** command. PR1409767
- The configuration database might not be unlocked automatically if the related user session is disconnected when the commit operation in progress. PR1410322
- Slow SNMP on entityMIB during subscribers load test. PR1411062
- A steady increase of the Packet Forwarding Engine heap memory utilization may happen when PPPoE subscribers are flapping. PR1411389
- Parity error might cause FPC alarm. PR1411610
- **JTASK_SCHED_SLIP** error might be observed on VRR platform during NTP synchronization. PR1411679
- FPC CPU may not be displayed correctly. PR1412314
- Junos PCC may reject PCUpdate/PCCreate message if there is metric type other than type 2. PR1412659
- J-Flow: To reduce maximum flow table size when using **flex-flow-sizing**. PR1413513
- The services load balance might not be effective for AMS if the hash key under the **forwarding-options** hierarchy is configured. PR1414109
- NPC might not apply the configured resource-monitor thresholds after NPC restart. PR1414650

- ICMP MTU exceeded error generated from the Packet Forwarding Engine does not reach the expected source. PR1415130
- Port speed change and scaled aggregated Ethernet configuration can lead to MQSS errors and subsequent card crash. PR1415183
- The IRB interface might flap after committing configuration change on any interface. PR1415284
- Some IPsec tunnels might fail to pass traffic after GRES on an MX Series platform. PR1417170
- The IPv6 neighbor might become unreachable after the primary link goes down in a VPLS multihoming scenario. PR1417209
- An IPv4 packet with a zero checksum may not be translated to IPv6 packet properly in a NAT64 scenario. PR1417215
- Some subscribers might be offline when doing GRES or daemon restart. PR1417574
- Observed zero tunnel statistics on the **soft-gre** tunnel. PR1417666
- CGNAT with MS-MPC card doesn't account for AP-P out of port errors or generate a syslog message when this condition is met. PR1418128
- JDI-RCT:M/Mx: Observed rpd core files when the rpd is restarted by the user or when the logical system is deactivated. PR1418192
- MX-GX+ services are not synchronized with the backup Routing Engine with GRES/NSR enabled. PR1418594
- A PPP session under negotiation might be terminated if another PPPoE client with the same session ID. PR1419500
- **ROUTING_LOOP_DETECTED** subcode is not generated under PATHERR_RECV code when a strict path loop is created for LSP event telemetry notifications. PR1420763
- On MX Series routers, PTP is phase aligned but TE/cTE is not correct. PR1420809
- The FPC CPU might be overused if channelized interfaces are configured. PR1420983
- An interface might go to downstate on QFX10000 platform. PR1421075
- MX Series LNS might fail to forward the traffic on the subscriber access route. PR1421314
- MX Series Virtual Chassis: VCP port reports MTU value 9152 in the ICMP MTU exceeded message while the VCP port MTU is set to 9148. PR1421629
- The ps access interface is not marked as ccc down on standby/non-designated PE device. PR1421648
- **RPT_REG_SERVICES:RPM** syslogs are not getting generated after deactivating aggregate interface. PR1421934
- The changed value of **remote-gateway** does not take effect when the router acts as an initiator of an IPsec-VPN tunnel. PR1421977
- RSI bloat due to VM host-based log collection. PR1422354

- Packet Forwarding Engine wedge may be observed after running the **show forwarding-options load-balance ...** command. PR1422464
- The allocation of MAC address may fall out of the MAC address pool on an MX204 platform. PR1422679
- SFP-T/SX/LX is not working with QSA adapter in MX10003. PR1422808
- Incorrect PIC mode on MX1RU when PIC mode is changed to default mode. PR1423215
- While committing a large configuration the **error: mustd trace init failed** error is seen. PR1423229
- The **set forwarding-options enhanced-hash-key symmetric** command is not effective on MX10003. PR1423288
- Configuration commit might fail when the file system gets into full state. PR1423500
- vMX RIOT Process Panic with lu_reorder_send_packet_postproc messages. PR1423575
- IP packet drop might be seen in a Layer 2 circuit scenario. PR1423628
- Traffic is dropped after FPC reboot with aggregated Ethernet member links deactivated by the remote device. PR1423707
- MPC10: crash seen @ **Ktree alloc (jnh_dfw_instance_add (filter_index=< optimized out>) at ../../../../../src/pfe/common/applications/dfw/dfw_iff.c:1030 with inline + scale prefix filter**PR1423709
- On MX204 optics "SFP-1GE-FE-E-T", I2C read errors are seen when an SFP-T is inserted into a disabled state port. PR1423858
- PTP asymmetry change needs ptp bouncing. PR1423860
- The bbe-smgd process might crash after executing the **show system subscriber-management route prefix <>** command. PR1424054
- The port configured for 1-Gbps speed flaps after Routing Engine switchover. PR1424120
- The interface configured with 1G speed on JNP10K-LC2101 cannot come up. PR1424125
- gNMI errors message update. PR1424128
- The system does not reboot or halt as configuration when encountering the disk error. PR1424187
- On vMX, continuous disk error logs on vCP Console (requesting switchover due to disk failure on ada1). PR1424771
- The rpd keeps crashing after changing configuration. PR1424819
- The jdhcpd might consume 100% CPU and crash if **dhcp-security** is configured. PR1425206
- Interface with FEC disabled might flap after Routing Engine mastership switchover. PR1425211
- The rpd will crash continuously if MD5 authentication on any protocols is used along with the master password. PR1425231
- Soft-gre tunnel route lost after reboot/GRES or upgrade in WAG scenario. PR1425237

- The mspmand process might crash and restart with a mspmand core file created after doing a commit change to deactivate and activate the service set.PR1425405
- Following log message appears continuously on MX204 router: **fru_is_present: out of range slot 0 for**. PR1425411
- 100% CPU on route monitor of static routes after the client is disconnected from prpd server. PR1425559
- All interfaces creation failed after NSSU. PR1425716
- IFL targeting: 18,000 phantom distributed interfaces are displayed for aggregated Ethernet interface with the targeted distribution enabled on it, when there are no active subscribers. PR1426157
- Interfaces may be down after device reboots. PR1426349
- PEMs lose DC output power load sharing after the PEM is switched off and on on MX Series platforms. PR1426350
- The host-bound traffic might be dropped after performing change configuration related to **prefix-list**. PR1426539
- Some CFM and BFD sessions might flap while collecting MPLS statistics. PR1426727
- show lldp neighbors interface does not display all interface information. PR1426793
- The decoding of telemetry data at the collector may not be proper if configuring the sensors. PR1426871
- Traffic loss might be seen when multiple IPsec tunnels are established with the remote peer. PR1426975
- Traffic may not flow through MACsec interface even after an unsupported cipher-suite is removed. PR1427294
- ENTITY MIB has incorrect containedIn values for some fixed MPCs with built-in PICs PR1427305
- Rebooting or halting a Virtual Chassis member might cause the RTG link to be down for 30 seconds. PR1427500
- When broadband edge PPPoE and DHCP subscribers coming up over Junos fusion satellite ports are active, **commit full** and **commit synchornization full** commands fail. PR1427647
- When installing YANG package without the **proxy-xml** statement, the CLI environment did not work well. PR1427726
- The subscriber IP route may get stuck in **bbe-smgd** if the subscriber IP address is the same as the local IP address. PR1428428
- Incorrect display of MAC/MAC+IP and count values, after setting **global-mac-limit** and **global-mac-ip-limit**. PR1428572
- Incorrect normalization on **routing-instance** where an interface includes a VLAN ID range. PR1428623
- The PTSP subscribers are stuck in the configured state after being rejected by the RADIUS server. PR1428688
- Incorrect IGMP statistics for dynamic PPP interfaces. PR1428822

- Fabric drops might be seen on MX10003 platform when two FPCs come online together. PR1428854
- Incorrect IGMP interface counter for dynamic PPP interfaces. PR1429018
- The emitted XML is INVALID is thrown for show virtual-network-functions. PR1429090
- L2TP subscriber and MPLS pseudowire subscriber volume accounting statistics value remain unchanged post ISSU. PR1429692
- [PRPD][RPSD]:rpsd daemon is not getting killed when simultaneous to toggling rpd 'force-64-bit', rpsd core file is seen 10 minutes later. PR1429770
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs. PR1429821
- The aggregated Ethernet interface does not come up after rebooting the FPC/device although the physical member link is up. PR1429917
- Cmerror Op Set log message missing for bringup jspec command-based error simulation in Junos Evolved. PR1430300
- Configuration is prevented from being applied on MX Series routers in a subscriber scenario. PR1430360
- The firewall filters might not be attached on the interfaces after doing some changes. PR1430385
- Performance degradation is observed for about 20 seconds after the fabric board on MX10008 or 100016 is taken offline. PR1430739
- Disabling DAC QSFP port may not work on MX204, MX10003, or EX9251. PR1430921
- Traceoptions file is exceeding configured file size limit and the file keeps on growing. PR1431033
- Inline LSQ might not work when it is configured on the same FPC where MIC-3D-16CHE1-T1 is slotted. PR1431069
- Error might be observed when using a script to load the configuration. PR1431198
- Destination unreachable counter was incrementing without receiving traffic. PR1431384
- During the stress tests, the bbe-smgd process might crash on the backup Routing Engine when performing GRES. PR1431455
- The bbe-smgd might crash if subscribers are trying to log in or log out and a configuration commit activity happens at the same time. PR1431459
- Subscribers coming from new physical interfaces might not log in in due to the 512 entries limit in the **subscriber-limit** table. PR1431566
- **SIB Link Error** is detected on a specific Packet Forwarding Engine might cause complete service impact. PR1431592
- Allow installation of three identical framed routes in the same routing instance. PR1431891
- On an MX10003, the PEM not present alarm is raised when the minimum required PEMs exist in the system. PR1431926

- MPC10E-15C(2xPFE)/OldMidPlane: The MPC10E-15C is in offline state forever due to unreachable destinations after the Packet Forwarding Engine PFE2 is powered off. PR1432019
- Dual stack subscriber accounting statistics are not baselined when one stack logs out. PR1432163
- Traffic might be sent on the standby link of an aggregated Ethernet bundle and get lost with LACP fast-failover enabled. PR1432449
- Change to in-use parameterized filter prefix-list could result in bbe-smgd core files on the backup Routing Engine. PR1432655
- Output traffic statistics may be incorrect with Routing Engine-generated traffic. PR1432724
- After deleting the CLI configuration **chassis license bandwidth**, the bandwidth value is not defaulting to maximum bandwidth value. PR1433157
- A few entries specific to **show dynamic-tunnels database** output are not getting populated while testing the functionality after both PICs are taken offline and then one PIC is brought online. PR1433247
- Traffic will be dropped if **sa-multicast** is in the configuration. PR1433306
- The gNMI 'set' RPC with 'replace' field does not work and the mgd-api crashes. PR1433378
- RSI and RSI brief should not include **show route forwarding-table** when Tomcat enabled. PR1433440
- Junos Telemetry Interface-firewall: Collected service statistics all 0 after ISSU for MPC2. PR1433589
- Lawful intercept for subscriber traffic is not programmed in Packet Forwarding Engine if it's activated by Access-Accept. PR1433911
- MX URLF: URL case sensitivity support. PR1434004
- Incorrect PLUGGABLE ID 17 on MX10003-LC2103. PR1434183
- rpd core files during the route flash when the policy is removed. PR1434243
- Packet Forwarding Engine memory leak might be seen if MLPPP links are flapped. PR1434980
- Error message for **show system resource-monitor** and **show system resource-cleanup** is error: **command is not valid on the qfx5220-32cd**. PR1435136
- Traffic drop when session key rolls over between primary and fallback for more than 10 times. PR1435277
- DHCPv6 Advertise to client might use incorrect destination MAC address. PR1435694
- Total number of packets mirrored after adding the DTCP trigger and enabling DTCP is not in the expected range while verifying traffic on the mirror port after DTCP drop policy is enabled. PR1435736
- MPC7/8/9/MX10003 MPC/EX9200-12QS/EX9200-40XS line card might crash in a scaling setup. PR1435744
- The mc-ae interface may get stuck in waiting state after a device reboot. PR1435874
- ifHCInOctets counter on aggregated Ethernet interface shows the zero value when snmp mib walk is executed. PR1436201

- A few static PPP subscribers are stuck in init state permanently and the **Failed to create client session, err=SDB data corrupted** error is seen. PR1436350
- JDI_MMX_REGRESSIONS : Router is not reachable after downgrade from Junos OS Release 18.2-20190513.0 to 18.2R2.6. PR1436832
- MPC10: Micro-BFD sessions do not come up in centralized mode. PR1436937
- Schema XSDs are missing objects/commands from 19.1. PR1437469
- CI-PR: Ping is failing on logical interfaces with dual tag. PR1437608
- The CPU utilization on a daemon might be around 100% or the backup Routing Engine might crash in race conditions. PR1437762
- LNS router might send the router-advertisement packet with NULL source link-layer option field. PR1437847
- The chassisd might crash after enabling hash-key.PR1437855
- (SEEN ONLY ON LEGACY IMAGE) ISSU is failing from Junos OS Release 19.1R1 legacy Junos OS release images. PR1438144
- Subscriber flows might not be synchronized between aggregated Ethernet members on MX Series Virtual Chassis platforms. PR1438621
- MPC10: drop_unknown_dmac counters are not reported (my-mac check failed exceptions are not repoted). PR1438761
- Carrier-grade NAT logs are not received by the syslog server over TCP-based-syslog when data traffic is sent at 10,000 sessions/sec. PR1438928
- FPC on Virtual Chassis backup router might reboot in an MX Series Virtual Chassis scenario. PR1439170
- Firewall: Interface-specific filters do not take effect on MPC10E line cards. PR1439327
- The **vlan all interface all** combination is not working as expected under VSTP. PR1439583
- The bbe-smgd core files is seen after restarted. PR1439905
- PRPD Flexible Tunnel Profile queries do no return DMAC when set to all zeros by client. PR1439940
- CoS related errors are seen and subscribers could not get service. PR1440381
- DHCP offer packets toward IRB over LT interface getting dropped in DHCP relay environment. PR1440696
- MX MPC11 gNMI: DUT not exporting firewall sensor information. PR1440817
- The Layer 2 dynamic VLANs miss when an interface is added to or removed from an aggregated Ethernet bundle. PR1440872
- On PTX Series or QFX Series devices, aggregated Ethernet outgoing traffic might be dropped after changes are made to the aggregated Ethernet. PR1441772

- LINX:SNMP trap comes twice for FRU removal in MX10000 one trap with FRU name as FPC: JNP10K-LC2101 and second with FRU name as FPC @ 1/*/*. PR1441857
- The packets originating from the IRB interface might be dropped in a VPLS scenario. PR1442121
- The chassisd is unable to power off a faulty FPC after Routing Engine switchover, which leads to chassisd restart loop. PR1442138
- The operational status of the interface in HW and SW might be out of synchronization in EVPN setup with arp-proxy feature enabled. PR1442310
- EVENT UpDown interface logs are partially collected in syslog messages. PR1442542
- Different formats of the B4 addresses may be observed in the **SERVICES_PORT_BLOCK_ALLOC/RELEASE/ACTIVE** log messages. PR1442552
- A few Path Computation Element Protocol (PCEP) logs are marked as ERROR even though they are not. Now severity of those logs are corrected as INFO. PR1442598
- DHCPv6 client might fail to get an IP address. PR1442867
- The kmd process may crash and restart with a kmd core file created if IP of NAT mapping address for IPsec-VPN remote peer is changed. PR1444183
- Inline-keepalive might stop working for LNS subscribers if the **routing-services** statement is enabled. PR1444696
- MX:EAPoL: **Macsec sessions are down with unicast EAPOL destination address**. PR1445052
- Access route might be stuck in bbe-smgd and rpd not cleared. PR1445155
- The CPCDD process continuously generates core files and then the process stops, in ServicesMgr::ServicesManager::cpcddSmdInterface::processInputMsg. PR1445382
- ECMP-FRR may not work for BGP multipath ECMP routes. PR1445391
- Detached LACP member link gets LACP State as enabled in Packet Forwarding Engine when switchover occurs because of device reboot. PR1445428
- The 1-Gbps interface on MX204 might stay down after the device is rebooted. PR1445508
- Lawful Intercept on LAC access interface might not work as expected due to MTU check failure. PR1445637
- Junos OS Release 19.2 group level use of wildcard <*>. PR1445651
- The mspmand process might crash if URL filtering is configured and one blacklisted domain name is a substring of another blacklisted domain name in URL filter database file. PR1445751
- The jdhcpd process may crash after issuing the **show access-security router-advertisement-guard** command is issued. PR1446034
- [TopGun] rpd core file observed: **task_block_verify(task_io_hook_block, hook),jtask_jthr_endpoint_internal_sanity ,jtask_jthr_endpoint_sanity,**. PR1446320

- Accurate statistics might not include packets forwarded during the last two seconds before subscriber termination. PR1446546
- NAT service set in certain scale might fail to get programmed. PR1446931
- PCEP: PCE-initiated SR LSP in the first PCE is tearing down when PCInitiate LSP is brought up and brought down in the second PCE. PR1448665
- DCD CPU spike seen after a Junos OS upgrade from Junos OS Release 14.2 to Release 16.1. PR1448858
- PR-1444575-fix-test: FPC rebooted when PIC 0 is taken offline. PR1449067
- The DHCP relay feature might not work as expected with **helpers bootp** configured. PR1449201
- Increase in the maximum value of **delegation-cleanup-timeout**. PR1449468
- The burst size is not updated when the static traffic control profile is used by the dynamic profile. PR1451033
- [MX Series] Error: Dropped packets are seen on MQ/XM-based MPCs, although there is no traffic flowing through the system. PR1451958
- MX10003: MACsec framing errors are seen whenever the sequence number exceeds 2 power 32 with XPN (Extended Packet Numbering). PR1452851
- PTP can go out of sync due to l2ald hwdb access failure. PR1453531
- [PFE] [grpc] grpc updates on_change not working when performing delete operations. PR1459038

Infrastructure

- The traffic to the NLB server may not be forwarded if the NLB cluster works on multicast mode. PR1411549
- Increase in Junos OS image size for Junos OS Release 19.1R1. PR1423139
- The operations on console might not work if the **system ports console log-out-on-disconnect** statement is configured. PR1433224

Interfaces and Chassis

- In a VPLS multihoming scenario, the CFM packets can be forwarded over a backup, standby, non-designated, or CCC-down link, which might cause the traffic to be endlessly looped. PR1253542
- Issue reported in MX Series Virtual Chassis wes the error message **?CHASSISD_CONFIG_ACCESS_ERROR: pic_parse_ifname: Check fpc rnage failed** flooding with LACP-enabled aggregated Ethernet interfaces. PR1349277
- LFM sessions might flap during unified ISSU. PR1377761
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. PR1402606
- Unrelated aggregated Ethernet interfaces might go down if committing configuration changes. PR1409535

- MX Series Virtual Chassis unified ISSU is not supported when Redundant LT (RLT) is configured. PR1411729
- Invalid speed value on an interface might cause other interface configuration loss. PR1421857
- The syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** upon LFM-related configuration commit on aggregated Ethernet interfaces. PR1423586
- The demultiplexer interfaces will be down after the MTU of the underlying et-interface is changed. PR1424770
- The cfmd might crash on DPCE. PR1424912
- The logical interfaces in EVPN routing instances might flap after committing configurations. PR1425339
- Upgrade from releases before Junos OS Release 17.4R1 results in cfmd core files. PR1425804
- CFM message flooding. PR1427868
- The vrrpd process might crash after VRRP sessions are deleted for several times. PR1429906
- The NCP session might be brought down after IPCP Configure-Reject is sent. PR1431038
- VRRP mastership might flap when the tracked route is deleted or the tracked interface goes down. PR1432361
- jppd No termination Ack for a LACP termination request RFC 1661. PR1433489
- Mixed link-speed aggregated Ethernet bundle could not add new subinterface successfully. PR1437929
- Targeted distribution for static demultiplexing interface over aggregate ether interfaces does not take correct lacp link status into consideration when choosing primary and backup links. PR1439257
- The cfmd process might crash after a restart on Junos OS Release 17.1R1 and later. PR1443353
- MX960: Validation failing while upgrading to Junos OS Release 19.2R1.8 with the error **Maximum 64 AE IFLs can be configured for mixed-rate AE</message>**. PR1445040
- The OAM CCM messages are sent with single-tagged VLAN even when configuring with two VLANs. PR1445926
- MX Series Virtual Chassis on MX10003: Not able to connect to newly installed Routing Engine from other Routing Engines in an MX Series Virtual Chassis. PR1446418
- Initiating a Routing Engine switchover on a VRRP backup router through a CLI command could cause the VRRP state for ae-bundle interfaces to transition to the master state even though the **protocols vrrp delegate-processing ae-irb** statement is configured, then very shortly afterward to backup again. PR1447028
- l2ald failed to update conposite next hops. PR1447693

Layer 2 Features

- Q-in-Q might be malfunctioning if VLAN ID lists are configured. PR1395312
- The rpd crashes after iw0 interface is configured under a VPLS instance. PR1406472

- Broadcast traffics may be discarded in a VPLS local-switching scenario. PR1416228
- VPLS neighbors might stay in the down state after VLAN ID configuration changes. PR1428862
- After disabling and enabling the aggregate interface, the next-hop of the CE-facing aggregate interface may be in a wrong state. PR1436714

