

# Release Notes: Junos<sup>®</sup> OS Release 19.3R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

22 April 2021

<b>Contents</b>	<b>Introduction   10</b>
	<b>Junos OS Release Notes for ACX Series   10</b>
	<b>What's New   11</b>
	What's New in Release 19.3R3   11
	What's New in Release 19.3R2   11
	What's New in Release 19.3R1-S1   11
	What's New in Release 19.3R1   12
	<b>What's Changed   19</b>
	What's Changed in Release 19.3R3-S1   20
	What's Changed in Release 19.3R3   20
	What's Changed in Release 19.3R2-S6   20
	What's Changed in Release 19.3R2   21
	What's Changed in Release 19.3R1   21
	<b>Known Limitations   23</b>
	General Routing   23
	<b>Open Issues   25</b>
	General Routing   25
	Platform and Infrastructure   27
	Virtual Chassis   27

**Resolved Issues | 27****Resolved Issues: 19.3R3 | 28****Resolved Issues: 19.3R2 | 30****Resolved Issues: 19.3R1 | 31****Documentation Updates | 34****Migration, Upgrade, and Downgrade Instructions | 35****Upgrade and Downgrade Support Policy for Junos OS Releases | 35****Junos OS Release Notes for EX Series Switches | 37****What's New | 37****What's New in 19.3R3 | 38****What's New in 19.3R2 | 38****What's New in 19.3R1 | 39****What's Changed | 47****What's Changed in 19.3R3 | 48****What's Changed in Release 19.3R2-S6 | 49****What's Changed in 19.3R2 | 50****What's Changed in 19.3R1 | 50****Known Limitations | 52****EVPN | 52****Infrastructure | 52****Platform and Infrastructure | 52****Open Issues | 53****Authentication and Access Control | 54****Infrastructure | 54****Interfaces and Chassis | 54****Layer 2 Features | 54****Network Management and Monitoring | 54****Platform and Infrastructure | 54****Virtual Chassis | 56****Resolved Issues | 56****Resolved Issues: 19.3R3 | 57****Resolved Issues: 19.3R2 | 62****Resolved Issues: 19.3R1 | 64****Layer 2 Ethernet Services | 64**

Network Management and Monitoring	65
Platform and Infrastructure	65
Routing Protocols	68
Subscriber Access Management	69
User Interface and Configuration	69
Virtual Chassis	69
VPNs	69
Documentation Updates	69
Migration, Upgrade, and Downgrade Instructions	70
Upgrade and Downgrade Support Policy for Junos OS Releases	70
Junos OS Release Notes for Junos Fusion Enterprise	71
What's New	72
What's New in Release 19.3R3	72
What's New in Release 19.3R2	72
What's New in Release 19.3R1	72
What's Changed	73
Known Limitations	73
Open Issues	74
Junos Fusion for Enterprise	74
Resolved Issues	75
Resolved Issues: 19.3R3	75
Resolved Issues: 19.3R2	75
Resolved Issues: 19.3R1	75
Documentation Updates	76
Migration, Upgrade, and Downgrade Instructions	76
Basic Procedure for Upgrading Junos OS on an Aggregation Device	77
Upgrading an Aggregation Device with Redundant Routing Engines	78
Preparing the Switch for Satellite Device Conversion	79
Converting a Satellite Device to a Standalone Switch	80
Upgrade and Downgrade Support Policy for Junos OS Releases	80
Downgrading from Junos OS	81

## Junos OS Release Notes for Junos Fusion Provider Edge | 82

### What's New | 82

What's New in Release 19.3R3 | 83

What's New in Release 19.3R2 | 83

What's New in Release 19.3R1 | 83

### What's Changed | 83

### Known Limitations | 84

Junos Fusion Provider Edge | 84

### Open Issues | 85

Junos Fusion Provider Edge | 85

### Resolved Issues | 86

Resolved Issues: 19.3R3 | 86

Resolved Issues: 19.3R2 | 86

Resolved Issues: 19.3R1 | 86

### Documentation Updates | 87

### Migration, Upgrade, and Downgrade Instructions | 87

Basic Procedure for Upgrading an Aggregation Device | 88

Upgrading an Aggregation Device with Redundant Routing Engines | 90

Preparing the Switch for Satellite Device Conversion | 91

Converting a Satellite Device to a Standalone Device | 92

Upgrading an Aggregation Device | 95

Upgrade and Downgrade Support Policy for Junos OS Releases | 95

Downgrading from Junos OS Release 19.3 | 95

## Junos OS Release Notes for MX Series 5G Universal Routing Platform | 96

### What's New | 97

What's New in Release 19.3R3 | 97

What's New in Release 19.3R2 | 97

What's New in Release 19.3R1 | 116

### What's Changed | 135

What's Changed in Release 19.3R3-S1 | 136

What's Changed in Release 19.3R3 | 136

What's Changed in Release 19.3R2-S6 | 139

What's Changed in Release 19.3R2-S5 | 140

What's Changed in Release 19.3R2 | 140

What's Changed in Release 19.3R1	142
Known Limitations	146
General Routing	146
Infrastructure	149
Interfaces and Chassis	149
MPLS	150
Platform and Infrastructure	150
Routing Protocols	151
Open Issues	151
EVPN	152
Forwarding and Sampling	153
General Routing	153
Infrastructure	161
Interfaces and Chassis	161
Layer 2 Ethernet Services	162
MPLS	162
Platform and Infrastructure	163
Routing Protocols	164
Services Applications	166
Subscriber Access Management	166
User Interface and Configuration	166
VPNs	166
Resolved Issues	167
Resolved Issues: 19.3R3	167
Resolved Issues: 19.3R2	186
Resolved Issues: 19.3R1	195
Documentation Updates	214
Migration, Upgrade, and Downgrade Instructions	215
Basic Procedure for Upgrading to Release 19.3	216
Procedure to Upgrade to FreeBSD 11.x based Junos OS	216
Procedure to Upgrade to FreeBSD 6.x based Junos OS	219
Upgrade and Downgrade Support Policy for Junos OS Releases	220
Upgrading a Router with Redundant Routing Engines	221
Downgrading from Release 19.3	221

## Junos OS Release Notes for NFX Series | 222

### What's New | 222

Release 19.3R3 New and Changed Features | 223

Release 19.3R2 New and Changed Features | 223

Release 19.3R1 New and Changed Features | 223

### What's Changed | 224

What's Changed in Release 19.3R2-S6 | 224

### Known Limitations | 225

Interfaces | 225

Platform and Infrastructure | 225

Virtual Network Functions (VNFs) | 226

### Open Issues | 226

Interfaces | 226

High Availability (HA) | 227

Platform and Infrastructure | 227

Virtual Network Functions (VNFs) | 227

### Resolved Issues | 228

Resolved Issues: 19.3R3 | 228

Resolved Issues: 19.3R2 | 229

Resolved Issues: 19.3R1 | 230

### Documentation Updates | 232

### Migration, Upgrade, and Downgrade Instructions | 232

Upgrade and Downgrade Support Policy for Junos OS Releases | 232

Basic Procedure for Upgrading to Release 19.3 | 233

## Junos OS Release Notes for PTX Series Packet Transport Routers | 234

### What's New | 235

What's New in Release 19.3R3 | 235

What's New in Release 19.3R2 | 235

What's New in Release 19.3R1 | 236

### What's Changed | 243

What's Changed in Release 19.3R3-S1 | 244

What's Changed in Release 19.3R3 | 244

Juniper Extension Toolkit (JET) | 245

What's Changed in Release 19.3R2-S6 | 245

What's Changed in Release 19.3R2	246
What's Changed in Release 19.3R1	246
Known Limitations	248
General Routing	248
Interfaces and Chassis	249
Open Issues	249
General Routing	249
Routing Protocols	250
Resolved Issues	251
Resolved Issues: Release 19.3R3	251
Resolved Issues: Release 19.3R2	254
Resolved Issues: Release 19.3R1	255
Documentation Updates	257
Migration, Upgrade, and Downgrade Instructions	258
Basic Procedure for Upgrading to Release 19.3	258
Upgrade and Downgrade Support Policy for Junos OS Releases	261
Upgrading a Router with Redundant Routing Engines	262
Junos OS Release Notes for the QFX Series	263
What's New	263
What's New in Release 19.3R3	264
What's New in Release 19.3R2	264
What's New in Release 19.3R1	264
What's Changed	272
What's Changed in Release 19.3R3-S1	273
What's Changed in Release 19.3R3	273
What's Changed in Release 19.3R2-S6	275
What's Changed in Release 19.3R2	276
What's Changed in Release 19.3R1	276
Known Limitations	278
Class of Service (CoS)	279
EVPN	279
Platform and Infrastructure	279
Infrastructure	280
Layer 2 Features	280

Routing Protocols | **281**

Open Issues | **281**

EVPN | **282**

High Availability (HA) and Resiliency | **282**

Infrastructure | **282**

Interfaces and Chassis | **282**

Junos Fusion Provider Edge | **283**

Layer 2 Features | **283**

Platform and Infrastructure | **283**

Routing Protocols | **285**

Virtual Chassis | **286**

Resolved Issues | **286**

Resolved Issues: Release 19.3R3 | **287**

Resolved Issues: Release 19.3R2 | **295**

Resolved Issues: Release 19.3R1 | **299**

Documentation Updates | **306**

Migration, Upgrade, and Downgrade Instructions | **307**

Upgrading Software on QFX Series Switches | **307**

Installing the Software on QFX10002-60C Switches | **310**

Installing the Software on QFX10002 Switches | **310**

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 | **311**

Installing the Software on QFX10008 and QFX10016 Switches | **313**

Performing a Unified ISSU | **317**

Preparing the Switch for Software Installation | **318**

Upgrading the Software Using Unified ISSU | **318**

Upgrade and Downgrade Support Policy for Junos OS Releases | **320**

Junos OS Release Notes for SRX Series | **321**

What's New | **322**

What's New in Release 19.3R3 | **322**

What's New in Release 19.3R2 | **322**



What's New in Release 19.3R1	322
What's Changed	330
What's Changed in Release 19.3R3	331
What's Changed in Release 19.3R2-S6	332
What's Changed in Release 19.3R2	333
What's Changed in Release 19.3R1	333
Known Limitations	335
Authentication and Access Control	336
J-Web	336
Logical Systems and Tenant Systems	337
Platform and Infrastructure	337
VPNs	337
Open Issues	338
Flow-Based and Packet-Based Processing	338
Intrusion Detection and Prevention (IDP)	338
J-Web	339
Routing Policy and Firewall Filters	339
VPNs	339
Resolved Issues	340
Resolved Issues: Release 19.3R3	340
Resolved Issues: Release 19.3R2	346
Resolved Issues: Release 19.3R1	349
Documentation Updates	357
Migration, Upgrade, and Downgrade Instructions	357
Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases	358
Upgrading Using ISSU	359
Licensing	359
Finding More Information	360
Documentation Feedback	360
Requesting Technical Support	361
Self-Help Online Tools and Resources	361
Opening a Case with JTAC	362
Revision History	362

# Introduction

Junos OS runs on the following Juniper Networks<sup>®</sup> 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 19.3R3 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.

## Junos OS Release Notes for ACX Series

### IN THIS SECTION

- [What's New | 11](#)
- [What's Changed | 19](#)
- [Known Limitations | 23](#)
- [Open Issues | 25](#)
- [Resolved Issues | 27](#)
- [Documentation Updates | 34](#)
- [Migration, Upgrade, and Downgrade Instructions | 35](#)

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- [What's New in Release 19.3R3 | 11](#)
- [What's New in Release 19.3R2 | 11](#)
- [What's New in Release 19.3R1-S1 | 11](#)
- [What's New in Release 19.3R1 | 12](#)

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

### What's New in Release 19.3R3

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 19.3R3.

### 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 Release 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 leaves 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



Platforms	Routing Engine Supported
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
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 | 19](#)

[Known Limitations | 23](#)

[Open Issues | 25](#)

[Resolved Issues | 27](#)

[Documentation Updates | 34](#)

[Migration, Upgrade, and Downgrade Instructions | 35](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 19.3R3-S1 | 20](#)
- [What's Changed in Release 19.3R3 | 20](#)
- [What's Changed in Release 19.3R2-S6 | 20](#)
- [What's Changed in Release 19.3R2 | 21](#)
- [What's Changed in Release 19.3R1 | 21](#)

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

## What's Changed in Release 19.3R3-S1

### *Routing Protocols*

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.

## What's Changed in Release 19.3R3

### *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](#).]

### *Juniper Extension Toolkit (JET)*

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

See [traceoptions \(Services\)](#).

## 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

### *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.

### *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.

- **XML RPC equivalent included for the `show bgp output-scheduler | display xml rpc` CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, we have included an XML RPC equivalent for the `show bgp output-scheduler | display xml rpc` CLI command. In Junos OS releases before Release 19.3R1, the `show bgp output-scheduler | display xml rpc` CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler.](#)]

### **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 | 11](#)

[Known Limitations | 23](#)

[Open Issues | 25](#)

[Resolved Issues | 27](#)

[Documentation Updates | 34](#)

[Migration, Upgrade, and Downgrade Instructions | 35](#)

## Known Limitations

### IN THIS SECTION

- [General Routing | 23](#)

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

- PTP IP packets queued to the best -effort queue instead of to the network control queue by default. [PR1361315](#)
- Filter for sensors (Optical and OTN) does not work. [PR1371996](#)
- For the et interface, the FEC and FRR alarms are not raised. [PR1371997](#)
- The clock class to QL level mapping must be changed through CLI on the ACX2100 router with network option-1. [PR1384968](#)
- The clock class must be changed to quality level mapping on the ACX2100 router with network option-2. [PR1386916](#)
- The PTP slave gets stuck in the **Initialize** state intermittently due to the ARP resolution issue when PTP and interfaces flaps continuously. [PR1410746](#)
- For the ACX6360-OR routers, use the PTX10001 router documentation and not the ACX Series router documentation. [PR1417258](#)
- Hardware-based fragmentation or reassembly is not supported. Software-based fragmentation rates goes extremely slow depending on CPU load. [PR1419371](#)
- Oversized frames are accounted twice (input packet and oversized frames) in the **extensive** command. [PR1425748](#)
- Error messages might be observed if the optics is unplugged during the EEPROM read. This is expected and does not impact any functionality. [PR1429016](#)
- On the ACX5448-D interfaces, the pps and bps for the aggregated Ethernet logical interfaces do not get displayed during traffic flow. [PR1429590](#)
- Incorrect queue statistics are observed for PTPOE multicast packets. [PR1429628](#)
- On the ACX5448-D interfaces, the Jabber frames counter does not increment correctly. [PR1429923](#)

- On the ACX5448-D interfaces, the input bytes value displayed in the **show interfaces extensive** command output is not at par with that of the older ACX Series devices. [PR1430108](#)
- On the ACX5448-D interfaces, the interface LEDs glow after the system halts. [PR1430129](#)
- On the ACX5448-D interfaces, the following error message is observed:  
**rio\_mic\_cfp2\_ifd\_ot\_eth\_media\_conf port command handler failed.** [PR1430355](#)
- On the ACX5448-D interfaces, the MTU error count in the output interface does not increment when frames are dropped due to less MTU output. [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](#)
- The l2cpd process generates the core file continuously after loading the 19.3 DCB image. [PR1431866](#)
- On the ACX5448-D interface, when configured over Layer 2 VPN with control-channel set, BFD gets to the **Down** state. [PR1432854](#)
- In case of Dot1P, CFI rewrite based on TC or DP classification is not possible on the ACX5448, ACX-710, and ACX-R6675 Series routers. [PR1435966](#)
- 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 router takes lesser time. 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. [PR1439960](#)
- The hold timer expiry is common across all platforms. It is not specific to the ACX5448 Series of routers. [PR1439980](#)
- With an asymmetric network connection (for example, a 10GbE MACsec port connected to a 10GbE channelized port), high and asymmetric T1 and T4 time errors lead to a high two-way time error. [PR1440140](#)
- By default, the management interface speed is always displayed as 1000 Mbps in Junos OS command output. [PR1440675](#)
- TC performance with 1-Gigabit Ethernet enables MACsec. [PR1441388](#)
- On the ACX5448-M interface, the Jitter tolerance test fails over MACsec ports. [PR1447296](#)
- IGMPv3 rate of join supported in the ACX5448 routers or an ACX5448-D interface is around 900 joins per second. [PR1448146](#)
- The aggregated Ethernet interface speed is displayed incorrectly when 1-Gigabit Ethernet member links are a part of the aggregated Ethernet interface. [PR1449887](#)
- Input packet rejects counter when the MAC address is not supported. [PR1463981](#)

SEE ALSO



---

[What's Changed | 19](#)


---

[Open Issues | 25](#)


---

[Resolved Issues | 27](#)


---

[Documentation Updates | 34](#)


---

[Migration, Upgrade, and Downgrade Instructions | 35](#)


---

## Open Issues

### IN THIS SECTION

- [General Routing | 25](#)
- [Platform and Infrastructure | 27](#)
- [Virtual Chassis | 27](#)

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

- Loopback status is not displayed for ot interfaces on CLI, which is available from vty only. [PR1358017](#)
- The **show class-of-service interface** command does not show the classifier information. [PR1522941](#)
- The Layer 2 VPN might flap and the CE device facing interface cannot restore the TX optical laser power even if the Layer 2 VPN is in the **Up** status under the asynchronous-notification. [PR1486181](#)
- The routes update might fail upon the HMC memory issue and traffic impact might be seen. [PR1515092](#)
- The SD threshold is set above the SF threshold. [PR1376869](#)
- On the ACX6360-OR router, enhancement is needed for the FRR BER threshold SNMP support. [PR1383303](#)
- Tx power cannot be configured using the + sign. [PR1383980](#)
- Traffic gets lost for more than 50 minutes under certain conditions. [PR1387858](#)
- The circuit cross-connect logs are not compressed after rotation. [PR1398511](#)
- The **jnxIfOtnOperState** trap notification is sent for all ot interfaces. [PR1406758](#)

- The Layer 2 rewrite does not occur on the UNI side for the packet arriving from the MPLS core. [PR1414414](#)
- Clock class value is incorrect in the default data when the slave interface is down in the PTP-OC device. [PR1416421](#)
- On the ACX5448 router, ZTP proceeds to load the image even when there is a platform string mismatch with the platform and software image names on the ZTP server. [PR1418313](#)
- On the ACX5448 router, the firewall filter works in a different way if it is untagged traffic. [PR1423214](#)
- The fpc process might crash with em2 configuration during initialization. [PR1429212](#)
- DHCP clients are not able to scale to 96000. [PR1432849](#)
- Logical interfaces or protocols take time to come up. [PR1433426](#)
- Forwarding occurs when using non-existing SSM map source address in IGMPv3 instead of pruning. [PR1435648](#)
- 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 percent of the mean latency introduced by the each phy which is of around 2.5us). [PR1437175](#)
- On the ACX5448-D interfaces, the following error message is observed when an interface is brought up: `fpc0 otnpic_otn_frr_set pic='et-0/2/0' link=0 IFD=et-0/2/0 error=7`. [PR1438039](#)
- Memory leakage is observed while running the longevity check. [PR1438358](#)
- On the ACX5448-D interfaces, the logical interface statistics are not displayed correctly. [PR1439124](#)
- Using an unnumbered interface for ECMP results in a few next hop being nonresponsive in the KRT queue. [PR1442622](#)
- Recovery of Junos OS volume from OAM becomes nonresponsive indefinitely. [PR1446512](#)
- On the ACX5448-X or ACX5448-M interfaces, the drop profile maximum threshold might not reach to its limit when the packet size is other than 1000 bytes. [PR1448418](#)
- oper-state for the et interface does not transition from the **Init** state to the **Normal** state. [PR1449937](#)
- While checking the BFD functionality over a Layer 2 VPN client, the BFD session does not come up. [PR1483014](#)
- On the ACX5048 routers, traffic loss is observed during the unified ISSU upgrade. [PR1483959](#)
- The CLI command configures **announce-interval** to **-3** so that the announce messages rate is set to 8 pps. [PR1507782](#)
- On the ACX5048 or ACX5096 line of routers, the LACP control packets might get dropped due to high CPU utilization. [PR1493518](#)

Platform and Infrastructure

- On the ACX5448-M interface, the CFM remote MEP does not come up after configuration. [PR1460555](#)

Virtual Chassis

- On the ACX5000 router, the following false positive parity error message is observed:  
\_soc\_mem\_array\_sbusdma\_read. [PR1276970](#)

SEE ALSO

<a href="#">What's New</a>	<a href="#">  11</a>
<a href="#">What's Changed</a>	<a href="#">  19</a>
<a href="#">Known Limitations</a>	<a href="#">  23</a>
<a href="#">Resolved Issues</a>	<a href="#">  27</a>
<a href="#">Documentation Updates</a>	<a href="#">  34</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  35</a>

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.3R3](#) | [28](#)
- [Resolved Issues: 19.3R2](#) | [30](#)
- [Resolved Issues: 19.3R1](#) | [31](#)