Layer 2 Ethernet Services

- LACP PDU may be looped toward peer MC-LAG nodes. PR1379022
- jdhcpd becomes aware about some of the existing configuration only after 'commit full' or jdhcpd restart. PR1419437
- Change the ND6 next hops to reject next hops after l2 interfaces get disassociated with ipv6 entries. PR1419809
- JDI-RCT:BBE:DHCP subscribers on nondefault routing instance went down after ISSU. PR1420982
- jdhcpd daemon might crash during continuous stress test. PR1421569
- MX:LACP:- Error message **fpc3 user.err aftd-trio: [bt] #1 JnhHandle::** is logged. PR1424106
- The DHCP DECLINE packets are not forwarded to DHCP server when **forward-only** is set within **dhcp-reply**. PR1429456
- DHCP request may get dropped in DHCP relay scenario. PR1435039
- The **dhcp-relay** statement might not work on MX10008 or MX10016 platforms. PR1447323
- DHCPv6 authentication via RADIUS server might fail as a result of the missing VSA option 26-207. PR1448100
- PPPoE holding DHCPv6 prefix causes DHCPv6 binding failure due to a duplicate prefix. PR1453464

MPLS

- SR-TE path does not install with "0" explicit NULL as the innermost label. PR1287354
- The rpd may restart after an MPLS LSP flap if **no-cspf** and **fast-reroute** are configured in an LSR ingress router. PR1368177
- DSCP bit marking of LSP self-ping is not compliant with RFC7746. PR1371486
- The rpd might crash in BGP-LU with egress protection while committing configuration changes. PR1412829
- RSVP LSP might get stuck in down state in OSPF multiarea topology. PR1417931
- Traffic might be dropped due to LDP label corruption after Routing Engine switchover. PR1420103
- Bad length for sub-TLV 34 (RFC 8287) in an MPLS echo request. PR1422093
- LDP might not update the LDP ingress route metric when inet.3 route flash happens before inet.0. PR1422645
- The dynamic bypass RSVP LSP tears down when it is being used to protect an LDP LSP. PR1425824

- The **ping mpls sweep** command stops working and makes the CLI irresponsive. PR1426016
- M/Mx: continuous generation of rpd core files at **l2ckt_alloc_label , l2ckt_standby_assign_label , l2ckt_intf_change_process in new backup during GRES in MX2010 box**. PR1427539
- Traffic loss might be observed after changing configuration under **protocols mpls** in an LDP tunneling scenario. PR1428081
- The LDP might withdraw a label for an FEC after the IGP route is inactive in inet.0. PR1428843
- When MBB for P2MP LSP fails, it is stuck in old path. PR1429114
- MPLS ingress LSPs for LDP link protection do not come up after MPLS is disabled and then enabled. PR1432138
- Restart Routing might result in RPD core files while GRES and NSR are enabled. PR1433857
- The P2MP LSP branch traffic might be dropped for a while when the sender PE device is doing switchover. PR1435014
- The rpd might crash after the **ping mpls ldp** command is executed. PR1436373
- The LDP route and LDP output label are not displayed in the inet.3 table and LDP database, respectively, if **ospf rib-group** is enabled. PR1442135
- Backup LSP signaling after if NP bypass is an inter-area LSP using loose-hop expansion. PR1442789
- P2MP LSP might get stuck in the down state after link flaps. PR1444111
- Silent dropping of traffic is likely if two consecutive PLRs along the LSP perform local repair simultaneously under certain misconfigured conditions. PR1445994
- The transit packets might be dropped if an LSP is added or changed on an MX Series or PTX Series device. PR1447170
- rpd crashed and generated core files at ted_delete_abstract_hop (instance=0x75d33c0, hop_name=< optimized out>) during abstract-hop testing. PR1448769
- The LDP route timer is reset when committing unrelated configuration changes. PR1451157

Network Address Translation (NAT)

- The nsd process might crash when SNMP query deterministic NAT pool information. PR1436775

Network Management and Monitoring

- The SNMP query may not get data in a scaled Layer 2 circuit environment. PR1413352
- MX10000 reports jail socket errors. PR1442176

Platform and Infrastructure

- DDoS violation for LLDP, MVRP, provider MVRP, and dot1x is incorrectly reported as LACP DDoS violation. PR1409626
- Error logs might be observed after performing unified ISSU. PR1412463

- Distributed multicast forwarding to the subscriber interface may not work. PR1416415
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10 (based system). PR1417321
- The ARP request might not be replied to although **proxy-arp** is configured. PR1422148
- The native VLAN ID of packets might fail to be removed when leaving out. PR1424174
- The policer bandwidth might be wrong for the aggregate interface after the **shared-bandwidth-policer** statement is enabled. PR1427936
- The BGP session with hold-time 6 seconds or less flaps after the backup Routing Engine is pulled out ungracefully. PR1428518
- With CNH for 6PE, MPLS EXP rewrite rule for non-VPN IPv4 over MPLS traffic might not work. PR1430878
- Pre-fragmented ICMP IPv4 packets might fail to arrive at the destination. PR1432506
- Enabling the sensor /junos/system/linecard/qmon/ causes continuous **ppe_error_interrupt** errors. PR1434198
- Traffic from the same physical interface cannot be forwarded. PR1434933
- The device might not be accessible after the upgrade. PR1435173
- BR for MAP-E does not return ICMP Type=3/Code=4 when over MTU sized packet comes with DF bit. PR1435362
- MAP-E encapsulation or de-encapsulation with specific parameters might work incorrectly. PR1435697
- The **/var/db/scripts** directory might be deleted after **request system zeroize** is executed. PR1436773
- The BGP session might flap after Routing Engine switchover is done simultaneously on both BGP peer devices in scaled BGP session setup. PR1437257
- The next-hop MAC address in the output of the **show route forwarding-table** command might be incorrect. PR1437302
- A certain combination of **allow-commands**/**deny-commands** does not work properly after Junos OS Release 18.4R1. PR1438269
- The inner IPv4 packet might get fragmented using the same size as the configured mtu-v6, which is used for the MAP-E softwire tunnel in the MAP-E configuration. PR1440286
- When host-bound packets are received in MAP-E BR router, service interface statistics counter shows incorrect number of bytes. PR1443204
- Packets drop due to missing destination MAC in the Packet Forwarding Engine. PR1445191
- Python op scripts executed as user "nobody" if started from a NETCONF session, not as the logged-in user, resulting in failing PyEZ connection to the device. PR1445917

Routing Policy and Firewall Filters

- The **route-filter-list** configuration with noncontinuous match might not work as expected after being updated. PR1419731
- Policy matching RD changes next hop of the routes that do not carry the RD. PR1433615

Routing Protocols

- When the prefix limit is reached, increasing maximum-prefixes does not take effect. PR1323765
- The rpd crashes due to assert in **bgp_io_write_user_handler_int()**. PR1351639
- Qualified next hop of static route might not be withdrawn when BFD is down. PR1367424
- Routing Engine-based micro-BFD packets do not go out with the configured source IP address when the interface is in logical-system. PR1370463
- The rpd might crash under a rare condition if GR helper mode is triggered. PR1382892
- BGP sessions might keep flapping on the backup Routing Engine if **proxy-macip-advertisement** is configured on an IRB interface for EVPN-VXLAN. PR1387720
- In rare cases rpd might crash after Routing Engine switchover when BGP multipath and Layer 3 VPN **vrf-table-label** are configured. PR1389337
- Processing a large-scale AS-path regex causes the flap of the route protocols to flap. PR1396344
- IGMP join through PPPoE sub is not propogated to the upstream PIM. PR1407202
- BFD link-failure detection of the broken path is delayed when IGP link-state update is received from the same peer through an alternative path. PR1410021
- BGP stuck in Idle (Close in progress) state after rpd is started on the peer. PR1412538
- The Layer3VPN link protection doesn't work after flapping the CE facing interface. PR1412667
- The CPU utilization of the rpd process is stuck at 100% if BGP multipath is configured. PR1414021
- Route information might be inconsistent between RIB and OSPF database when using the OSPF LFA feature. PR1416720
- A memory leak in rpd might be seen if source packet routing is enabled for IS-IS protocol. PR1419800
- BFD crash after GRES was done @ **__assert (func=0x831a40e "bfdd_link_session", file=0x831a24a "../../../../../../src/junos/usr.sbin/bfdd/bfdd_session.c**". PR1420694
- IPv6 IS-IS routes might be deleted and not be reinstalled when MTU is changed under the logical interface level for family inet6. PR1420776
- The rpd might crash if **no-propagate-ttl** is configured in a BGP multipath scenario. PR1425173
- The multicast traffic might be dropped when proxy mode is used for **igmp-snooping**. PR1425621
- The rpd might crash in a PIM scenario with **auto-rp** enabled. PR1426711

- The rpd might crash while removing multicast routes that do not have an associated (S,G) state or enabling the **accept-remote-source** statement on PIM upstream interface. PR1426921
- The rpd might crash while handling the withdrawal of an imported VRF route. PR1427147
- MVPN traffic might be lost for around 30 seconds during Routing Engine switchover. PR1427720
- The rpd generates core files due to improper handling of graceful restart stale routes. PR1427987
- The rpd might crash with OSPF overload configuration. PR1429765
- The next-hop of an IPv6 route remains empty when a new IS-IS link comes up. PR1430581
- The BGP configuration statement **multipath multiple-as** does not work in specific scenarios. PR1430899
- IPv6 aggregate routes are hidden. PR1431227
- The **show isis adjacency extensive** command output does not display state transition details. PR1432398
- In BFD and GR enabled scenario, BFD DOWN packets are not being sent immediately after BFD failure. PR1432440
- PP-LFA is not working on the penultimate hop, thereby causing micro-loops. PR1432615
- The **request system core-dump routing** command is not supported in cRPD. PR1433349
- PIM-SM join message might be delayed when MSDP is enabled. PR1433625
- With SR enabled, 6PE next-hop is not installed. PR1435298
- Clearing BGP neighbors takes much longer to delete routes. PR1435466
- The rpd might crash during the best path changes in BGP-L3VPN with multipath and **no-vrf-propagate-ttl** enabled. PR1436465
- BGP route next-hop can be incorrect in some scenarios with PIC edge configuration. PR1437108
- The backup Routing Engine might go out of synchronization if BGP sessions are cleared on the master Routing Engine. PR1439620
- Removing SSH Protocol version 1 from configuration. PR1440476
- RIP routes are discarded by the Juniper Networks device when the next-hop field in the RIPv2 response packet contains a subnet broadcast address. PR1441452
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. PR1443507
- The rpd might crash in an OSPF scenario due to invalid memory access. PR1445078
- BRP: RPC call is missing for **show bgp output-scheduler**. PR1445854
- The BGP route prefixes are not being advertised to the peer. PR1446383
- The as-external route may not work in OSPF overload scenario for a VRF instance. PR1446437
- The rpd might crash when the policy applied to the MoFRR is deleted. PR1446472
- The rpd crashes and the configuration fails when you try to commit configuration. PR1447595

- Intra-router PPMD[RE] to PPMAN[FPC] connection could be closed if the session timeout is greater than 3 seconds in either direction. PR1448670
- rpd core files at **rt_nhn_tree_stop,rt_table_tree_free_family, bgp_sync_free_tsp** after deactivating protocols. PR1457358

Services Applications

- ms- used for IPsec PIC is listed in **show services ha detail** as standby. It is a cosmetic issue. PR1383898
- **SPD_CONN_OPEN_FAILURE: spd_svc_set_summary_query: unable to open connection to si-0/0/0 (No route to host)**. PR1397259
- [technology/subscriber_services/jl2tpd] [all] RPT BBE Regressions : ERA value does not match with the configured values while verify new ERA settings are reflected in messages log. PR1410783
- In subscriber with L2TP scenario, subscribers are stuck in init state forever. PR1425919
- Some problems might be seen if the client negotiates LCP with no ppp-options to LAC. PR1426164
- The kmd process may crash when DPD timeout for some IKEv2 SAs happens. PR1434521
- Traffic might be dropped in an IPsec VPN scenario when the VPN peer is behind a NAT device. PR1435182
- The output of "show subscriber user-name" on LTS shows only one session instead of two. PR1446572

Software Installation and Upgrade

- JSU might be deactivated from FPC after the device is power cycled. PR1429392

Subscriber Access Management

- Authd telemetry: Linked pool head attribute is incorrect for single pools. PR1413293
- CoA-NACK is not sent when performing negative COA Request tests by sending an incorrect session ID. PR1418144
- Address allocation issue with linked pools when using **linked-pool-aggregation**. PR1426244
- RADIUS authentication server might always be marked as DEAD. PR1429528
- Incorrect Acct-Session-Time : Acct-Session-Time is not zero, although no start event occurred. PR1433251
- Test aaa ppp, output enhancement. PR1444438
- Subscriber login fails when the PCRF server is unreachable. PR1449064

User Interface and Configuration

- Junos fusion: **show chassis hardware satellite** command is not available in 17.3 Junos versions. PR1388252
- commit reject ae0.0 vlan-id-list + routing-instance vlan-id (but does not reject for vlan-range). PR1427278

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. PR1282875
- The deletion of (S,G) entry might be skipped after the PIM join timeout. PR1417344
- The rpd might crash if Layer 2 circuit or local switching connections flap continuously. PR1418870
- JDI-RCT: Permanent traffic loss is seen on next-generation MVPN selective tunnels after Routing Engine switchover (one-time). PR1420006
- The rpd process might crash and core files generated when an MPLS ping command is executed on Layer 2 circuit. PR1425828
- MVPN using PIM Dense mode does not prune the OIF when PIM prune is received. PR1425876
- The resumed multicast traffic for certain groups might be stopped in overlapping MVPN scenario. PR1441099
- Result of **show task replication** shows MVPN as "InProgress" when the active master Routing Engine is plugged out and NSR is enabled. PR1441292
- Memory leak might happen if PIM messages are received over an MDT (mt- interface) in a Draft-Rosen MVPN scenario. PR1442054

SEE ALSO

What's New | **84**

What's Changed | **121**

Known Limitations | **128**

Open Issues | **132**

Documentation Updates | **180**

Migration, Upgrade, and Downgrade Instructions | **181**

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for MX Series.

SEE ALSO

What's New | **84**

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 18.3R1 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 19.3

> **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-x86-32-19.3R1.9-signed.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-19.3R1.9-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-x86-32-19.3R1.x-limited.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-19.3R1.9-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 19.3R1, 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 19.3 **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-19.3R1.9-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-19.3R1.9-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 19.3 **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

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before

or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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 19.3

To downgrade from Release 19.3 to another supported release, follow the procedure for upgrading, but replace the 19.3 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

- What's New | 84
- What's Changed | 121
- Known Limitations | 128
- Open Issues | 132
- Resolved Issues | 152

Documentation Updates | 180

Junos OS Release Notes for NFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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 devices.

Release 19.3R2 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 19.3R2.

Release 19.3R1 New and Changed Features

Application Security

- **Application quality of experience (AppQoE) (NFX150 and NFX250 NextGen)**—Starting in Junos OS Release 19.3R1, Application Quality of Experience (AppQoE) enforces the configuration limit for overlay paths, metric profiles, probe parameters, and SLA rules per profile when you configure application-specific SLA rules and associate the SLA rules to an APBR profile. If you configure more parameters than the allowed limit, an error message is displayed after you commit the configuration.

  [See Understanding AppQoE Configuration Limits.]

- **Application path selection based on link preference and priority (NFX150 and NFX250 NextGen)**—Starting in Junos OS Release 19.3R1, you can configure application quality of experience (AppQoE) to select the application path based on the link priority and the link type when multiple paths that meet the SLA requirements are available.

  [See Understanding Application Path Selection Based on Link Preference and Priority.]

Interfaces

- **Dual virtual function support (NFX150)**—Starting in Junos OS Release 19.3R1, you can configure and map a maximum of two L3 interfaces to a single physical port.

SEE ALSO

What's Changed | **190**

Known Limitations | **191**

Open Issues | **192**

Resolved Issues | **194**

Documentation Updates | **196**

Migration, Upgrade, and Downgrade Instructions | **197**

What's Changed

IN THIS SECTION

See what changed in this release for NFX Series.

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

SEE ALSO

New and Changed Features | **188**

Known Limitations | **191**

Open Issues | **192**

Resolved Issues | **194**

Documentation Updates | **196**

Migration, Upgrade, and Downgrade Instructions | **197**

Known Limitations

IN THIS SECTION

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

Interfaces

- On NFX150 devices, the link does not come up if a 1-gigabit SFP transceiver is connected from heth-0-4 and heth-0-5 to a peer device. As a workaround, disable the auto-negotiation for the interface connected to the NFX150 on the remote device. PR1428020

Platform and Infrastructure

On NFX250 NextGen devices, the error message **pki-service subsystem is not responding to management requests** appears when you issue the **show security pki ca-certificate** command. PR1435639

SEE ALSO

New and Changed Features | 188

What's Changed | 190

Open Issues | 192

Resolved Issues | 194

Documentation Updates | 196

Migration, Upgrade, and Downgrade Instructions | 197

Open Issues

IN THIS SECTION

Learn about open issues in this release for NFX Series. 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 NFX150 chassis cluster, the host logs updated in the system log messages might not show the correct timestamp. As a workaround, convert the UTC timestamp to the local time zone. PR1394778

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
- When a DHCP server assigns a conflicting IP address to the NFX Series device interfaces, the NFX Series device does not send a **DHCP DECLINE** message in response. PR1398935
- If you plug an unsupported SFP-T transceiver into an NFX150 device and reboot the device, the FPC1 WAN port does not come online. PR1411851
- When the interface configuration has the encapsulation **flexible-ethernet-services** enabled on a 10-Gigabit Ethernet interface, traffic is dropped. PR1425927

Platform and Infrastructure

- On NFX250 NextGen devices, jumbo frames are not supported through OVS. PR1420630
- On NFX150 devices, random RPM probe losses are noticed if the probe packets are fragmented because the data-size more than the **inet** MTU. PR1447082

- On NFX150 devices, the following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. PR1315605
- On an NFX Series device running Junos OS Release 19.3R1, a srxpfe core may be seen when you attempt to configure, reconfigure, or delete the dual VF mappings on the device. PR1458452

Virtual Network Functions (VNFs)

- After you create or delete a VNF on NFX150 and NFX250 NextGen devices, the **request virtual-network-functions console *vnf-name*** command gives an error that the VNF domain is not found. VNFs are reachable through SSH in this state. PR1433204
- 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
- On NFX150-S1 devices, configuring the VNF with the **no-default-interfaces** option and disabling the internal management interface (eth0) with the **link disable** command might not disable the interface. Hence, the liveliness status remains alive even if the link is configured to be disabled. PR1442064
- 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. As a workaround, reboot the device, which resets the MAC address to the default value. PR1448595
- On NFX150 devices, when you need to change the vmhost mappings of a particular NIC or NICs, you must delete the existing vmhost mapping and commit the configuration. Now you can configure the new mappings for the respective NICs. You cannot change the NIC vmhost mappings in the same commit to delete and add a new mapping to the heth NICs. PR1459885

SEE ALSO

New and Changed Features | **188**

What's Changed | **190**

Known Limitations | **191**

Resolved Issues | **194**

Documentation Updates | **196**

Migration, Upgrade, and Downgrade Instructions | **197**

Resolved Issues

IN THIS SECTION

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

Resolved Issues: 19.3R2

Interfaces

- When applying firewall filters on lo0.0 on an NFX250 NextGen device, FPC0 disappears. PR1448246
- On NFX150 devices, the **show security dynamic-address** command does not work for port 3. PR1448594

Platform and Infrastructure

- On NFX Series devices, the HTTP traffic flow is created with a different routing instance when an APBR profile is configured with category and application in the same profile. PR1447757
- Informational log message **LIBCOS_COS_RETRIEVE_FROM_PVIDB: feature cos_fc_defaults num elems 4 rc 0** is displayed on the console when you commit after you configure AppQoS rule set. PR1457328
- Version compare in PHC might fail, causing the PHC to download the same image. PR1453535
- The REST service might become nonresponsive when the REST API receives a number of continuous HTTP requests. PR1449987

Virtual Network Functions (VNFs)

NFX250 devices do not allow “jdm” (case-insensitive) as a VNF name. You can use “jdm” as part of the name. For example, jdm123, abcJDM, abcJDM123 are valid VNF names, but jdm, JDM, Jdm, JDm are not valid VNF names. PR1463963

Resolved Issues: 19.3R1

Class-of-Service (CoS)

- In the NFX Series device configuration, traffic is being sent to the incorrect queue when configuring CoS with **forwarding-classes** class vs. queue. The **forwarding-classes** class is not supported and is hidden. As a workaround, use **forwarding-classes** queue when you configure CoS. PR1436408

Interfaces

- On NFX250 devices with XDSL SFP transceiver used on the fiber ports, the status of the transceiver is displayed under the **ADSL Status** field in the output of the **show interfaces *int-name*** command. But whenever a user hot-swaps an XDSL SFP with another XDSL SFP on the same port, the **ADSL Status** field is not displayed in the output of the **show interfaces** command. PR1408597
- Starting in Junos OS Release 19.2R1, when you transition NFX150 devices from a PPPoE configuration to a non-PPPoE configuration in a non-promiscuous mode, the interface hangs without any traffic flow. PR1409475
- The limit on maximum OVS interfaces is restored to the originally defined limit 25 for backward compatibility. As a workaround, reduce the number of OVS interfaces in the configuration to 20 or less. PR1439950
- On NFX150 and NFX250 NextGen devices, cross-connect stays down even if all linked interfaces are up. PR1443465
- On NFX150 devices, whenever you need to change the vmhost mappings of particular NICs, you must delete the existing VM-host mapping and commit the configuration so that the existing mapping is cleared. Then you can configure the new mappings for the respective heth NICs. Changing the NIC VM-host mappings in the same commit, which will delete and then add a new mapping to the heth NICs, is not supported. PR1450147

Platform and Infrastructure

- On an NFX250 device, the console is not accessible and JDM stops working. These issues occur because the libvirtd process stops responding. PR1341772
- On an NFX250 device, if the **idle-time out** parameter for a user login class on JDM is configured in minutes, the system considers the configured idle timeout value in seconds. The user is logged out based on the idle timeout value in *seconds*. PR1435310

Protocols

- On NFX150 devices, SNMP does not work for the following commands:
  - **show snmp mib walk jnxIpSecTunMonOutEncryptedBytes**
  - **show snmp mib walk jnxIpSecTunMonOutEncryptedPkts**
  - **show snmp mib walk jnxIpSecTunMonInDecryptedBytes**
  - **show snmp mib walk jnxIpSecTunMonInDecryptedPkts**
  - **show snmp mib walk jnxIpSecTunMonLocalGwAddr**
  - **show snmp mib walk jnxIpSecTunMonLocalGwAddrType**

PR1386894

Virtual Network Functions (VNFs)

- When you issue the **show virtual-network-functions *vnf-name*** command, the system creates a defunct process due to the presence of popen() calls and pclose() calls that do not match. This issue is fixed in Junos OS Release 15.1X53-D497 onward by ensuring that pclose() calls match the popen() calls. PR1415210
- With a VNF running, when MTU is configured, the KVM crashes and the VNF goes down. PR1417103
- On NFX150 devices, FPC0 may not be online after an upgrade and a device reboot is required. PR1430803
- When you run the **show chassis fpc** or **show chassis fpc details** command, the **Temperature** field in the command output is displayed as **Testing**. PR1433221
- On NFX150 devices with VNFs configured, when the VNF interfaces are moved from the default OVS bridge to a custom OVS bridge, duplicate VNF host entries are present in the **/etc/hosts** file on JDM. PR1434679
- When you downgrade from Junos OS Release 19.2 to Junos OS Release 18.4, the **show virtual-network-functions *vnf-name*** command does not display the VNF information. PR1437547

SEE ALSO

New and Changed Features | **188**

What's Changed | **190**

Known Limitations | **191**

Open Issues | **192**

Documentation Updates | **196**

Migration, Upgrade, and Downgrade Instructions | **197**

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for NFX Series.