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: 19.3R3

### General Routing

- On the ACX5000 router, the following error message is observed: **MacDrainTimeOut and bcm\_port\_update failed: Internal error.** [PR1284590](#)
- The **gether-options** command are enabled again under the interface hierarchy. [PR1430009](#)
- On the ACX5448-D interface, while performing repeated power-off or power-on of the device, SMBUS transaction timeout is observed. [PR1463745](#)
- High CPU utilization of the fxpc process might be observed with the **class-of-service** changes on interfaces. [PR1407098](#)
- Policer discarded count is displayed incorrectly to the enq count of the interface queue. But the traffic behavior is as expected. [PR1414887](#)
- On the ACX2200 routers, the drift messages, which is a PTP hybrid (PTP + Synchronous Ethernet), are observed. [PR1426910](#)
- The input bytes value displayed in the **show interfaces extensive** command output is not at par with that of the older ACX Series routers. [PR1430108](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- On the ACX5448 routers, the encapsulation **flexible-vlan-tagging** is not supported with the MPLS family. [PR1445046](#)
- On the ACX5448-M routers, fans might not run at the correct speed. [PR1448884](#)
- On the ACX5048 routers, the SNMP polling is stalled after the link is flapped or the SFP transceiver is replaced. The following error message is observed:  
**ACX\_COS\_HALP(acx\_cos\_gport\_sched\_set\_strict\_priority:987): Failed to detach logs.** [PR1455722](#)
- The rpd process might crash if a BGP route is resolved over the same prefix protocol next hop in an **inet.3** table that has both RSVP and LDP routes. [PR1458595](#)
- On enabling the local loopback, the 10-Gigabit Ethernet interface goes down. [PR1460715](#)
- The ARP resolution issue might be observed while bringing up the aggregated Ethernet interface after reboot. [PR1461485](#)
- Memory leakage on the l2cpd process might lead to l2cpd crash. [PR1469635](#)
- RED drop might be seen on an interface even when there is no congestion. [PR1470619](#)
- On the ACX5048 routers, the egress queue statistics are not applicable to the aggregated Ethernet interfaces. [PR1472467](#)
- FPC might continuously crash after deactivating or activating the loopback filter or rebooting the system after configuring the loopback filter. [PR1477740](#)
- The dcpfe process might crash when disabling or enabling MACsec. [PR1479710](#)

- Links might not come up when the 100-Gigabit Ethernet interface is channelized into 4x25-Gigabit Ethernet interfaces. [PR1479733](#)
- RPD memory chunk size is optimized for the ACX Series routers to reduce the memory foot print. [PR1481151](#)
- On the ACX5448 routers, the following logs are reported: **dnx\_nh\_mpls\_tunnel\_install**. [PR1482529](#)
- On the ACX5448 router, the fpc process might crash. [PR1485315](#)
- On the ACX5448 routers, the Layer 2 VPN with the **ethernet-ccc input-vlan-map/output-vlan-map** interface silently discard the traffic. [PR1485444](#)
- The queue statistics are not as expected after configuring the IFD and logical interface shaping with the transmit rate and scheduler-map. [PR1488935](#)
- During speed mismatch, the QSFP28 or QSFP+ optics or cables might not work. [PR1494600](#)
- Outbound SSH connection flap or memory leakage issue might be observed while pushing the configuration to the ephemeral database with a high rate. [PR1497575](#)
- On the ACX500 router, the SFW sessions might not get updated on the MS interfaces. [PR1505089](#)
- The PIC slot might get shut down in less than 240 seconds due to the over-temperature start time being handled incorrectly. [PR1506938](#)
- BFD flaps with the following error message after a random time interval: **ACX\_OAM\_CFG\_FAILED: ACX Error (oam):dnx\_bfd\_l3\_egress\_create : Unable to create egress object**. [PR1513644](#)
- The l2ald process crashes during the stability test with traffic on a scaled setup. [PR1517074](#)
- On the ACX5448 Chassis, the MAC-address and label MAC-address might not match. [PR1489034](#)

### ***Interfaces and Chassis***

- The MC-AE interface displays status as **Unknown** if the subinterface is added as part of the VLAN on the peer MC-AE node. [PR1479012](#)

### ***Layer 2 Ethernet Services***

- Member links states might be asynchronized on a connection between the PE and CE devices in the EVPN active/active scenario. [PR1463791](#)

### ***MPLS***

- The BGP session might flap between two directly connected BGP peers because of the wrong TCP MSS (maximum segment size) in use. [PR1493431](#)

### ***Routing Protocols***

- The BGP route target family might prevent the route reflector from reflecting the Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

### ***VPNs***

- The Layer 2 circuit neighbor might get stuck in the **RD** state at one end of the MG-LAG peer. [PR1498040](#)

## **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 I2circuit 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 gigheter 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 a 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 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. user@host# 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 theoretical 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 Vitesse 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

<a href="#">What's New   11</a>
<a href="#">What's Changed   19</a>
<a href="#">Known Limitations   23</a>
<a href="#">Open Issues   25</a>
<a href="#">Documentation Updates   34</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   35</a>

**Documentation Updates**

There are no errata or changes in Junos OS Release 19.3R3 documentation for the ACX Series.

SEE ALSO

What's New   11
What's Changed   19
Known Limitations   23
Open Issues   25
Resolved Issues   27
Migration, Upgrade, and Downgrade Instructions   35

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 35

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 | 11](#)

---

[What's Changed | 19](#)

---

[Known Limitations | 23](#)

---

[Open Issues | 25](#)

---

[Resolved Issues | 27](#)

---

[Documentation Updates | 34](#)

# Junos OS Release Notes for EX Series Switches

## IN THIS SECTION

- What's New | 37
- What's Changed | 47
- Known Limitations | 52
- Open Issues | 53
- Resolved Issues | 56
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 70

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- What's New in 19.3R3 | 38
- What's New in 19.3R2 | 38
- What's New in 19.3R1 | 39

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 19.3R3: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

### **What's New in 19.3R3**

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.3R3.

### **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/rtssock/`

[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

<a href="#">What's Changed   47</a>
<a href="#">Known Limitations   52</a>
<a href="#">Open Issues   53</a>
<a href="#">Resolved Issues   56</a>
<a href="#">Documentation Updates   69</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   70</a>

# What's Changed

IN THIS SECTION

- [What's Changed in 19.3R3 | 48](#)
- [What's Changed in Release 19.3R2-S6 | 49](#)
- [What's Changed in 19.3R2 | 50](#)
- [What's Changed in 19.3R1 | 50](#)

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

## What's Changed in 19.3R3

### General Routing

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.
- **Command to view summary information for resource monitor (MX Series routers and EX9200 line of switches)**—You can use the **show system resource-monitor** command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

See [[show system resource-monitor and](#) ].

- **New commit check for MC-LAG (MX Series)**— We've introduced a new commit check to check the values assigned to the redundancy group identification number on the MC-AE interface ( **redundancy-group-id** ) and ICCP peer ( **redundancy-group-id-list** ) when you configure multichassis aggregation groups (MC-LAGs). If the values are different, the system reports a commit check error. In previous releases, if the configured values were different, the l2ald process would crash.

[See [iccp](#) and [mc-ae](#).]



### *Juniper Extension Toolkit (JET)*

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

See [traceoptions \(Services\)](#).

### *Multicast*

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 19.3R3, we provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... n packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). Before Junos OS Release 19.3R3, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** do not provide valid statistics for multicast traffic at Layer 2.

[See [show multicast snooping route](#).]

### *Routing Protocols*

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the **Isdist.0** and **Isdist.1** routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the **Isdist.0** and **Isdist.1** routing tables as part of node characteristics and advertised them as the router ID.

## **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

### *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.
- **XML RPC equivalent included for the show bgp output-scheduler | display xml rpc CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc**

CLI command. In Junos OS releases before Release 19.3R1, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler](#).]

**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

<a href="#">What's New   37</a>
<a href="#">Known Limitations   52</a>
<a href="#">Open Issues   53</a>
<a href="#">Resolved Issues   56</a>
<a href="#">Documentation Updates   69</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   70</a>

## Known Limitations

### IN THIS SECTION

- [EVPN | 52](#)
- [Infrastructure | 52](#)
- [Platform and Infrastructure | 52](#)

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](#)

### Infrastructure

- On an EX4300 switch, the CLI configuration **set chassis routing-engine on-disk-failure disk-failure-action (reboot | halt)** is not supported. [PR1450093](#)

### Platform and Infrastructure

- With the 288,000 MAC scale, the Routing Engine command **show ethernet-switching table summary** output displays the learned scale entries after a delay of around 60 seconds. [PR1367538](#)
- During software upgrade to a more recent 19.3 images image, the system hangs after the **request system software add /var/tmp/<image.gz>** command is issued. Device could be recovered by power cycling the device. [PR1405629](#)
- When the box is loaded and unloaded with the MACsec configuration multiple times with operations made continuously, Layer 3 connectivity lost, the system stops, followed by a reboot to resume operation. [PR1416499](#)
- On deactivating and activating POE, the POE interface draws more power (as per the display) for quite some time. [PR1431647](#)
- Filters are installed only during route add if there is enough space. If the filter fails because of the non-availability of TCAM space, those routes might not be processed for filter add later when space becomes available. [PR1419926](#)

- The **set class-of-service shared-buffer ingress buffer-partition lossless-headroom percent 0** command is not supported when in Virtual Chassis (VC), as the VCP ports should have some heardroom to support PFC. Configuration will be rejected at the hardware layer with a log message. [PR1448377](#)
- Even with the fix for PR1463015, the **Failed to complete DFE tuning** syslog might appear. This message has no functional impact and can be ignored. [PR1473280](#)

SEE ALSO

<a href="#">What's New   37</a>
<a href="#">What's Changed   47</a>
<a href="#">Open Issues   53</a>
<a href="#">Resolved Issues   56</a>
<a href="#">Documentation Updates   69</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   70</a>

## Open Issues

IN THIS SECTION

- [Authentication and Access Control | 54](#)
- [Infrastructure | 54](#)
- [Interfaces and Chassis | 54](#)
- [Layer 2 Features | 54](#)
- [Network Management and Monitoring | 54](#)
- [Platform and Infrastructure | 54](#)
- [Virtual Chassis | 56](#)

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, configure the path using the "set protocols dot1x ssl-certificate-path PATHNAME" command, if the default pathname is not "/var/tmp/". [PR1431086](#)

## Infrastructure

- On EX Series switches, if you are configuring a large number of firewall filters on some interfaces, the FPC might crash and generate core files. [PR1434927](#)
- EX 9251: "IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) error message is observed continuously in AD with base configurations. [PR1485038](#)

## Interfaces and Chassis

- After GRES, the VSTP port cost on aggregated Ethernet interfaces might get changed, leading to a topology change. [PR1174213](#)

## Layer 2 Features

- GARPs were being sent whenever there is a MAC (FDB) operation (add or delete). This is now updated to send GARP when the interface is UP and the Layer3 interface is attached to the VLAN. [PR1192520](#)

## Network Management and Monitoring

- On EX Series switches except EX4300, EX4600, and EX9200, when RTG (Redundant Trunk Group) switchovers are done, the /var/log/shadow.log or /var/log/shadow\_debug.log is rotated. This might cause the PFE process to crash. [PR1233050](#)

## Platform and Infrastructure

- ARP queue limit has been changed from 100 pps to 3000 pps. [PR1165757](#)
- On an EX2300 switch, the output of the **show chassis routing-engine** command might display an incorrect value of "Router rebooted after a normal shutdown" for the last reboot reason field. [PR1331264](#)
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)
- On an EX9208 switch, a few xe- interfaces go down with the error message "if\_msg\_ifd\_cmd\_tlv\_decode ifd xe-0/0/0 #190 down with ASIC Error". [PR1377840](#)

- There is a possibility of seeing multiple reconnect logs, JTASK\_IO\_CONNECT\_FAILED, during the device initialization. There is no functionality impact because of these messages. These messages can be ignored. [PR1408995](#)
- On QFX5110 and QFX5120 platforms, unicast RPF check in strict mode might not work properly. [PR1417546](#)
- Minimal traffic loss of ~100 pps is seen on EX9208 when the packets are sent between FPCs. This is due to random drops happening in the fabric. Amount of drop varies on the line rate and occurs less frequently. `router@host> show pfe statistics traffic |match dropSoftware input control plane drops : 0Software input high drops : 0Software input medium drops : 0Software input low drops : 0Software output drops : 0Hardware input drops : 179892Info cell drops : 0Fabric drops : 71.`[PR1429714](#)
- BUM traffic rate limiting is done after removing Ethernet headers. L1 Tx rate on ingress interface: 1G Tx rate with headers: 865 Mbps Rx rate on the egress interface: 800M L1 Rx rate on egress interface: 925Mbps Storm control functionalities in MX-L card is achieved by policer and hence the below mentioned policer inaccuracy is applicable for storm control feature as well. Since XM sprays packets to 4 different LUs, each LU will be processing packets of varying sizes. XM does not do strict round-robin, so even if all the incoming packets were to be of exact same sizes (which is not a practical scenario), each LU will still be loaded differently, hence there will be some periods where some LUs policing limit may reach sooner than the others (either due to processing more packets or due to processing larger packets). Hence, it is possible that some LUs, which see the policing limit reached sooner may drop the packet or color them differently that might result into eventual drop while the other LUs could queue the packets for transmission, we could see this behavior within a single flow as well. Hence, the policer functionality can be unpredictable at times. In an extreme case, a packet flow might be sent to a single LU and the policer result is one-fourth of what is expected. Because the policer functionality, in general, may not work correctly, we will see the impact on all the policing features - for example, input-policer, three-color-policer (srTCM, trTCM), output-policer. [PR1442842](#)
- On an EX9214 switch, if the MACsec-enabled link flaps after reboot, the error "errorlib\_set\_error\_log(): err\_id(-1718026239)" is observed. [PR1448368](#)
- In overall commit time, the evaluation of mustd constraints is taking 2 seconds more than usual. This is because the persist-group-inheritance feature has been made a default feature in the latest Junos OS releases. Eventually, this feature helps improve the subsequent commit times for scaled configurations significantly. The persist-group-inheritance feature is useful in customer scenarios where groups and nested groups are used extensively. In those scenarios, the group inheritance paths are not built every time; thus, subsequent commits are faster. [PR1457939](#)
- The image size grows over a period of time and subsequently storage might be insufficient to install images when upgrading EX2300 and EX3400 platforms. The upgrade might fail with the error message **not enough space to unpack**. [PR1464808](#)

## Virtual Chassis

- EX4650: **kldload: an error occurred while loading the module** is displayed during booting. [PR1527170](#)

### SEE ALSO

[What's New | 37](#)

[What's Changed | 47](#)

[Known Limitations | 52](#)

[Resolved Issues | 56](#)

[Documentation Updates | 69](#)

[Migration, Upgrade, and Downgrade Instructions | 70](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 19.3R3 | 57](#)
- [Resolved Issues: 19.3R2 | 62](#)
- [Resolved Issues: 19.3R1 | 64](#)
- [Layer 2 Ethernet Services | 64](#)
- [Network Management and Monitoring | 65](#)
- [Platform and Infrastructure | 65](#)
- [Routing Protocols | 68](#)
- [Subscriber Access Management | 69](#)
- [User Interface and Configuration | 69](#)
- [Virtual Chassis | 69](#)
- [VPNs | 69](#)

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: 19.3R3

### *Authentication and Access Control*

- EX4600 and EX4300—MAC entry missing in Ethernet Switching table for MAC-RADIUS client in server\_fail scenario when tagged is sent for the second client. [PR1462479](#)
- Authentication failure might occur on the captive portal. [PR1504818](#)
- The Junos OS event **DOT1XD\_AUTH\_SESSION\_DELETED** might not be triggered when the 802.1X (dot1x) authentication session is terminated. [PR1512724](#)
- The dot1x client will not be moved to the hold state when the authenticated PVLAN is deleted. [PR1516341](#)

### *Class of Service (CoS)*

- Shaping does not work after the reboot if **shaping-rate** is configured. [PR1432078](#)
- The traffic is placed in the network control queue on the extended port even if it comes in with a different DSCP marking. [PR1433252](#)

### *EVPN*

- rpd might crash after changing the EVPN-related configuration. [PR1467309](#)
- The ESI of IRB interfaces does not update after autonomous-system number change if the interface is down. [PR1482790](#)
- The l2ald memory leak might be observed in any EVPN scenario. [PR1498023](#)
- The VXLAN function might be broken due to a timing issue after the change in PR1495098. [PR1502357](#)
- Unable to create a new VTEP interface. [PR1520078](#)

### *Forwarding and Sampling*

- Type 1 ESI/AD route might not be generated locally on EVPN PE in the all-active mode. [PR1464778](#)

### *Infrastructure*

- EX4300—CLI configuration **on-disk-failure** is not supported. [PR1450093](#)
- l2ald and eventd are hogging 100% after the **clear ethernet-switching table** command is issued. [PR1452738](#)
- EX Series switches might not come up properly upon reboot. [PR1454950](#)
- Error messages related to soft reset of port due to queue buffers being stuck could be seen on an EX4600-EX4300 Virtual Chassis. [PR1462106](#)
- Continuous dcpfe error messages and eventd process hogged might be seen in an EX2300 Virtual Chassis scenario. [PR1474808](#)

- Kernel core file might be observed while deactivating the daemon on the EX2300 and EX3400 platforms. [PR1483644](#)
- fxpc might crash when configuring scaled configuration with 4093 VLANs. [PR1493121](#)

### ***Interfaces and Chassis***

- The syslog `scchassisd[ ]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x` is observed after an MX Series Virtual Chassis local or global switchover. [PR1428254](#)
- The MC-LAG configuration-consistency ICL-configuration might fail after committing some changes. [PR1459201](#)
- Executing commit might result in hangup due to the stuck dcd process. [PR1470622](#)

### ***Junos Fusion Enterprise***

- Observed error `jnh_dot1br_ktree_entry_create(1098): Entry Already Exists`. Observing duplicate ECID values for cluster/extended ports on member ports of same cluster. [PR1408947](#)
- SDPD core file found @ `vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry`. [PR1454335](#)
- Loop detection might not work on extended ports in Junos Fusion scenarios. [PR1460209](#)

### ***Junos Fusion Satellite Software***

- Temperature sensor alarm is seen on an EX4300 switch in a Junos Fusion scenario. [PR1466324](#)

### ***Layer 2 Features***

- The LLDP function might fail when a Juniper device connects to a non-Juniper one. [PR1462171](#)
- `fxpc.core` might be seen when committing the configuration altogether, for example, after the reboot. [PR1467763](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)

### ***Layer 2 Ethernet Services***

- Member links state might be asynchronized on a connection between Provider Edge and Customer Edge devices in an EVPN active/active scenario. [PR1463791](#)
- Issues with DHCPv6 Relay processing Confirm and Reply packets. [PR1496220](#)

## MPLS

- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP MSS (maximum segment size) in use. [PR1493431](#)

## Platform and Infrastructure

- EX3400-VC—unable to upgrade `/usr/libexec/ui/package: /var/tmp/mchassis-install.tgz: no such file.` [PR1440122](#)
- The laser Tx might be enabled while the interface is disabled. [PR1445626](#)
- NSSU causes a traffic loss again after the backup to master transitions. [PR1448607](#)
- On certain MPC line cards, cm errors need to be reclassified. [PR1449427](#)
- The IRB traffic might drop after mastership switchover. [PR1453025](#)
- The OSPF neighbor might go down when the mDNS/PTP traffic is received at a rate greater than 1400 pps. [PR1459210](#)
- Traffic loss may be observed for longer than 20 seconds when performing NSSU on an EX4300 VC. [PR1461983](#)
- IGMP reports are dropped with mixed enterprise and SP configuration styles on an EX4300 switch. [PR1466075](#)
- The switch might not be able to learn MAC address with dot1x and interface-mac-limit configured. [PR1470424](#)
- EX4300: Input firewall filter attached to isolated or community VLANs does not match the dot1p bits on the VLAN header. [PR1478240](#)
- MAC learning under bridge-domain stops after the MC-LAG interface flaps. [PR1488251](#)
- The traffic destined to VRRP VIP might be dropped after the IRB interface is disabled on the initial VRRP master [PR1491348](#)
- IPv6 neighbor solicitation packets might be dropped in a transit device [PR1493212](#)
- Packets get dropped when next hop is IRB over the LT- interface. [PR1494594](#)
- The NSSU upgrade might fail on EX4300 switches due to a storage issue in the `/var/tmp` directory [PR1494963](#)
- Traffic loss might be seen with framing errors or runts if MACsec is configured on an EX4300 switch. [PR1502726](#)
- VC split after network topology change. [PR1427075](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping [PR1431355](#)
- IRB over VTEP unicast traffic might get dropped on EX9200 and MX platforms [PR1436924](#)
- The MAC Pause frames will be incrementing on Receive direction if the half duplex mode on 10Mbps or 10 MBps speed is configured [PR1452209](#)

- A firewall filter might not be able to be applied in a particular VC or VCF member as the TCAM space is running out [PR1455177](#)
- Link up delay and traffic drop might be seen on mixed SP Layer 2 and Layer 3 and EP Layer 2 type configurations. [PR1456336](#)
- MAC addresses learned on RTG may not be aged out after aging time [PR1461293](#)
- RTG link faces nearly a 20-second down time during backup node rebooting [PR1461554](#)
- The VC function may be broken after upgrading on EX2300/EX3400 platforms [PR1463635](#)
- EX Series (ELS) / QFX Series—There are some commands to disable a MAC learning, and some of them were not working. [PR1464797](#)
- dhcpcd might consume high CPU and no further subscribers can be brought up if more than 4000 dhcp-relay clients in the MAC-MOVE scenario. [PR1465277](#)
- fxpc might crash after mastership election on on the EX2300 and EX3400 platforms. [PR1465526](#)
- The broadcast and multicast traffic might be dropped over IRB or LAG interface in in a QFX Series or EX Series Virtual Chassis scenario. [PR1466423](#)
- The MAC move message may have an incorrect "from" interface when the MAC moves rapidly. [PR1467459](#)
- FPCs might get disconnected from an EX3400 VC briefly after reboot/upgrade. [PR1467707](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI [PR1468435](#)
- FPC may be down when configuring VXLAN routing. [PR1468736](#)
- Traffic loss might be seen with framing errors or runts if MACsec is configured on the EX4600 and QFX5100 platforms. [PR1469663](#)
- The SSH session closes while checking for the **show configuration | display set** command for both local and non-local users. [PR1470695](#)
- An EX3400 switch is advertising only 100m when configured the speed 100m with autonegotiation enabled. [PR1471931](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- CoS 802.1p bit rewrite might not happen in Q-in-Q mode. [PR1472350](#)
- DSCP marking might not work as expected if the fixed classifiers are applied to interfaces on the QFX5000 and EX4600 platforms. [PR1472771](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)
- The dhcpcd process may crash in a Junos Fusion environment. [PR1478375](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- TFTP installation from loader prompt may not succeed on the EX Series devices [PR1480348](#)

- ARP request packets for unknown host might get dropped in remote PE in an EVPN-VXLAN scenario. [PR1480776](#)
- On an EX2300 switch, SNMP traps are not generated when MAC address limit threshold is reached. [PR1482709](#)
- Incorrect "frame length" of 132 bytes might be captured in a packet header. [PR1487876](#)
- DHCP binding might fail when the PVLAN is configured with a firewall to block or allow certain IPv4 packets. [PR1490689](#)
- Traffic loss could be seen under the MC-LAG scenario on a QFX5120 or EX4650 switch. [PR1494507](#)
- Outbound SSH connection flap or memory leak issue might be observed during pushing the configuration to the ephemeral DB with high rate. [PR1497575](#)
- Traffic might get dropped if an AE member interface is deleted or added or if a SFP of the AE member interface is unplugged or plugged. [PR1497993](#)
- The OSPF might be stuck in the "ExStart" state if it is configured across the master and the standby FPCs of the EX4300-48MP VC [PR1498903](#)
- Continuous traffic loss might happen during GRES. [PR1500783](#)
- LLDP neighborhood might be unable to set up on QFX5000 platforms [PR1504354](#)
- LLDP might not work when PVLAN is configured on EX Series switches or QFX Series Virtual Chassis. [PR1511073](#)
- "dot1x" memory leak is observed. [PR1515972](#)
- MPPE-Send/Recv-key attribute is not extracted correctly by dot1xd [PR1522469](#)
- "Drops" and "Dropped packets" counters in the output by "show interface extensive" are double counting. [PR1525373](#)

### **Routing Protocols**

- Host-destined packets with filter log action might not reach the RE if log/syslog is enabled [PR1379718](#)
- QFX5100—BGP v4/v6 convergence & RIB install/delete time degraded in Junos OS Release 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The "other querier present interval" timer cannot be changed in an IGMP/MLD snooping scenario [PR1461590](#)
- LACP interface might not go to "collecting and distributing" state if along with the knob of "force-up" [PR1484523](#)
- FPC may go to the "NotPrsnt" state after upgrading with non-tvp image in VC/VCF setup [PR1485612](#)
- The BGP route-target family might prevent the RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

- The firewall filter could not work in certain conditions under a VC setup [PR1497133](#)
- Packet loss is seen for stream bLock:irb\_lacp\_tr\_ospf while verifying traffic from access to core network for IPv4 [PR1520059](#)

#### ***User Interface and Configuration***

- The **umount: unmount of /.mount/var/val/chroot/packages/mnt/jweb-ex32-d2cf6f6b failed: Device busy** message is seen when Junos OS is upgraded with the validate option. [PR1478291](#)
- J-Web does not display the correct Flow-control status on EX Series devices. [PR1520246](#)

#### ***Virtual Chassis***

- Disabling one of the VCP ports might result in another VCP port flapping [PR1469257](#)

### **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 wrongly 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 EX9251 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 jdncpd 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 : I2ald 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

<a href="#">What's New   37</a>
<a href="#">What's Changed   47</a>
<a href="#">Known Limitations   52</a>
<a href="#">Open Issues   53</a>
<a href="#">Documentation Updates   69</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   70</a>

Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 documentation for the EX Series switches.

SEE ALSO

<a href="#">What's New   37</a>
<a href="#">What's Changed   47</a>
<a href="#">Known Limitations   52</a>

---

[Open Issues | 53](#)

---

[Resolved Issues | 56](#)

---

[Migration, Upgrade, and Downgrade Instructions | 70](#)

---

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 70](#)

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

---

[What's New | 37](#)

---

What's Changed	47
Known Limitations	52
Open Issues	53
Resolved Issues	56
Documentation Updates	69

## Junos OS Release Notes for Junos Fusion Enterprise

### IN THIS SECTION

- What's New | 72
- What's Changed | 73
- Known Limitations | 73
- Open Issues | 74
- Resolved Issues | 75
- Documentation Updates | 76
- Migration, Upgrade, and Downgrade Instructions | 76

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- [What's New in Release 19.3R3 | 72](#)
- [What's New in Release 19.3R2 | 72](#)
- [What's New in Release 19.3R1 | 72](#)

Learn about new features introduced in the Junos OS main and maintenance releases for Junos fusion for enterprise.

**NOTE:** For more information about the Junos Fusion Enterprise features, see the [Junos Fusion Enterprise User Guide](#).

### What's New in Release 19.3R3

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 19.3R3.

### What's New in Release 19.3R2

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

### What's New in Release 19.3R1

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 19.3R1.

### SEE ALSO

---

[What's Changed | 73](#)

---

[Known Limitations | 73](#)

---

[Known Issues | 74](#)

---



---

[Resolved Issues | 75](#)

---

[Documentation Updates | 76](#)

---

[Migration, Upgrade, and Downgrade Instructions | 76](#)

---

## 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 19.3R3 for Junos fusion for enterprise.

### SEE ALSO

---

[What's New | 72](#)

---

[Known Limitations | 73](#)

---

[Known Issues | 74](#)

---

[Resolved Issues | 75](#)

---

[Documentation Updates | 76](#)

---

[Migration, Upgrade, and Downgrade Instructions | 76](#)

---

## Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 19.3R3 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

---

[What's New | 72](#)

---

[What's Changed | 73](#)

---

[Known Issues | 74](#)

---

[Resolved Issues | 75](#)

---

[Documentation Updates | 76](#)

---

[Migration, Upgrade, and Downgrade Instructions | 76](#)

---

## Open Issues

### IN THIS SECTION

- [Junos Fusion for Enterprise | 74](#)

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 10GbE fiber port is using 1-gigabit Ethernet SFP optics, autonegotiation is enabled by default. To bring up the satellite device, BCM recommends to disable the autonegotiation for PHY84756 ports. [PR1420343](#)
- In a Junos fusion for enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, the loop-detect feature does not work for ports 0–23, because 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 a single interface as the ICL instead of a LAG. [PR1447873](#)
- Traffic drop seen from the satellite device to the aggregation device on L3 unicast OSPF routes. [PR1450373](#)

### SEE ALSO

---

[What's New | 72](#)

---

[What's Changed | 73](#)

---

[Known Limitations | 73](#)

---

[Resolved Issues | 75](#)

---

[Documentation Updates | 76](#)

---

[Migration, Upgrade, and Downgrade Instructions | 76](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 19.3R3 | 75](#)
- [Resolved Issues: 19.3R2 | 75](#)
- [Resolved Issues: 19.3R1 | 75](#)

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

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

### Resolved Issues: 19.3R3

- The sdpd process generates a core file at `vfpc_all_eports_deletion_complete` `vfpc_dampen_fpc_timer_expiry`. [PR1454335](#)
- Loop detection might not work on extended ports in Junos fusion scenarios. [PR1460209](#)
- The temperature sensor alarm is seen on EX4300 in a Junos fusion scenario. [PR1466324](#)

### Resolved Issues: 19.3R2

There are no resolved issues in Release 19.3R2.

### Resolved Issues: 19.3R1

- Traffic might be dropped in a Junos fusion for 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 | 72](#)

<a href="#">What's Changed</a>	<a href="#">73</a>
<a href="#">Known Limitations</a>	<a href="#">73</a>
<a href="#">Known Issues</a>	<a href="#">74</a>
<a href="#">Documentation Updates</a>	<a href="#">76</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">76</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 for documentation for Junos fusion for enterprise.

### SEE ALSO

<a href="#">What's New</a>	<a href="#">72</a>
<a href="#">What's Changed</a>	<a href="#">73</a>
<a href="#">Known Limitations</a>	<a href="#">73</a>
<a href="#">Known Issues</a>	<a href="#">74</a>
<a href="#">Resolved Issues</a>	<a href="#">75</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">76</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | [77](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | [78](#)
- [Preparing the Switch for Satellite Device Conversion](#) | [79](#)
- [Converting a Satellite Device to a Standalone Switch](#) | [80](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [80](#)
- [Downgrading from Junos OS](#) | [81](#)

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

[What's New | 72](#)

[What's Changed | 73](#)

[Known Limitations | 73](#)

[Known Issues | 74](#)

[Resolved Issues | 75](#)

[Documentation Updates | 76](#)

# Junos OS Release Notes for Junos Fusion Provider Edge

## IN THIS SECTION

- [What's New | 82](#)
- [What's Changed | 83](#)
- [Known Limitations | 84](#)
- [Open Issues | 85](#)
- [Resolved Issues | 86](#)
- [Documentation Updates | 87](#)
- [Migration, Upgrade, and Downgrade Instructions | 87](#)

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- [What's New in Release 19.3R3 | 83](#)
- [What's New in Release 19.3R2 | 83](#)
- [What's New in Release 19.3R1 | 83](#)

Learn about new features introduced in the Junos OS main and maintenance releases for Junos Fusion Provider Edge.

**What's New in Release 19.3R3**

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.3R3.

**What's New in Release 19.3R2**

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

**What's New in Release 19.3R1**

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.3R1.

SEE ALSO

<a href="#">What's Changed   83</a>
<a href="#">Known Limitations   84</a>
<a href="#">Open Issues   85</a>
<a href="#">Resolved Issues   86</a>
<a href="#">Documentation Updates   87</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   87</a>

**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 19.3R3 for Junos fusion for provider edge.

SEE ALSO

<a href="#">What's New   82</a>
<a href="#">Known Limitations   84</a>
<a href="#">Open Issues   85</a>
<a href="#">Resolved Issues   86</a>
<a href="#">Documentation Updates   87</a>

# Known Limitations

## IN THIS SECTION

- [Junos Fusion Provider Edge | 84](#)

Learn about known limitations in this release 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.

## 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

<a href="#">What's New   82</a>
<a href="#">What's Changed   83</a>
<a href="#">Open Issues   85</a>
<a href="#">Resolved Issues   86</a>
<a href="#">Documentation Updates   87</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   87</a>

## Open Issues

### IN THIS SECTION

- [Junos Fusion Provider Edge | 85](#)

This section lists the Open Issues in hardware and software in Junos OS Release 19.3R3 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 | 82](#)

---

[What's Changed | 83](#)

---

[Known Limitations | 84](#)

---

[Resolved Issues | 86](#)

---

[Documentation Updates | 87](#)

---

[Migration, Upgrade, and Downgrade Instructions | 87](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 19.3R3 | 86](#)
- [Resolved Issues: 19.3R2 | 86](#)
- [Resolved Issues: 19.3R1 | 86](#)

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: 19.3R3

#### *Junos Fusion Provider Edge*

- The statistics of extended ports on a satellite device cluster might show incorrect values from the aggregation device. [PR1490101](#)

### Resolved Issues: 19.3R2

There are no fixed issues in Junos OS 19.3R2 for Junos fusion for provider edge.

### 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

<a href="#">What's New   82</a>
<a href="#">What's Changed   83</a>
<a href="#">Known Limitations   84</a>
<a href="#">Open Issues   85</a>
<a href="#">Documentation Updates   87</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   87</a>

## Documentation Updates

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

SEE ALSO

<a href="#">What's New   82</a>
<a href="#">What's Changed   83</a>
<a href="#">Known Limitations   84</a>
<a href="#">Open Issues   85</a>
<a href="#">Resolved Issues   86</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   87</a>

## Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 88](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 90](#)
- [Preparing the Switch for Satellite Device Conversion | 91](#)
- [Converting a Satellite Device to a Standalone Device | 92](#)
- [Upgrading an Aggregation Device | 95](#)

- Upgrade and Downgrade Support Policy for Junos OS Releases | 95
- Downgrading from Junos OS Release 19.3 | 95

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 **bundle** 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 19.3R3 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/jinstall64-19.3R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot  
source/jinstall-19.3R3.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/jinstall64-19.3R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-19.3R3.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 19.3R3 **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 19.3R3, 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 | 82](#)

[What's Changed | 83](#)

[Known Limitations | 84](#)

[Open Issues | 85](#)

[Resolved Issues | 86](#)

[Documentation Updates | 87](#)

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

#### IN THIS SECTION

- [What's New | 97](#)
- [What's Changed | 135](#)
- [Known Limitations | 146](#)
- [Open Issues | 151](#)
- [Resolved Issues | 167](#)
- [Documentation Updates | 214](#)
- [Migration, Upgrade, and Downgrade Instructions | 215](#)

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- [What's New in Release 19.3R3 | 97](#)
- [What's New in Release 19.3R2 | 97](#)
- [What's New in Release 19.3R1 | 116](#)

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

### What's New in Release 19.3R3

There are no new and changed features for Junos OS 19.3R3 Release.

### 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 99 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 100](#):

**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	MPC2E-3D
	RE-S-1800X4-32G-UB	MPC2-3D-NG
		MPC3E and MPC3E-3D-NG
		MPC4E-3D
		MPC-3D-16XGE
SCBE2	RE-S-1800X4-16G-UPG-BB	MPC2E-3D
	RE-S-1800X4-32G-UB	MPC2-3D-NG
	RE-S-X6-64G-UB	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 default 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 QSFP-4x10-Gigabit Ethernet optics and 4x10-Gigabit Ethernet breakout cables.
- Port 0 of every PIC in 40-Gigabit Ethernet mode using QSFP 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.

[See [Understanding Junos Node Slicing](#).]

- **Support for link fault management (MX2K-MPC11E)**—Starting in Junos OS Release 19.3R2, you can configure IEEE 802.3ah link fault management (LFM) for MX2K-MPC11E on MX2010 and MX2020 routers. You can also configure the following supported LFM features:
  - Discovery and link monitoring
  - Distributed LFM
  - Remote fault detection and remote loopback

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

### *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]` hierarchy 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 2 VPN

- **Support for VPLS (MX series with MPC10 line card)**—Starting in Junos OS 19.3R2, you can configure VPLS on the MPC 10 line card in a MX Series router.

[See [Introduction to VPLS](#) and [VPLS Configuration Overview](#).]

### 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

- **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 Routers](#)[Understanding 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 support 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 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, 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 on 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

- **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](#).]

- **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 QSFP-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](#)].



### *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/rtssock/`

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

- **Physical interface operational status sensor (intf-exp) support on JTI (MX960, MX2010, and MX2020)**—Starting with Junos OS Release 19.3R1, Junos telemetry interface (JTI) uses the sensor `intf-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)(mLDP inband and PIM signaling). [See [Multicast Overview](#).]

### Management

- **Error recovery, fault handling, and resiliency support for MX2K-MPC11E (MX2010 and MX2020)**—Starting in Junos OS Release 19.3R1, the MX2010 and MX2020 routers with MX2K-MPC11E line cards support error recovery, fault handling, and software resiliency. The MX2K-MPC11E line cards support detecting errors, reporting them through alarms, and triggering resultant actions. To view application level errors, use the **show trace node fpc<#> application fabspoked-pfe-redbull** command. To check the status of the card, use the **show chassis fpc pic-status** command. Use the **show chassis errors active** command to view the fault details and the **show system alarm** command to view the alarm details.

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

- **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 are 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 support 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.](#)]

[See [Enabling 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 software-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/urllf_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

Platforms	Routing Engine Supported
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
QFX10002-60C	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 | 135](#)

[Known Limitations | 146](#)

[Open Issues | 151](#)

[Resolved Issues | 167](#)

[Documentation Updates | 214](#)

[Migration, Upgrade, and Downgrade Instructions | 215](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 19.3R3-S1 | 136](#)
- [What's Changed in Release 19.3R3 | 136](#)
- [What's Changed in Release 19.3R2-S6 | 139](#)
- [What's Changed in Release 19.3R2-S5 | 140](#)
- [What's Changed in Release 19.3R2 | 140](#)
- [What's Changed in Release 19.3R1 | 142](#)

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

## What's Changed in Release 19.3R3-S1

### General Routing

- **New commit check for MC-LAG (MX Series)**— We've introduced a new commit check to check the values assigned to the redundancy group identification number on the MC-AE interface ( **redundancy-group-id** ) and ICCP peer ( **redundancy-group-id-list** ) when you configure multichassis aggregation groups (MC-LAGs). If the values are different, the system reports a commit check error. In previous releases, if the configured values were different, the l2ald process would crash.

[See [iccp and .](#)]

### Routing Protocols

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

### VPN

- **Changes to Junos XML operational RPC request tag names (MX480)**—Starting in Junos OS Release, we've updated the Junos XML request tag name for the below operational RPCs. The changes include: `<get-security-associations-information>` is changed to `<get-re-security-associations-information>` and `<get-ike-security-associations-information>` is changed to `<get-re-ike-security-associations-information>`.

[See [Junos XML API Operational Developer Reference.](#)]

## What's Changed in Release 19.3R3

### General Routing

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.
- **Precision Time Protocol (PTP) interface configuration (MX2020, MX2010, MX480, MX960, and MX240)**— Remove the aggregated Ethernet interface association and upgrade the device when configuring PTP interface.
- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection**



**protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

- **Command to view summary information for resource monitor (MX Series routers and EX9200 line of switches)**—You can use the **show system resource-monitor** command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

See [[show system resource-monitor and](#) ].

## Infrastructure

- **Change in support for interface-transmit-statistics statement (MX Series)**—You cannot configure aggregated Ethernet interfaces to capture and report the actual transmitted load statistics by using the **interface-transmit-statistics** statement. Aggregated Ethernet interfaces do not support reporting of the transmitted load statistics. The **interface-transmit-statistics** statement is not supported in the aggregated Ethernet interfaces hierarchy. In earlier releases, the **interface-transmit-statistics** statement was available in the aggregated Ethernet interfaces hierarchy but not supported.

See [\[interface-transmit-statistics\]](#).

## Interfaces and Chassis

- **Change in error severity (MX960, MX240, MX2020, MX480, MX2008, and MX2010)**—Starting in Junos OS Release 19.3R3, we've reduced the severity of the CRC error 'XR2CHIP\_ASIC\_JGCI\_FATAL\_CRC\_ERROR' from Fatal to Major. Earlier, these errors caused the line card to be reset, if the **interasic-linkerror-recovery-enable** statement was configured. Now, these errors only disable the Packet Forwarding Engines that are affected. With this change, the **interasic-linkerror-recovery-enable** statement has no effect when this error occurs because of the reduced severity.

**NOTE:** This behavior change is applicable to the following line cards only: MPC5E, MPC6 MPC7, MPC8, and MPC9.

## Juniper Extension Toolkit (JET)

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

See [traceoptions \(Services\)](#).

## Network Management and Monitoring

- **Change in Startup Notification after GRES (MX Series routers)**—In Junos OS Release 19.3R3, 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.
- **Enhancement to the show snmp mib command**—Starting in Junos OS Release 19.3R3, a new option, **hex**, is supported to display the SNMP object values in the hexadecimal format. In earlier releases, the **show snmp mib** command displays the SNMP object values in ASCII and decimal format only.

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

### *Services Applications*

- **New option for configuring delay in IPsec SA installation**—In Junos OS Release 19.3R3, you can configure the **natt-install-interval seconds** option under the **[edit services ipsec-vpn rule rule-name term term-name then dynamic]** hierarchy to specify the duration of delay in installing the IPsec SA in a NAT-T scenario soon after the IPsec SA negotiation is complete. The default value is 0 seconds.

### *Subscriber Management and Services*

- **Improved tunnel session limits display (MX Series)**—Starting in Junos OS Release 19.3R3, the **show services l2tp tunnel extensive** command displays the configured value for maximum tunnel sessions. On both the LAC and the LNS, this value is the minimum from the global chassis value, the tunnel profile value, and the value of the Juniper Networks VSA, Tunnel-Max-Sessions (26–33). On the LNS, the configured host profile value is also considered.

In earlier releases, the command displayed the value 512,000 on the LAC and the configured host profile value on the LNS.

[See [Limiting the Number of L2TP Sessions Allowed by the LAC or LNS](#).]