SEE ALSO

New and Changed Features | **188**

What's Changed | **190**

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.

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information on EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html.

Basic Procedure for Upgrading to Release 19.3

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 19.3R2:

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 Packet Transport Routers

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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 the Junos OS main and maintenance releases for PTX Series routers.

What's New in 19.3R2

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 19.3R2.

What's New in 19.3R1

Class of Service

- **Support for explicit-null packet classification using the EXP value from MPLS explicit-null labels (PTX10002)**—The default classification for explicit-null packets is based on the payload (IPv4 or IPv6 DSCP bits). Starting with Junos OS Release 19.3R1, PTX10002 routers with third-generation FPCs (FPC3) support the CLI option, **[explicit-null-cos *inet*|*inet6*]** at the **[edit forwarding-options]** hierarchy level, that makes the packet classification based on the MPLS EXP value rather than on the payload, thus preserving the MPLS classification of the packet.

  [See explicit-null-cos.]

Interfaces and Chassis

- **FTIs with support for UDP encapsulation (PTX Series)**—Starting in Junos OS Release 19.3R1, you can configure flexible tunnel interfaces (FTIs) on the PTX Series routers/QFX Series switches, which provide support for static UDP tunnels only.

  With the UDP tunnels-over-FTI feature, you can benefit from better traffic distribution over ECMP, that is achieved by the UDP source port derived from the hash value of the inner payload. In addition to this, the other benefits of this feature include, shortened interface hop counts, smooth IGP domain separation, and reduced operational complexity.

  [See Flexible Tunnel Interfaces Overview.]

- **VLAN tag manipulation: pop, push, and swap (PTX1000, PTX10001, PTX10002, PTX10003, PTX10008, and PTX10016)**—Starting in Junos OS Release 19.3R1, you can configure your VLAN circuit cross-connect

(CCC) logical interface on a Layer 2 circuit to handle dual-tag and single-tag packets. You can also use the **l2circuit-control-passthrough** statement at the **[edit forwarding-options]** hierarchy level to enable passthrough of certain Ethertype/DMAC-matched frames over the Layer 2 circuit after successful VLAN tag manipulation on the VLAN CCC logical interface. The VLAN CCC logical interface can be on a single Ethernet interface or on an aggregated Ethernet interface.

[See Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough.]

- **QSFP28 DWDM2 transceivers for PTX5000**—Starting in Junos OS Release 19.3R1, PTX5000 routers with FPC type 3 line cards support QSFP28 DWDM2 transceivers. Use the existing show commands such as **show chassis pic** and **show chassis hardware** to view the inventory details of the transceivers.

  [See show chassis pic and show chassis hardware.]

Junos OS XML, API, and Scripting

- **IPv6 support in Python automation scripts (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.3R1, devices running Junos OS with upgraded FreeBSD support using IPv6 in:
  - Python automation scripts, including commit, event, op, and SNMP scripts
  - Juniper Extension Toolkit (JET) scripts
  - YANG action scripts

  IPv6 support enables Python scripts to establish connections and perform operations using IPv6 addresses.

  [See IPv6 Support in Python Automation Scripts.]

Junos Telemetry Interface

- **JTI support extended for Junos kernel GRES and RTSOCK (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos Telemetry Interface (JTI) extends support for streaming Junos kernel Graceful Routing Engine Switchover (GRES) and Routing Socket (RTSOCK) information using remote procedure call (gRPC) services. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel GRES and RTSOCK information:
  - /junos/chassis/gres/
  - /junos/kernel/rtsock/

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

- **gNMI support for Routing Engine statistics for JTI (MX960, MX2010, MX2020, PTX5000, PTX1000, and PTX10000)**—Junos OS Release 19.3R1 and supports the Junos telemetry interface (JTI) export of Routing Engine sensors using gRPC Management Interface (gNMI). gNMI is a protocol for configuration and retrieval of state information. Both streaming and ON_CHANGE export is supported using gNMI.

Export the following statistics using gNMI:

- Network discovery, ARP table state (resource path **/arp-information/**)
- Network discovery, NDP table state (resource paths **/nd6-information/** and **/ipv6-ra/**)

To provision the sensor to export data through gNMI, use the Subscribe RPC defined in the gnmi.proto to specify request parameters. Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **JTI support extended for Junos kernel LAG, NSR, and TCP (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos Telemetry Interface (JTI) extends support for streaming Junos kernel Link Aggregation Group (LAG), Non-Stop Routing (NSR) Junos Socket Replication (JSR), and Transport Control Protocol (TCP) information using remote procedure call (gRPC) services. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel LAG, NSR, and TCP information:

  - /junos/chassis/aggregated-devices/
  - /junos/routing-options/nonstop-routing/
  - /junos/kernel/tcpip/tcp/

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

- **JTI support extended for Junos kernel IPv4 and IPv6 (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel IPv4 and IPv6 information using remote procedure call (gRPC) services. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel IPv4 and IPv6 information:

  - /junos/kernel/tcpip/arp/ — Address Resolution Protocol cache
  - /junos/kernel/tcpip/ndp/ — Neighbor Discovery Protocol cache
  - /junos/kernel/tcpip/netisr/ — NETISR network queues
  - /junos/kernel/tcpip/nhdix/ — Nexthop index space exhaustion
  - /junos/kernel/tcpip/rtb/ — Route tables

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

- **JTI support extended for Junos kernel IP multicast, tunnels, TNP, and VPLS (EX9200, EX9251, EX9253, MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX10008, PTX10016, PTX3000 with RE-PTX-X8-64G, PTX5000 with RE-PTX-X8-64G)**—Starting in Junos OS Release 19.3R1, Junos telemetry interface (JTI) extends support for streaming Junos kernel IP multicast, tunnels, Trivial Network Protocol (TNP), and Virtual Private LAN Service (VPLS) information using remote procedure call (gRPC) services. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

  You can use the following base resource paths for exporting kernel IP multicast, tunnels, TNP, and VPLS information:

  - /junos/kernel/multicast/
  - /junos/kernel/tunnel/
  - /junos/kernel/tnp/
  - /junos/kernel/vpls/

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

Management

- **OpenConfig AAA data model support (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 19.3R1 supports the configuration leafs specified in the OpenConfig AAA data model. Mapping the OpenConfig AAA configuration to the Junos AAA configuration using the following yang files in the data model makes this support possible:

  - openconfig-aaa.yang
  - openconfig-aaa-types.yang
  - openconfig-aaa-tacacs.yang
  - openconfig-aaa-radius.yang

  The configuration model supporting the OpenConfig data model includes:

  - A translation script (.py / .slax) that maps each config leaf in the OpenConfig Schema to one or more config leafs in the JUNOS Schema.
  - A deviation file (.yang) that specifies how much the implementation deviates from the vendor-neutral model.

[See Mapping OpenConfig AAA Commands to Junos Configuration.]

MPLS

- **BGP multipath scaling to 128-way (PTX Series)**—The maximum number of supported ECMP paths is increased to 128. Use the **maximum-ecmp** command to configure the maximum limit for ECMP next hops. This provides more flexibility to load-balance on networks with high-volume traffic.

  [See maximum-ecmp].

Network Management and Monitoring

- **sFlow performance improvements (PTX10002)**—Starting in Junos OS Release 19.3R1, the following improvements have been made to sFlow for PTX10002 routers:

  - You can configure forwarding class and DSCP values per collector.
  - You can configure IPv6 addresses for the **source-ip** and **agent-id**.

  [See Understanding How to Use sFlow Technology for Network Monitoring.]

Routing Policy and Firewall filters

- **Filter-based GRE encapsulation (PTX10002)** Starting with Junos OS Release 19.3R1, for PTX10002 routers running third-generation line cards (the ExpressPlus chipset), you can use **tunnel-end-point** commands to enable line-rate, filter-based, GRE tunneling of IPv4 and IPv6 payloads across IPv4 networks.

  This GRE encapsulation is not supported for logical systems, or for MPLS traffic, and the route lookup for GRE encapsulated traffic is supported on the default routing instance only.

  [See tunnel-end-point and Components of Filter-Based Tunneling Across IPv4 Networks.]

- **IPv6 packet flow labels as load-balancing hash key (PTX10002)**—Starting in Junos OS Release 19.3R1, you can configure IPv6 packet flow labels for hash calculations on PTX10002 routers. Additionally, if you want the load balancing to be based on the flow label of the IPv6 header, include the **ipv6-flow-label** configuration statement at the **[edit forwarding-options hash-key family inet6 layer-3]** hierarchy level.

  The use of the flow labels enhances load balancers operating on IP packets and TCP sessions, commonly known as Layer 3/4 load balancers. With this approach, you improve the performance of most types of Layer 3/4 load balancers, especially for traffic including multiple IPv6 extension headers and for fragmented packets.

  [See ipv6-flow-label.]

Routing Protocols

- **Support for color mode in segment routing traffic engineering using BGP (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports color-only mode corresponding to color bits 01 and supports the steering fallback mechanism (in a limited manner) when color bits as set to 01 as described in IETF DRAFT-SPRING-SRTE. Use the **extended-nexthop-color** CLI configuration option to set color bits to 01 to enable color-only mode. Fall back to color-only SRTE policies is also supported and can be configured independently by configuring an import policy at the headend.

[See Understanding Ingress Peer Traffic Engineering for BGP SPRING.]

- **Support for OSPF TI-LFA back paths for Segment Routing (PTX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports creation of OSPF topology-independent TI-LFA backup paths where the prefix SID is learned from a segment routing mapping server advertisement when the PLR and mapping server are both in the same OSPF area.

  [See Configuring Topology-Independent Loop-Free Alternate with Segment Routing for OSPF.]

- **Support for 64 BGP add-path routes (PTX10002-60C)**—Starting in Junos OS Release 19.3R1, support is extended to 64 BGP add-path routes. Currently, Junos OS supports six add-path routes and BGP can advertise up to 20 add-path routes through policy configuration. If you enable advertisement of multiple paths to a destination or if you increase the add-path prefix policy send count, BGP can now advertise up to 64 add-path routes.

  To advertise all add-paths, up to 64 add-paths or only equal-cost paths, include the **path-selection-mode** statement at the **[edit protocols bgp group *group-name* family *name* addpath send]** hierarchy level. You cannot enable both **multipath** and **path-selection-mode** at the same time.

  To advertise a second best path as a backup path in addition to the multiple ECMP paths include the **include-backup-path *backup_path_name*** statement at the **[edit protocols bgp group *group-name* family *name* addpath send]]** hierarchy level.

  [See path-selection-mode.]

  [See include-backup-path.]

Security

- **Support of l2circuit-control-passthrough statement on PTX10002-60C**—Starting in Junos OS Release 19.3R1, you can enable the pass-through of the following Layer 2 protocols on a Layer 2 circuit by configuring the **l2circuit-control-passthrough** statement:

  - Link Aggregation Control Protocol (LACP)
  - Link Layer Discovery Protocol (LLDP)
  - OAM link fault management (LFM)
  - OAM connectivity fault management (CFM)

  Prior to this release, these Layer 2 protocols were classified as the control packets on the PTX10002-60C router, and were not routed on a Layer 2 circuit.

  [See l2circuit-control-passthrough.]

Software Installation and Upgrade

- **Zero touch provisioning using WAN interfaces (PTX1000)**—Starting in Junos OS Release 19.3R1, in Zero Touch Provisioning (ZTP), you can either use WAN interfaces or management interfaces, to automatically download and install the appropriate software and the configuration file on your device during the bootstrap process.

ZTP automatically configures WAN interfaces based on the optics type, and then connects your device to the Dynamic Host Configuration Protocol (DHCP) server to perform the bootstrap process.

[See Zero Touch Provisioning.]

- **Migration of Linux kernel version**—Starting in Junos OS Release 19.3R1, the following devices support the Wind River Linux 9 (WRL9) kernel version:

Platforms	Routing Engine Supported
ACX5448-D	RE-ACX-5448
MX240, MX480, and MX960	RE-S-X6-64G
MX2020 and MX2010	REMX2K-X8-64G
MX204	RE-S-1600x8
MX10003	RE-S-1600x8
MX2008	RE-MX2008-X8-64G
MX10016	RE X10
MX10008	RE X10
PTX5000	RE-PTX-X8-64G
PTX3000	RCBPTX
PTX10016	RE-PTX-2X00x4/RE X10
PTX10008	RE-PTX-2X00x4/RE X10
PTX1000	RE-PTX1000
PTX10002-XX	RE-PTX10002-60C
EX9208	RE-S-EX9200-2X00x6
EX9251	EX9251-RE
EX9253	EX9253-RE
EX9204	RE-S-EX9200-2X00x6

Platforms	Routing Engine Supported
EX9214	RE-S-EX9200-2X00x6
QFX10002	RE-QFX10002-60C
QFX10008	RE-QFX10008
QFX10016	RE-QFX10016

Starting in Junos OS Release 19.3R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following routers:

- MX Series—MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008
- PTX Series—PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

If you perform a software upgrade on a router with i40e NVM version earlier than 6.01, the upgrade fails and the following error message is displayed:

ERROR: i40e NVM firmware is not compatible ,please upgrade i40e NVM before installing this package

ERROR: Aborting the installation

ERROR: Upgrade failed

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

SEE ALSO

What's Changed | **208**

Known Limitations | **211**

Open Issues | **213**

Resolved Issues | **215**

Documentation Updates | **219**

Migration, Upgrade, and Downgrade Instructions | **219**

What's Changed

IN THIS SECTION

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

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

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

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

  See Configuring RFC-Compliant NETCONF Sessions.

What's Changed in 19.3R2

There are no changes in behavior and syntax for PTX Series in Junos OS Release 19.3R2.

What's Changed in 19.3R1

General Routing

- **User confirmation prompt for configuring the suboptions of request vmhost commands (MX Series and PTX series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm (with a **yes** or a **no**) whether you want to configure the suboptions also.
  - **request vmhost reboot**
  - **request vmhost poweroff**
  - **request vmhost halt**

  In earlier Junos OS releases, the confirmation prompt is available only for the main options.

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Release 19.3R1 and later, PTX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

  In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.
- **Monitoring information available only in trace log (PTX Series)**—In Junos OS Release 19.3R1 and later, the Ethernet link fault management daemon (lfmd) in the peer router stops monitoring the locally occurred errors until unified ISSU completes. You can view the monitoring-related details only through the trace log file.

Junos OS XML, API, and Scripting

- **Range defined for confirm-timeout value in NETCONF and Junos XML protocol sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, the value for the **<confirm-timeout>** element in the Junos XML protocol **<commit-configuration>** operation must be in the range 1 through 65,535 minutes, and the value for the **<confirm-timeout>** element in the NETCONF **<commit>** operation must be in the range 1 through 4,294,967,295 seconds. In earlier releases, the range is determined by the minimum and maximum value of its unsigned integer data type.

Software Defined Networking

- **Increase in the maximum value of delegation-cleanup-timeout (PTX Series)**—You can now configure a maximum of *2147483647* seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

  With the increase in maximum value of **delegation-cleanup-timeout** from *600* to *2147483647* seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

  [See delegation-cleanup-timeout.]

System Logging

- **Preventing system instability during core file generation (PTX Series)**—Starting with Release 19.3R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

What's New | **200**

Known Limitations | **211**

Open Issues | **213**

Resolved Issues | **215**

Documentation Updates | **219**

Migration, Upgrade, and Downgrade Instructions | **219**

Known Limitations

IN THIS SECTION

Learn about known limitations in this release 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

- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error**. PR1268678
- The Routing Engine boots from the secondary disk when you:

  Press the reset button on the RCB front panel, while the Routing Engine is booting up before Junos OS reboots.

  Upgrade the software by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network.

  Upgrade the BIOS and it fails.

  Reboot the system and it hangs before Junos OS reboots. As a workaround, interrupt the boot process to select the primary disk. PR1344342
- If output firewall filter is configured with the "syslog" or log option, the host interface might be wedged on PTX1000, PTX5000, and PTX10000. The change in this PR is to add the warning but it does not prevent the problem that causes the host interface stop sending packets. This condition might occur if the following conditions are met:

  The packet that is hitting the filter term should be less than 128 bytes

Output firewall filter has syslog, log or port-mirror and accept action.

Sample configuration for V4 and V6:

set interfaces <interface name> unit family inet filter output <filter-V4>

set firewall family inet filter <filter-V4> term 1 then log

set firewall family inet filter <filter-V4> term 1 then accept

set interfaces <interface name> <unit> family inet6 filter output <filter-V6>

set firewall family inet6 filter <filter-V6> term 1 then log

set firewall family inet6 filter <filter-V6> term 1 then accept PR1354580

- The **request vmhost power-off** command does not actually power off the system in the latest releases. It only does a reboot and the system comes back up. PR1393061
- One hundred percent traffic loss is seen on all streams from PTX10001 to MX240. PR1435069
- PTX10000/QFX10000 devices will learn source MAC information even when the traffic is explicitly dropped through the Ethernet-switching filter. This is because a learning event is triggered in the Source Lookup block of the ASIC that occurs before the filter rule is executed. Hence the learning event can't be avoided. The learning event generated on PTX Series or QFX Series devices is not dependent on forwarding decisions taken in subsequent stages of the ASIC pipeline. PR1436377
- Due to an issue in the BIOS:QFXS_SFP_00.32_02.01 version, when the watchdog is killed, the device will not reboot. PR1441963
- The **clear interface statistics** command is taking much longer time to execute than expected. PR1447851
- The local-loopback test is failing with gigether-options. PR1458814
- Traffic failure is seen with gcm-aes-xpn-128 cipher on performing event. PR1460254

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after upgrade. This is because of the old version of **/var/db/cfm.db**. PR1281073

SEE ALSO

What's New | **200**

What's Changed | **208**

Open Issues | **213**

Resolved Issues | **215**

Open Issues

IN THIS SECTION

Learn about open issues in this release 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

- The routing information base (RIB, that is routing table) and forwarding information base (FIB, that is forwarding table) might get out of sync in a very large-scale network due to kernel routing table queue being stuck. The KRT queue is used by the rpd process to send forwarding information messages to Packet Forwarding Engines. With the stuck state, the queue can get into a state where no more messages can be sent to the Packet Forwarding Engines. This issue is applicable only to Junos OS Release 16.1 through Junos OS Release 17.3. PR1315212
- Users might not be able to stop the ZTP bootstrap, when a PTX10016 or PTX10008 router with many line cards is powered ON with the factory-default configuration. PR1369959
- The **rx_power** value streamed to the telemetry server is the raw value (mW) returned directly from the transceiver driver. The Junos CLI value has been transformed in the transportd daemon into different units: Rx input total power(0.01dBm). PR1411023
- Traffic loss for more than 15 seconds is seen when 50 percent of ae links are brought down by restarting multiple FPCs. PR1412578
- VTY command **show filter index < number> counter** shows values as zero at 28-02-HOSTBOUND_NDP_DISCARD_TERM on PTX5000 platform. Basically, the counter does not increase for NDP packets. The issue is seen only with **show filter index**, which is a debug tool in VTY. This issue has no impact on NDP functionality for user traffic. There are no issues with NDP functionality, and DDoS for NDP is also working. PR1420057
- The firewall counter for the lo0 interface might not increase. PR1420560

- Due to a Broadcom firmware issue, under some error conditions, the Broadcom retimer will assert pmd-lock bit with a “closed eye”, in order to recover from this condition, Junos code will issue the **force pmd-lock** command every time you try to bring up a link in "DOWN" state. PR1428307
- An em2 interface configuration is causing FPC to crash during initialization and the FPC does not come online. After deleting the em2 configuration and restarting, the router FPC comes online. PR1429212
- Interface statistics are not getting updated with port-mirroring. PR1431607
- More packet loss is seen after unified ISSU with InterAS Layer 3 VPN OptionB configuration. PR1435578
- Due to an issue in the BIOS:QFXS_SFP_00.32_02.01 version, when the watchdog is killed, the device will not reboot. PR1441963
- When BGP labeled-unicast (BGP-LU) IPv6 session is configured on PTX10000 platform and more than 65,000 IPv6 labeled-unicast routes are received on this session, the F-label might be exhausted because chained composite next hops for ingress labeled-bgp LSPs are not supported on this platform. The F-label exhaustion could cause the kernel routing table (KRT) queue to be stuck with the error **ENOMEM -- Cannot allocate memory**, which could cause routes to be missing in the forwarding table. PR1442760
- On PTX5000 and PTX3000 router with 15x100G and 96x10G PIC, the interface bcm8238 line side amplitude setting is incorrect, and it might cause optic reliability issues. PR1453217
- The local-loopback test is failing with gigether-options. PR1458814
- After injecting the errors on FPC, all interrupts are not recorded. PR1459367
- The ae interface egress or output statistics will not be in sync with the aggregated members (ports) egress/output statistics. PR1459633
- Layer 3 traffic fragmentation is failing without DF bit in Junos OS Release 19.3R1. PR1459738
- Statistics comparison between CLI and Junos telemetry interface for queue is failing because the buffers are showing incorrect values. PR1460246
- Traffic failure with gcm-aes-xpn-128 cipher is observed on performing event. PR1460254
- The following message is seen: **core-argus-bng-reg-node-cda-zh.0.tgz @ __assert_fail_base, __GI___assert_fail, zephyr_filter_regs_pf2_3_v2_beta_node_mem_bank_set, zephyr_beta_node_write, Cda::GrpcUnaryApi, Cda::AsicServerGrpc::runloop(), Cda::AsicServerGrpc::run()**. PR1467741

Routing Protocols

- With BFD configured on the ae interface, if you disable or enable the ae interface, afterward that ae interface and Bidirectional Forwarding Detection (BFD) session might not come up. PR1354409

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in Junos OS main and maintenance releases for PTX Series.