## **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*

- **IPv6 statistics added in aggregate Ethernet interfaces (MX2K-MPC11E line cards)**—In Junos OS Release 19.3R2 and later, fixed an issue wherein IPv6 statistics were not being added for aggregate Ethernet Interfaces.

An example showing IPv6 statistics can be seen here:

```
user@device# run show interfaces ael extensive
```

```
Physical interface: ael, Enabled, Physical link is Up
  Interface index: 304, SNMP ifIndex: 830, Generation: 342
  Link-level type: Flexible-Ethernet, MTU: 16000, Speed: 200Gbps, BPDU Error:
None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:00:ab:cd:ef, Hardware address: ab:cc:dd:34:af:8f
  Last flapped   : 2020-11-04 14:09:06 IST (00:21:16 ago)
  Statistics last cleared: Never
```

```

Traffic statistics:
  Input  bytes   :           229134933812           16068257904 bps
  Output bytes   :           229139043862           16067232936 bps
  Input  packets:           228202271             2000539 pps
  Output packets:           228206299             2000411 pps
IPv6 transit statistics:
  Input  bytes   :           229133203684
  Output bytes   :           229137314396
  Input  packets:           228188259
  Output packets:           228192346
Dropped traffic statistics due to STP State:
  Input  bytes   :           0
  Output bytes   :           0
  Input  packets:           0
  Output packets:           0
MAC statistics:
  Broadcast packets          Receive          Transmit
  Multicast packets          13603             13522

```

- **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](#).]

### Services Applications

- **Update to CLI option for configuring the version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 and the latest available version**—In Junos OS Release 19.3R2, the **version-3** option under the **[edit services softwire softwire-concentrator map-e]** hierarchy is optional. You can choose to configure the version number to distinguish between the currently supported version of the Internet draft draft-ietf-softwire-map-03 and the latest available version. In the earlier Junos OS releases, if you do not configure the **version-3** option, the configuration results in an error.

[See [map-e](#).]

### 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.
- **XML RPC equivalent included for the show bgp output-scheduler | display xml rpc CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.3R1, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler](#).]

#### *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 software 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 software 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 software 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 | 97](#)

[Known Limitations | 146](#)

[Open Issues | 151](#)

[Resolved Issues | 167](#)

[Documentation Updates | 214](#)

[Migration, Upgrade, and Downgrade Instructions | 215](#)

## Known Limitations

### IN THIS SECTION

- General Routing | [146](#)
- Infrastructure | [149](#)
- Interfaces and Chassis | [149](#)
- MPLS | [150](#)
- Platform and Infrastructure | [150](#)
- Routing Protocols | [151](#)

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](#)
- In an MPC8 line card, enabling both the bandwidth and the **flex-flow-sizing** statements might result in J-Flow service getting disabled because the memory requested by the **flex-flow-sizing** statement is not allocated or denied. [PR1413513](#)
- 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](#)
- MX Series routers report Routing Engine and FPC policer violation when DDoS violation occurs. [PR1439427](#)
- 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](#)

- When an xe- interface working in 1-Gigabit mode is added to a member link of an aggregated Ethernet interface, the speed of the aggregated Ethernet interface is incorrectly shown as 10 Gbps. There is no functional impact. This is a display issue. [PR1449887](#)
- 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](#)
- When an MX2K-MPC11E boots, the harmless errors which do not have a functional impact are seen. [PR1457655](#)
- We cannot change pic-mode to port-mode within 20 seconds as MPC11E needs more than 15 seconds to bring the PIC online. [PR1458617](#)
- 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](#)
- Changes to RLT with PS anchored over it is not recommended. For more information, refer to: [https://www.juniper.net/documentation/en\\_US/junos/topics/reference/configuration-statement/anchor-point-edit-interfaces-ps0.html](https://www.juniper.net/documentation/en_US/junos/topics/reference/configuration-statement/anchor-point-edit-interfaces-ps0.html). [PR1460898](#), [PR1460910](#)
- When untagged frames are received on the 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 an interface on MPC11E line card, in some instances, when the interface extensive output is displayed for the first time, the local statistics of the logical interface are also shown in the transit statistics. This is a transient issue and will correct itself in the next command display. [PR1467181](#)

- When a connection between aftman and aft-ulcd is dropped, aft-ulcd will exit. As a result, FPC reboots. [PR1467246](#)
- 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](#)
- The following error message might appear: **Failed to complete DFE tuning**. This error message has no functional impact and can be ignored. [PR1473280](#)
- The aftd hogs on executing the command to clear the VPLS table and MACs are not learned for less than 5 minutes. [PR1473334](#)
- Firewall filter counters do not get the expected value for **family inet firewall** in the input direction on an AF interface because share-bw-Policer is not yet supported on AF. [PR1474389](#)

## 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](#)
- In scenarios when an MPLS LSP is using the whole or closer to all available bandwidth reservation on links, LSP reestablishment (P2MP LSP, for instance on deactivating and activating the associated routing-instance) might take more than around 30 seconds. The issue occurs because the time needed for the network to converge after LSP down/up events. Bandwidth resource is not released during the timing window so bandwidth is still occupied by the LSP instance going down, as more bandwidth is not available. LSP fails to reestablish rapidly going for a second try after the predefined timeout. Therefore, we observe the delay. [PR1509791](#)

## Platform and Infrastructure

- On all platforms running Junos OS, execution of Python scripts through enhanced automation does not work on veriexec images. [PR1334425](#)
- 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](#)
- Salt execution module junos.install\_os fails to remove the tmp image file `__salt.tmp.ecc1ME` on target minion after image installation. [PR1432123](#)
- On the MX Series platform with Protocol Independent Multicast (PIM) implemented and the number of IGMP groups exceeding 15,000, join message (S,G) might not be created after GRES. [PR1457166](#)
- NPC core file is generated at `trinity_rt_iff_attach,pfe_ifl_family_attach,ifrt_ifl_family_adder,ifrt_ifl_family_add_vector,ifrt_command_handler`. [PR1461892](#)
- Unknown unicast filter applied in an EVPN routing instance blocks the unexpected traffic. [PR1472511](#)

## Routing Protocols

- [BGP] When scaling RIB to 80 million after FPC restart, cannot scale on backup Routing Engine. [PR1444073](#)
- Users that replace simple VLAN interfaces with PS over RLT might notice an increase in FPC CPU usage. This is in keeping with the increased processing and resources needed to support these types of interfaces, which are similar in this regard to that of an aggregated Ethernet interface. [PR1396925](#)
- Packet drop and CPU spike on the Routing Engine might be seen in certain conditions if labeled-unicast protection is enabled for a CsC-VRF peer. [PR1456260](#)

### SEE ALSO

<a href="#">What's New</a>		<a href="#">97</a>
<a href="#">What's Changed</a>		<a href="#">135</a>
<a href="#">Open Issues</a>		<a href="#">151</a>
<a href="#">Resolved Issues</a>		<a href="#">167</a>
<a href="#">Documentation Updates</a>		<a href="#">214</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>		<a href="#">215</a>

## Open Issues

### IN THIS SECTION

- [EVPN](#) | [152](#)
- [Forwarding and Sampling](#) | [153](#)
- [General Routing](#) | [153](#)
- [Infrastructure](#) | [161](#)
- [Interfaces and Chassis](#) | [161](#)
- [Layer 2 Ethernet Services](#) | [162](#)
- [MPLS](#) | [162](#)
- [Platform and Infrastructure](#) | [163](#)
- [Routing Protocols](#) | [164](#)
- [Services Applications](#) | [166](#)
- [Subscriber Access Management](#) | [166](#)

- User Interface and Configuration | 166
- VPNs | 166

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

- A few duplicate packets might be seen in an active-active EVPN scenario when the remote PE device sends a packet with an IM label due to MAC not learned on the remote PE device, but learned on the A/A local PE device. The nondesignated forwarder sends the IM-labeled encapsulated packet to the PE-CE interface after MAC lookup instead of dropping the packet, which causes duplicate packets to be seen on the CE side. [PR1245316](#)
- With Junos OS Release 19.3R1, VXLAN OAM host-bound packets are not throttled with DDoS policers. [PR1435228](#)
- When DHCP is used with EVPN, the Layer 2 learning daemon adds a destination route to the kernel with a "permanent remote" flag, while the DHCP process adds a destination route with a "permanent" flag. There could be a race condition where the Layer 2 learning destination route is overwritten by the DHCP route, causing the remote flag to be deleted. This subsequently leads the ARP route to age out in the kernel. To ensure that DHCP routes are not added to the kernel, the "forward-only" statement must be configured under forwarding-options dhcp-relay. [PR1439568](#)
- Three BGP replication flaps are seen on the new master Routing Engine after GRES. The route synchronization issue is also seen between the Routing Engines without GRES. [PR1441925](#)
- In an EVPN-VXLAN core isolation scenario, the server is multihomed to the leaf devices through LACP interfaces. If graceful restart is enabled, when you reboot the system or restart routing on the leaf device, the core isolation does not work. If you reboot the system, the issue results in the leaf device silently dropping the traffic sent from the server during the time window between LACP coming up and BGP coming up. If you restart routing, there might be no traffic drop because of the graceful restart. [PR1461795](#)
- The rpd crashes upon changing CoS configuration. The following error is seen:  
**COSD\_CLASS\_8021P\_UNSUPPORTED: ieee-802.1 classifier is not valid on interface ge-0/0/1.0.**  
[PR1530991](#)



## Forwarding and Sampling

- The skip-service configuration does not work with IPv6 NDP negotiation or ping. [PR1074853](#)
- When an IPv4 prefix is added to a prefix list referenced by an IPv6 firewall filter, then the log message **Prefix-List [Block-Host] in Filter [Protect\_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized** is not seen. [PR1395923](#)
- When GRES is triggered by SSD hardware failure, the syslog error of **rpd[2191]: krt\_flow\_dfwd\_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out** might be seen. [PR1397171](#)
- A few error messages are observed in the dfwd process at the **[edit interfaces ae2 unit 0 family inet]** hierarchy level, filter the index for referenced filter **input\_ipv4\_ngn\_filter** is not defined. [PR1433146](#)
- After upgrading to Junos OS Release 19.2R2-S1.4, the following false-positive commit warning is seen even though there is no configuration change under 'forwarding-options vxlan-overlay-load-balance': # commit and-quit re0: [edit] 'forwarding-options' warning: vxlan-overlay-load-balance configuration for forwarding options has been changed. A system reboot is mandatory. Please reboot \*ALL\* routing engines NOW. Continuing without a reboot might result in unexpected system behavior. configuration check succeeds re1: configuration check succeeds commit complete re0: commit complete Exiting configuration mode. [PR1459833](#)
- Subscriber statistics are not supported for V44 EP ports configuration. When RADIUS accounting is enabled for the subscriber interface, statistics retrieval has issues. Statistics will not be reported and subscriber logout can take 5 minutes. [PR1505409](#)

## General Routing

- The fxp0 is marked "Dest-route-down" because of specific operations such as disabling and enabling operations. [PR1052725](#)
- If a Layer 3 interface is receiving a GRE-encapsulated packet and the interface has the filter "family any" with action as "mirror" and the filter "family inet" with action as "decapsulate gre" attached at ingress, 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 the GRE headers), the workaround is to deactivate or disable the "decapsulate gre" action of the filter. [PR1090854](#)
- On an MX104 platform, when using "snmpbulkget" or "snmpbulkwalk" (for example, used by the SNMP server) on a chassisd-related component (for example, jnxOperatingEntry), a hardware limitation might result in high CPU usage and slow response by the chassisd process, which may also lead to query timeout on the SNMP client. In addition, the issue may not be seen when using SNMP query for interface statistics. For avoiding the issue, either of the two options listed below could be regarded as a workaround: Option 1—Use snmpget or snmpwalk instead of snmpbulkget or snmpbulkwalk. Include "-t 30" options

when doing SNMP query, for example, "snmpget -v2c -c XX -t 30". Option 2—Use "-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 the 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 synchronization of the configuration 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](#)
- The mgd process might crash when executing the **get-configuration** RPC several times. [PR1247465](#)
- If a VM host snapshot is taken on the alternate disk and there is no further VM host software image upgrade, the expectation is that if the current VM host image gets corrupted, the system will boot from the alternate disk so that the user can recover the primary disk to restore the state. However, if the host root file system is corrupted, the node boots with the previous VM host software instead of booting from the alternate disk. [PR1281554](#)
- In some MX Series deployments running Junos OS, random syslog messages are observed for FPCs: **fpcx ppe\_img\_ucode\_redistribute Failed to evict needed instr to GUMEM - xxx left**. These messages are not an issue and might not have a service impact. These messages will be addressed as information level messages. On a Packet Forwarding Engine, there are dedicated UMEM and shared GUMEM memory blocks. This informational message indicates some evicting events between UMEM and GUMEM and can be safely ignored. [PR1298161](#)
- On MX Series routers, the routing information base (RIB, that is routing table) and forwarding information base (FIB, that is forwarding table) might not synchronize in a large-scale network because of 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 releases from 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 an ingress rewrite premise. Without this statement, things work just fine. [PR1318984](#)
- With regard to FPC restarts or Virtual Chassis splits, the design of the MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections and hence the reactions to failure situations might not be handled in a graceful way. For example, TCP connection timeout because of jlock hog crossing the boundary value (5 seconds) can cause bad consequences in MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling the marker infrastructure in an MX Series Virtual Chassis setup. Currently, there is no workaround on enabling the marker as it was causing a lot of issues in MX Series Virtual Chassis when we tried to enable it. [PR1332765](#)
- The first packet pertaining to the J-Flow Packet Forwarding Engine sensor in UDP mode is missing after the line card reboots on an 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 long, the backup Routing Engine might not be able to synchronize with the master Routing Engine and will not be ready for GRES. [PR1346300](#)

- The backup Routing Engine might crash after GRES occurs more than 10 times. [PR1348806](#)
- In some cases, online insertion and removal (OIR) of a MIC on an FPC can lead to traffic destined to the FPC to be silently dropped or discarded. The only way to recover from this is to restart the FPC. The issue will not be seen if you use the corresponding CLI commands to take the MIC offline and bring it back online. [PR1350103](#)
- Packet Forwarding Engine error messages "localtp\_offload\_tx\_errcheck: failed to send packet 11 times in last one second" may be seen on a Junos Node Slicing deployment. No known impact. [PR1359149](#)
- 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% 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 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 also show an alarm. [PR1360426](#)
- When an FPC is booting up (either during unified ISSU or router reboot, or FPC restart), I2C timeout errors for an SFP transceiver are seen and the I2C action is not completed because the device is busy. When the FPC is up, all the I2C transactions to the device are normal, 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 transactions from the 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, link not coming up or the MIC itself not booting up on restart. I2C Failed device: group 0xa0 address 0x70 Failed to enable PCA9548(0x70):grp(0xa0)->channel(0) mic\_sfp\_select\_link:MIC(0/0) - Failed to enable PCA9548 channel, PCA9548 unit:0, channel ID: 0, SFP link: 0 mic\_sfp\_id\_read: Failed to select link 0. The only way to recover from these failures is to detect and replace the faulty SFP/SFP+ transceiver plugged into the GMIC2 ports. [PR1375674](#)
- On MX480, a few xe- interfaces go down with the error message `if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error`. [PR1377840](#)
- The virtio throughput remains the same for multi-queue and single-queue deployments. [PR1389338](#)
- The MX104 router has the following limitations in error management: The show chassis fpc error CLI command is not available for MX104 in the 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 Disable PFE for major errors, Junos OS does not disable its only 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](#)
- Deleting the **active-lease-query** configuration may sometimes lead to the generation of core files, or the TCP connection may remain active. [PR1415990](#)
- The `show services hybrid-access sessions`, `show services hybrid-access statistics`, and `show services hybrid-access tunnels` commands display values of zero for hybrid access gateway traffic statistics even when traffic is active in the gateway sessions and the tunnels. [PR1419529](#)
- If Hypertext Transfer Protocol (HTTP) Header Enrichment function is used, the traffic throughput decreases when traffic passes through Header Enrichment. [PR1420894](#)
- For ALGs with out-to-in sessions, if the data sessions come from an IP address that is different from the IP address available in the control sessions with the NAT rule matching, such ALGs should have the match condition for the destination-address as any and not a specific IP address, or you must add all possible IP addresses from where the data sessions for the ALGs can come. [PR1421555](#)
- On an MX204 platform, the allocation of MAC address for the second PIC in the FPC may fall out of the MAC address pool, which may further cause a MAC conflict in the network. [PR1422679](#)
- Junos OS Releases 19.1 and later support RFC 8231/8281 compliance by default. However, if the controller is not compliant with RFC 8231/8281, backward compatibility can be configured to fall back to pre-RFC 8231/8281 behavior. [PR1423894](#)
- A new configuration statement, **nh-detail**, is added in **show route detail**, and **show route extensive**, which display next-hop detail in extensive format with proper and formatted XML tags. [PR1424930](#)
- On MPC10, error message **failed, Return code: 500** is seen with baseline. [PR1431552](#)
- JSD generates a core file when aggressively subscribing and unsubscribing both gRPC and gNMI subscriptions from multiple sessions. [PR1433744](#)
- On MX Series platforms, if the clock frequency is slowly changing on CB0 (slow drift), the clock source for MPC-3D-16XGE-SFPP might not be changed to CB1, which will cause interfaces on it to go down and remain in the downstate. [PR1433948](#)
- In gRIBI, programmed routes references 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>). 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::<grp-id> and nexthop ID <FC01> through an actual indirect IPv4 gateway address. This is a Junos OS limitation. [PR1434050](#)
- Packet Forwarding Engine major error **QSS\_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-15xQSFP: L2overGRE is not supported in Junos OS Release 19.3R1. Though the configuration gets committed, the feature will not work. [PR1435855](#)
- Traffic errors do not get policed as expected after being locally switched for VLAN 100 and 101, while verifying the selective local-switching functionality with 4000 VLANs. [PR1436343](#)
- The Routing Engine interprets any input from the console port as interrupts. Depending on the frequency, console noise impacts the Routing Engine interruption handling to different extents, even with the current mechanism. When the interrupt frequency is too high for the Routing Engine to handle, the impact might vary from the line card reboot (partial impact) to the Routing Engine reboot (chassis-wide impact). [PR1436386](#)
- QFX Series routers report Routing Engine and FPC policer violation when DDoS violation occurs. [PR1439427](#)
- 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, the telemetry value corresponding to -infinity should be equivalent to IEEE 754, which is a single precision float and the 32-bit 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](#)
- On the MPC7E, MPC8E, and MPC9E cards, egress stream flush failure and silent dropping of traffic might occur in a rare occasion for a repeatedly flapping link. [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](#)
- Establishing a BGP session over the GRE tunnel fails when the router receives the BGP packets encapsulated as GRE and uses the firewall filter action to de-capsulate the GRE header. [PR1443238](#)
- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card at FRS. [PR1444516](#)
- The **show ddos-protection protocols arp statistics |display xml** command does not show APR violation packets and are not incremented. [PR1449968](#)
- Error message are observed during loading of the RLI configurations. [PR1451213](#)
- Error messages **mpc0 aftd-trio[17006]: [Error] IF:Failed to setup bridge family - invalid IFL, proto:Bridge iffIndex:27** is seen. [PR1451259](#)
- When you use the **replace pattern** command to replace the name in the apply-group, the mgd crashes. [PR1452136](#)

- Physical interface policer are not supported in Junos OS Release 19.3 for MPC11. [PR1452963](#)
- SDK upgrade from Junos OS Release 5.10 to 6.5.16 is done in this release for the following ACX platforms ACX1000, ACX1100, ACX2000, ACX2100, ACX4000, and ACX500. [PR1454144](#)
- With FPC restart done in presence of dynamic tunnels configuration, traffic might not resume sometimes. [PR1454325](#)
- When you edit a command and run the command from CLI command history, the timestamp might not appear. [PR1454387](#)
- Behavior has been modified to display the correct protocol number instead of 255 whenever unknown protocol is encountered. [PR1454792](#)
- 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 space GUI for both the MX10008 and MX10016 devices. [PR1455383](#)
- SmiHelperd process is not initialized in Junos OS PPC releases. [PR1455667](#)
- With **logical-system** configuration, filter-based GRE encapsulation is not working. [PR1456762](#)
- Effective FIB download rate of an MPC11E is lower than that of MPC10E by 30 percent because MPC11E has 8 targeted Packet Forwarding Engines while MPC10E has 2-3 targeted Packet Forwarding Engines. [PR1456816](#)
- Syslog/link-down action on low-light through **set interfaces <x> optics-options** is not supported on MPC11E. [PR1456894](#)
- 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](#)
- 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](#)
- NAT performance is impacted with remote syslog enabled. [PR1460211](#)
- The **show ancp subscriber detail** command displays incorrect values for line attributes. [PR1460812](#)
- 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](#)
- A BFD session might flap when it is moving to aggressive interval after coming up with slow/nonaggressive interval. This issue is mainly seen in a scaled setup. [PR1462775](#)
- Backport jemalloc profiling CLI support to all releases where jemalloc is present. [PR1463368](#)

- The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline/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](#)
- On MX2000 platforms with MPC11E, if Link Aggregation Control Protocol (LACP) link protection is configured, when the Physical Interface Card (PIC) gets offline on which the active link is located, the failover time might be more than 2 seconds. Due to this issue, more traffic loss might happen than expected. [PR1464652](#)
- Commit script does not apply changes in private mode unless a commit full is performed. [PR1465171](#)
- A BFD session might flap when it is moving to aggressive interval after coming up with slow/non-aggressive interval. This issue is mainly seen in a scaled setup. [PR1465285](#)
- **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](#)
- 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, aggIfName:ps1.0 memberIfName: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, aggIfName:ps1.0 memberIfName:lt-5/0/0.32767. [PR1466531](#)
- When a connection between aftman and aft-ulcd is dropped, aft-ulcd will exit. As a result, FPC reboots. [PR1467246](#)
- Not able to get the service sessions when NAT64 is configured with destination-prefix length 32 [PR1468058](#)
- VPLS statistics are not showing up in 8000 scaled VPLS scenarios after subsequent deactivation and activation of interfaces for the first 10 minutes after which the statistics are updated correctly. There is no functionality impact and issue is seen with only in 8000 scaled scenarios. Complete design flow study needs to be done for resolution. Will be fixed in the upcoming release. [PR1468535](#)
- This issue observed with BGP rib-sharding and update-threading configuration which has SCALE VRFs and below steps are performed. 1. Delete all VRFs. 2. Rollback or reconfigure the same VRFs immediately. The issue can also be seen if the above steps are performed directly or indirectly such as loading configuration which doesn't have VRFs and rollback immediately. [PR1469873](#)



- An occasional afd core is observed on setups with GNF configuration after rebooting another line card such as MPC8 within the same GNF. [PR1473403](#)
- After doing Routing Engine switchover on Bsys , the AF interface on peer router shows status as down (instead of being up) with reason Packet Forwarding Engine is down on the GNF with MPC11E. [PR1473555](#)
- Creating more than 500 PS interfaces on MPC11 takes a long time. [PR1473559](#)
- IS-IS and micro BFD sessions are not coming up during baseline for MPC10E card. [PR1474146](#)
- On rebooting the external server, the SNMP values configured within /etc/snmp/snmpd.conf at the server get overwritten with content from the JDM SNMP stanza. Trap configuration changes get completely removed. JDM restart or stop/start does not change the host /etc/snmp/snmpd.conf file, only system reboot of the server does. [PR1474349](#)
- Changing framing modes on a CHE1T1 MIC between E1 and T1 on a MPC3E NG HQoS line card will cause the PIC to go offline. [PR1474449](#)
- If the LDP is scaled to 480,000 over four interface, the mirror subsystem times out and goes down. [PR1474965](#)
- When specific hardware failure conditions occur in MX2000 platforms, fabric healing attempts to auto-heal the fault location in three phases to prevent traffic get silently dropped and discarded. In such fault conditions, fabric healing process in last phase-3 might not be able to decide which FPC slot should be marked faulty and off-line all the FPCs in the system reporting fabric destination time out. [PR1482124](#)
- The next-generation services MX Series SPC3 services card can exhibit inconsistent behavior when the vmhost image is installed on the next-generation Routing Engine (NG-RE): RE-S-X6-64G-UB. To use the NG-RE, a workaround is available. See the "Problem" and "Workaround" section for more detail. Other Routing Engines that are compatible with next-generation services do not experience this problem. These Routing Engines are RE-S-1800X4-16G-UPG-BB, RE-S-1800X4-32G-UB, RE-S-1800X4-16G-UPG-BB, RE-S-1800X4-32G-UB. [PR1482334](#)
- Support for upgrading PSM's firmware on the MX2000 platforms. [PR1489967](#)
- Viewing a large file from vFPC console using cat application might hog the console. [PR1493805](#)
- Auto-configure stacked vlan ranges" is not handled on vMX platform and VLANs are not programmed on driver causing the issue. [PR1503538](#)
- On all Junos OS platforms with the JTI (Juniper Telemetry Interface) configured, the rpd may crash when there is telemetry streaming is in progress and meanwhile there is a network churn. This is a timing issue, and the rpd will recover automatically. [PR1505425](#)
- In an EVPN scenario with VRRPv6 is used, the Ethernet source MAC address might be used for IPv6 mac-ip binding when the NA is sent from VRRPv6 master. As this unexpected behavior is triggered on regular intervals, it causes the entries to keep refreshing in the EVPN database because NS from VRRPv6 master changes the mac-ip binding. This impacts the traffic. [PR1505976](#)



- A 10-Gigabit Ethernet interface configured with WAN-PHY framing might flap continuously if the hold-down timer is set to 0 (which is the default). This is not applicable to an interface with the default framing LAN-PHY. [PR1508794](#)
- Disruptive switchover (no GRES or NSR configured) can lead to stale Periodic Packet Management (PPM) entries programmed on the new master Routing Engine, if both GRES and NSR are activated after disruptive switchover and then a GRES switchover is performed, BFD sessions might flap continuously. [PR1518106](#)
- Constant messages are flooding in log **summit\_pic\_port\_profile\_isvalid: VALID Port profile**. These messages are cosmetic messages and do not have any service impact. [PR1464879](#)
- With NAT/Stateful-firewall/TCP tickle (enable by default) configured on MS-MPC/MS-MIC, the vmcore crash sometimes along with mspmand process might crash when large-scale traffic flows (for example, million flows) are processed by it. [PR1482400](#)
- In Junos Telemetry Interface (JTI) with gRPC Network Management Interface (gNMI) sensor scenario, when the telemetry data of transceivers diagnostics is exported. The Kernel might not handle it correctly. It might delay exporting the specific telemetry data located all leaves under gNMI sensor /components/. [PR1504733](#)
- Issue occurs only with specific RLT interface configuration followed by FPC restart. In this PR, GRES was done when FPC was doing the interface cleanup after FPC restart. Amount of time taken for interface cleanup on FPC depends on scale. If GRES is done before FPC comes to clean state, backup Routing Engine might not be in sync with master that may cause interface states inconsistent after GRES. [PR1518543](#)
- On all Junos OS platforms that support Open vSwitch database (OVSDB), the vgd core file might get generated when the OVSDB server is restarted. The vgd daemon restart after the core might cause traffic impact. This issue happens when OVSDB server is disconnected and the device (switch/router) sends some updates events to server. [PR1518807](#)

## 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](#)
- **IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151)** error message is observed continuously in AD with base configurations. [PR1485038](#)

## Interfaces and Chassis

- After GRES, the 1-Gigabit Ethernet changes to 10-Gigabit Ethernet. [PR1326316](#)
- In an L2TP scenario when the MX Series router functions as an L2TP tunnel switch (LTS), there is a memory leak in the 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 a large-scale subscriber environment, the two Routing Engines generate core files when changing aggregated Ethernet member link configuration. [PR1375638](#)
- The SFP index in Packet Forwarding Engine starts at 1, while the port numbering starts at 0. This causes confusion in the log analysis. [PR1412040](#)
- 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 restarts (**restart chassis-control**) reboot both Routing Engines (**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](#)

## Layer 2 Ethernet Services

- DHCPV6\_LEASEQUERY counter might not be as expected in the **show dhcpv6 relay statistics** command output. [PR1521227](#)
- With access-profile being configured under the **forwarding-options dhcp-relay** or **system services dhcp-local-server** statement, there might be memory leak in jdncpd when DHCP request is processed. The leaked memory size for each processing will be the number of bytes equal to the length of the access profile. [PR1525052](#)

## MPLS

- Root XML tag in the output is changed from **rsvp-pop-and-fwd-info** to **rsvp-pop-and-fwd-information** to be consistent with the XML tag convention. [PR1365940](#)
- 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](#)

## 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, the 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.
- An accuracy issue occurs with three-color policers of both types 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](#)
- In Junos OS EVPN/MPLS, an MPC core file is triggered after deleting the MAC on bridge domains in random scenarios. The core file is generated when **MAC delete request** comes from the Routing Engine to the Packet Forwarding Engine for an EVPN instance in scenarios where L2 route information on Packet Forwarding Engine is in corrupt state. [PR1388373](#)
- On MX Series routers with MPCs, unicast traffic might drop when the destination is reachable over an integrated routing and bridging (IRB) interface and a label-switched interface (LSI) with two next hops. [PR1420626](#)
- On the EX9208 and MX480 devices, traffic loss is observed if the ingress and egress ports are on different FPCs. [PR1429714](#)
- For the bridge domains configured under an EVPN instance, ARP suppression is enabled by default. This enables the EVPN to proxy the ARP, and reduces the flooding of ARP in the EVPN networks. As a result, storm control does not affect the ARP packets on the ports under such bridge domains. [PR1438326](#)
- Arrival rates are not seen at the system level when **global-disable fpc** is configured on QFX Series. [PR1438367](#)
- The error message **LUCHIP(0) PPE\_0 Errors sync xtxn error** is continuously flooding after the unified ISSU on MX104. [PR1446800](#)
- 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 after a manual GNF Routing Engine mastership switchover attempt with the **request chassis routing-engine master [acquire|release|switch]** CLI command from either GNF Routing Engine CLI. [PR1456565](#)
- On the MX Series platform with Protocol Independent Multicast (PIM) implemented and the number of IGMP groups exceeding 15,000, the join message (S,G) might not be created after the GRES. [PR1457166](#)
- The CFM REMOTE MEP does not come up after configuration or if the state remains in Start state. [PR1460555](#)

- In NTP with the boot-server scenario, when the router or switch boots, the NTP daemon sends an ntpdate request to poll the configured NTP boot-server to determine the local date and time. If the ntpdate is not activated correctly while the device is booting, the ntpdate might not work successfully. Then, some cosmetic error messages of time synchronization might be seen, but there is no impact with time update because the NTP daemon will update the time eventually. [PR1463622](#)
- With multiple different fixed sized traffic streams configured at 1,000,000 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 instruction of an MPC11E line card, small varying packet drops for every iteration on ae1 on disabling ae0. The drops might vary from 200 to certain 1000 frames. [PR1464549](#)
- The TCP connection is still established even after the interface is deleted/deactivated/disabled. **show agent sensors** still shows the subscription. [PR1477790](#)
- Traffic loss is observed after ISSU, while enable/disable and activate/deactivate the interface. [PR1493723](#)

## 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](#)
- When interoperating with other vendors in a draft-rosen multicast VPN, by default the Junos OS attaches a route target to multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities will be excluded from propagating if the BGP route target filtering is enabled on a device running Junos OS. Note that **draft-rosen-idr-rtc-no-rt** has been created in IETF to document this issue and carry the proposed fix through standards. [PR993870](#)
- Certain BGP traceoption flags (for example, open, update, and keepalive) might result in (trace) logging of debugging messages that do not fall within the specified traceoption category, which results in some unwanted BGP debug messages being logged to the BGP traceoption file. [PR1252294](#)
- LLDP and OSPF are in the in-sync state because an IGP interface is down with LDP synchronization enabled for OSPF. According to the current analysis, the 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](#)
- In rare cases, RIP replication might fail as a result of performing NSR Routing Engine switchovers when the system is not NSR ready. [PR1310149](#)
- BGP I/O threading was added in Junos OS Release 16.1R1 whereby BGP writes were batched to improve efficiency. This might sometimes lead to some latency in sending BGP update while reacting to certain network events. [PR1332301](#)

- SCP command with routing option (-JU) is not supported. [PR1364825](#)
- At scale, a GNF with PS over RLT and multiple MPCs might show BFD flap at recovery. [PR1386574](#)
- Performance improvement with the **addpath-optimization** statement configured will vary across releases because of variability of baseline convergences without the configuration statement. [PR1395684](#)
- When **auto-export** is configured on the devices running Junos OS in two VPN routing and forwarding (VRF) instances, the routes get exported from or 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 due to the infinite loop of route flashing in the VRF instances. [PR1402140](#)
- During NSR initial state replication on a scaled setup, while BGP state replication is still ongoing, the BGP task replication might get marked as completed. This is because BGP replication is triggered and controlled through the backup Routing Engine. You must check the output of the **show bgp replication** command to confirm whether replication has actually completed. This corner case scenario is valid only on a scaled setup and during initial state synchronization. [PR1404470](#)
- 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 the route distinguisher and then reactivate. [PR1433913](#)
- On the MX2010 Series routers, the BFD session on IS-IS step up flaps during the ISSU - FRU upgrade stage. [PR1453705](#)
- With NSR enabled, the current BGP design supports 3000 BGP IPv6 peers or 8000 BGP IPv4 peers. If you try to bring up more than 3000 BGP IPv6 sessions or more than 8000 BGP IPv4 sessions, the rpd might crash. [PR1461436](#)
- If the **wide-metrics-only** statement is enabled for any IS-IS level and a metric configured on the IS-IS enabled interface for that level has ASCII representation in decimal more than 6 characters long. This interface metric for that level will be merged with 'priority' field value in the output of **show isis interface detail**. [PR1482983](#)
- Improve BGP-LS NLRI handling with regard to 'BGP-LS identifier TLV' both at BGP-LS receiving speaker and at a router performing TED export into BGP-LS. [PR1521258](#)
- On a Junos OS platform, when flood-group is configured on any interface and after IS-IS adjacency flap, the device sends self-generated LSPs only without incrementing the flapped LSP updates received from the neighbor. As a result, the LSP database do not sync. [PR1526447](#)

## 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

- Subscriber username with space in between is not displayed by [show network-access aaa subscribers statistics username "<>"] command. [PR1518016](#)

## User Interface and Configuration

- When the nested apply-groups applied is getting deleted, the logical interface under the nested groups is not removed. [PR1427962](#)

## VPNs

- In the MVPN environment with the SPT-only option, if the source or receiver is connected directly to the c-rp PE device and the MVPN data packets arrive at the c-rpce PE device before its transition to SPT, the MVPN data packets might be dropped. [PR1223434](#)
- P2MP LSP replication to the backup Routing Engine is not proper. [PR1453900](#)
- MPC10E: NG-MVPN GTM for remote IPv6 source is not working. [PR1454163](#)
- The LSP might stay down if you configure both the virtual-tunnel (VT) interface and vrf-table-label in an MVPN scenario. In this case, VT is preferred over LSI. Later when the VT interfaces are deleted, there is no notification to MVPN, indicating that LSI is still available. Hence traffic loss might be seen. [PR1474830](#)

## SEE ALSO

[What's New | 97](#)

[What's Changed | 135](#)

[Known Limitations | 146](#)

[Resolved Issues | 167](#)

[Documentation Updates | 214](#)

[Migration, Upgrade, and Downgrade Instructions | 215](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 19.3R3 | 167](#)
- [Resolved Issues: 19.3R2 | 186](#)
- [Resolved Issues: 19.3R1 | 195](#)

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: 19.3R3

#### *Application Layer Gateways (ALGs)*

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

#### *Class of Service (CoS)*

- MX Series generated OAM/CFM LTR message sent with a different priority than the incoming OAM/CFM LTM message. [PR1466473](#)
- The following syslog errors are seen: **error cosd[10290]: LIBCOS\_COS\_ATTRIBUTE\_RETRIEVE\_FAILED: FAILED to retrieve cos field (cos\_fc\_defaults\_0\_fc\_no\_loss)**. [PR1470252](#)
- Unexpected traffic loss might be discovered in certain conditions under fusion scenario. [PR1472083](#)
- When **policy-map** configuration is set, MX10008 and MX100016 might generate cosd core file after executing **commit** or **commit check** commands. [PR1475508](#)
- MX Series with line cards using MPC1-Q/MPC2-Q might report memory errors. [PR1500250](#)

#### *EVPN*

- The rpd might crash after changing EVPN related configuration. [PR1467309](#)
- Remote MAC address present in EVPN database might be unreachable. [PR1477140](#)
- Deleting an L2 logical interface generates an error if not deleted first from EVPN. [PR1482774](#)
- The ESI of IRB interfaces does not get updated after autonomous-system number change if the interface is down. [PR1482790](#)

- Dead next hops might flood in a rare scenario after remote PE devices are bounced. [PR1484296](#)
- The ARP entry is gone from kernel after adding and deleting the virtual-gateway-address. [PR1485377](#)
- The rpd core file might be generated when Routing Engine switchover is performed after disabling BGP protocol globally. [PR1490953](#)
- VXLAN bridge domain might lose VTEP logical interface after restarting chassisd. [PR1495098](#)
- The l2ald memory leak might be observed in any EVPN scenario. [PR1498023](#)
- The VXLAN function might be broken because of a timing issue. [PR1502357](#)
- The MAC address of the LT interface might not be installed in the EVPN database. [PR1503657](#)
- A functionality breakage is observed when **proxy-macip-advertisement** statement is configured for EVPN-MPLS. [PR1506343](#)
- Unable to create a new VTEP interface. [PR1520078](#)

### ***Forwarding and Sampling***

- The following syslog errors are seen: **pfed: rtslib: ERROR received async message with no handler: 28.** [PR1458008](#)
- A problem with statistics on some interfaces of a router might be observed after FPC or PIC reboot. [PR1458143](#)
- Type 1 ESI/AD route might not be generated locally on EVPN PE device in the all-active mode. [PR1464778](#)
- An output bandwidth-percent policer with logical-bandwidth-policer applied to an aggregated Ethernet bundle along with an output-traffic-control-profile has incorrect effective policing rate. [PR1466698](#)
- Traffic might not be forwarded into the right queue but the default queue when VPLS traffic has three or more VLAN tags with VLAN priority 5. [PR1473093](#)
- The filter might not be installed if the "policy-map xx" is present under it. [PR1478964](#)
- Use of UTC timestamp in flat-file-accounting files when profile configured. [PR1509467](#)

### ***General Routing***

- UID might not be released properly in some scenarios after service session deactivation. [PR1188434](#)
- IPv6 dynamic subscribers might be unable to login from Junos OS Release 18.2R3 and onwards. [PR1468414](#)
- DHCP relay with forward-only might fail to send OFFER messages when DHCP client is terminated on logical tunnel interface. [PR1471161](#)
- Dynamic-profile for VPLS-PW pseudowire is incorrectly reporting Dynamic Static Subscriber Base Feature license alarm. [PR1473412](#)
- Subnet information might be corrupted if it is passed by a RADIUS server. [PR1474097](#)



- Syslog error messages **PFEIFD: Could not decode media address with length 0** might be generated by Packet Forwarding Engine. [PR1341610](#)
- `show security group-vpn member ipsec security-associations detail | display xml` is not in expected format. [PR1349963](#)
- The following errors are seen **Failed to get xfchip which\_xfchip 7, CMTY\_FM\_HSL2: [cmt\_y\_fm\_hsl2\_adc\_xbar\_set\_cell\_limit\_per\_link] xbar limit set failed pb\_slot[3], pfe[1], sfb\_slot[7], sub\_plane[2], link\_index[83]** during unified ISSU from Junos OS Release 17.3 to 18.2 in MX2010 box. [PR1388076](#)
- The severity of the following error is reduced from fatal to major: **XR2CHIP\_ASIC\_JGCI\_FATAL\_CRC\_ERROR**. [PR1390333](#)
- NAPT66 pool split is not supported with AMS. Therefore, commit fails with V6 pool in AMS. [PR1396634](#)
- The non-existent subscribers might appear in the output of the **show system resource-monitor subscribers-limit chassis extensive** statement. [PR1409767](#)
- Egress monitored traffic is not mirrored to destination for analyzers on MX Series router. [PR1411871](#)
- "FPC x Voltage Tolerance Exceeded" alarm raised and cleared upon bootup of JNP10K-LC2101. [PR1415671](#)
- "System\_id" of an old master Routing Engine is reported by FPCs even after GRES. [PR1417366](#)
- Resetting the Playback Engine logs are seen on MPC5E. [PR1420335](#)
- PF Core Voltage is not set according to the required e-fuse value and remains at the default value (0.9V) on JNP10008-SF and JNP10016-SF. [PR1420864](#)
- FPC crash may be observed after GRES when committing changes in firewall filter with "next term" statements in subscriber scenario. [PR1421541](#)
- PTP might not work on MX104 if phy-timestamping is enabled. [PR1421811](#)
- RPD scheduler slips could be seen upon executing **show route resolution extensive 0.0.0.0/0 | no-more** CLI command if the number of routes in the system is big (several millions). [PR1425515](#)
- Observing 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](#)
- Disable PTP and show warning when hyper mode is configured. [PR1429527](#)
- The following syslog error is seen **dfw\_abstract\_issu\_stats\_counters\_restore:2222 Failed to find Index = 4613734?** during unified ISSU with 19.3I-20190409\_dev\_common.0.2212. [PR1429879](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- ZF interrupts [MAJOR] Out-of-range Dest PFE INTR for Gnt seen during MPC6 or MPC9 line card bringup. [PR1436148](#)
- Control logical interface is not created by default for LLDP. [PR1436327](#)