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

Resolved Issues: 19.3R2

Forwarding and Sampling

- PFED core files are seen and MIB2D is reported as slow peer due to a Packet Forwarding Engine accounting issue, original PR1405091. PR1452363

General Routing

- The ae interface is not coming up with LACP enabled over the CCC circuit between R0 and R3. PR1424553
- BCM FW needs to be upgraded to DE2E. PR1445473
- Changing the hostname will trigger lsp on -change notification, not an adjacency on-change notification. Also, currently IS-IS is sending hostname instead of system-id in OC paths. PR1449837
- Interfaces might flap forever after deleting the interface disable configuration. PR1450263
- Firewall filter applied at interface level is not working when entropy level is present in certain scenarios. PR1452716

- The FPC might crash when the severity of error is modified. PR1453871
- Grpc updates on_change is not working when performing delete operations. PR1459038
- Traffic gets dropped silently upon interface flap after DRD autorecovery. PR1459698
- The **forwarding** option is missed in routing instance type. PR1460181
- Hardware failure in CB2-PTX causes traffic interruption. PR1460992
- IPV6 ping is not working between CE to CE in Layer 3 VPN network. PR1466659

Interfaces and Chassis

- After member interface flapping, the ae interface remains down on 5X100GE DWDM CFP2-ACO PIC. PR1429279

Platform and Infrastructure

- REST API process will get nonresponsive when a number of requests come at a high rate. PR1449987

Routing Protocols

- The rpd might crash with SR-TE configuration change. PR1442952
- Layer 3 VPN PE-CE link protection exhibits unexpected behavior on MX2000/PTX10000 platforms. PR1447601

Resolved Issues: 19.3R1

General Routing

- agentd sensor transmits multiple interface telemetry statistics per FPC slot. PR1392880
- Confirmation message is missing when you issue **request vmhost reboot re**. PR1397912
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. PR1399683
- ZTP upgrade might fail if there are more than one 10-Gigabit Ethernet interfaces connected to the DHCP server. PR1404832
- On PTX3000 and PTX5000, the backup CB's chassis environment status keeps 'Testing' after backup CB removal or reinsertion. PR1405181
- **openconfig-network-instance:network-instances** support for IS-IS must be hidden unless supported. PR1408151
- Junos OS PCC might reject PCUpdate or PCCreate message if there is a metric type other than type 2. PR1412659
- The core file **fpc3-sevfpc.elf.0.tgz** is generated after GRES in PTX3000. PR1415145
- Support for 140e NVM firmware upgrade through CLI for PTX10002-60c. PR1418909
- An interface might go down on the PTX10000 platform. PR1421075

- Virtual Chassis might become unstable and FXPC might generate core files when there are a lot of configured filter entries. PR1422132
- Packet Forwarding Engine wedge might occur after issuing the **show forwarding-options load-balance** command. PR1422464
- 4x10G interfaces on PTX3000 and PTX5000 third-generation FPCs might not come up after frequent flapping for a long duration of time. PR1422535
- While committing a huge configuration the following error **error: mustd trace init failed** is seen. PR1423229
- Traffic is dropped after FPC reboot with an aggregated Ethernet member links deactivated by the remote device. PR1423707
- **per-interface-per-member-link** statement is hidden for PTX5000 FPC. PR1425372
- The host-bound traffic might be dropped after performing change configuration related to prefix-list. PR1426539
- A specific interface on the P3-15-U-QSFP28 PIC card remains down until another interface comes up. PR1427733
- An interface on port 7, 9, 17, 19, 27, or 29 might go down on 30-port 100-Gigabit or 40-Gigabit line cards. PR1427883
- On PTX10000, when an interface is configured with jumbo frames support (for example, MTU = 9216), the effective MTU size for locally sourced traffic is 24 bytes less than the expected value. PR1428094
- **show chassis environment** shows input0 and input1. PR1428690
- Inline J-Flow might cause a major error on the chip. PR1429419
- IPFIX flow timestamp does not match the with NTP-synchronized system time. PR1431498
- **SIB Link Error** detected on a specific Packet Forwarding Engine might cause complete service impact. PR1431592
- The scaled filter might drop packets with the **flt.Dispatcher.flt_err** error on the PTX Series routers. PR1433648
- IPv6 neighbor solicitation packets are dropped on PTX Series routers. PR1434567
- On PTX10002, **No chassis alarm** is raised when a PEM is removed or power off to PEM. PR1439198
- The PTX10K-LC1104 and PTX10K-LC1105 line cards might continuously crash when an inline flow monitoring is configured. PR1439956
- Interfaces on PTX Series routes might not come up after FPC restart or port flap. PR1442159
- BCM FW needs to be upgraded to DE2E. PR1445473
- Receipt of a malformed packet for J-Flow sampling might create an FPC core file. PR1445585

- Firewall filter applied at the interface level is not working when entropy level is present in certain scenarios. PR1452716
- The jdhcpd process might crash after issuing the **show access-security router-advertisement-guard** command. PR1446034
- Egress sampling for sflow might stop working for more than 8 interfaces on PTX Series platforms. PR1448778
- GRPC updates on_change does not work when performing delete operations. PR1459038

Infrastructure

- The CLI command **request system recover oam-volume** might fail on PTX Series routers. PR1425003
- An unsupported package warning after the system is upgraded. PR1427344

Interfaces and Chassis

- Syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** on committing LFM-related configuration is commit on aggregated Ethernet interfaces. PR1423586
- Some ports on PTX Series routers might remain down after the FPC or device at the remote side is rebooted. PR1429315

Layer 2 Ethernet Services

- DHCP request might get dropped in DHCP relay scenario. PR1435039

MPLS

- The optimization timer is being updated in an incorrect manner in the code path. Due to this, a particular check fails when the exponential increase function is called. PR1416948
- The dynamic bypass RSVP LSP tears down when being used to protect an LDP LSP. PR1425824
- The transit packets might be dropped if an LSP is added or changed on a PTX Series router. PR1447170

Routing Protocols

- Routing Engine-based micro-BFD packets do not go out with the configured source IP address when the interface is in the logical system. PR1370463
- PTX Series routers cannot intercept a PIM BSR message. PR1419124
- The rpd might crash with **ospf overload** configuration. PR1429765

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. PR1282875
- Memory leak might happen if PIM messages received over an MDT (mt- interface) in Draft-Rosen MVPN scenario. PR1442054

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for the PTX Series.

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 19.3

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 19.3R2:

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. Click 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 comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-19.3R2.9.tgz
```

Customers in the Eurasian Customs Union (currently comprised 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-x86-64-19.3R2.9-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.
- Starting in Junos OS Release 19.3R2, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following PTX Series routers:
  - PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

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

NOTE: After you install a Junos OS Release 19.2**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

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases

provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://support.juniper.net/support/eol/software/junos/.

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 | 200
- What's Changed | 208
- Known Limitations | 211
- Open Issues | 213
- Resolved Issues | 215
- Documentation Updates | 219

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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 19.3R2: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5200-32CD, 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 19.3R2

There are no new features or enhancements to existing features for QFX Series Switches in Junos OS Release 19.3R2.

What's New in 19.3R1

Hardware

- **JNP-SFPP-10GE-T transceivers (QFX10000-60S-6Q line card)**—Starting in Junos OS Release 19.3R1, the QFX10000-60S-6Q line card supports the JNP-SFPP-10GE-T transceivers.

  > **NOTE:** The JNP-SFPP-10GE-T SFP+ operates in multi-rate speeds of 100/1000/10G BASE-T. If the attached device advertises only 100/1000 BASE-T speeds, the SFP+ transceiver sets the line rate to the agreed autonegotiation rates. However, the transceiver considers the link to be a 10Gbps link which might cause one of the following unexpected behaviors:
  >
  > - Packet loss occurs because of the difference in actual versus provisioned speeds.
  > - The SFP+ tranceiver link goes down reflecting the xe-*a*/*b*/*c* state to be down.

EVPN

- **Selective multicast forwarding and SMET support in EVPN-VXLAN (QFX5110 and QFX5120 switches)**—Starting in Junos OS Release 19.3R1, Junos OS supports selective multicast Ethernet forwarding in an EVPN-VXLAN network. IGMP snooping enabled devices on a bridge domain monitor and selectively forward traffic from the access interface to the core. Devices that support selective multicast Ethernet forwarding do not send multicast traffic to all devices. Instead, they replicate and forward multicast traffic only to the devices that indicate an interest. This feature is supported on a spine-and-leaf topology where the network can consist of a mix of devices that support selective multicast Ethernet and those that do not support this feature.

  [See Selective Multicast Forwarding.]

- **BPDU protection in EVPN-VXLAN (QFX5100, QFX5110, and QFX5200)**—Starting in Junos OS Release 19.3R1, you can enable BPDU protection to avoid network outages due to STP, MSTP, and RSTP miscalculations. Without BPDU protection, STP, MSTP, and RSTP BPDUs are not recognized and are flooded as unknown Layer 2 packets on the VXLAN interfaces. With BPDU protection, when a BPDU is received on an edge port in an EVPN-VXLAN environment, the edge port is disabled, and it stops forwarding all traffic. You can also configure BPDU protection to drop BPDU traffic but have all other traffic forwarded on interfaces without having to configure a spanning-tree protocol.

  - To enable BPDU protection on an edge port with RSTP on access and leaf devices:

    set protocols rstp interface *interface-name* edge

    set protocols rstp bpdu-block-on-edge

  - To enable BPDU protection without a spanning-tree protocol configured on access and leaf devices:

    set protocols layer2-control bpdu-block interface *interface-name*

  - To enable BPDU protection without a spanning- tree protocol but still forward other traffic on access and leaf devices:

    set protocols layer2-control bpdu-block interface *interface-name* drop

Forwarding and Sampling

- **Customizing hashing parameters and shared-buffer alpha values for better load balancing (EX4650 and QFX5120 switches)**—These switches achieve load balancing through use of a hashing algorithm, which determines how to forward traffic over LAG bundles or to next-hop devices when ECMP is enabled. The hashing algorithm makes hashing decisions based on values in various packet fields. Starting with Junos OS Release 19.3R1, you can explicitly configure some hashing parameters to make hashing more efficient. The shared-buffer pool is a global memory space that all ports on the switch share dynamically as they need buffers. The switch uses the shared-buffer pool to absorb traffic bursts after the dedicated-buffer pool is exhausted. The shared-buffer pool threshold is dynamically calculated based on a factor called alpha. Also starting with Junos OS Release 19.3R1, you can specify the alpha, or dynamic threshold, value to determine the change threshold of shared buffer pools for both ingress and egress buffer partitions.

  To specify hashing parameters:

  `user@switch#` **set forwarding-options enhanced-hash-key hash-parameters (ecmp | lag)**

  To specify a threshold value for a particular queue:

  `user@switch#` **set class-of-service shared-buffer (ingress|egress) buffer-partition *buffer* dynamic-threshold *value***

  [See hash-parameters and buffer-partition].

Interfaces and Chassis

- **FTIs with support for UDP encapsulation (QFX Series)**—Starting in Junos OS Release 19.3R1, you can configure flexible tunnel interfaces (FTIs) on the PTX Series routers/QFX Series switches, which provide support for static UDP tunnels only.

  With the UDP tunnels-over-FTI feature, you can benefit from better traffic distribution over ECMP, that is achieved by the UDP source port derived from the hash value of the inner payload. In addition to this, the other benefits of this feature include, shortened interface hop counts, smooth IGP domain separation, and reduced operational complexity.

  [See Flexible Tunnel Interfaces Overview.]

- **Gigabit Ethernet Optics for the QFX5110**—Starting in Junos OS Release 19.3R1, QFX5110 switches support these optics:

  - SFP-GE10KT15R13
  - SFP-GE10KT13R15
  - SFP-GE40KT13R15
  - SFP-GE40KT15R13
  - EX-SFP-GE10KT15R13
  - EX-SFP-GE10KT13R15
  - EX-SFP-GE40KT13R15
  - EX-SFP-GE40KT15R13

  See the [Hardware Compatibility Tool].

- **Host route generation support for ARP and Neighbor Discovery Protocol (NDP) (QFX5100)**—Starting in Release 19.3R1, Junos OS supports host route generation for devices connected to QFX5100 switches in a data center. When you enable this feature on an interface for IPv4 or IPv6, host routes are created in the routing table for each device present in ARP (IPv4) and NDP (IPv6). These host routes can be exported to routing protocols to be advertised to the network by matching the new policy qualifier **l2-learned-host-routing** statement.

  You can configure the **host-route-generation** statement under the **[edit interfaces *name* unit *name* family inet/inet6]** hierarchy, on each interface and for each address family.

  > **NOTE:** Host route generation is disabled by default.

- **Proactive ARP detection (QFX5110 and QFX5120)**—Starting with Junos OS Release 19.3R1, you can check the reachability of connected devices (within an IP subnet range) on a specified interface. To enable proactive ARP detection, configure the **proactive-arp-detection** statement at the **[edit system arp]** hierarchy level. After enabling the **proactive-arp-detection** statement, you can set the ARP

configurations at the interface level by the setting the: **host-discovery *address-range***, **ageing-time-out *seconds***, and **discovery-time-interval *seconds*** options at the [edit interfaces *interface-name* family inet address *ip-address*] hierarchy level. Likewise, you can delete the ARP configuration settings by using the **delete interfaces *interface-name* unit *unit*** family inet address *ip-address* host-discovery *address-range* command.

[See proactive-arp-detection.]

- **QFX5120 supports JNP-SFPP-10GE-T**—Starting in Junos OS Release 19.3R1, QFX5120 switches support the new copper 10GBASE-T SFP+ transceiver (JNP-SFPP-10GE-T), which provides a speed of 10 Gbps. Use the existing show commands such as **show chassis pic** and **show chassis hardware** to view the details of the transceivers.

  > **NOTE:** In case a device with a different interface speed (that is, 1 Gbps or 100 Mbps) is connected on the other side of the wire, the interface on the Juniper device does not come up.

  [See show chassis pic and show chassis hardware.]

Junos OS XML, API, and Scripting

- **IPv6 support in Python automation scripts (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.3R1, devices running Junos OS with upgraded FreeBSD support using IPv6 in:
  - Python automation scripts, including commit, event, op, and SNMP scripts
  - Juniper Extension Toolkit (JET) scripts
  - YANG action scripts

  IPv6 support enables Python scripts to establish connections and perform operations using IPv6 addresses.

  [See IPv6 Support in Python Automation Scripts.]

Junos Telemetry Interface

- **JTI support for interface burst monitoring (QFX5220-128C and QFX5220-32CD)**—Junos OS Evolved Release 19.3R1supports interface burst monitoring on Junos telemetry interface (JTI) to monitor physical interfaces for bursts. Use interface burst monitoring to help troubleshoot problems, make decisions, and adjust resources as needed.

  Exported statistics report:
  - Peak bytes
  - The time peak bytes are detected
  - The direction (transmit or receive)

You can export interface burst statistics from the Juniper device to an outside collector by including the sensor **/junos/system/linecard/bmon-sw/** in a subscription using remote procedure call (gRPC) services.

To provision the sensor to export data through gRPC services, use the **telemetrySubcribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support JTI.

> **NOTE:** This feature does not detect microbursts.

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

Management

- **OpenConfig AAA data model support (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 19.3R1 supports the configuration leafs specified in the OpenConfig AAA data model. Mapping the OpenConfig AAA configuration to the Junos AAA configuration using the following YANG files in the data model makes this support possible:

  - openconfig-aaa.yang
  - openconfig-aaa-types.yang
  - openconfig-aaa-tacacs.yang
  - openconfig-aaa-radius.yang

  The configuration model supporting the OpenConfig data model includes:

  - A translation script (**.py / .slax**) that maps each configuration leaf in the OpenConfig schema to one or more configuration leafs in the Junos OS Schema.
  - A deviation file (**.yang**) that specifies how much the implementation deviates from the vendor-neutral model.

[See Mapping OpenConfig AAA Commands to Junos Configuration.]

Multicast

- **MLDv1, MLDv2, and MLD snooping (EX4650 and QFX5120-48Y switches and Virtual Chassis)**—Starting in Junos OS Release 19.3R1, you can configure Multicast Listener Discovery (MLD) version 1 (MLDv1), MLD version 2 (MLDv2), and MLD snooping on EX4650 and QFX5120-48Y switches and Virtual Chassis. With MLD snooping enabled, these switches or Virtual Chassis replicate and forward IPv6 traffic for a multicast group only to the interfaces in a VLAN with listeners who joined the group, rather than flooding to all interfaces in the VLAN.

  [See Examples: Configuring MLD and Understanding MLD Snooping.]

Routing Protocols

- **RIPng routing protocol supported (EX4650 and QFX5120 switches)** Starting with Junos OS Release 19.3R1, EX4650 and QFX5120 switches support the RIPng routing protocol.

  [See Basic RIPng Configuration.]

- **Support for color mode in segment routing traffic engineering using BGP (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.3R1, Junos OS supports color-only mode corresponding to color bits 01 and supports the steering fallback mechanism (in a limited manner) when color bits as set to 01 as described in IETF DRAFT-SPRING-SRTE. Use the **extended-nexthop-color** CLI configuration option to set color bits to 01 to enable color-only mode. Fall back to color-only SRTE policies is also supported and can be configured independently by configuring an import policy at the headend.

  [See Understanding Ingress Peer Traffic Engineering for BGP SPRING.]

Routing Protocols and Firewall Filters

- **Support for IPv6 Filter-Based Forwarding (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 19.3R1, you can use stateless firewall filters in conjunction with filters and routing instances to control how IPv6 traffic travels in a network on EX4650 and QFX5120 switches. This is called IPv6 filter-based forwarding. To set up this feature, you define a filtering term that matches incoming packets based on the source or destination address and then specify the routing instance to send packets to. You can use filter-based forwarding to route specific types of traffic through a firewall or security device before the traffic continues on its path. You can also use it to give certain types of traffic preferential treatment or to improve load balancing of switch traffic.

  [See Firewall Filter Match Conditions for IPv6 Traffic and Filter-Based Forwarding Overview.]

Services Applications

- **Support for real-time performance monitoring or RPM (QFX5120)** —Starting in Junos OS Release 19.3R1, you can configure active probes to track and monitor traffic across the network and to investigate network problems on QFX5120 switches.

  You can use RPM in the following ways:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.
- Set thresholds to trigger SNMP traps when values are exceeded.

  You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test.
- Determine automatically whether a path exists between a host router or switch and its configured BGP neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

[See Understanding Real-Time Performance Monitoring on Switches.]

Software Installation and Upgrade

- **Migration of Linux kernel version**—Starting in Junos OS Release 19.3R1, the following devices support the Wind River Linux 9 (WRL9) kernel version:

Platforms	Routing Engine Supported
ACX5448-D	RE-ACX-5448
MX240, MX480, and MX960	RE-S-X6-64G
MX2020 and MX2010	REMX2K-X8-64G
MX204	RE-S-1600x8
MX10003	RE-S-1600x8
MX2008	RE-MX2008-X8-64G
MX10016	RE X10
MX10008	RE X10
PTX5000	RE-PTX-X8-64G
PTX3000	RCBPTX
PTX10016	RE-PTX-2X00x4/RE X10
PTX10008	RE-PTX-2X00x4/RE X10
PTX1000	RE-PTX1000

Platforms	Routing Engine Supported
PTX10002-XX	RE-PTX10002-60C
EX9208	RE-S-EX9200-2X00x6
EX9251	EX9251-RE
EX9253	EX9253-RE
EX9204	RE-S-EX9200-2X00x6
EX9214	RE-S-EX9200-2X00x6
QFX10002	RE-QFX10002-60C
QFX10008	RE-QFX10008
QFX10016	RE-QFX10016

Starting in Junos OS Release 19.3R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following routers:

- MX Series—MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008
- PTX Series—PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

If you perform a software upgrade on a router with i40e NVM version earlier than 6.01, the upgrade fails and the following error message is displayed:

ERROR: i40e NVM firmware is not compatible ,please upgrade i40e NVM before installing this package

ERROR: Aborting the installation

ERROR: Upgrade failed

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

Virtual Chassis

- **Virtual Chassis support (EX4650 and QFX5120-48Y switches)**—Starting in Junos OS Release 19.3R1, you can interconnect two EX4650 or two QFX5120-48Y switches into a Virtual Chassis, which operates as one logical device managed as a single chassis.
  - Member switches must be two EX4650 or two QFX5120 switches (no mixed mode).
  - Both member switches take the Routing Engine role with one as master and one as backup.
  - You can use any of the 100-Gbps QSFP28 or 40-Gbps QSFP+ ports on the front panel (ports 48 through 55) as Virtual Chassis ports (VCPs) to connect the member switches.

- You can run nonstop software upgrade (NSSU) to update the Junos OS release on both member switches with minimal traffic disruption during the upgrade.
- EX4650 and QFX5120 Virtual Chassis support the same protocols and features as the standalone switches in Junos OS Release 19.3R1 except for the following:
  - IEEE 802.1X authentication
  - EVPN-VXLAN (QFX5120)
  - Layer 2 port security features, DHCP, and DHCP snooping
  - Junos telemetry interface (JTI)
  - MPLS
  - Multichassis link aggregation (MC-LAG)
  - Redundant trunk groups (RTG)
  - Priority-based flow control (PFC)

Configuration parameters and operation are the same as for other non-mixed EX Series and QFX Series Virtual Chassis.

[See Virtual Chassis Overview for Switches.]

SEE ALSO

What's Changed

IN THIS SECTION

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

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

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

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

  See Configuring RFC-Compliant NETCONF Sessions.

Changes in Behavior and Syntax: 19.3R2

Network Management and Monitoring

- **entPhysicalTable fetched on QFX10002**—In Junos OS Release 19.3R2, the MIB data for entPhysicalTable will be fetched on a QFX10002-72Q or QFX10002-36Q switch.

  [See SNMP Explorer.]

Changes in Behavior and Syntax: 19.3R1

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Release 19.3R1 and later, QFX Series switches support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

  In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

- **Logical interfaces created along with physical interfaces by default (QFX10000 and QFX5000 switches)**—On the QFX10000 line of switches, logical interfaces are created along with the physical et-, sxe-, xe-, and channelized xe- interfaces. In earlier releases, only physical interfaces are created.

  On the QFX5000 line of switches, by default, logical interfaces are created on channelized xe- interfaces. In earlier releases, logical interfaces are not created by default on channelized xe- interfaces (xe-0/0/0:1, xe-0/0/0:2, and so on), but they are created on et-, sxe-, and nonchannelized xe- interfaces.

Junos OS XML, API, and Scripting

- **Range defined for confirm-timeout value in NETCONF and Junos XML protocol sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, the value for the **<confirm-timeout>** element in the Junos XML protocol **<commit-configuration>** operation must be in the range 1 through 65,535 minutes, and the value for the **<confirm-timeout>** element in the NETCONF **<commit>** operation must be in the range 1 through 4,294,967,295 seconds. In earlier releases, the range is determined by the minimum and maximum value of its unsigned integer data type.

Layer 2 Features

- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—From Junos OS Release 19.3R1, the configuration statement **input-native-vlan-push** at the **[edit interfaces *interface-name*]** hierarchy level is introduced. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement **input-vlan-map** with a **push** operation is configured.

  [See input-native-vlan-push.]

Services Applications

- **Commit check for incomplete tunnel encapsulation configuration on flexible tunnel interface (FTI)**—Tunnel encapsulation configuration is mandatory for FTI interfaces. In Junos OS Release 19.3R1, when you try to commit any incomplete tunnel encapsulation configuration on an FTI, the CLI displays a commit error message.

Software Defined Networking

- **Increase in the maximum value of delegation-cleanup-timeout (QFX Series)**—You can now configure a maximum of *2147483647* seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

  With the increase in maximum value of **delegation-cleanup-timeout** from *600* to *2147483647* seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

  [See delegation-cleanup-timeout.]

System Logging

- **Preventing system instability during core file generation (QFX Series)**—Starting with Release 19.3R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

What's New | 224

Known Limitations | 237

Open Issues | 240

Resolved Issues | 244

Documentation Updates | 256

Migration, Upgrade, and Downgrade Instructions | 257

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for QFX Series. 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 PFC feature is not be supported on QFX5120/EX4650 2-member VC currently due to BCM limitation. PR1431895

EVPN

- When a VLAN uses an IRB interface as the routing interface, the **vlan-id** parameter must be set to "none" to ensure proper traffic routing. This issue is platform-independent. PR1287557

General Routing

- The chip has VLAN-based logical interface statistics. For a given logical interface, both IPv4 and IPv6 packets use the same VLAN, so both v4 and v6 are counted together in the statistics. There is no way to separately count them. Hence, **IPv6 transit statistics** is always 0. However, the total transit statistics (IPv4 + IPv6) will be displayed under Transit statistics. PR1327811
- After installing the Junos OS Release 14.1X53-D51 on an EX4300, xe- interfaces are not seen. PR1336416
- On the QFX5100, if a scaled configuration involving a LAG interface, more that 3000 VLANs, and corresponding next hops is removed and a new configuration involving a LAG interface is applied at the same time, the new configuration might not take effect until the previous configuration has been deleted. During this time, FXPC might consume high CPU resources. No other system impact is observed. PR1363896
- For USB installation if the USB storage device is not removed from device after a USB upgrade, the system might not come up and the system might reboot continuously. The Customer needs to manually change the boot sequence from BIOS menu to select boot from SSD. For PXE installation, the system boots twice from PXE before booting from SSD, and this increases boot time. PR1404717
- Packets of size greater than the MTU of a GRE interface are not fragmented. PR1420803
- During software validation Junos OS mounts the new image and validates the configuration against the new image. Since the TVP-based QFX Series platforms (QFX5000 and QFX10000 are already mounting the maximum 4 disks during normal execution it cannot mount the extra disk for this purpose. Thus QFX currently does not support configuration validation during upgrade on QFX5000 which is why the syntax error appears when the image installation is triggered with "validation". PR1421378
- VLAN is not deleted in the hardware on IRB disable leading to ARP getting refreshed even though IRB is disabled. PR1421382
- The chassisd core file is generated at **fpc_sfxpc_la_ng_show_hw ui_sfxpc_show_hardware ms_parse_substring**. PR1434188
- On QFX5110-32Q running Junos OS 18.1R1 and later, due to a platform limitation, the channelization of the ports should follow the following design recommendations:

- With 100-gigabit transceivers connected in the port range 28–31, only ports 0–19 can be channelized in default system-mode.
- If a 40-gigabit transceiver is connected in any of the 100G supported ports, only ports in the range 1–18 can be channelized in default system mode.
- If all 32 ports have 40-gigabit transceivers connected, only ports in the range 1–18 can be channelized in default system-mode.
- In non-oversubscribed mode, all the valid ports (that is, 0–23) can be channelized as expected.

PR1438319

- The **set class-of-service shared-buffer ingress buffer-partition lossless-headroom percent 0** is not supported when in a Virtual Chassis, as the VCP ports should have some headroom to support PFC. The configuration is rejected at the hardware layer with a log message. PR1448377

Infrastructure

- CRON core file is generated when the statement **cron_popen child_process do_command** is executed. PR1434152

Layer 2 Features

- The **Targeted-broadcast forward-only** command does not broadcast the traffic. PR1359031

Routing Protocols

- Targeted broadcast functionality with VXLAN is not supported yet on QFX5000 platforms. In a non VXLAN scenario, bcast dest IP look up results in a next hop with the destination MAC address of all 0xffs and gives the class ID for IFP to match and action to redirect to IPMC with VLAN membership check. In case of a VxLAN, **l3 egress intf**, **egr l3 next hop**, and **ingress l3 entry** creations are failing. PR1397086
- When IRACL v6 and loopback v6 entries are present, delete and rollback of loopback v6 takes time to re-program the entries in hardware. This is because loopback v6 has high priority in the same IRACL groups and the existing IRACL v6 entries have to be reshuffled in the hardware. PR1428087

SEE ALSO

What's New | 224

What's Changed | 233

Open Issues

IN THIS SECTION

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

EVPN

- OVSDB-managed QFX5100 or QFX5110 is encapsulating VXLAN traffic and sending to wrong destination mac when multiple remote VTEPs are in the same subnet and reachable via irb interface in stretched vlan. This issue is planned to be resolved on the QFX5110 but will not be resolved on the QFX5100. Resolution is still pending for the QFX5110. PR1424698

General Routing

- On the QFX10002-60C, the filter operation with log action only supports Layer 2, IPv4, and IPv6 protocols. The following message is seen in firewall logs: Protocol 0 not recognized.. PR1325437
- Backup Routing Engine might crash after the GRES switchover more than 10 times. PR1348806

- QFX10000 platform drops the Aruba wireless Access Point (AP) heartbeat packets. As a result, the Aruba wireless AP cannot work. PR1352805
- On the QFX10000 line of switches, with EVPN-VXLAN, the following error is seen: expr_nh_fwd_get_egress_install_mask:nh type Indirect of nh_id: # is invalid. PR1367121
- User is not able to stop the ZTP bootstrap process, when the QFX10016 and QFX10008 switches with more than 10 line cards are powered-on with factory-default configuration. PR1369959
- Junos on White Box : USB upgrade of NOS image is not supported. PR1373900
- 100 Gbps VCP links might go down (become unavailable) after the line card member of a QFX5200 Virtual Chassis is renumbered. PR1374655
- Intermittent traffic loss is observed with RTG streams while flapping the RTG Primary interface. PR1388082
- The **show chassis fpc** command displays an incorrect amount of available memory on a QFX10000 FPCs. PR1394978
- On QFX5000 platforms with a scaled setup of the aggregated Ethernet bundles and VLANs, if Link Aggregation Control Protocol (LACP) is enabled, and there are scaled configuration change (for example, delete 4000 VLANs or VxLAN and apply them again) some interfaces of the aggregated Ethernet bundle might go to the detached state. Because of this issue, the running routing protocols (for example, LACP and BGP) might get down over the affected aggregated Ethernet bundles. PR1406691