- IRB over VTEP unicast traffic might get dropped on EX9200 or MX Series platforms. [PR1436924](#)
- (SEEN ONLY ON LEGACY IMAGE) ISSU is failing from Junos OS Release 19.1R1 legacy Junos OS release images. [PR1438144](#)
- FPC might reboot if jlock hog occurs on all Junos VM-based platforms. [PR1439906](#)
- The EX Series device ports might stay in up state even if the EX4600/QFX5100 device is rebooted. [PR1441035](#)
- The interface might go into admin down state after FPC restart with PTP configuration enabled. [PR1442665](#)
- MPC9E is not going to offline due to unreachable destinations in phase 3 stage. [PR1443803](#)
- Irregular traffic drop might be seen incoming traffic from MPC3E and exiting to MPC10E. [PR1445649](#)
- bbe-smgd core file is generated on the backup Routing Engine in bbe\_ifd\_add\_vlan (ifd=0x8c3e835, ifl=0xcaf59f18) at ../../../../src/junos/usr.sbin/bbe-svcs/smd/infra/bbe\_ifd.c:6374. [PR1447493](#)
- CLI prompt configured as part of [edit system login class] and [edit system login user] is not taking into effect with subsequent login post-commit. [PR1447497](#)
- FEC statistics are not getting reset after changing FEC mode. [PR1449088](#)
- IPv6 throughput numbers for NAT with HTTP traffic is not at par with IPv4. [PR1449435](#)
- Changing the hostname triggers an on-change notification, not an adjacency on-change notification. Also, currently IS-IS is sending the hostname instead of the system ID in OC paths. [PR1449837](#)
- RE-MX2008-X8-128G secure BIOS version mismatch alarms. [PR1450424](#)
- JNP10K-LC2101 FPC generates the "Voltage Tolerance Exceeded" major alarm for EACHIP 2V5 sensors. [PR1451011](#)
- Main chassisd thread at a JNS GNF could experience stalls upon GNF SNMP polling for hardware-related OIDs. [PR1451215](#)
- Support for drop flows when packet is dropped. [PR1451921](#)
- Interfaces shut down by 'disable-pfe' action might not be up using MIC offline/online command. [PR1453433](#)
- FPC/PFE crash may happen with ATM MIC installed in the FPC. [PR1453893](#)
- The 100G interfaces might not come up again after going down on MPC3E-NG. [PR1454595](#)
- Multiple daemons may crash on committing configuration changes related to groups. [PR1455960](#)
- 4x1GE using QSFP28 optics - Continuous logging in chassisd when 'speed 1g' is configured - pic\_get\_nports\_inst and ch\_fru\_db\_key. [PR1456253](#)
- Queue data might be missing from path /interfaces/interface/state. [PR1456275](#)
- All the IPsec tunnels might be cleared when the clear command is executed for only one IPsec tunnel with specified service-set name. [PR1456749](#)

- The LCMD will not restart automatically after JSU installation. [PR1457304](#)
- Traffic silently drops or get discarded or MPC crash might be seen on MPC10E during firewall filter terms changing on it. [PR1458499](#)
- The rpd crash might be seen if a BGP route is resolved over the same prefix protocol next hop in an inet.3 table that has both RSVP and LDP routes. [PR1458595](#)
- The traffic might be stuck on MS-MPC/MS-MIC with sessions receiving a huge number of affinity packets. [PR1459306](#)
- Inline S-BFD packets are dropped on MPC6E MIC1/PIC1 ports 0-11. [PR1459529](#)
- Subscriber statistics can be broken after ISSU. [PR1459961](#)
- Multiple leaf devices and prefixes missing when LLDP neighbor added after the streaming started at global level [PR1460347](#)
- In an MVPN instance, traffic drops on multicast receivers within the range of 0.1 to 0.9 percent. [PR1460471](#)
- Support of del\_path for the LLDP neighbor change at various levels. [PR1460621](#)
- IPv6 prefix might be hidden when received over IPv4 BGP session. [PR1460786](#)
- Explicit Deletion Notification (del\_path) not received when LLDP neighbor is lost as a result of disabling local interface on the DuT via CLI (gNMI). [PR1461236](#)
- Traffic might be impacted because the fabric hardening is stuck. [PR1461356](#)
- On MPC10E, more number of output packets are seen than expected when ping function is performed. [PR1461593](#)
- The rpd might crash after committing the dynamic-tunnel-anchor-pfe configuration. [PR1461980](#)
- The rpd process might crash if **show v4ov6-tunnels information anti-spoof-ip ip** is executed. [PR1462047](#)
- Traffic drop might be seen in a scaled scenario with VRRP sessions configured on aggregated Ethernet interfaces. [PR1462310](#)
- The statistics of the subscribers might be incorrect when accounting interval statistics are not populated on MX204. [PR1462325](#)
- MAC address learning might not work as expected in VPLS scenario on MX Series platforms with MPC10 and MPC11. [PR1463078](#)
- The Routing Engine switchover might not be triggered when the master CB clock failure. [PR1463169](#)
- MVPN traffic might be dropped after performing switchover. [PR1463302](#)
- The subscribers might not pass traffic after doing some changes to the dynamic-profiles filter. [PR1463420](#)
- The MPC2E-NG/MPC3E-NG card with specific MIC might crash after a high rate of interface flaps. [PR1463859](#)

- The mspmand might crash when stateful firewall and RPC ALG used on MX Series platforms with MS-MIC/MS-MPC. [PR1464020](#)
- The IPoE subscriber route installation may fail. [PR1464344](#)
- The bbe-smgd core file is generated (0x000000000088488c in bbe\_autoconf\_delete\_vlan\_session\_only (session\_id=918) at `../../../../src/junos/usr.sbin/bbe-svcs/smd/plugins/autoconf/bbe_autoconf_plugin.c:3115`). [PR1464371](#)
- The CPU utilization on mgd daemon might be stuck at 100% after the NETCONF session is interrupted by flapping interface. [PR1464439](#)
- The MS-MIC might not work when it is used on specific MPC. [PR1464477](#)
- The LLDP packets might get discarded on all platforms running Junos OS. [PR1464553](#)
- The jdhcpd might consume a high CPU and no further subscribers can be brought up if there are more than 4000 dhcp-relay clients in the MAC-MOVE scenario. [PR1465277](#)
- The physical interface of aggregated Ethernet might take time to come up after disable/enable. [PR1465302](#)
- Junos OS: MX Series: PFE crash on MPC7, MPC8, and MPC9 upon receipt of small fragments requiring reassembly (CVE-2020-1649). [PR1465490](#)
- The bbe-smgd process generates core files on backup Routing Engine. [PR1466118](#)
- ICMP error messages are still unreceived after enabling the **enable-asymmetric-traffic-processing** statement. [PR1466135](#)
- The PPPoE subscribers get stuck because the PPPoE inline keepalives won't work properly. [PR1467125](#)
- A few of the DHCP INFORM packets specific to a particular VLAN might be taking the wrong resolve queue. [PR1467182](#)
- MPC11E: The DOM MIB alarm for the channelized 10G interface is not showing any alarm for LF/RF. [PR1467446](#)
- L2 wholesale not forwarding all client requests with stacked VLAN. [PR1467468](#)
- Hotswap between MPC11E and MPC9 MPC8, and MPC6 is not supported. [PR1467725](#)
- The rpd process might stop after several changes to the flow-spec routes. [PR1467838](#)
- Crypto code can cause high CPU utilization. [PR1467874](#)
- Daemons might not be started if commit is executed after commit check. [PR1468119](#)
- Benign logs might appear in Junos OS Release 19.3R2 when switching between configurations using load-override with GRES and commit-synchronize. [PR1468234](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)

- MX10003/MX204: Temperature sensor name for PEMs under show chassis environment pem is incorrect. [PR1468642](#)
- The rpd crash might be seen if BGP sharding is enabled. [PR1468676](#)
- The tcp-log connections fail to reconnect and are stuck in "Reconnect-In-Progress" state. [PR1469575](#)
- nsd core file is generated after committing the configuration successfully, if destination NAT rule matches '0.0.0.0/X' destination address. [PR1469613](#)
- Memory leak on l2cpd process might lead to l2cpd crash. [PR1469635](#)
- A hierarchical scheduler should not be configured on a "ps" interface. [PR1470049](#)
- On MPC11E, some of the 10-Gbps interface states may not be cleaned up correctly when performing GRES with invalid profile configuration. [PR1470153](#)
- Traffic might be not policed properly on MPC10E/MPC11E interfaces. [PR1470629](#)
- SNMP interface-mib stops working for PPPoE clients. [PR1470664](#)
- Multiple FreeBSD vulnerabilities are fixed in Junos OS. (CVE-2018-6916, CVE-2018-6918) [PR1470693](#)
- Unable to setup 26M sessions (NAPT44) at 900Kpps/s. [PR1470833](#)
- Syslog message **fpcX user.notice logrotate: ALERT exited abnormally with [1]** appears at 04:02:01. [PR1471006](#)
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card may cause major alarms on other FPCs in the system. [PR1471372](#)
- The clksyncd crash might be seen when PTP over AE is configured on MX104 platform. [PR1471466](#)
- MPC11E: MTU errors count captured in the show pfe statistics traffic do not match exactly with the actual count of frames dropped. Some incorrect cumulative value might be shown. [PR1471554](#)
- ARP suppression (default enabled) in EVPN not working on MX10008/MX10016 line cards. [PR1471679](#)
- The pccd core and PCEP session flaps might be seen in PCE-initiated or PCE-delegated LSP scenario. [PR1472051](#)
- Junos VMhost upgrades might fail on MX240/MX480/MX960 platforms with NG-RE installed and USF mode enabled. [PR1472287](#)
- The chassis alarm might appear on BSYS: **RE0 to one or many FPCs is via em1: Backup RE**. [PR1472313](#)
- Service accounting statistics don't get updated after making change to firewall filters. [PR1472334](#)
- The kernel may crash and vmcore may be observed after configuration change is committed. [PR1472519](#)
- Performing back-to-back rpd restarts might cause rpd to crash. [PR1472643](#)
- Active errors counts are not increasing for I2C in SYNCE cards. [PR1472660](#)
- Manually configured ERO on NS controller lost when PCEP session bounced. [PR1472825](#)

- SDB goes down very frequently if the "reauthenticate lease-renewal" configuration is enabled for DHCP. [PR1473063](#)
- Enhanced-mode ISSU may fail for MPC7E/MPC8E/MPC9E as the card might get wrongly marked as offline due to "Bad voltage". [PR1473722](#)
- Junos OS: MX Series: Services card might restart when DNS filtering is enabled (CVE-2020-1645). [PR1474056](#)
- Ingress multicast replication does not work with GRES configuration. [PR1474094](#)
- MX150 core files are not seen under **show system core-dumps**. [PR1474118](#)
- Junos OS: MX Series: PFE crash on MPC7/8/9 upon receipt of large packets requiring fragmentation (CVE-2020-1655). [PR1474154](#)
- An MPC11 crash might occur on MX2K platform using multi-dimensional advanced scale configuration having Inline Keep Alive Sessions [PR1474160](#)
- QSA adapter Lane 0 port might also be brought down when disabling one of the other lanes. [PR1474231](#)
- A newly added LAG member interface might forward traffic even though its micro-BFD session is down. [PR1474300](#)
- **show services sessions** and **show services sessions extensive output** not displaying member interface of the AMS where the session got landed. It is displaying only AMS interface name. [PR1474313](#)
- clksyncd core files after GRES. [PR1474987](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- The RADIUS accounting updates of service session have bogus statistics data. [PR1475729](#)
- When xSTP protocols are enabled on interface all, it might run on 'vlan-tagging'/'flexible-vlan-tagging' L3 interfaces and lead to blocking of SXE interface. [PR1475854](#)
- Traffic drop might be observed while performing ISSU on MX2020/MX2010/MX960 platforms. [PR1476505](#)
- The bbe-mibd might crash on the MX Series platform in subscriber environment. [PR1476596](#)
- The host-generated packets may get dropped at the other end. [PR1476764](#)
- Traffic loss may be observed to the LNS subscribers in case the "routing-service" configuration is enabled under the dynamic-profile. [PR1476786](#)
- In a NAT-T scenario, the IKE version 2 IPsec tunnel might flap if the tunnel initiator is not behind NAT. [PR1477483](#)
- Observed messages "spmb0 cmt\_y\_sfb\_temp\_check: sfb[0] is powered OFF" and "spmb0 cmt\_y\_sfb\_voltage\_check\_one: sfb[0] is powered OFF" are flooding even though SFBs are online in MX2010. [PR1477924](#)

- Error log messages: **chassisd[7836]: %DAEMON-3-CHASSISD\_IOCTL\_FAILURE: acb\_get\_fpga\_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device)** after every commit. [PR1477941](#)
- Packet Forwarding Engine might be disabled due to major error on MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9. [PR1478028](#)
- "show evpn statistics instance" command gets stuck on multihomed scenario. [PR1478157](#)
- The ukern-platformd process might crash on MX2000 platforms with MPC11. [PR1478243](#)
- FPC memory leak might happen after executing "show pfe route". [PR1478279](#)
- Traffic drop might be seen if interface-specific filter is configured. [PR1478358](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- The protocol MTU may not be changed on It- interface from the default value. [PR1478822](#)
- TCP-log sessions might be in Established state but no logs get sent out to the syslog server. [PR1478972](#)
- PPPoE subscribers might not all log in after FPC restart. [PR1479099](#)
- The SCBE3 fabric plane gets into check state in MX-VC. [PR1479363](#)
- After kmd restart, IPsec SA comes up but traffic fails for some time in certain scenarios. [PR1480692](#)
- 100-Gbps interface may randomly fail to come up after maintenance operations. [PR1481054](#)
- The rpd may crash when executing "show route protocol l2-learned-host-routing" or "show route protocol rift" CLI command on a router. [PR1481953](#)
- Logging in some PPPoE subscribers through aggregated Ethernet interface might cause the device to reboot. [PR1482431](#)
- Traffic impact might be seen when policy-multipath is configured without LDP on SPRING-TE scenario. [PR1483585](#)
- The downstream IPv4 packet greater than BR MTU are getting dropped in MAP-E. [PR1483984](#)
- Packet loss might be observed after device is rebooted or l2ald is restarted in EVPN-MPLS scenario. [PR1484468](#)
- ARP entry might not be created in the EVPN-MPLS environment. [PR1484721](#)
- The logical tunnel interface might not work on MPC10 line card. [PR1484751](#)
- Interface input error counters are not increasing on MX150. [PR1485706](#)
- krt-nexthop-ack-timeout may not automatically be picked up on rpd start or restart. [PR1485800](#)
- Kernel core might be seen if deleting an ifstate. [PR1486161](#)
- Command completion help text for LLDP-MED co-ordinate configuration statement contains spelling errors. [PR1486327](#)
- The aftd might crash. [PR1487416](#)

- Incorrect 'frame length' of 132 bytes might be captured in packet header. [PR1487876](#)
- Add support for PSM firmware upgrade on MX2000. [PR1488575](#)
- With MX-SPC3 service card, NAT might not be processed on order as setup. [PR1489581](#)
- With 4 Member AMS used in the service-set, commit check should fail when /30 subnet address is used as NAT pool IP. [PR1489885](#)
- Prolonged flow-control might happen on MX Series platforms with MS-MPC/MS-MIC. [PR1489942](#)
- Error syslog message "Failed to connect to the agentx master agent (/var/agentx/master): Unknown host (/var/agentx/master) (No such file or directory)" is continuously being generated with DNS-sinkhole. [PR1490487](#)
- When NAT/SFW rule configured with application-set with multiple applications having different TCP inactivity-timeout, sessions are not getting TCP inactivity-timeout according to the configured application order. [PR1491036](#)
- ISSU to Junos OS Release 19.4R1 or later is not supported on NG-MPC cards. [PR1491337](#)
- Multiple deactivating/activating security traceoptions along with clear single NAPT44 session might result in the flowd crash. [PR1491540](#)
- MS-MIC is down after loading some releases in MX-VC scenario. [PR1491628](#)
- FPCs might stay down or restart when swapping MPC7/8/9 with MPC10/11 or vice versa in the same slot. [PR1491968](#)
- User-configured MTU might be ignored after the ISSU upgrade using "request vmhost software in-service-upgrade". [PR1491970](#)
- MX10008 SNMP table entPhysicalTable does not match the PICs displayed **show chassis hardware**. [PR1492996](#)
- MPC10 and MPC11 line card might crash if the interface is configured with firewall filter referencing shared-bandwidth policer. [PR1493084](#)
- One port on MIC-3D-2XGE-XFP flapping might cause the other port on the same MIC to flap. [PR1493699](#)
- Frequency traceable flag reflects the incorrect status and get toggled unexpectedly. [PR1493743](#)
- In node slicing setup after GRES has been done, radius interim updates may not carry actual statistics. [PR1494637](#)
- The B4 might not able to establish the softwire with AFTR [PR1496211](#)
- **PFEIFD: Could not decode media address with length 0** error messages generated by Packet Forwarding Engine when subscribers come up over a pseudowire interface. [PR1496265](#)
- Outbound SSH connection flap or memory leak issue might be observed while pushing configuration to ephemeral DB with high rate. [PR1497575](#), [PR1508324](#)
- Port numbers logged in ALG syslog are incorrect. [PR1497713](#)



- Subscribers may be disconnected after one of the AE participating FPCs comes online in a Junos node slicing scenario. [PR1498024](#)
- SNMP polling does not show correct PSM jnxOperatingState when one of the PSM inputs failed. [PR1498538](#)
- The rpd might crash when multiple VRFs with 'IFLs link-protection' are deleted at a single time. [PR1498992](#)
- MX2020/MX2010 continuously logging "pem\_tiny\_power\_remaining:" in chassisd log. [PR1501108](#)
- The chassisd process may get stuck. [PR1502118](#)
- The packets from non existing source on GRE/UDP designated tunnel might be accepted on MX Series platforms. [PR1503421](#)
- "show bridge statistics" will not display the statistics information for pseudowire subscriber interfaces. [PR1504409](#)
- Fan speed might toggle between full and normal on MX960 with enhanced FRU. [PR1504867](#)
- The l2cpd crash might be seen if adding/deleting ERP configuration and then restarting l2cpd. [PR1505710](#)
- GnmiJuniperTelemetryHeader incompatibility introduced in Junos OS Release 19.3. [PR1507999](#)
- The host-generated packets might get dropped if the **force-control-packets-on-transit-path** statement is configured. [PR1509790](#)
- The disabled QSFP transceiver might fail to be turned on. [PR1510994](#)
- Static subscribers are logged out after creating a unit under demux0 interface. [PR1511745](#)
- Memory leak on l2ald might be seen when adding/deleting the routing-instances/bridge-domains configuration. [PR1512802](#)
- The wavelength configured using CLI might not be set on SFP+-10G-T-DWDM-ZR optics when the optics is used on MPC7E line card. [PR1513321](#)
- Modifying the segment list of the SR LSP might not work. [PR1513583](#)
- Subscribers might not be able to bind again after performing back-to-back GRES followed by FPC restart. [PR1514154](#)
- MPC7E with QSFP installed may get rebooted when 'show mtip-chmac <1|2> registers' vty command is executed. [PR1517202](#)
- PADI packets are dropped when encapsulation VPLS is set along with dynamic profiles accepted protocols configured as pppoe. [PR1523902](#)
- pkid process might crash at bn\_i2c (pval=0x1d, cont=0x0, putype=0xffffcce8, it=0xc8c848b8 < BIGNUM\_it>) at ../../../../src/crypto/openssl/crypto/asn1/x\_bignum.c:127. [PR1471878](#)
- The kmd might crash in specific simultaneous rekey scenario on MX Series routers. [PR1474797](#)
- Error logs **ac200\_dcfp2\_pm\_get\_info: ifd 0xb3d090a8, info 0x196cc518, ac200\_dcfp2 0x0** being generated periodically. [PR1487070](#)

- Previous configuration might still take effect after "rollback rescue" is performed. [PR1489575](#)
- The heap memory leak might be seen on MPC10/MPC11 line cards. [PR1499631](#)
- Slow FPC Heap memory leak observed after extensive subscriber login or logout. [PR1508291](#)
- Used-service-unit of the CCR-U has output-bytes counter zero. [PR1516728](#)
- Adding and removing an aggregated Ethernet member link might cause PPPoE subscribers session and traffic to get drop. [PR1525585](#)

### ***High Availability (HA) and Resiliency***

- The unified ISSU might fail on MX204 or MX10003 Virtual Chassis or MX Series with MPC7E, MPC8E, or MPC9E on Virtual Chassis with an error message. [PR1480561](#)

### ***Infrastructure***

- An alarm might be seen if the PEM's serial number starts with "1F1". [PR1398128](#)
- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. [PR1462986](#)
- The scheduled tasks might not be executed if "cron" daemon goes down without restarting automatically. [PR1463802](#)
- SNMP poling might return unexpectedly high value of ifHCOutOctets counter for physical interface when any jnxDom OID is processed at the same time. [PR1508442](#)

### ***Interfaces and Chassis***

- Restarting chassisd with GRES disabled might cause FPC to restart and some demux interfaces deletion. [PR1337069](#)
- The syslog `scchassisd[]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x` is observed after MX Series Virtual Chassis local or global switchover. [PR1428254](#)
- Decoupling of L2 logical interface from bridge EVPN configuration. [PR1438172](#)
- When the logical interface is associated to a routing instance inside an LR is removed from the routing instance, the logical interface is not added to the default routing instance. [PR1444131](#)
- Interface descriptions might be missing under logical systems CLI. [PR1449673](#)
- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. [PR1459201](#)
- The EOAM CFM primary-vid functionality does not work if the **enhanced-cfm-mode** is enabled. [PR1465608](#)
- The IPv6 local statistics are counted against the IPv6 transit traffic statistics as well. [PR1467236](#)
- `vrpv3mibs` are not working on QFX Series platform to poll VRRPv6 related objects. [PR1467649](#)

- **MIC Error code: 0x1b0002** alarm might not be cleared for MIC on MPC5E when the voltage has returned to normal. [PR1467712](#)
- When ESI is configured on a physical interface, and if a logical interface is disabled under the physical interface, a traffic drop is seen. [PR1467855](#)
- On performing the **commit** command, the dcd process might hang up. [PR1470622](#)
- The **sonet-options** configuration stanza is disabled for xe interface working in wan-phy mode. [PR1472439](#)
- Commit error is not thrown when a member link is added to multiple aggregation group with different interface-specific options. [PR1475634](#)
- The interface on MIC3-100G-DWDM might be going down after performing an interface flap. [PR1475777](#)
- MC-AE interface displays unknown status if the subinterface is added as part of the VLAN on the peer MC-AE node. [PR1479012](#)
- The **vrrp-inherit-from** change operation leads to a packet loss when traffic is forwarded to the VIP gateway. [PR1489425](#)
- Unexpected dual VRRP backup state might happen after performing two subsequent Routing Engine switchovers with **track priority-hold-time** configured. [PR1506747](#)

#### **J-Web**

- Security vulnerability in J-Web and web-based (HTTP/HTTPS) services. [PR1499280](#)

#### **Juniper Extension Toolkit (JET)**

- JET certificate request utility is broken in JET sb. [PR1473892](#)

#### **Layer 2 Ethernet Services**

- No alarm is generated when FPC is connected to master Routing Engine through backup RE/CB. [PR1461387](#)
- Member links state might be asynchronized on a connection between PE and CE devices in EVPN active/active scenario. [PR1463791](#)
- The unified ISSU might fail during subscriber in-flight login. [PR1465964](#)
- Telemetry data for relay/bindings/binding-state-v4relay-binding and relay/bindings/binding-state-v4relay-bound is not correct. [PR1475248](#)
- For MX204 platform, the Vendor-ID is set as MX10001 in factory-default configuration and DHCP client messages. [PR1488771](#)
- Issues with DHCPv6 relay processing confirm and reply packets. [PR1496220](#)
- The MC-LAG might become down after disabling and then enabling the force-up. [PR1500758](#)
- The aggregated Ethernet interface sometimes might not come up after the switch is rebooted. [PR1505523](#)
- DHCPv6 lease query is not as expected while verifying the DHCPv6 server statistics. [PR1506418](#)

- The **show dhcp relay statistics** command displays **DHCPLEASEUNASSIGNED**, instead of **DHCPLEASEUNASSIGNED** which has spelling error. [PR1512239](#)
- The **show dhcpv6 relay statistics** command should display **DHCPV6\_LEASEQUERY\_REPLY** instead of **DHCPV6\_LEASEQUERY\_REPL** for messages sent. [PR1512246](#)
- The memory leak in **jdhcpd** might be seen if **access-profile** is configured at the **[dhcp-relay]** or **[dhcp-local-server]** hierarchy level. [PR1525052](#)

## MPLS

- The FPC might be stuck in Ready state after applying a configuration change that will remove RSVP and trigger FPC restart. [PR1359087](#)
- Traffic loss might happen if P2MP with NSR is enabled. [PR1434522](#)
- Traffic might silently get dropped or discarded after the LSP protection link when third-party transit router goes down. [PR1439251](#)
- The traffic might silently get dropped or discarded after the LACP timeout. [PR1452866](#)
- P2MP LSP might flap after VT interface in MVPN routing instance is reconfigured. [PR1454987](#)
- Static MPLS LSP Label might not get installed in mpls.0 after link flap. [PR1457432](#)
- The rpd crash might be observed with traceoption enabled in MPLS. [PR1457681](#)
- RSVP interface bandwidth calculation might be incorrect when RSVP subscription percentage is configured under RSVP interface. [PR1458527](#)
- The device might use the local-computed path for the PCE-controlled LSPs after link/node failure. [PR1465902](#)
- The rpd might crash in PCEP for the RSVP-TE scenario. [PR1467278](#)
- Fast reroute detour next hop down event might cause primary LSP down in particular scenario. [PR1469567](#)
- The rpd process might crash during shutdown. [PR1471191](#)
- The rpd crash might be seen after doing some commit operations, which could affect RSVP ingress routes. [PR1471281](#)
- Errors continuously flood in backup Routing Engine **JTASK\_IO\_CONNECT\_FAILED: RPDtm./var/run/rpdtmd\_control: Connecting to 128.0,255.255,255.255,0.0.0.0,0.0.0.0, failed: No such file or directory**. [PR1473846](#)
- LDP/BFD sessions are not coming up in scaled setup. [PR1474204](#)
- RSVP LSPs might not come up in scaled network with a high number of LSPs if NSR is used on transit router. [PR1476773](#)
- Kernel crash and device restart might happen. [PR1478806](#)

- The rpd process loads CPU 100 percent and the rpd core files are seen on the backup Routing Engine. [PR1479249](#)
- High CPU utilization for rpd process might be seen if RSVP is implemented. [PR1490163](#)
- The rpd might crash when flapping BGP with FEC 129 VPWS enabled. [PR1490952](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP maximum segment size (MSS) in use. [PR1493431](#)
- The rpd might crash in a rare condition under SR-TE scenario. [PR1493721](#)
- The rpd process might generate core files while performing unified ISSU from Junos OS Release 19.3R2 to Junos OS Release 20.1R1 or later. [PR1493969](#)
- The rpd might crash when SNMP polling is done using OID `jnxMplsTeP2mpTunnelDestTable`. [PR1497641](#)
- The traffic loss might happen if ISSU is performed when P2MP is configured for an LSP. [PR1500615](#)
- CSPF job might get stalled for new or existing LSP in high-scale LSP setup. [PR1502993](#)
- The rpd process might crash with RSVP configured in a rare timing case. [PR1505834](#)
- Activating or deactivating the LDP-sync under OSPF might cause LDP neighborship to go down and stay down. [PR1509578](#)
- The rpd might crash after upgrading to Junos OS Release 18.1 and later. [PR1517018](#)
- SNMP trap is sent with incorrect OID `jnxSpSvcSetZoneEntered`. [PR1517667](#)

### ***Network Management and Monitoring***

- The SNMP coldStart trap might be seen after Routing Engine switchover. [PR1461839](#)
- SNMPv3 informs might not work properly after rebooting. [PR1497841](#)

### ***Platform and Infrastructure***

- jcrypto syslog help package and events are not packaged even when error message is compiled. [PR1290089](#)
- CI-PR:'core.vmx.mpc0' seen at 5 0x096327d5 in l2alm\_sync\_entry\_in\_pfes (context=0xd92e7b28, sync\_info=0xd92e7a78) at ../../../../src/pfe/common/applications/l2alm/l2alm\_common\_hw\_api.c:1727 [PR1430440](#)
- With CNH enabled, the MPLS CoS rewrite does not work for 6PE traffic. [PR1436872](#)
- On certain MPC line cards, cm errors need to be reclassified. [PR1449427](#)
- Traffic loss might be seen in case of Ethernet frame padding with VLAN. [PR1452261](#)
- The MPC might drop packets after enabling firewall fast lookup filter. [PR1454257](#)
- DDoS-protection does not stop logging when remote tracing is enabled. [PR1459605](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)

- On MX204 platform, Packet Forwarding Engine errors might occur when incoming GRE tunnel fragments 1) get sampled and 2) undergo inline reassembly. [PR1463718](#)
- EVPN-VXLAN Type-5 tunnel might not work properly on MX Series platforms. [PR1466602](#)
- All the subscriber services might be unavailable on vBNG running on MX150 or vMX running in payg mode. [PR1467368](#)
- The Layer-2 traffic over aggregated Ethernet interfaces sent from one member to another is corrupted on MX Series Virtual Chassis setup. [PR1467764](#)
- Kernel crash (vmcore) or FPC crash due to mbuf leak (CVE-2020-1653). [PR1468183](#)
- JNH memory leaks might be seen after CFM session flap for LSI/VT interfaces. [PR1468663](#)
- The switch might not be able to learn MAC address with dot1x and **interface-mac-limit** configured. [PR1470424](#)
- SSH login might hang and the TACAAS + server closes the connection without sending any authentication failure response. [PR1478959](#)
- The convergence time for MVPN fast upstream failover might be more than 50 ms. [PR1478981](#)
- **show system buffer** command displays all zero in the MX104 chassis and it looks like cosmetic issue as there is no service impact reported. [PR1484689](#)
- MAC learning under bridge-domain stops after MC-LAG interface flap. [PR1488251](#)
- MAC malformation might occur in a rare scenario under MX-VC setup. [PR1491091](#)
- In node slicing setup, MPLS TTL might be set to zero when the packet goes through AF interface configured with CCC family. [PR1492639](#)
- Packets get dropped when next hop is IRB over It- interface. [PR1494594](#)
- Routing Engine crash might be seen when a large number of next hops are quickly deleted and re-added in large ARP/ND scale scenario. [PR1496429](#)
- Traffic to VRRP virtual IP/MAC might be dropped when ingress queuing is enabled. [PR1501014](#)
- Python or SLAX script might not be executed. [PR1501746](#)
- Traffic originated from another subnet is sent out with 0x8100 instead of 0x88a8. [PR1502867](#)
- MPCs might crash when there is a change on routes learned on IRB interface configured in VPLS/EVPN instances. [PR1503947](#)
- Traffic loss is seen in certain conditions under MC-LAG setup. [PR1505465](#)
- During route table object fetch failure, FPC might crash. [PR1513509](#)
- Traffic from IRB interface towards LSI interface gets dropped with adaptive or per-packet load balancing. [PR1458825](#)
- Traceroute monitor with MTR version v.69 shows a false 10 percent loss. [PR1493824](#)

### ***Routing Policy and Firewall Filters***

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

### ***Routing Protocols***

- The BGP sessions might flap with statement **keep none** configured when a VRF is deleted. [PR1439560](#)
- The CPU utilization on rpd spins at 100 percent after the same external BGP route is learned in different VRF tables. [PR1442902](#)
- PIM RPF selection for the specific multicast group might get incorrectly applied to other multicast groups. [PR1443056](#)
- The rpd crash might be seen after configuring OSPF NSSA area-range and summaries. [PR1444728](#)
- The TI-LFA backup path for adj-sid is broken in OSPF. [PR1452118](#)
- SSH login might fail if a user account exists in both the local database and RADIUS/TACACS+. [PR1454177](#)
- MoFRR with MLDP inband signaling is not working. [PR1454199](#)
- The "other querier present interval" timer cannot be changed in an IGMP/MLD snooping scenario. [PR1461590](#)
- The IS-IS IPv6 routes might flap when there is an unrelated commit under protocol stanza. [PR1463650](#)
- The rpd might crash if IPv4 routes are programmed with IPv6 next hop through JET APIs. [PR1465190](#)
- BGP peers might flap if a small hold-time parameter value is set. [PR1466709](#)
- The configured BGP damping policy might not take effect after BGP is disabled and then enabled followed by **commit**. [PR1466734](#)
- The BGP multipath does not work if multitopology routing is used. [PR1467091](#)
- The rpd might crash after configuring **independent-domain** under the master routing-instance. [PR1469317](#)
- The mcsnoopd might crash when the STP moves the mrouter port to the blocked state. [PR1470183](#)
- The BFD client session might flap when removing BFD configuration from the peer end (from other vendor) of the BFD session. [PR1470603](#)
- The rpd might stop when both instance-import and instance-export policies contain the as-path-prepend action. [PR1471968](#)
- The rpd process might crash with BGP multipath and damping configured. [PR1472671](#)
- Removing cluster from BGP group might cause prolonged convergence time. [PR1473351](#)
- Adjacency SID might be missed and not be advertised to peer/controller/BMP monitor in BGP-LS NLRI. [PR1473362](#)
- SFTP does not connect properly. The following error is displayed: **Received message too long**. [PR1475255](#)
- The rpd process might crash with BGP multipath and route is withdrawn occasionally. [PR1481589](#)

- The rpd process crashes due to specific BGP UPDATE packets (CVE-2020-1644). [PR1481641](#)
- Multicast traffic loss could be seen in certain conditions while enabling **IGMP snooping** under EVPN-VXLAN ERB scenario. [PR1481987](#)
- BGP multipath traffic might not fully load-balance for a while after adding a new path for load sharing. [PR1482209](#)
- The rpd might crash after BGP peer flapping. [PR1482551](#)
- RIPv2 could malfunction when changing interface type from P2MP to broadcast. [PR1483181](#)
- The rpd crashes if the same neighbor is set in different RIP groups. [PR1485009](#)
- There might be rpd memory leak in a certain looped MSDP scenario. [PR1485206](#)
- The BGP-LU routes do not have the label when BGP sharding is used. [PR1485422](#)
- The rpd crashes if BGP LLGR with RIB sharding and traceoptions for **graceful-restart** configured. [PR1486703](#)
- The rpd might crash when performing GRES with MSDP configured. [PR1487636](#)
- High CPU utilization might be observed when the outgoing BGP updates are send slowly. [PR1487691](#)
- The rpd process might generate core files after **always-compare-med** is configured for BGP path-selection. [PR1487893](#)
- BGP RIB sharding feature cannot be run on a system with a single CPU. [PR1488357](#)
- The rpd crashes when OSPF neighbors are reset. [PR1489637](#)
- The BGP route-target family might prevent route reflector from reflecting L2VPN and L3VPN routes. [PR1492743](#)
- The rpd might crash on PTX Series and QFX10000 due to rpd resolver problem of indirect next hop. [PR1494005](#)
- The static route in inet6.0 or inet6.3 RIB might fail to be deleted. [PR1495477](#)
- Receipt of certain genuine BGP packets from any BGP Speaker causes rpd to crash. (CVE-2020-1640). [PR1497721](#)
- The rpd might crash if the import policy is changed to accept more routes that exceed the teardown function threshold. [PR1499977](#)
- The rpd process crash might occur in a multicast scenario with BGP configured. [PR1501722](#)
- Change in x-path output for value "input-updates" in **show bgp neighbors**. [PR1504399](#)
- On all Junos OS dual Routing Engine GRES/NSR enabled routers, the rpd process might crash on a new master Routing Engine if Routing Engine switchover occurs right after massive routing-instances deletion. [PR1507638](#)
- The rpd crash might occur due to RIP updates being sent on an interface in down state. [PR1508814](#)



- The rpd might crash on backup Routing Engine if BGP (standby) received a route from a peer which is rejected due to invalid target community. [PR1508888](#)
- The rpd might crash after deleting and re-adding a BGP neighbor. [PR1517498](#)
- The rpd process might crash if there is a huge number of SA messages in MSDP scenario. [PR1517910](#)

### ***Services Applications***

- The flow-tap add function may not work as expected on MX Series platforms. [PR1472109](#)
- MX L2TP LTS fails to forward the agentCircuitId and agentRemotId AVP toward the LNS. [PR1472775](#)
- The kmd might crash due to the incorrect IKE SA establishment after the remote peer's NAT mapping address has been changed. [PR1477181](#)
- The FPC might crash with npc core file if the service interface is configured under service-set in USF mode. [PR1502527](#)
- Output of "show services l2tp tunnel extensive" doesn't show configured session limit. [PR1503436](#)

### ***Subscriber Access Management***

- No volume attribute is in the Accounting Stop for Service session when the Activated Services session is configured. [PR1470434](#)
- The subinterfaces might be missing in the NAS Port ID. [PR1472045](#)
- The authd might crash after ISSU from 18.3- to 18.4+. [PR1473159](#)
- Some address relevant fields are missing when executing "test aaa ppp" command. [PR1474180](#)
- Syslog messages: pfe\_tcp\_listener\_open\_timeout: Peer info msg not received from addr: 0x60000080. Socket 0xfffff804ad23c2e0 closed. [PR1474687](#)
- The delete request of a specified service session via CoA could fail. [PR1479486](#)
- The CoA request may not be processed if it includes "proxy-state" attribute. [PR1479697](#)
- "mac-address" CLI option is hidden under "access profile <profile-name> radius options calling-station-id-format" statement. [PR1480119](#)
- The authd logs events might not be sent to syslog host when destination-override is used. [PR1489339](#)
- The LTS is incorrectly sending the access-request with the Tunnel-Assignment-ID, which is not compliant with RFC 2868. [PR1502274](#)
- On subscriber termination, CCR-T is sent to the PCRF server reporting the same. But this does not contain the usage monitoring information. [PR1517507](#)

### ***User Interface and Configuration***

- The version information under configuration is changed from Junos OS Release 19.1. [PR1457602](#)
- In MX J-Web page might not get redirected to login once session is expired with idle timeout. [PR1459888](#)

## VPNs

- The core file is seen due to a double free of a label. The issue happens in a BGP-based VPLS setup where BGP has RR configuration. Because of the route reflector configuration, the BGP-VPLS label routes are exported into the bgp.l2vpn table. [PR1379621](#)
- Traffic loss might be observed when the inter-AS NG-MVPN VRF is disabled on one of the ASBRs. [PR1460480](#)
- The l2circuit displays MM status, which might cause traffic loss. [PR1462583](#)
- The l2circuit connections might be stuck in OL state after changing the l2circuit community and flapping the primary LSP path. [PR1464194](#)
- The rpd might crash when "link-protection" is added/deleted from LSP for MVPN ingress replication selective provider tunnel. [PR1469028](#)
- The l2circuit neighbor might be stuck in RD state at one end of MG-LAG peer. [PR1498040](#)
- MPLS label manager could allow configuration of a duplicated VPLS static label. [PR1503282](#)
- The rpd might crash after removing the last interface configured under the l2circuit neighbor. [PR1511783](#)
- The rpd might crash when deleting l2circuit configuration in a specific sequence. [PR1512834](#)

## 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 mismatch. [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 inconsistency 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 pfd 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 jdncpd 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](#)
- **mshpmand[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 is 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 jdncpd 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 synchronizaton 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 ESI. [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 jdncpd 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 was the error message **?CHASSISD\_CONFIG\_ACCESS\_ERROR: pic\_parse\_ifname: Check fpc rname 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 composite 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 I2 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 unresponsive. [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 software 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 propagated 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 PPM[D][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 | 97](#)

[What's Changed | 135](#)

[Known Limitations | 146](#)

[Open Issues | 151](#)

[Documentation Updates | 214](#)

[Migration, Upgrade, and Downgrade Instructions | 215](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 documentation for MX Series.

## SEE ALSO

[What's New | 97](#)

[What's Changed | 135](#)
[Known Limitations | 146](#)
[Open Issues | 151](#)
[Resolved Issues | 167](#)
[Migration, Upgrade, and Downgrade Instructions | 215](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.3 | 216](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 216](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 219](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 220](#)
- [Upgrading a Router with Redundant Routing Engines | 221](#)
- [Downgrading from Release 19.3 | 221](#)

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-x86-32-19.3R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.3R3.9-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-x86-32-19.3R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.3R3.9-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-19.3R3.9-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-19.3R3.9-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

<a href="#">What's New   97</a>
<a href="#">What's Changed   135</a>
<a href="#">Known Limitations   146</a>
<a href="#">Open Issues   151</a>
<a href="#">Resolved Issues   167</a>

# Junos OS Release Notes for NFX Series

## IN THIS SECTION

- What's New | 222
- What's Changed | 224
- Known Limitations | 225
- Open Issues | 226
- Resolved Issues | 228
- Documentation Updates | 232
- Migration, Upgrade, and Downgrade Instructions | 232

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os)

## What's New

## IN THIS SECTION

- Release 19.3R3 New and Changed Features | 223
- Release 19.3R2 New and Changed Features | 223
- Release 19.3R1 New and Changed Features | 223

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

Release 19.3R3 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 19.3R3.

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

<a href="#">What's Changed   224</a>
<a href="#">Known Limitations   225</a>
<a href="#">Open Issues   226</a>
<a href="#">Resolved Issues   228</a>
<a href="#">Documentation Updates   232</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   232</a>

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 19.3R2-S6 | 224](#)

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 | 222](#)

---

[Known Limitations | 225](#)

---

[Open Issues | 226](#)

---

[Resolved Issues | 228](#)

---

[Documentation Updates | 232](#)

---

[Migration, Upgrade, and Downgrade Instructions | 232](#)



## Known Limitations

### IN THIS SECTION

- [Interfaces | 225](#)
- [Platform and Infrastructure | 225](#)
- [Virtual Network Functions \(VNFs\) | 226](#)

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](#)
- The Routing Engine boots from the secondary disk when you:
  - Press the reset button on the RCB front panel, while the RE 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](#)

- With SRX1500 device used as HUB and NFX350 device as SPOKE, **IPSEC replay-errors** are seen with HTTP traffic when the AppQoE passive probing is enabled. As a workaround, use SRX4200 as HUB. [PR1461068](#)

## 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](#)

### SEE ALSO

[New and Changed Features | 222](#)

[What's Changed | 224](#)

[Open Issues | 226](#)

[Resolved Issues | 228](#)

[Documentation Updates | 232](#)

[Migration, Upgrade, and Downgrade Instructions | 232](#)

## Open Issues

### IN THIS SECTION

- [Interfaces | 226](#)
- [High Availability \(HA\) | 227](#)
- [Platform and Infrastructure | 227](#)
- [Virtual Network Functions \(VNFs\) | 227](#)

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.