- There is a possibility of seeing multiple reconnect logs, **JTASK_IO_CONNECT_FAILED**, during the device initialization. There is no functionality impact due to these messages. PR1408995
- On QFX5110 and QFX5120, unicast RPF check in strict mode does not work properly. PR1417546
- On QFX10000 devices, if the analyzer is configured to a mirror traffic of an input aggregated Ethernet interface and a new member is added to the same aggregated Ethernet interface bundle. In this case, the analyzer might not sample packets that flow through the newly added child interface. PR1417694
- Persistent MAC learning is not expected as per TC. PR1422446
- The **show ptp lock-status** command is not supported on qfx5110-48s-4c device from 19.3DCB. PR1426863
- Power cycling while ISSU/ISSR is in progress does not seems to be a valid test scenario. The problem occurs because of ISSU specific sysctl and nvram variables are left with intermediate state values and those should be cleared. However, if such abnormal event occurs, while ISSU/ISSR is in progress, and the system enters in to problematic state. As a workaround, use the following commands to clear ISSU/ISSR specific sysctl / nvram variables from their intermediate state:
  - SYSCTL CLEAR : -**RE:0% sysctl hw.re.tissu=0 RE:0% sysctl hw.re.issu_state=0 RE:0% sysctl hw.lc.issuboot=0**
  - SYSCTL CLEAR : - **RE:0% sysctl hw.re.tissu=0 RE:0% sysctl hw.re.issu_state=0 RE:0% sysctl hw.lc.issuboot=0 NVRAM CLEAR COMMANDS : ---------- RE:0% nvram setenv**

hw.lc.issuboot 0 RE:0% nvram getenv hw.lc.issuboot hw.lc.issuboot=0 RE:0% nvram setenv hw.re.tissu 0 RE:0% nvram getenv hw.re.tissu hw.re.tissu=0

Then the system should be able to recover, provided sudden power failure must not damage anything beyond the ISSU. PR1427563

- The ISSU is not supported on QFX5200 switches and fails from Junos OS Release 17.2X75-D4(x) to Junos OS Release 19.2R1. PR1438690
- The ISSU fails and will not be supported for QFX5200 from Junos OS Release 17.2X75-D4(x) to Junos OS Release 19.2R1. PR1440288
- On QFX5100-VC, VGD process hogs the CPU without **switch-options vtep-source-interface lo0.0** configuration. PR1454014
- On QFX5200-32c-32q a vmcore occurs at **...../.amd/svl-engdata1vs1/occamdev/build/freebsd/stable_11/20190614.234225 __ci_fbsd_builder_stable_11.0.269d466/src/sys/kern/kern_shutdown.c:313** after upgrading from Junos OS Release 18.3 and later to Junos OS Release 19.3R1. PR1455851
- Fan display in show chasis environment is not proper PR1457896
- On QFX5100, when a ISSU is performed with Layer 3 protocols configured then the traffic loss of 0.8 seconds is observed.. PR1459701
- The statement **show forwarding-options enhanced-hash-key** is not supported for QFX10000 platform in Junos OS Release 19.3R1 PR1462519
- On QFX10000, the show command **show forwarding-options enhanced-hash-key** for hashd function does not give the correct output. PR1462563
- Due to a firmware issue on the power supplies (PEMs) of the switch, the routing engine may spontaneously misread the status registers of a power supply. This produces erroneous messages of PEM not present. Although the power supply is present and can deliver power, the system may then deactivate the power supply believing it not to be present. A flashing LED on the PEM may accompany this situation. The system continues to function on the remaining power supply and recovery of non-working power supply requires a full reboot. Errors PR1465183

High Availability (HA) and Resiliency

- The message **kernel: GENCFG: op 51 (AE bias) failed; err 255 (Undefined)** is seen in syslog. These messages don't have any functionality impact. PR1416004

Infrastructure

- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)** messages are seen during FTP. PR1315605

- QFX5100 goes to db prompt when rpd process is restarted using a CLI. PR1372810

Interfaces and Chassis

- You might notice the flooding of ARP reply unicast packets as a result of an ARP request sent for the device's VRRP MAC address. This should not cause major issues. The ARP reply which is flooded in the VLAN by the device has the correct DMAC of the originator of the ARP request. In other words, the ARP reply is flooded but with the correct unicast DMAC. The ARP reply is not broadcasted.

  PR1454764

Layer 2 Features

- In case of QFX5000 Virtual Chassis/VCF setups, when **IGMP-snooping** is enabled, multicast traffic is forwarded based on IGMP joins/reports. But, when IGMP report is timed out, traffic is dropped. This occurs only in case of QFX5000 Virtual Chassis/VCF. PR1431893

MPLS

- The dcpfe core file is generated at **../../../../../../../../../src/pfe/common/applications/nh/hal/nh_db.c** on multiple DUT's while verifying MPLS profile configuration. PR1457356

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. PR1054798

Routing Protocols

- On QFX10000 switches, VRRP is not converging in an MC-LAG environment, because the default routing instance, lo0.0, has been moved to a user-defined routing instance. As a workaround, do not move lo0.0 to the user-defined routing instance. Use a different lo0 logical interface, such as lo0.1 or lo0.2 in the user-defined routing instance. PR1274204
- When 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**. PR1407175
- On QFX5100, the BGP IPv4/IPv6 convergence and RIB install/delete time degraded in Junos OS Releases 19.1R1, 19.2R1, and 19.3R1. PR1414121

- On QFX5110 **MCLAG L2_L3_INTF_OPS_ERROR** error messages are seen after rebooting the node. PR
- On QFX5100 , when a unified ISSU is performed, a traffic loss of 15–20 seconds is observed. PR1449581

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in Junos OS main and maintenance releases for QFX Series.

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

Resolved Issues: 19.3R2

Class of Service (CoS)

- Without this fix, **show cos scheds-per-pfe** and **show cos pfe-scheduler-ifds** Packet Forwarding Engine CLI causes the forwarding plan to restart on QFX10008 switches.PR1452013

EVPN

- ARP and IPv6 neighbor entries cannot be cleared when they are learned from EVPN multi-home ESI. The following commands do not clear ARP and IPv6 neighbor entries when they are learned from EVPN

multi-home ESI. **clear ethernet-switching evpn arp-table clear ethernet-switching evpn nd-table clear ethernet-switching mac-ip-table** PR1446957

- When there is a VXLAN with VLAN ID of 2 on a QFX5100, ARP does not get resolved. PR1453865
- Under EVPN multihoming mode, if ARP Request or Neighbor Solicitation (NS) message encapsulated in Dual Tagged VLAN arrives at the DF (designated forwarder) which might send it back to the local segment as it was, that might cause a loop and at last, overwhelms the device. However, it does not happen with normal broadcast traffic. BDF (backup designated forwarder) does not have this behavior. PR1459830

Forwarding and Sampling

- Commit failure with error might be seen and the dfwd crashes when applying a firewall filter with action **then traffic-class** or **then dscp** to an interface. PR1452435

General Routing

- On QFX5100 platforms, LR4 QSFPs might take longer to come up than others (up to 15 minutes). This is an intermittent occurrence. PR1337340
- mib2d generates a core file in mib2d_write_snmpidx at snmpidx_sync.c on both ADs while bringing. PR1354452
- On QFX5110, the interface FEC counter does not work though FEC function has been supported. PR1382803
- The **show chassis errors active detail** command does not support QFK5000 platform. It is hidden and taken care in other opened scopes. PR1386255
- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped. But, as the port is not started, it does not disable Tx. PR1411015
- When IPv4 and IPv6 are programmed at the same time, most of the IPv6 routes are not installed due to the hardware route table getting full. PR1412873
- On all Junos OS platforms with channelizing ports on FPCs, if 40G ports that are channelized to 10G ports already (eg:xe-2/0/16:0) are being channelized to 10G again, they might get incorrectly channelized. PR1423496
- The dcpfe/Packet Forwarding Engine might not start on AS7816-64X and QFX5K platform devices which results in all the interfaces going down. PR1426737
- When configuring the **global-mac-limit** or **global-map-ip-limit** lower than the number of currently learned MAC/MAC-IP entries, the total number of learned MAC/MAC-IP entries might be more than the configured limit. PR1428572
- On QFX10k platforms, the dcpfe might crash on all line cards if VTEP flap or next-hop deletion happens in scaled environment. PR1431735
- When you plug in unsupported SFP-T module to MPC/DPC/FPC, the line card might crash. PR1432809

- When NSSU is done from 18.1R3 to any forward image on QFX5100-VC with LACP link protection configuration, there might be around 5 minutes traffic loss. Traffic loss is not seen during NSSU if link protection configuration is not present. PR1435519
- dcpfe core file generation is observed with sxe interface and NSSU is performed. PR1435963
- On QFX platforms, the FPC might crash if both the AE (Aggregate Ethernet) bundle flapping on the local device and the configuration change on the peer device which can cause the interface down occur at the same time. PR1437295
- On QFX10002 or QFX10008 or QFX10016 Series platforms with MACsec feature enabled, the BGP neighborship might not be established. PR1438143
- After upgrading Junos OS Release 19.1, port LED shows different from the previous. PR1438359
- On QFX10008 or QFX10016 platform, xSTP recognizes 1G SFP-T optic interface as LAN type link even if it is in full-duplex mode. This might cause the xSTP to converge slowly. As a workaround, configure the xSTP link type from LAN to Pt-Pt (Point To Point) using the command **set protocols <vstp/rstp/mstp> <vlan X> interface <interface name> mode point-to-point**. PR1439095
- When lacp is configured with link protection and force-up on local, and peer is configured with link protection, disabling the active member on peer device causes the LACP MUX state to be stuck in attached state. Issue is not seen if link protection is not configured on the peer device. PR1439268
- EX4600 VC might split if you replace the Virtual Chassis-port connection from SR4 or other fiber QSFP with DAC cable. PR1440062
- There is a IPC sequence issue when Virtual Chassis member rebooted in aggregated interface. After rebooting Virtual Chassis member, RE kernel inject mac entry to fpc. Because of IPC sequence issue, RE added mac entry, originally source mac entry, is added to fpc as remote mac entry. And entry is never be aged out because it is remote entry. PR1440574
- On QFX5000 Series platforms(except for the QFX5100) in the VXLAN scenario, the traffic is not classified properly on the UNI interface which has multiple VLANs configured. PR1445960
- When unicast arp request is received by EX3400/QFX5100 switch and it is configured with "set switch-options no-arp-trap option", the arp request may not be replied. This has been fixed and unicast ARP request will be replied even with "set switch-options no-arp-trap option" configuration. PR1448071
- On QFX series platform, the sFlow sample packets would stop on one Aggregated Ethernet (AE) member link if ingress sFlow is configured on the member link. This would cause inaccurate monitoring on network traffic. PR1449568
- On QFX5000 platform when dual vlan tag feature is configured on physical interface a warning/error is not received. Since this feature is supported only on logical and sub interfaces, it does not work when configured on a physical interface. PR1450455
- On QFX5120 Series switches (function as PE devices), all the L3 tunneling (for example, IP in IP, GRE, VXLAN) packets hit the wrong routing table while receiving the encapsulated packets on the L3VPN MPLS PE-CE interface. This can result in a black hole issue. PR1451032

- If OVSDB is enabled on the device, in a rare case, vgd (VTEP gateway daemon) core file might be seen when a tunnel is deleted twice. It might cause OVSDB to not work properly. PR1452149
- On QFX10000 platforms, DHCP offer packet with unicast flag set gets dropped if anycast IP is used in a VXLAN multi-homed setup. PR1452870
- The VLAN specific parameters might not be used if configuring VLAN all option and VLAN specific configuration. PR1453505
- On QFX5100 or QFX5110 or QFX5120 or QFX5200 or QFX5210 Series platforms with an EVPN/VXLAN scenario, the classifier might not be applied to the interface successfully and all the traffic flows in the best-effort queue. PR1453512
- **show chassis led** status outputs might not be proper along with some port status. PR1453821
- On QFX5100-VC VGD process hogs the CPU without **switch-options vtep-source-interface lo0.0** configuration. PR1454014
- EVPN-VXLAN: MAC+IP Count may be shown as Zero in the output of CLI **show ethernet-switching global-information** on Xellent (QFX10002-60c). PR1454603
- On the QFX5210 platform, the laser emits from the 10G SFP+ port even though the interface is disabled or the device is rebooted. It will cause the peer's interface to be up and might impact traffic. PR1456742
- QFX series switches generate SNMP trap for high temperature after upgrading to any of the affected Junos software. This is due to a temperature threshold value being set incorrectly in the software, SNMP false trap related to temperature gets generated and results in "over temperature" logs. PR1457456
- In EVPN-VXLAN with retaining S-VLAN Tags and C-VLAN tags scenario, both S-VLAN and C-VLAN tags are treated as the data of a packet when it is transported. When a dual-tagged ARP packet arrives at ingress PE, the device could only recognize either untagged ARP packet or single tag ARP packet, and if it is not, the device assumes that it is not an ARP packet. Since ARP resolution fails, all subsequent communication does not happen. PR1458206
- On QFX5000 platforms dhcp6 security with LDRA option not supported, When ldra is configured, ldra filter from punt packets to the host path conflicts with system default dhcpv6 relay filter. Therefore, packets are not punted to host path. PR1459499
- On QFX10008 or QFX10016 Series platforms, the "forwarding" option for routing-instance type configuration is missed, it might impact the function related to this configuration. For example, FBF does not work because of the missing configuration. PR1460181
- "entPhysicalTable" MIB does not fetch expected data on QFX10002-72Q / 36Q platforms PR1462582

Interfaces and Chassis

- VRRP-V6 state flaps with init and idle states after configuring vlan-tagging. PR1445370

Layer 2 Features

- When QFX5100 is initialized, in rare condition, if storm control is configured on the interface, it might not work as expected. The traffic levels are not monitored and the unknown unicast packets are not dropped. PR1354889
- In ERPS topologies, after failure recovery or reboot, some nodes might not converge to IDLE state and their interfaces might remain in discarding state. This is caused because two STP instances might get created, resulting in two STG groups. Traffic loss might be caused. PR1431262
- On QFX5100 or QFX5110 or EX4600 platforms, if copper base SFP-T is used, the MAC/ARP learning might not work for the SFP-T. PR1437577
- On QFX5100 or QFX5110 or QFX5120 or QFX5200 or QFX5210 Series platforms with load-balance configuration, the uneven traffic distribution might be seen on the link aggregation group (LAG) interfaces. PR1455161

Platform and Infrastructure

- When several continuous HTTP requests received through REST API, the REST service might get non-responsive. PR1449987

Routing Protocols

- Due to Bad Chip ID, fxpc core file generation is encountered once during reboot of device, later it recovers by itself with no other issues. PR1432023
- When applying a firewall filter, which has a modifier to change the DSCP value of a packet, to an IRB interface, the action modifier has no effect. PR1441444
- With **protocol igmp-snooping** configured, if some receiver joins/leaves a group, few seconds of traffic drop might be seen on the existing receivers. PR1457228
- On edge-routed bridging (ERB) EVPN-VXLAN multihoming designs with QFX5110 and QFX5120 switches work as Layer 3 gateways, in some rare condition, when one of the switches acting as L3 gateway comes up after reboot, the egress interface in Packet Forwarding Engine for some end-host might not be updated to the correct next-hop interface in the hardware on that gateway. This issue causes traffic disruption for the affected end host PR1460688
- When deleting IRB on QFX-5110, IRB does not get removed from Packet Forwarding Engine and will blackhole traffic to irb mac address. PR1463092

User Interface and Configuration

- QFX5100 might be unable to commit baseline configuration after zeroize {master:0}[edit] root# commit check Mar 26 05:50:48 mustd: UI_FILE_OPERATION_FAILED: File /var/run/db/enable-process.data doesn't exist Mar 26 05:50:48 mgd[1938]: UI_FILE_OPERATION_FAILED: Failed to open /var/run/db/enable-process.data+ file error: Failed to open /var/run/db/enable-process.data+ file error: configuration check-out failed: daemon file propagation failed. PR1426341

Resolved Issues: 19.3R1

Class of Service (CoS)

- On QFX10008, FPC0 generates a core file after running the Packet Forward Engine command **show cos sched-usage**. PR1449645

EVPN

- The rpd process might crash with EVPN type-3 route churn. PR1394803
- Multicast MAC address might be learned in the Ethernet switching table on QFX5000 or QFX10000 platforms with EVPN-VXLAN configured. PR1420764
- The device might proxy the ARP probe packets in an EVPN environment. PR1427109
- Asynchronous between ARP table and Ethernet switching table happens if EVPN ESI link flap multiple times. PR1435306
- Configuring ESI on a single-homed 25G port might not work. PR1438227
- MAC and IP addresses routes are not consistent. PR1441464
- A change in VLAN configuration is seen when l2ald restarted resulting in kernel sync issues due and impact forwarding. PR1450832
- When there is a VxLAN with a vlan-id of 2 on a QFX5100, ARP will not get resolved. PR1453865

General Routing

- [SIRT]Certain QFX and EX Series devices are vulnerable to 'Etherleak' memory disclosure in Ethernet padding data (CVE-2017-2304) PR1063645
- The 1G copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms PR1286709
- Port LEDs do not work on QFX5100 in QFX5110-QFX5100 mixed mode virtual chassis PR1317750
- QFX10002-60C: Commit should deny when mixed L2 and L3/L4 match conditions are configured on a L2 filter. PR1326715
- When powering off an individual FPC the other FPC PFE might go offline too PR1344395
- Interface flap 100GBASE-LR4 seen during a unified ISSU. PR1353415
- QFX5120/EX4650 : Convergence delay between PE1 and P router link is more than expected delay value. PR1364244

- Traffic spikes generated by IPFIX might be seen on QFX10002 PR1365864
- Error logs seen when channelization is deleted in AS7816-64X product. [Err] 0:_pm4x25_line_side_phymod_interface_get: [Wed Jun 13 08:22:45.845 LOG: Err] ERROR: u=0 p=81 interface type 16 not supported by internal SERDES for this speed 50000 PR1366137
- The backup member switch might fail to become the master switch after switchover on QFX5100/QFX5200/EX4600 Virtual Chassis platform PR1372521
- New CLI knob to enable copying of Open vSwitch Database (OVSDB) to RAM on Virtual Chassis backup RE instead of SSD PR1382522
- Static default route with next-table inet.0 does not work PR1383419
- CLI "show chassis errors active detail" not supported for QFK5K platforms. PR1386255
- QFX5110 - Fan LED turns Amber randomly PR1398349
- CPU Interrupt process high due to intr{swi4: clock (0)} on qfx5100-48t-6q running a "QFX 5e Series" image and 18.x code PR1398632
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay PR1399683
- QSFP-100GBASE-SR4/LR4 might take a long time to come up after disabling interface or reboot PR1402127
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface PR1403528
- Executing command "request system configuration rescue save" may fail with error messages PR1405189
- DHCP Not working for some clients in dual AD fusion setup on EP ports. PR1405495
- Ping over loopback might not work over TYPE 5 tunnel on QFX10000 platforms PR1405786
- QFX5120 : In VxLAN-EVPN configuration , transition from collapsed to non-collapsed L2/L3 GW and vice versa needs switch reload PR1405956
- QFX5200/5100 might not be able to send out control plane traffic to the peering device PR1406242
- QFX10002 showing error fpc0 prds_ptc_clear_all_pulse_and_samples: prds_ptc_clear_all_pulse_and_samples PE 4 PTC 2: after clearing sample, sample still valid 1 PR1407095
- After upgrading junos to 18.1R2 QFX10k send packet without innner vlan tag PR1407347
- MAC address movement might not happen in Flexible Ethernet Services mode when family inet/inet6 and vlan-bridge are configured on the same ifd PR1408230
- Fan failure alarms might be seen on QFX5100-96S after upgrade to 17.3R1 PR1408380
- LLDP memory leak when ieee dcbx packet is received in auto-neg mode followed by another dcbx packet with none of ieee_dcbx tlvs present. PR1410239
- EX2300-24P,error message: dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family PR1410717

- Storm control not shutting down mc-ae interface PR1411338
- FPC CPU may not be displayed correctly PR1412314
- Junos PCC may reject PCUpdate/PCCreate message if there is metric type other than type 2 PR1412659
- QFX5K : Intermittently chassis alarms not raised after power-cycle of the device PR1413981
- QFX5K: EVPN / VxLAN: Mutlicast NH limit is 4K PR1414213
- VC Ports using DAC may not establish link on QFX5200 PR1414492
- Two instances of Junos are running after Junos upgrade to 18.1R3-S3.7 PR1416585
- Mac learning might not happen on trunk mode interface in EVPN/MPLS scenario PR1416987
- Traffic loss might be seen on the ae interface on QFX10000 platforms PR1418396
- Traffic loss might be seen after NSSU operation PR1418889
- Rebooting QFX5200-48Y using "request system reboot" doesn't take physical links offline immediately PR1419465
- libvirtMib_suba core seen during installation PR1419536
- The 100G PSM4 optics connected ports go down randomly during the repeated power cycle PR1419826
- Ping fails over Type-5 tunnel on IRB interfaces under EVPN-VXLAN scenario PR1420785
- An interface may go to downstate on QFX10000/PTX10000 platform PR1421075
- QFX5120-32C: DHCP binding on client might fail when QFX5120-32C acting as DHCP server, this is seen only for channelized port PR1421110
- BFD might stuck in slow mode on QFX10002/QFX10008/QFX100016 platform PR1422789
- QFX5100-48T 10G interface might be auto-negotiated at 1G speed instead of 10G PR1422958
- The interface can not get up when the remote-connected interface only supports 100M in QFX5100 VC setup PR1423171
- IPv6 multicast traffic received on one VC member might be dropped when egressing on other VC member if MLD snooping is enabled. PR1423310
- ON QFX5120-32C , BUM traffic coming over IRB underlay interface gets dropped on destination vtep in PIM based VxLAN. PR1423705
- Traffic is dropped after FPC reboot with AE member links deactivated by remote device. PR1423707
- The J-Flow export might fail when channelization is configured on FPC QFX10000-30C. PR1423761
- A ping over EVPN type-5 route to QFX10000 does not work. PR1423928
- All interfaces will be down and the dcpfe will get crash if SFP-T is inserted on QFX5210. PR1424090
- IPv6 communication issue might be seen after passing through QFX10002-60C platforms. PR1424244
- QFX5120 QSFP-100G-PSM4 become undetected and come back up as channelized interfaces PR1424647

- All interfaces creation failed after NSSU. PR1425716
- The host-bound traffic might be dropped after performing change configuration related to prefix-list. PR1426539
- QFX5210: Received LLDP frames on em0 not displaying in LLDP neighbor output. PR1426753
- Heap memory leak might be seen on QFX10000 platforms. PR1427090
- CRC errors can be seen when other manufacturer device is connected to QFX10000 with QSFP-100GBASE-LR4-T2 optics. PR1427093
- Rebooting or stopping Virtual Chassis member might cause 30 seconds down on RTG link. PR1427500
- QFX5100-VCF rollback for uncommitted configuration takes 1 hour. PR1427632
- The dcpfe process might crash and restart in MC-LAG scenario when the ARP/NDP next-hop is changed. PR1427994
- Interface with optic "QSFP-100GBASE-ER4L" is not coming up in Junos Is Release 18.3R1-S2.1. PR1428113
- Licenses used flag for ovsdb on **show system license** will not be flagged even though ovsdb is configured and working. PR1428207
- Incorrect display of MAC/MAC+IP and count values, after setting **global-mac-limit** and **global-mac-ip-limit**. PR1428572
- Show chassis environment shows Input0 and Input1. PR1428690
- L2ALD generates a core file when number of VXLAN HW IFBDS exceeds the maximum limit of 16382. PR1428936
- On QFX10008, after Routing Engine switchover, led status is not set for missing fan tray. PR1429309
- When forward-only is set within dhcp-reply, dhcp declines are not forwarded to server. PR1429456
- DHCP-relay might not work in an EVPN-VxLAN scenario. PR1429506
- DHCP-relay might not work in an EVPN-VxLAN scenario. PR1429536
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs. PR1429821
- Interface on QFX does not come up after the transceiver is replaced with one having different speed. PR1430115
- In a collapsed VGA4 script ping on shared ESI R6 to R7 IRB address is failing. PR1430327
- The firewall filters might not be attached on the interfaces after doing some changes. PR1430385
- Traffic impact might be seen on QFX10000 platforms with **interface hold-down timer** configured. PR1430722
- On QFX Series platforms the validation of meta data files failed on hypervisor. PR1431111

- **SIB Link Error** detected on a specific Packet Forwarding Engine might cause complete service impact. PR1431592
- The dcpfe might crash on all line cards on QFX10000 in scaled setup. PR1431735
- All ingress traffic might be dropped on 100m fixed speed port with no-auto-negotiation enabled. PR1431885
- The optical power of interface may gradually reduce the optical power for almost 3 mins after issuing **request system reboot at now** on QFX5110 and QFX5120. PR1431900
- L2 traffic drop on QFX10000 with interface MTU lower than 270 bytes. PR1431902
- Outer VLAN tag may not be pushed in the egress VXLAN traffic towards the host for QinQ scenario PR1432703
- L3 filters applied to PVLAN IRB interface may not work after ISSU PR1434941
- SIB/FPC Link Error alarms might be observed on QFX10K due to a single CRC PR1435705
- The mc-ae interface may get stuck in waiting state in dual mc-ae scenario PR1435874
- DHCP discover packets sent to IP addresses in the same subnet as irb interface cause the QFX5110 to send bogus traffic out of dhcp-snooping enabled interfaces PR1436436
- Unknown SNMP trap (1.3.6.1.4.1.2636.3.69.1.0.0.1) sent on QFX5110 restart PR1436968
- QFX5110, QFX5200, QFX5210 There is no jnxFruOK SNMP trap message when only the Power cable is disconnected and connected back. PR1437709
- The DHCP Snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it PR1438351
- Interfaces configured with flexible-vlan-tagging might loss connectivity PR1439073
- DHCPv6 relay binding is not up while verifying the DHCP Snooping along with DHCPv6 Relay PR1439844
- Traffic drop seen on disable/enable MC-LAG. PR1440732
- From interface match condition with IRB over AE interface not working. PR1441230
- QFX5110 - L2 & L3 IFL on IFD - flexible-ethernet-services - VXLAN passing over L2 ifd breaks, L3 P2P communication. PR1441690
- The interface's operational status in HW and SW might be out of sync in EVPN setup with arp-proxy feature enabled. PR1442310
- Flow control does not work as expected on 100G interface of QFX5110. PR1442522
- The PMTUD might not work for both IPv4 and IPv6 if the ingress L3 interface is an IRB. PR1442587
- DHCPv6 Client might fail to get an IP address. PR1442867
- When a line-card is rebooted, the MC-LAG might not get programmed after the line-card comes back online. PR1444100
- QFX5200 : Observing "DCBCM[bcore_init]: ioctl call failed ret:0" failure message when changing UFT profile in FPC logs PR1445855

- On QFX10008 traffic impact might be seen when the JSRV interface is used. PR1445939
- Traffic Discarded for only specified VLAN in IPACL_VXLAN filters PR1446489
- Long IPv6 address are not displayed fully on ipv6 neighbor table. PR1447115
- Unicast arp requests are not replied with no-arp-trap option. PR1448071
- Rebooting QFX5120-48Y using "request system reboot" doesn't take physical links offline immediately PR1448102
- On QFX5120, the incoming layer 3 encapsulated packets are dropped on L3VPN MPLS PE-CE interface. PR1451032
- vgd core file might be generated on any platforms supporting OVSDB. PR1452149
- DHCP offer packet with unicast flag set gets dropped by 10k in a vxlan multi-homed (ESI) setup using anycast IP PR1452870
- QFX10002-60c: EVPN-VXLAN: MAC+IP Count is shown as Zero PR1454603

Interfaces and Chassis

- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. PR1402606
- The logical interfaces in EVPN routing instances might flap after committing configurations. PR1425339
- An ARP entry is not learned at one of mc-lag device at QFX10000. PR1449806

Layer 2 Ethernet Services

- LACP PDU might be looped towards peer MC-LAG nodes. PR1379022

Layer 2 Features

- On QFX Series switches the error message **Failed with error (-7) while deleting the trunk 1 on the device 0**. PR1393276
- QinQ might be malfunctioning if **vlan-id-lists** are configured. PR1395312
- On all QFX5000, symmetric hashing can be done with the hashing options Broadcom provides, though it cannot be enabled and stored in the Junos OS configuration. PR1397229
- On QFX Series EVPN-VXLAN, the unicast IPv6 NS message gets flooded on L3GW. Both IPv4 and IPv6 traffic gets dropped on L2SW. PR1405814
- **IGMP-snooping** on EVPN-VXLAN might impact OSPF hello packets flooding after VTEP leaf reboot. PR1406502
- QFX5110 VC generates DDOS messages of different protocols on inserting a 1G/10G SFP or forming VCP connection. PR1410649
- Stale entries might be observed in a layer 3 VXLAN gateway scenario. PR1423368

- The FXPC might continually crash when firewall filter is applied on a logical unit of a DSC interface. PR1428350
- JTASK and multimove depth failed errors seen after HALT. PR1434687
- Transit DHCPv6 packets might be dropped on QFX5100 and QFX5200 platforms. PR1436415
- QFX5000 switches not properly hashing MPLS transit traffic from VXLAN to L2 LAG. PR1448488

MPLS

- Traffic loss might be observed after changing the configuration **protocols mpls** in ldp-tunneling scenario. PR1428081
- In QFX5110, the l2circuit traffic might be silently dropped or discarded at EVPN SPINE/MPLS LSP TRANSIT device if VXLAN access interface flaps on remote PE node. PR1435504
- Packet loss is seen with **ECMP resilient-hash** enabled on QFX Series platforms. PR1442033

Routing Protocols

- Some storm control error logs might be seen on QFX Series platforms. PR1355607
- Host destined packets with filter log action might not reach to the Routing Engine if log/syslog is enabled. PR1379718
- The IRB transit traffic might not be counted for EVPN/VXLAN traffic. PR1383680
- AUTONEG errors and flush operation failed error, seen after power cycle of the device. PR1394866
- On QFX5110, the firewall filter applied on VxLAN mapped VLAN is not supported on EVPN-VXLAN scenario. PR1398237
- The same traffic flow might be forwarded to different ECMP next hops on QFX5000 platforms. PR1422324
- The traffic loss might start after deleting IRB logical interface. PR1424284
- The rpd process generates a core file due to improper handling of Graceful Restart stale routes. PR1427987
- BGP statement **multipath multiple-as** does not work in specific scenario. PR1430899
- BGP session might go into down status once the traffic flow starts. PR1431259
- Ping fails over Type-5 tunnel on IRB interfaces under EVPN-VXLAN scenario. PR1433918
- The IPv4 fragmented packets might be broken if PTP transparent clock is configured. PR1437943
- The bandwidth value of the DDOS-protection might cause the packets loss after the device reboot. PR1440847
- One of the downstream interfaces flapped and the traffic through xe-2/0/38 broken interface. PR1441402
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. PR1443507
- QFX5110 MCLAG: L2_L3_INTF_OPS_ERROR messages seen after node reboot. PR1435314

- PIM (S,G) joins can cause MSDP to incorrectly announce source active messages in some cases. PR1443713
- The QFX5120 might drop the tunnel encapsulated packets if it acts as a transit device. PR1447128
- Loopback address exported into other VRF instance might not work on QFX Series platforms. PR1449410
- MPLS LDP might still use stale MAC of the neighbor even the LDP neighbor's MAC changes. PR1451217
- Few seconds of traffic drop might be seen towards the existing receivers when another receiver joins/leaves. PR1457228

User Interface and Configuration

- QFX5100 were unable to commit baseline configuration after zeroize. PR1426341

SEE ALSO

What's New | **224**

What's Changed | **233**

Known Limitations | **237**

Open Issues | **240**

Documentation Updates | **256**

Migration, Upgrade, and Downgrade Instructions | **257**

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for the QFX Series.

SEE ALSO

What's New | **224**

What's Changed | **233**

Known Limitations | **237**

Open Issues | **240**

Resolved Issues | **244**

Migration, Upgrade, and Downgrade Instructions | **257**

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 **19.3** 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 19.2 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-x86-64-19.3-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 19.3**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-x86-64-19.3R1.9.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-x86-64-19.3R1.9.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-x86-64-19.3R1.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-x86-64-19.3R1.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.

> **CAUTION:** 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-x86-64-19.3R1.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-x86-64-19.3R1.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 268
- Upgrading the Software Using Unified ISSU on page 268

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-x86-64-19.3R1.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

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases

provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html.

SEE ALSO

- What's New | 224
- What's Changed | 233
- Known Limitations | 237
- Open Issues | 240
- Resolved Issues | 244
- Documentation Updates | 256

Junos OS Release Notes for SRX Series

IN THIS SECTION

These release notes accompany Junos OS Release 19.3R2 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.

Release 19.3R2 New and Changed Features

There are no new features in Junos OS Release 19.3R2 for the SRX Series devices.

Release 19.3R1 New and Changed Features

Application Security

- **Operational commands for SSL sessions (SRX Series and vSRX)**—In Junos OS Release 19.3R1, we've introduced new operational mode CLI commands to monitor and troubleshoot SSL-related issues.

  You can use the new **show** commands to display information and statistics related to SSL configurations, sessions, counters, and logs. You can also use the output of the CLI commands to understand the issue and plan the required next steps accordingly.

  [See Troubleshooting SSL Proxy.]

- **DSCP support in APBR rule (SRX Series and vSRX)**—Starting in Junos OS Release 19.3R1, you can use a differentiated services Code Point (DSCP) value in an APBR rule as a match criteria to perform advanced policy-based routing. You can configure the DSCP value in addition to the other matching criteria of the APBR rule such as dynamic application and dynamic application group.

  By configuring the DSCP value in an APBR rule, you can extend the APBR service to the encrypted traffic or to the traffic with the DSCP markings.

[See Advanced Policy-Based Routing.]

- **User-defined ICAP request header extension (SRX Series)**—Starting in Junos OS Release 19.3R1, Internet Content Adaptation Protocol (ICAP) redirect adds **X-Client-IP**, **X-Server-IP**, **X-Authenticated-User**, and **X-Authenticated-Groups** header extensions in an ICAP message to provide information about the source of the encapsulated HTTP message.

  [See ICAP Service Redirect.]

Chassis Clustering

- **Dedicated fabric ports support (SRX4600)**—Starting in Junos OS Release 19.3R1, you can use the built-in dedicated fabric ports as fabric link ports in chassis cluster mode.

  [See Understanding Chassis Cluster Slot Numbering and Physical Port and Logical Interface Naming, SRX Series Chassis Cluster Configuration Overview, and Chassis Cluster Control Plane Interfaces.]

Flow-Based and Packet-Based Processing

- **Express Path (SRX4600)**—Starting in Junos OS Release 19.3R1, SRX4600 devices support Express Path (formerly known as services offloading) functionality. The Express path support is already available on SRX5000 line devices.

  Express Path considerably reduces packet-processing latency.

  [See Express Path]

General Packet Radio Switching (GPRS)

- **Validate IP address in GTP messages to prevent security threats (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—In Junos OS Release 19.3R1, we've aligned with the GSMA FS.20 standards, which enables you to configure IP addresses in an IP group list. You can prevent a variety of attacks by validating the IP addresses of incoming and outgoing packets in GTP messages against the IP addresses configured in the IP group list.

  [See Understand Validation of IP Address in GTP Messages.]

Hardware

- Starting with Junos OS Release 19.3R1, the following hardware is available to enhance the performance and scalability of the SRX5000 line of devices:
  - SRX5K-IOC4-10G (IOC4): SRX5K-IOC4-10G is a fourth-generation fixed-configuration I/O card with two Packet Forwarding Engines that provide 400 Gbps line rate with 40x10GbE interfaces.
  - SRX5K-IOC4-MRAT (IOC4): SRX5K-IOC4-MRAT is a fourth-generation fixed-configuration I/O card with two Packet Forwarding Engines that provide 480 Gbps (240 Gbps per PFE) line rate with 48x10GbE, 12x40GbE, or 4x100GbE interface options.
  - SRX5K-SCB4 (SCB4): The SCB4 is an enhanced Switch Control Board that provides improved fabric performance and bandwidth capabilities for high-capacity line cards using the ZF-based switch fabric.

The SCB4 enables 480 Gbps throughput per SCB and can be configured with intra chassis and inter chassis redundancy.

- SRX5K-RE3-128G (RE3): The RE3 for the SRX5000 line is, based on the Intel Haswell-EP CPU with six core processors running at 2.0 GHz and 128 GB of DDR4 memory. It provides increased control plane performance and scalability along with virtualization features in SRX5000 line chassis.

For more information about the new hardware support and interoperability, see Cards Supported on SRX5400, SRX5600, and SRX5800 Services Gateways.

J-Web

- **Support for line cards (SRX5000 line of devices)**—Starting in Junos OS Release 19.3R1, J-Web supports IOC4 and RE3 line cards for the SRX5000 line of devices and SCB4 line cards for SRX5600 and SRX5800 devices.

  [See Dashboard Overview, Monitor Ports, and About the Ports Page.]

- **New J-Web Launch Pad (SRX Series)**—Starting in Junos OS Release 19.3R1, after you successfully log in to the J-Web user interface, the J-Web launch pad appears. The launch pad provides a quick view of system identification details, active users, and interface status.

  [See Explore J-Web.]

- **Improved Setup wizard (SRX Series)**—Starting in Junos OS Release 19.3R1, you can configure device and users, time and DNS servers, management interface, zones and interfaces, and security policies using the Setup wizard in the factory default settings to get a fully functional device. If you do not want to perform the initial configuration, you can click **Skip** in the Setup wizard. You can then select Configure > Setup Wizard on the J-Web menu and perform the initial configuration.

  [See Start J-Web and Configure Setup Wizard.]

- **Simplified Juniper Sky ATP enrollment process (SRX Series)**—Starting in Junos OS Release 19.3R1, you can enroll your device to Juniper Sky ATP directly through J-Web. You no longer need to switch between the Juniper Sky ATP portal and J-Web to fetch the enrollment URL and new registrations.

  [See Enroll Your Device with Juniper Sky ATP.]

- **Improved Dashboard widget categories (SRX Series)**—Starting in Junos OS Release 19.3R1, you can choose any one of the following categories in the J-Web dashboard to view supported widgets on your device:

  - All Widgets
  - Applications
  - Devices
  - Security

The dashlet data is refreshed every minute by default. You cannot manually configure the refresh interval of the dashlet. If the data is not aged in the cache, data loads from the cache during the dashlet refresh. If the data is aged, it is retrieved from the device during the next refresh interval cycle.

[See Dashboard Overview.]

- **UTM enhancements (SRX Series)**—Starting in Junos OS Release 19.3R1, the following UTM (Configure > Security Services > UTM) pages are refreshed for a seamless experience:
  - Web Filtering
  - Category Update
  - Antispam Profiles
  - Custom Objects

  [See About the Web Filtering Page, About the Category Update Page, About the Antispam Page, and About the Custom Objects Page.]

Logical Systems and Tenant Systems

- **Secure wire support for user logical system (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)**—Junos OS Release 19.3R1 extends support for secure wire (on root logical systems) to user logical systems. You can forward traffic that arrives on a specific interface to another interface without modifying any received frames on the user logical systems.

  [See Secure Wire for Logical Systems.]

- **User firewall support in customized mode for logical systems and tenant systems (SRX Series)**—Starting in Junos OS Release 19.3R1, a customized model through integrated Juniper Identity Management Service (JIMS) with active mode improves the user firewall authentication process. In this model, the logical system and tenant system extract the authentication entries from JIMS servers configured at the root level based on the logical system and tenant system names.

  [See Understanding Integrated User Firewall support in a Logical System, and Firewall Authentication for Tenant Systems.]

- **Application quality of services support for logical systems and tenant systems (SRX Series)**—Starting in Junos OS Release 19.3R1, logical systems and tenant systems support application quality of services (AppQoS). You can configure a default AppQoS rule set to manage conflicts in the logical systems or tenant systems if multiple security policies match the traffic.

  [See AppQoS for Logical Systems, and AppQoS for Tenant Systems.]

Network Address Translation (NAT)

- **Support for NAT features in PMI mode (SRX5000 devices with SRX5K-SPC3 card, SRX4200, SRX4100, and vSRX)**—Starting in Junos OS Release 19.3R1, you can configure all NAT features in PowerMode IPsec (PMI) mode. Configuration and operational commands for NAT remain the same for both PMI and regular mode. You can configure source NAT, destination NAT, and static NAT for both IPv4 and IPv6 traffic in PMI mode. NAT64 is not supported in PMI mode. However, NAT64 works properly in normal mode, when PMI is enabled.

  See [Introduction to NAT and Improving IPsec Performance with PowerMode IPsec.]

Network Management and Monitoring

- **Improved on-box reporting performance (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Release 19.3R1, Junos OS stores logs in multiple tables instead of a single table in a database file. Each table contains the timestamp of the oldest and latest logs. When you initiate a query based on the start and end time, the local log management daemon (llmd process) finds the latest table to generate reports.

  [See Understanding On-Box Logging and Reporting.]

- **Packet capture from operational mode (SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.3R1, you can capture packets from operational mode without committing the configurations. You can define the packet filter to trace only a certain type of traffic, such as logical interface, protocol, source IP address prefix, source port, destination IP address prefix, and destination port. In addition, you can modify the filename, file type, file size, and capture size of the packet capture output.

  [See Packet Capture from Operational Mode.]

Platform and Infrastructure

- **New SCB, IOC, and Routing Engine improve performance and scalability (SRX5400, SRX5600, SRX5800)**—In Junos OS Release 19.3R1, we've introduced the following new hardware to enhance the performance and scalability of our SRX5000 line of devices:

  - **Switch Control Board SCB4 (model number: SRX5K-SCB4)**—Supports high traffic capacity, and provides greater link speeds, fabric capacity, and improved services. The SCB4 is supported only on SRX5600 and SRX5800 devices.
  - **I/O card IOC4 (model numbers: SRX5K-IOC4-MRAT and SRX5K-IOC4-10G)**—Enhances processing speed, provides line rates of up to 480 Gbps per slot, and supports Media Access Control Security (MACsec).
  - **Routing Engine (model number: SRX5K-RE3-128G)**—Supports higher CPU speed, 128-GB RAM, a trusted platform module (TPM), and increased processing capacity.

  The IOC4 can interoperate with the SCB3, SCB4, SPC2, SPC3, IOC2, IOC3, IOC4, and the Routing Engines SRX5K-RE-1800X4 and SRX5K-RE3-128G. However:

  - The SCB4 can interoperate with all of these components except the SCB3.

- The Routing Engine SRX5K-RE3-128G can interoperate with all of these components except the SRX5K-RE-1800X4.

You cannot use any of these components in a chassis that has the Switch Control Board SCB2 installed. For more information about the new hardware interoperability, see Cards Supported on SRX5400, SRX5600, and SRX5800 Services Gateways.

With the new hardware installed, the SRX5000 line of devices support the firewall and advanced security services—such as application security, unified threat management (UTM), intrusion prevention system (IPS)—and all other software features that they supported before this release, except the following:

- Layer 2 Ethernet switching mode
- Port mirroring

For the complete list of features supported on the SRX5000 line of devices, see Feature Explorer.

[See Chassis Cluster Control Plane Interfaces and show chassis hardware (View).]

Routing Protocols

- **Support for nondefault routing instance for outbound SSH (MX Series and SRX Series)**—Starting in Junos OS Release 19.3R1, you can specify the name of the routing instance on which the outbound SSH connectivity needs to be established using the **routing-instance** statement at the **[edit system services outbound-ssh]** hierarchy level. If you do not specify a routing instance, your device will establish the outbound SSH connection using the default routing table.

  [See outbound-ssh, Configuring Outbound SSH Service.]

Security

- **High Availability (HA) synchronization of address name resolving cache (SRX Series and vSRX)**—Starting in Junos OS Release 19.3R1, the policy DNS cache memory is synchronized into a single local DNS cache file on the HA active node and is copied to the HA backup node. This process suppresses Domain Name System (DNS) queries and responses during Network Security Process (NSD) restart. In releases before Junos OS Release 19.3R1, a few system resources become a bottleneck when a large number of DNS queries and responses are sent and received at the same time. During this period, security policies use empty source and destination addresses. Therefore, the new pass-through traffic is blocked as no policy can be matched, and flow sessions cannot be established.

  [See High Availability (HA) Synchronization of Address Name Resolving Cache.]

- **Support for bundle feeds in dynamic address groups (SRX Series and vSRX)**—Starting in Junos OS Release 19.3R1, you can configure bundle feeds for dynamic address groups in a security policy. You can download a single **.tgz** file from the server and extract it into multiple child feed files. Each individual file corresponds to one feed. Individual dynamic-addresses reference the feed inside the bundle file.

  You can update IP addresses, IP prefixes, or IP ranges contained in a dynamic address entry periodically by downloading an external feed. SRX Series devices periodically initiate a connection to the feed server to download and update the IP lists that contain the updated dynamic addresses.

You can configure the **url** option for the feed server by using the **set security dynamic-address feed-server *feed-server-name*** at the **[edit]** hierarchy level.

[See Dynamic Address Groups in Security Policies.]

Juniper Sky ATP

- **Juniper Sky ATP block files with unknown verdict and send user notification**—Starting in Junos OS Release 19.3, for advanced anti-malware policies, you can block a file when the verdict is **unknown**. You can also send a user notification when a file is blocked. We've introduced the following new commands: **set services advanced-anti-malware policy p1 http file-verdict-unknown (block|permit)** and **set services advanced-anti-malware policy p1 http client-notify (message|file|redirect-URL)**.

  See set services anti-malware policy and request services advanced-anti-malware redirect-file.

- **Juniper Sky ATP onboarding changes**—Starting in Junos OS Release 19.3, you can use an alternative onboarding procedure to perform all enrollment steps using the CLI on the SRX Series device without having to access the Sky ATP Web Portal. Run the **request services advanced-anti-malware enroll** command on the SRX Series device to begin the process. Both the original enrollment process that obtains an op script from the Web Portal and the new CLI-only enroll process are valid procedures. Use either one.

  See Enroll the SRX Series Device using the Enroll Command.

Subscriber Management and Services

- **Diameter S6a authentication (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 19.3R1, you can configure the diameter-based authentication S6a application on SRX series devices at [edit access] hierarchy. The MME uses S6a application to retrieve authentication information from Home Subscriber Server (HSS).

  [See Configuring S6a and s6a.]

Virtual Routing

- **VRF-Group in L3VPN traffic (SRX Series and vSRX)**—Starting in Junos OS Release 19.3R1, to support mid-stream routing, VRF undergoes changes for processing a session among a group of MPLS VRF instances in an L3VPN MPLS network. These VRF instances which are logically part of a given L3VPN traffic are grouped and this is a VRF-Group. The VRF-Groups allows the session to switch from one MPLS VRF to another MPLS VRF

  VRF-Group supports the following features:

  - Overlapping in VPN session
  - VRF-Group Policy
  - VRF-Group NAT
  - VRF-Group ALG

[See Security Policy for Controlling Traffic for VRF Routing-Instance

SEE ALSO

What's Changed

IN THIS SECTION

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

What's Changed in Release 19.3R2-S6

Network Management and Monitoring

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

  See ssh (NETCONF).

- **Changes to commit RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—When you configure the **rfc-compliant** statement at the **edit system services netconf** hierarchy level, the NETCONF server's response for **commit** operations includes the following changes:
  - If a successful **commit** operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the **<source-daemon>** element as a child of the **<error-info>** element instead of the **<rpc-error>** element.
  - If you also configure the **flatten-commit-results** statement at the **edit system services netconf** hierarchy level, the NETCONF server suppresses any **<commit-results>** XML subtree in the response and only emits an **<ok>** or **<rpc-error>** element.

  See Configuring RFC-Compliant NETCONF Sessions.

Release 19.3R2 Changes in Behavior and Syntax

Interfaces and Chassis

- **Change in output of show interfaces (SRX300, SRX320, SRX340, SRX345, SRX550M)**—Starting in Junos OS Release 19.3R2, the output of the **show interfaces** command on the SRX300 line of devices and on the SRX550M, no longer displays **vlan** as the value of the **Physical interface** field. On these devices, the value of the **Physical interface** field in the command output appears as **irb** instead of **vlan**.

Network Management and Monitoring

- **Change in jnxJsFlowMIB statistics display (SRX Series)**—Starting in Junos OS Release 19.3R2, in a chassis cluster, you can see the statistics on all SPUs of both nodes using the **show snmp mib walk jnxJsFlowMIB** command. In earlier releases, you can see the statistics only on local SPUs.

  [See SNMP MIB Explorer.]

Release 19.3R1 Changes in Behavior and Syntax

Application Security

- Starting in Junos OS Release 19.3R1, you can schedule automatic download of the application signature package in a new format. Use the YYYY-MM-DD.hh:mm format to configure the time to automatic download for application signatures. For example, the following statement sets the start time as 10 AM on June 30, 2019:

  ```
  user@host# set services application-identification download automatic start-time 2019-06-30.10:00:00
  ```

You can configure the automatic updates using the new format once you upgrade your previous Junos OS version to any of the above supported Junos OS version.

Authentication and Access Control

- **SSH protocol version v1 option deprecated from CLI (SRX Series)**—Starting in Junos OS Release 19.3R1, we've removed the nonsecure SSH protocol version 1 (**v1**) option from the [**edit system services ssh protocol-version**] hierarchy level. You can use the SSH protocol version 2 (**v2**) as the default option to remotely manage systems and applications. With the **v1** option deprecated, Junos OS is compatible with OpenSSH 7.4 and later versions.

  Junos OS releases earlier than Release 19.3R1, continue to support the **v1** option to remotely manage systems and applications.

  [See protocol-version.]

Junos OS XML API and Scripting

- **Range defined for confirm-timeout value in NETCONF and Junos XML protocol sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, the value for the **<confirm-timeout>** element in the Junos XML protocol **<commit-configuration>** operation must be in the range 1 through 65,535 minutes, and the value for the **<confirm-timeout>** element in the NETCONF **<commit>** operation must be in the range 1 through 4,294,967,295 seconds. In earlier releases, the range is determined by the minimum and maximum value of its unsigned integer data type.

Licensing

- Starting in Junos OS Release 19.3R1, the SNMP OID **jnxLicenseKeys** is deprecated.

  [See Licensing Guide.]

Network Management and Monitoring

- **Default system log messages (SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550M)**—Starting in Junos OS Release 19.3R1, we've changed the default mode for system log messages from event mode to stream mode.

  [See Understanding System Logging for Security Devices and mode (Security Log).]

System Logging

- **Preventing system instability during core file generation (SRX Series)**—Starting with Release 19.3R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

Unified Threat Management (UTM)

- **Support to adjust core allocation ratio of UTM onbox-AV**— Starting in Junos OS Release 19.3R1, to improve the throughput of low scan cost file such as doc file and big exe file, the on-box AV load flavor light ratio is changed from 1/3 to 1/4, and the onbox AV load flavor heavy ratio is changed from 2/3 to 1/2.

  [See Example: Configuring On-Device Antivirus Feature Profile.]

VPN

- **Power Mode IPsec (SRX Series)**—Starting in Junos OS Release 19.3R1, when you enable the Power Mode IPsec, the **show security flow statistic** and **show security flow session tunnel summary** commands does not count, or display the number of packets that are processed within the Power Mode IPsec.

  show security flow statistics

SEE ALSO

Known Limitations

IN THIS SECTION

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

J-Web

- After you generate the Default Trusted CA profile group under Certificate Management>Trusted Certificate Authority in J-Web, J-Web does not display the CA profile group local under Certificate Management>Certificate Authority Group page. PR1424131
- The CA profile group imported using J-Web does not populate the group on the Certificate Authority Group initial landing page grid, but all the CA profiles of a group are populated on the Trusted Certificate Authorities landing page. PR1426682

Logical Systems and Tenant Systems

- In case of logical systems, secure wire cannot work with the user firewall AD integrated solution together, because secure wire cannot support forwarding traffic between different logical systems. The user firewall AD integrated solution cannot probe client PCs, which are located at non-root logical systems. PR1436546

VPNs

- In the HA design for SRX Series devices, the **anti-replay** window is synced to the backup only when the total incoming packet count is an odd multiple of 128 packets. When a failover occurs, the **anti-replay** bitmap is not synchronized. Again, when the node comes back online, the SA is installed but the **anti-replay** bitmap is reset to 0 along with the in and out sequence number. PR1420521
- In a chassis cluster, ESP or AH packet sequence number is not synchronized to the backup node after the backup node is rebooted. PR1433424

SEE ALSO

What's New | **272**

What's Changed | **279**

Open Issues | **284**

Resolved Issues | **286**

Documentation Updates | **296**

Migration, Upgrade, and Downgrade Instructions | **297**

Open Issues

IN THIS SECTION

Learn about open issues in this release for SRX Series. 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

- Auto completion is not working on proxy terminator profile name. PR1424822
- On SRX Series device, MTU is wrongly calculated in a **gr0.0** or **st0.0** interface after a GRE or VPN route change. If the command **do-not-fragment** is configured and the packet is bigger than the MTU, the packet might be dropped. PR1462825
- Tail drop on all ports is observed when any switch side egress port get congested. PR1468430

Intrusion Detection and Prevention (IDP)

- Rogue **.gz** files in **/var/tmp/sec-download/** might fail offline **secpack** update. PR1466283

J-Web

- Due to **set chassis auto-image-upgrade** in factory configuration, from phone home page you are not able to skip to J-Web and get the error **Bootstrap is in progress, Can't Skip!!**. PR1420888
- 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

- Policy rules grid will be blank when you navigate to rules menu after creating shared objects or security profiles. The committed or discarded changes will not be visible in UI. As a workaround, re-click **Rules**, and menu will populate the rules of a policy grid or log out and then log in again to J-Web to view the committed or discarded changes. PR1460210
- When a dynamic application is created for an edited policy rule, the list of services will be blank when the **Services** tab is clicked and then the policy grid will be auto-refreshed. 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

Platform and Infrastructure

- On SRX5400, SRX5600, and SRX5800 devices with SPC3, it is possible that when multiple core files are generated in quick succession, the cold-sync-monitored status is displayed and cannot be removed even though cold-sync has finished. You must reboot the affected node to recover. PR1403000
- Cm errors on certain MPC line cards are classified as major which should be minor or non-fatal. If these errors are generated, it might get projected as a bad hardware condition and therefore trigger Packet and Forwarding Engine disable action. PR1449427

Routing Policy and Firewall Filters

- SSL reverse proxy feature must be used instead of SSL inspection feature because SSL inspection is being deprecated in favor of SSL reverse proxy. PR1450900

Routing Protocols

- On SRX Series device, the SSH login from automation tools to the Junos device is not successful when using authentication method password (not 'keyboard-interactive'). If the username is configured both as a local user and also on a remote RADIUS or TACACS server, using the **Juniper-Local-User-Name** attribute pointing to a different local username. PR1454177

VPNs

- On SRX Series devices, in case multiple traffic selectors are configured for a peer with IKEv2 reauthentication, only one traffic selector is rekeyed at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors are cleared without immediate rekeying. New negotiation of these traffic selectors is triggered through other mechanisms such as traffic or by peer. PR1287168
- On SRX5400, SRX5600, and SRX5800 devices, during in-service software upgrade (ISSU), the IPsec tunnels flap, causing a disruption of traffic. The IPsec tunnels recover automatically after the ISSU process is completed. PR1416334

- On the SRX5000 line of devices with an SPC3 cards, 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

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in 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.

Resolved Issues: 19.3R2

Application Layer Gateways (ALGs)

- Sometimes unexpected forwarding sessions appear for tenant ALG SIP traffic in cross tenant. PR1409748
- After Layer 3 HA enable, ALG H.323 group or resource cannot be synchronized to the peer node correctly. PR1456709

Authentication and Access Control

- The same source IP sessions are cleared when the IP entry is removed from UAC table. PR1457570

Application Security

- The AAMW diagnostic script gives incorrect error **Error: Platform does not support SkyATP: srx300**. PR1423378
- Unable to get more than 60 Gbps of AppQoS throughput. PR1439575

Chassis Clustering

- Hardware failure is seen on both nodes in **show chassis cluster status**. PR1452137
- On SRX Series devices with chassis cluster, the control link remains up even though the control link is actually down. PR1452488

Flow-Based and Packet-Based Processing

- Packet loss is caused by FPGA back pressure on SPC3. PR1429899
- VPN traffic fails after primary node reboot or power off. PR1433336
- On SRX4600 device, core file might be observed and SPM might be in present state. PR1436421
- While checking the flow session XML for source NAT under tenant, there is no value identifier for **tenant-name (< tenant>< /tenant>)**. PR1440652
- J-Flow version 5 stops working after changing input rate value. PR1446996
- AAWM policy rules for IMAP traffic sometimes might not get applied when passed through SRX Series device. PR1450904
- FTP data cannot pass through SRX320 4G wireless from FTP server to client. PR1451122
- Traffic forwarding on Q-in-Q port and VLAN tagging is not observed properly on R0. PR1451474
- The rpd process might stop and restart and an rpd core file is generated when committing the configuration. PR1451860
- The **peers** and **peers-synchronize** commands are removed from SRX Series devices. PR1456661
- Added some JP APN settings to default list in LTE mPIM. PR1457838
- Changing the **RESET** configuration button behavior on the SRX1500 does not work. PR1458323
- The **security flow traceoptions** fills in with RTSP ALG related information. PR1458578
- The **security-intelligence** CC feed does not block HTTPS traffic based on SNI. PR1460384
- The AAMWD process exceeds 85 percent RLIMIT_DATA limitation due to memory leak. PR1460619
- Fragmented traffic might get looped between the fab interface in Z mode. PR1465100
- HTTP block message stops working after SNI check for HTTPS session. PR1465626

Interfaces and Chassis

- SCB4 or SCB3 ZF or XF2 fabric plane retraining is needed after switching the fabric redundancy mode. PR1427119
- The fxp0 interface might redirect packet not destined to itself. PR1453154

J-Web

- J-Web fails to display the traffic log in event mode when stream mode host is configured. PR1448541
- Editing destination NAT rule in J-Web introduces a non-configured routing instance field. PR1461599

Network Address Translation (NAT)

- There is port endian issue in SPU messages between SPC3 and SPC2, which results in one redundant NAT binding being created in CP when one binding is allocated in SPC2 SPC. PR1450929

Network Management and Monitoring

- Control links are logically down on SRX Series device chassis cluster running Junos OS Release 12.3X48. PR1458314

Platform and Infrastructure

- The node0 stayed in secondary hold status for a long time but cannot change back to secondary status after manual failover in RG0. PR1421242
- On the SRX300 line of devices, interface LED does not work properly. PR1446035
- REST API process will get non-responsive when a number of requests come at a high rate. PR1449987

Routing Policy and Firewall Filters

- The NSD process might stop due to a memory corruption issue. PR1419983
- The NSD process might get stuck and cause problems. PR1458639

Services Applications

- On SRX Series devices, the **lcore-slave** core files are seen. PR1460035

Unified Threat Management (UTM)

- Blacklist compilation failures are reported. PR1418980
- The command **show security utm web-filtering status** now provides additional context when the status of EWF is **down**. PR1426748

VPNs

- The IKE and IPsec configuration under groups is not supported. PR1405840
- Old tunnel entries are also seen when new tunnel negotiation happens from peer device after change in IKE gateway configuration at peer side. PR1423821

- The P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. PR1432434
- The P1 or P2 SAs are deleted after RG0 failover. PR1433355
- IPsec SA in and out key sequence number update missing with cold-sync (secondary node reboot). PR1433424
- Sequence number is reset to zero while recovering SA after SPC3 or flowd stop or reboot. PR1433568
- The IKED stops on the SRX5000 line of devices with SPC3 when IPsec VPN or IKE is configured. PR1443560
- Sometimes old SAs are not deleted after rekey and the number of IPsec tunnels shows up more than the configured tunnels. PR1449296
- Traffic is not sent out through IPsec VPN after update to Junos OS Release 18.2 or later. PR1461793

Resolved Issues: 19.3R1

Application Layer Gateways (ALGs)

- The TCP reset packet is dropped when any TCP proxy-based feature and the **rst-invalidate-session** command are enabled simultaneously. PR1430685
- The H.323 connection might not be established when the H.323 packet passes SRX devices twice through different virtual routers. PR1436449
- Packet loss happens during cold synchronization from secondary node after rebooting. PR1448252

Application Security

- Automatic application-identification download stops after going over the year and reboot. PR1436265
- The flowd or srxpfe process might crash when advanced anti-malware service is used. PR1437270
- The applications that get declassified in the middle of a session are not identified properly. PR1437816
- The flowd process core files might be seen when the traffic hits AppQoS policy. PR1446080

Authentication and Access Control

- Support redirecting HTTP or HTTPS request to firewall Web authentication server with the server's domain name. PR1421725
- The CPU utilization of the uacd is high, about 100 percent, in the output of **show chassis routing-engine**. PR1424971

Chassis Clustering

- Mixed mode (SPC3 coexisting with SPC2 cards) high availability (HA) IP monitoring fails on the secondary node with **secondary arp entry not found** error. PR1407056
- Memory leaks might be seen on the jsqlsyncd process on SRX Series chassis clusters. PR1424884

- The flowd or srxpfe process might stop when SCCP or MGCP ALG works on SRX Series chassis clusters. PR1426722
- RG0 failover sometimes causes FPC offline/present status. PR1428312

Class of Service (CoS)

- Frequent issuance of the **show class-of-service spu statistics** command cause rtlogd busy. PR1438747

Flow-Based and Packet-Based Processing

- Password recovery menu does not appear on SRX Series device. PR1381653
- Invalid sessions timeout over 48 hours with stress TCP traffics in the backup node. PR1383139
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when PowerMode IPsec is enabled, the **show security flow statistics** and **show security flow session tunnel summary** commands do not count or display the number of packets processed within PowerMode IPsec, because these packets do not go through the regular flow path. PR1403037
- CPU is hitting 100 percent with fragmented traffic. PR1402471
- Throughput or latency performance of TCP traffic is dropped when TCP traffic is passing through from one logical system to another logical system. PR1403727
- While PMI is on, IPsec-encrypted statistics on the Routing Engine **show security ipsec statistics** are not working anymore for fragment packets. PR1411486
- The input and output bytes or bps statistic values might not be identical for the same size of packets. PR1415117
- None of the operational **web-proxy** command have **clear** support. PR1415753
- Force clearing client session from flow does not clean up proxy session. PR1415756
- Juniper Sky ATP does not escape the \ inside the username before the metadata is sent to the cloud. PR1416093
- The TCP session might not get cleared even after it reaches the timeout value. PR1416385
- TCP segmented client-side session fails to create transparent proxied relay session, and session stays idle. PR1417389
- The **show security flow session session-identifier <*sessID*>** command is not working if the session ID is bigger than 10M on SRX4600 platform. PR1423818
- The tunnel ID information is displayed in the flow session. PR1423889
- PIM neighbors might not come up on SRX Series chassis cluster. PR1425884
- When configuring a GRE tunnel (GRE-over-IPsec-tunnel) or an IPsec tunnel on an SRX Series device, the MTU of the tunnel interface is calculated incorrectly. PR1426607
- The IPsec traffic going through the SRX5000 line of devices with SPC2 cards installed causes high SPU CPU utilization. PR1427912

- The flowd process might stop on the SRX5000 line of devices. PR1430804
- SRX550M running Junos OS Release 18.4R1 shows PEM 1 output failure message, whereas with Junos OS Release 15.1X49 or Junos OS Release 18.1R3.3 it does not show any alarms. PR1433577
- Currently PMI doesn't support mirror-filter functionality. If there are any mirror filters configured, PMI flaps all of the traffic to the regular flow path. PR1434583
- Intermittent packets drop might be observed if IPsec is configured. PR1434757
- On SRX series, syslog severity level of **msg subtype is end of policy** is set to **error** although this message can be ignored. PR1435233
- The rtlogd process on the two Routing Engine HA nodes go into deadlock state when rtlogd on both nodes are busy with sending data to each other in the single thread context. PR1435352
- The second IPsec ESP tunnel might not be able to establish between two IPv6 IKE peers. PR1435687
- On an SRX4600 device, core file generation might be observed and SPM might be in present state. PR1436421
- The ipfd process might crash when SecIntel is used. PR1436455
- Packet reorder does not work when sending traffic over IPsec tunnel with session-affinity. PR1436720
- Member of dynamically created VLANs information is not displaying on show VLANs. PR1438153
- Security logs cannot be sent to the external syslog server through TCP. PR1438834
- Decryption traffic doesn’t take PMI path after IPsec rekey (initiated by peer) when loopback interface is configured as external interface. PR1438847
- The wmic process might stop and restart when using user firewall with Active Directory. PR1439538
- The IKE pass-through packet might be dropped after source NATed. PR1440605
- Performance improvements were made to Screens, which benefit multi-socket systems. PR1440677
- SPC2 wrongly forwarded packet to SPC3 core0 and core14. PR1441234
- The configured RPM probe server hardware timestamp does not respond with correct timestamp to the RPM client. PR1441743
- New CLI option to show only userful group infotmations for an Active Directrory user. PR1442567
- The flowd or srxpfe process might crash when processing fragmented packets. PR1443868
- Packet loss happens during cold sync from secondary node after rebooting. PR1447122
- LACP cannot work with the **encapsulation flexible-ethernet-services** configuration. PR1448161
- SPC3 talus FPGA stuck on 0x3D or 0x69 golden version. PR1448722
- FTP data cannot pass through SRX320 4G wireless from FTP server to client. PR1451122
- Traffic forwarding on Q-in-Q port and VLAN tagging is not observed properly on R0. PR1451474

Infrastructure

- Increase in Junos OS image size for Junos OS Release 19.1R1. PR1423139

Interfaces and Routing

- The fxp0 interface might redirect packet not destined itself. PR1453154

Installation and Upgrade

- SRX Series devices go into DB mode after USB installation. PR1390577
- SPMC version mismatch errors after Junos OS install using USB method. PR1437065

Interfaces and Chassis

- Both nodes in the SRX Series chassis cluster go into DB mode after downgrading to Junos OS Release 18.1. PR1407295
- The reth interfaces are now supported when configuring SSL decryption mirroring (**mirror-decrypt-traffic** interface). PR1415352
- Disabling the interface on the primary node causes traffic to get silently dropped through the secondary. PR1424705
- SCB4 or SCB3 ZF or XF2 fabric plane retraining is needed after switching the fabric redundancy mode. PR1427119
- MTU change after a CFM session is up can impact L2 Ethernet ping (loopback messages). If the new change is less than the value in the initial incarnation then L2 Ethernet ping would fail. PR1427589
- LFM remote loopback is not working as expected. PR1428780
- The LACP interface might flap if performing a failover. PR1429712

Intrusion Detection and Prevention (IDP)

- NSD fails to push security zone to the Packet Forwarding Engine after reboot, if there is an active IDP rule configured with FQDN. PR1420787

J-Web

- J-Web configuration change for an address set using the search function results in a commit error. PR1426321
- User unable to view GUI when logged in as read-only user. The user is presented with an empty page after login. PR1428520
- IRB interface is not available in the zone option of J-Web. PR1431428
- Launch pad is not loading in the foreground and not showing details for any widgets. PR1446802
- The idle-timeout for J-Web access doesn't work properly. PR1446990
- J-Web fails to display the traffic log in event mode when stream mode host is configured. PR1448541

Network Address Translation (NAT)

- RTSP resource session is not found during NAT64 static mapping. PR1443222

Network Management and Monitoring

- MIB OID **dot3StatsDuplexStatus** shows wrong status. PR1409979
- Partial traffic might get dropped on an existing LAG. PR1423989
- SNMPD might generate core files after restarting NSD process by **restart network-security gracefully**. PR1443675

Platform and Infrastructure

- Memory leak might occur on the data plane during composite next-hop installation failure. PR1391074
- On SRX4600 device, the 40-Gigabit Ethernet interface might flap continuously by MAC local fault. PR1397012
- The **show security flow session** command fails with error messages when SRX4600 has over a million routing entries. PR1408172
- On PEM 0 or PEM 1 or fan, I2C failure major alarm might be set and cleared multiple times. PR1413758
- Complete device outage might be seen when an SPU VM core file is generated. PR1417252
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10 (based system). PR1417321
- On SRX Series device, the flowd process might stop. PR1417658
- On SRX4600 devices, commit failed while configuring 2047 VLAN IDs on the reth interface. PR1420685
- SPC in slot1 of node0 remained in offline state for more than 1 hour after the cluster was upgraded from Junos OS Release 18.2R2-S1.3 to Junos OS Release 18.2X41.1. PR1423169
- Screen sync cookie causes 100 percent CPU utilization across all SPC3 cards of SRX5800, when packet rate is high. PR1425332
- The ipfd process might crash if the security intelligence feature is configured. PR1425366
- Alarms triggered due to high temperature when operating within expected temperatures. PR1425807
- The PICs might go offline and split-brain might be seen when interrupt storm happens on internal Ethernet interface em0 or em1. PR1429181
- REST API does not work properly. PR1430187
- Uneven distribution of CPU with high PPS on device. PR1430721
- Packet Forwarding Engine crashes might be seen on SRX1500 platform. PR1431380
- The false license alarm may be seen even if there is a valid license. PR1431609
- The kmd log shows resource temporarily unavailable repeatedly and VPNs might be down. PR1434137
- The interface using LACP flaps when the Routing Engine is busy. PR1435955

- CLI giving error as **usp_ipc_client_open: failed to connect to the server after 1 retries(61)** when SRX4100 or SRX4200 has large entries on RIB or FIB. PR1445791
- On the SRX300 line of devices, interface LED does not work properly. PR1446035
- IS-IS adjacencies between the GE link is not up. PR1446533

Routing Policy and Firewall Filters

- Memory leak in nsd causes configuration change to not take effect after a commit. PR1414319
- The flowd process stops on SRX Series devices while deleting a lot of policies from Junos Space. PR1419704
- A commit warning is now presented to the user when a traditional policy is placed below a unified policy. PR1420471
- The dynamic-address summary's IP entry count does not include IP entries in the root logical system. PR1422525
- After a new alarm is created, the NSD process fails to restart because subcomponents fail. PR1422738
- DNS cache entry does not time out from device even after TTL=0. PR1426186
- The ipfd generates a core file while scaling. PR1431861
- An SRX1500 device allows only a maximum of 256 policies with counting enabled. PR1435231
- Two ipfd processes appear in **ps** command and the process pauses. PR1444472

Unified Threat Management (UTM)

- Unable to achieve better Avira antivirus TP on SRX4600 as mbuf high watermark is reached. PR1419064
- When using unified policies, the base filter for certain UTM profiles might not be applied correctly. PR1424633
- The **custom-url-categories** configuration is now pushed correctly to the Packet Forwarding Engine under all circumstances. PR1426189
- Memory issue due to SSL proxy whitelist or whitelist URL category. PR1430277
- Replace the **bypass-on-dns-cache-miss** command with the **drop_on_dns_error** command in the Web proxy profile. If the **drop_on_dns_error** command is not set and DNS failure occurs for a session, that session passes through bypass mode. If the **drop_on_dns_error** command is set and DNS failure occurs for a session, that session is dropped by the Web proxy plug-in. PR1430425
- Adjust core allocation ratio for on-box antivirus. PR1431780

User Interface and Configuration

- Tenant system administrator cannot view its configuration with empty database message when using groups. PR1422036

VPNs

- Tunnel flapping is seen after doing RG0 failover. PR1357402
- With a large number of IPsec tunnels established, a few tunnels may fail during rekey negotiation if the SRX Series device initiates the rekey. PR1389607
- VPN tunnels may flap upon commiting changes in configuration groups on SRX Series devices. PR1390831
- Idle IPsec VPN tunnels without traffic and with ongoing DPD probes can be affected during RG0 failover. PR1405515
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when the SRX Series device is configured to initiate IKEv2 reauthentication when NAT traversal is active, occasionally reauthentication might fail. PR1414193
- The iked process does not handle cases and core files might be generated when a remote gateway address is configured as an IPv6 address while the local interface where the tunnel is anchored has an IPv4 address. PR1416081
- Group VPN IKE security associations cannot be established before RG0 failover. PR1419341
- SSL proxy did not correctly warn users about unsupported certificates. PR1419485
- The iked process might stop when IKE and IPsec SA rekey happens simultaneously. PR1420762
- The 4G network connection might not be established if LTE mPIM card is in use. PR1421418
- Tenant system administrator can change VLAN assignment beyond the allocated tenant system. PR1422058
- The **show security ike sa detail** command shows incorrect values in the IPsec security associations column. PR1423249
- IPsec packet throughput might be impacted if NAT-T is configured and the fragmentation operation of post fragment happens. PR1424937
- On SRX Series devices with SPC3, the device does not send IKE delete notification to the peer if the traffic selector configuration is changed. PR1426714
- The kmd process stops and generates a core file after running the **show security ipsec traffic-selector** command. PR1428029
- In SPC3 and SPC2 mixed mode, IPsec SA is not getting cleared by executing the **clear security ipsec sa** command. PR1428082
- On the SRX5000 line of devices with SPC3, with P2MP and IKEv1 configured, if negotiation fails on the peer device, then multiple IPsec SA entries are created on the device if the peer keeps triggering a new negotiation. PR1432852

- IPsec rekey triggers for when sequence number in AH and ESP packet is about to exhaust is not working. PR1433343
- On SRX Series devices, fragments exit VPN traffic earlier than required by ingress packet sizes. PR1435700
- The IPsec VPN traffic drop might be seen on SRX Series platforms with NAT-T scenario. PR1444730

SEE ALSO

What's New | 272

What's Changed | 279

Known Limitations | 282

Open Issues | 284

Documentation Updates | 296

Migration, Upgrade, and Downgrade Instructions | 297

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R2 documentation for the SRX Series.

SEE ALSO

What's New | 272

What's Changed | 279

Known Limitations | 282

Open Issues | 284

Resolved Issues | 286

Migration, Upgrade, and Downgrade Instructions | 297

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.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths. You can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 15.1X49, 17.3, 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from one Junos OS Release to the next release or one release after the next release. For example you can upgrade from Junos OS Release 15.1X49 to Release 17.3 or 17.4, Junos OS Release 17.4 to Release 18.1 or 18.2, and from Junos OS Release 18.1 to Release 18.2 or 18.3 and so on.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

For information about ISSU, see the Chassis Cluster User Guide for Security Devices.

SEE ALSO

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 is a framework, set of policies, and 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 on the list of supported products, see Juniper Flex Program.

Finding More Information

- **Feature Explorer**—Determine the features supported on MX Series, PTX Series, QFX Series devices. The Juniper Networks Feature Explorer is a Web-based app that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. https://pathfinder.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:

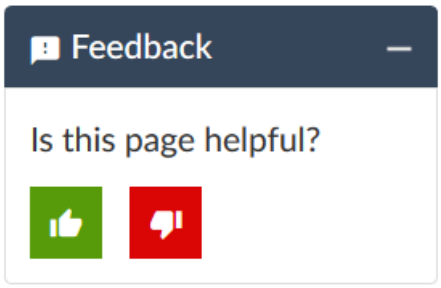

  - Click the thumbs-up icon if the information on the page was helpful to you.

- 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 J-Care or JNASC 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.

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: https://support.juniper.net/support/
- Search for known bugs: https://kb.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://support.juniper.net/support/downloads/

- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://forums.juniper.net
- Open a case online in the CSC Case Management tool: https://www.juniper.net/cm/

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at https://www.juniper.net/cm/.
- 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, visit us at https://support.juniper.net/support/requesting-support/.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to support@juniper.net. For documentation issues, fill out the bug report form located at https://www.juniper.net/documentation/feedback/.

Revision History

22 April 2021—Revision 18, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 17, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 October 2020—Revision 16, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 15, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 July 2020—Revision 14, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 June 2020—Revision 13, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 May 2020—Revision 12, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 March 2020—Revision 11, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 March 2020—Revision 10, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

5 March 2020—Revision 9, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 February 2020—Revision 8, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 January 2020—Revision 7, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 January 2020—Revision 6, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 January 2020—Revision 5, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 December 2019—Revision 4, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

5 December 2019—Revision 3, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 November 2019—Revision 2, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 November 2019—Revision 1, Junos OS Release 19.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 October 2019—Revision 4, Junos OS Release 19.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 October 2019—Revision 3, Junos OS Release 19.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 October 2019—Revision 2, Junos OS Release 19.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 September 2019—Revision 1, Junos OS Release 19.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.