### 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](#)
- On NFX150 devices, srxpfe core dump is observed when testing the ADSL interfaces. [PR1485384](#)

## High Availability (HA)

- On NFX250 devices, virtual port peer (VPP) functionality might not run on dual customer premises equipment (CPE) and occasionally on single CPE. This issue is noticed when a **redis** database file, **appendonly.aof** has stale entries before redis-server started after a boot. [PR1461238](#)

## 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](#)
- When you upgrade the NFX250 devices from Junos OS Release 15.1X53-D47.4 to Junos OS Release 18.4R1, it overwrites the existing **vsrx.xml** file to factory default file and the jdm subsystem becomes available. As a workaround, fix **vsrx.xml** file with the correct path to the **qcow2** file. [PR1456900](#)
- After you upgrade the NFX250 devices to Junos OS Release 18.4R3, login access to JDM via TACACS fails. [PR1504915](#)

## Virtual Network Functions (VNFs)

- 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 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](#)

SEE ALSO

[New and Changed Features | 222](#)

[What's Changed | 224](#)

[Known Limitations | 225](#)

[Resolved Issues | 228](#)

[Documentation Updates | 232](#)

[Migration, Upgrade, and Downgrade Instructions | 232](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 19.3R3 | 228](#)
- [Resolved Issues: 19.3R2 | 229](#)
- [Resolved Issues: 19.3R1 | 230](#)

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: 19.3R3

#### *Application Security*

- AppQOE is sending active prob packets for the deleted **active-probe-params** option. [PR1492208](#)

#### *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*

- On NFX Series devices, ping is not working between the cross-connected interfaces with interface deny-forwarding configuration. [PR1442173](#)
- On NFX150 devices, heth-0-4 and heth-0-5 ports do not detect traffic when you try to activate the ports by plugging or unplugging the cable. [PR1449278](#)
- On NFX Series devices, the static MAC address on VNF might get replaced by random MAC. [PR1458554](#)

- When traffic goes through vSRX3.0 devices, Packet Forwarding Engine does not come up and all interfaces are down. [PR1465132](#)
- On NFX Series devices, GRE tunnel interface (gr-1/0/0) may not appear if **clear-dont-fragment-bit** option is configured for the GRE interface. [PR1472029](#)
- On NFX350 devices, deleting and adding SXE interfaces move SXE interface to STP BLK state and the traffic drops on that interface. [PR1475854](#)

#### *Platform and Infrastructure*

- On NFX Series devices, process jdmd is not responsive because the **/etc/hosts** file gets corrupted. [PR1477151](#)
- On NFX250 devices, coredumps are observed while adding the second LAN subnet. [PR1490077](#)
- On NFX250-ATT-LS1 devices, LTE packages are lost after a software image upgrade from Junos OS Release 15.1X53-D497.1 to Junos OS Release 18.4R3.3. [PR1493711](#)
- The NFX250 devcies might fail to work after you upgrade the software image from Junos OS Release 15.1X53-D497.1 to Junos OS Release 18.4R3-S2. [PR1507165](#)
- On NFX150 devices, ZTP over LTE configuration commit fails for **operation=create** in xml operations configuration. [PR1511306](#)

#### *Virtual Network Functions (VNFs)*

- On NFX150-S1 devices, management ports are not disabled with the **link disable** command. [PR1442064](#)
- On NFX150 devices, **vmhost virtualization-options** command is not working as expected for heth to ge interface mapping. [PR1459885](#)
- On NFX Series devices, no error message is displayed if the commit fails when you configure **native-vlan-id** on an access VNF interface. [PR1438854](#)

### **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 jnxlpSecTunMonOutEncryptedBytes**
  - **show snmp mib walk jnxlpSecTunMonOutEncryptedPkts**
  - **show snmp mib walk jnxlpSecTunMonInDecryptedBytes**
  - **show snmp mib walk jnxlpSecTunMonInDecryptedPkts**
  - **show snmp mib walk jnxlpSecTunMonLocalGwAddr**
  - **show snmp mib walk jnxlpSecTunMonLocalGwAddrType**

[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 | 222](#)

[What's Changed | 224](#)

[Known Limitations | 225](#)

[Open Issues | 226](#)

---

[Documentation Updates | 232](#)

---

[Migration, Upgrade, and Downgrade Instructions | 232](#)

---

## Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 documentation for NFX Series.

### SEE ALSO

---

[New and Changed Features | 222](#)

---

[What's Changed | 224](#)

---

[Known Limitations | 225](#)

---

[Open Issues | 226](#)

---

[Resolved Issues | 228](#)

---

[Migration, Upgrade, and Downgrade Instructions | 232](#)

---

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 232](#)
- [Basic Procedure for Upgrading to Release 19.3 | 233](#)

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 19.3R3:

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

---

[New and Changed Features | 222](#)

---

[What's Changed | 224](#)

---

[Known Limitations | 225](#)

---

[Open Issues | 226](#)

---

[Resolved Issues | 228](#)

---

[Documentation Updates | 232](#)

## Junos OS Release Notes for PTX Series Packet Transport Routers

#### IN THIS SECTION

- [What's New | 235](#)
- [What's Changed | 243](#)
- [Known Limitations | 248](#)
- [Open Issues | 249](#)
- [Resolved Issues | 251](#)

- Documentation Updates | 257
- Migration, Upgrade, and Downgrade Instructions | 258

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- What's New in Release 19.3R3 | 235
- What's New in Release 19.3R2 | 235
- What's New in Release 19.3R1 | 236

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

### What's New in Release 19.3R3

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 19.3R3.

### What's New in Release 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 Release 19.3R1

### *Class of Service*

- **Support for class-based-forwarding (CBF) with 16 forwarding-classes support (PTX Series)**—Starting in Junos OS Release 19.3R1, we support CBF with 16 forwarding-classes.  
[See [Understanding 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/rtssock/

[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 are 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

Platforms	Routing Engine Supported
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
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 | 243](#)

[Known Limitations | 248](#)

[Open Issues | 249](#)

[Resolved Issues | 251](#)

[Documentation Updates | 257](#)

[Migration, Upgrade, and Downgrade Instructions | 258](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 19.3R3-S1 | 244](#)
- [What's Changed in Release 19.3R3 | 244](#)
- [Juniper Extension Toolkit \(JET\) | 245](#)
- [What's Changed in Release 19.3R2-S6 | 245](#)
- [What's Changed in Release 19.3R2 | 246](#)
- [What's Changed in Release 19.3R1 | 246](#)

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

## What's Changed in Release 19.3R3-S1

### *Routing Protocols*

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.

## What's Changed in Release 19.3R3

### *General Routing*

- **Trigger alarms when a PTX10008 or PTX10016 router has a mix of AC and DC power supplies**—If you insert a mix of AC and DC power supply units (PSU) into a PTX10008 or PTX10016 router, Junos OS raises an alarm to indicate that there is a mix of AC and DC power supplies in the router. To fix this alarm, you must ensure that the router has the same type of power supplies.

[See [Understanding Chassis Alarms.](#)]

- **Control-plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in Junos OS Release 19.3R3, the `arp-snoop` packet type option at the `[edit system ddos-protection protocols arp]` protocol group is renamed `arp`. This packet type option enables you to change default control-plane DDoS protection policer parameters for ARP traffic. After this change, the `edit system ddos-protection protocols arp` protocol group includes `aggregate`, `arp`, and `unclassified` packet type options.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\).](#)]

### *Interfaces and Chassis*

- **Change in fabric error handling behavior (PTX10008, PTX10016, PTX5000 routers [with FPC3-PTX-U2, FPC3-PTX-U3 FPCs], QFX10008, QFX10016, and QFX10002 switches)**—Starting in Junos OS Release 19.3R3, when the Packet Forwarding Engine encounters ECC errors or parity errors related to fabric that are fatal, major, or correctable minor errors, the interfaces on the Packet Forwarding Engine are disabled. You must reboot the FPC manually to recover from the error. If you still face an issue after rebooting the FPC, contact our Customer Support. In earlier releases, when the Packet Forwarding Engine encounters any error (fatal, major, minor\_correctable, minor\_transient, and info), the errors are incorrectly classified as info and, as a result, ignored.

### **Juniper Extension Toolkit (JET)**

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

[See [traceoptions \(Services\)](#).]

### **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

### *Interfaces and Chassis*

- **Updates to the show interfaces and show policer commands (PTX Series)**—Starting in Junos OS release 19.3R2, on PTX Series routers, when you issue the **show interfaces** command or the **show policer** command, the output does not display the default ARP policer (**\_default\_arp\_policer\_**). In earlier releases, when you issue the **show interfaces** command or the **show policer** command, the output displays the default ARP policer (**\_default\_arp\_policer\_**) even though PTX Series routers don't support that policer.

## What's Changed in Release 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 | 235](#)

[Known Limitations | 248](#)

[Open Issues | 249](#)

[Resolved Issues | 251](#)

[Documentation Updates | 257](#)

[Migration, Upgrade, and Downgrade Instructions | 258](#)

## Known Limitations

### IN THIS SECTION

- [General Routing | 248](#)
- [Interfaces and Chassis | 249](#)

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

- The following error message is observed during LC1101 reboot: **pechip\_cmerror\_set\_error:3113: Level: Major, cmerror\_code: 0x210613 (id=1555)**. [PR1268678](#)
- On the PTX10008 and PTX100016 routers, GRES takes more than 3 minutes to complete when shutdown is initiated by the internal **vmhost init 0** command. [PR1312065](#)
- The system might start from the second disk in certain scenarios. [PR1344342](#)
- Adding a warning while committing the host interface might stop sending packets on the PTX1000, PTX3000, PTX5000, and PTX10000 routers when using outbound firewall filter with the syslog option. [PR1354580](#)
- The **request vmhost power-off** command does not shut down the system. [PR1393061](#)
- 100 percent traffic loss is observed on all streams from the PTX10001 router to MX240 router. [PR1435069](#)
- On the PTX10000 routers, source MAC is learned even when the traffic drops as a result of using the Ethernet-switching filter. [PR1436377](#)
- For the PTX10002-60C router, watchdog does not work. [PR1441963](#)
- Trace observed during image upgrade from WRL6 to WRL9. [PR1442017](#)
- The **clear interface statistics** command takes much longer time to execute than expected. [PR1447851](#)
- Interface statistics might not get updated for the local-loopback test. [PR1458814](#)
- Traffic failure with the **gcm-aes-xpn-128** cipher observed on performing an event. [PR1460254](#)
- The **aftd** process gets hogged on executing the **clear vpls table** command and MAC entries are not learned for less than 5 minutes. [PR1473334](#)



## Interfaces and Chassis

- The cfmd process might continuously crash after the upgrade. [PR1281073](#)

### SEE ALSO

[What's New | 235](#)

[What's Changed | 243](#)

[Open Issues | 249](#)

[Resolved Issues | 251](#)

[Documentation Updates | 257](#)

[Migration, Upgrade, and Downgrade Instructions | 258](#)

## Open Issues

### IN THIS SECTION

- [General Routing | 249](#)
- [Routing Protocols | 250](#)

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

- When CFP2-DCO-T-WDM-1 is plugged to the PIC of the device, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)
- The RIB and FIB might be out of synchronization if the KRT asynchronous queue becomes nonresponsive. [PR1315212](#)
- The alarm action does not work for minor errors after changing the threshold to 1. [PR1345154](#)
- Unable to stop ZTP bootstrap when the PTX10016 or PTX10008 router is powered on with the factory-default configuration. [PR1369959](#)

- The value of **cur-rx-input-power** is different compared to the CLI output value of the OTN optics. [PR1411023](#)
- Traffic is lost for more than 15 seconds when 50 percent of the aggregated Ethernet interface links are brought down by restarting multiple FPCs. [PR1412578](#)
- On the PTX10008 and PTX5000 routers, the output of the **show filter index < number> counter** command shows value as zero at **28-02-HOSTBOUND\_NDP\_DISCARD\_TERM**. [PR1420057](#)
- On the PTX3000 routers, the firewall counter for lo0 does not increment. [PR1420560](#)
- The FPC might crash with the em2 configuration during initialization. [PR1429212](#)
- More packet loss observed after ISSU with the InterAS Layer 3 VPN OptionB configuration. [PR1435578](#)
- Memory leak observed while running the longevity check. [PR1438358](#)
- Watchdog does not work. [PR1441963](#)
- The KRT queue might become nonresponsive when more than 65,000 IPv6 labeled-unicast routes are received on the BGP-LU IPv6 session, which is configured on PTX10000 routers. [PR1442760](#)
- The **show snmp mib walk jnxContentsDescr** command output does not display the fan controllers. [PR1455640](#)
- After injecting the errors on the FPC, all interrupts do not get recorded. [PR1459367](#)
- Output statistic are not synced with the aggregated Ethernet interface from the member physical interface. [PR1459633](#)
- Layer 3 traffic fragmentation fails without DF bit with l3fwd configuration. [PR1459738](#)
- Statistics comparison between the CLI and JTI for a queue fails as the buffers display incorrect values. [PR1460246](#)
- Traffic fails with the **gcm-aes-xpn-128** cipher on performing an event. [PR1460254](#)
- LSP of the MPLS goes down with the MPLS TED. [PR1461832](#)

## Routing Protocols

- The aggregated Ethernet interface and BFD session remain down after the interface is disabled or enabled. [PR1354409](#)
- The **show dynamic-tunnels database** command does not show the statistics for the first time and fetches the traffic statistics the second time. [PR1445705](#)
- The rpd process might crash if there are more than 3000 BGP IPv6 peers or more than 8000 BGP IPv4 peers in the NSR setup. [PR1461436](#)

## SEE ALSO

<a href="#">What's New</a>	<a href="#">  235</a>
<a href="#">What's Changed</a>	<a href="#">  243</a>
<a href="#">Known Limitations</a>	<a href="#">  248</a>
<a href="#">Resolved Issues</a>	<a href="#">  251</a>
<a href="#">Documentation Updates</a>	<a href="#">  257</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  258</a>

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: Release 19.3R3](#) | [251](#)
- [Resolved Issues: Release 19.3R2](#) | [254](#)
- [Resolved Issues: Release 19.3R1](#) | [255](#)

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: Release 19.3R3

#### *General Routing*

- The PTX Series router interface remains down after maintenance. [PR1412126](#)
- Interface does not come up after the interface flaps and the FPC reboots. [PR1428307](#)
- Reclassification policy applied on the route prefixes might not work. [PR1430028](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- Upgrade fails due to communication failure between Junos-VM and the host OS. [PR1438219](#)
- FPC might reboot if jlock hog occurs on all Junos VM based platforms. [PR1439906](#)

- Changing the hostname triggers an on-change notification, not an adjacency on-change notification. [PR1449837](#)
- FPC generates the **Voltage Tolerance Exceeded** major alarm for each IP 2V5 sensor. [PR1451011](#)
- The 100-Gigabit Ethernet interface might not come up after flapping on the PTX5000 router. [PR1453217](#)
- Along with the 4x1-Gigabit Ethernet feature using the QSFP28 optics, continuous logging in the chassisd file is observed when speed 1-Gigabit Ethernet is configured with `pic_get_nports_inst` and `ch_fru_db_key`. [PR1456253](#)
- Interface statistics might not get updated for local-loopback test. [PR1458814](#)
- PTX1000 and PTX10002 routers might get drop or discarded after transient SIB or FPC voltage alarms. [PR1460406](#)
- Hardware failure in CB2-PTX causes traffic interruption. [PR1460992](#)
- On the PTX5000 routers, for the FPC3 line card, the **optics-options** syslog and link-down do not work as expected. [PR1461404](#)
- The sample, syslog, or log action in the output firewall filters for packets of size less than 128 bytes might cause an ASIC wedge (all packet loss). [PR1462634](#)
- On the PTX3000 or PTX5000 routers, PICs might restart if the QSFP optics get overheated. [PR1462987](#)
- On the PTX10000 routers, FPCs might restart during runtime. [PR1464119](#)
- EBUF parity interrupt is not observed. [PR1466532](#)
- IPv6 traffic might be dropped in a Layer 3 VPN network. [PR1466659](#)
- On the PTX5000 routers, the following Packet Forwarding Engine error logs might be reported when the IGMP packets get sampled: **prds\_packet\_classify\_notification: Failed to find fwd nh for flabel 48**. [PR1466995](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- Incorrect counter values are displayed for the arrival rate and peak rate for DDoS commands. [PR1470385](#)
- Traffic loops for pure Layer 2 packets coming over an EVPN tunnel with the destination MAC address matching the IRB MAC address. [PR1470990](#)
- A PTX5000 SIB3 might fail to come up in slot 0 and/or slot 8 when Routing Engine 1 is the master. [PR1471178](#)
- **input-vlan-map** or **output-vlan-map** might not work properly in the Layer 2 circuit local-switching scenario. [PR1474876](#)
- Sampling process on FPC might crash when the MPLS traffic is sampled. [PR1477445](#)
- Multicast routes add or delete events might cause adjacency and LSPs to go down. [PR1479789](#)
- Traffic impact might be observed when **policy-multipath** is configured without LDP on the SPRING-TE scenario. [PR1483585](#)

- FPC might crash when dealing with invalid next hops. [PR1484255](#)
- BFD sessions start to flap when the firewall filter in loopback0 is changed. [PR1491575](#)
- Pushing configuration to the ephemeral database with a high rate causes outbound SSH connection to flap or leak memory. [PR1497575](#)
- The following error message is observed: **PFE\_ERROR\_FAIL\_OPERATION: IFD et-1/0/8: RS credits failed to return: init=192 curr=193 chip=5.** [PR1502716](#)
- On the PTX10008 and PTX10016 routers, a few TCP-based application sessions might flap upon Routing Engine switchover or application sessions bouncing in the backup Routing Engine. [PR1503169](#)
- On the PTX3000 or PTX5000 router, unable to bring the ports up when plugging in the optic QSFP-100G-LR4-T2(740-061409). [PR1511492](#)
- Sampling with the **rate limiter** command enabled, crosses the sample rate 65535. [PR1525589](#)

### **Infrastructure**

- Slow response from SNMP might be observed after upgrading to Junos OS Release 19.2R1 and later. [PR1462986](#)

### **Layer 2 Ethernet Services**

- Member links state might be asynchronized on a connection between the PE and CE devices in an EVPN active/active scenario. [PR1463791](#)

### **MPLS**

- Kernel might crash and device might restart. [PR1478806](#)
- BGP session might keep flapping between two directly connected BGP peers because of the wrong TCP MSS (maximum segment size) in use. [PR1493431](#)
- The rpd process might crash in a rare condition in an SR-TE scenario. [PR1493721](#)
- SNMP trap is observed with incorrect OID jnxSpSvcSetZoneEntered. [PR1517667](#)

### **Routing Protocols**

- SSH login might fail if a user account exists in both the local database and RADIUS/TACACS+. [PR1454177](#)
- The **other querier present interval** timer cannot be changed in an IGMP or MLD snooping scenario. [PR1461590](#)
- The ppmmd process crashes after configuring S-BFD responder on the PTX Series routers with RE-DUO-2600. [PR1477525](#)
- The rpd process might crash with BGP multipath and routes withdrawn occasionally. [PR1481589](#)

- The BGP multipathed traffic might not fully load-balance for a while after adding a new path for load sharing. [PR1482209](#)
- The BGP route-target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

## Resolved Issues: Release 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: Release 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 jdncpd 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

[What's New | 235](#)

[What's Changed | 243](#)

[Known Limitations | 248](#)

[Open Issues | 249](#)

[Documentation Updates | 257](#)

[Migration, Upgrade, and Downgrade Instructions | 258](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 documentation for the PTX Series.

## SEE ALSO

- [What's New | 235](#)
- [What's Changed | 243](#)
- [Known Limitations | 248](#)
- [Open Issues | 249](#)
- [Resolved Issues | 251](#)
- [Migration, Upgrade, and Downgrade Instructions | 258](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.3 | 258](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 261](#)
- [Upgrading a Router with Redundant Routing Engines | 262](#)

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 **bundle** 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 19.3R3:

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-x86-64-19.3R3.9.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-x86-64-19.3R3.9-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 19.3R3, 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

<a href="#">What's New   235</a>
<a href="#">What's Changed   243</a>
<a href="#">Known Limitations   248</a>
<a href="#">Open Issues   249</a>
<a href="#">Resolved Issues   251</a>
<a href="#">Documentation Updates   257</a>

# Junos OS Release Notes for the QFX Series

## IN THIS SECTION

- What's New | 263
- What's Changed | 272
- Known Limitations | 278
- Open Issues | 281
- Resolved Issues | 286
- Documentation Updates | 306
- Migration, Upgrade, and Downgrade Instructions | 307

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- What's New in Release 19.3R3 | 264
- What's New in Release 19.3R2 | 264
- What's New in Release 19.3R1 | 264

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 19.3R3: 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 Release 19.3R3

There are no new features or enhancements to existing features for QFX Series switches in Junos OS Release 19.3R3.

## What's New in Release 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 Release 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+ transceiver 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 bpdv-block-on-edge
```

- To enable BPDU protection without a spanning-tree protocol configured on access and leaf devices:

```
set protocols layer2-control bpdv-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 bpdv-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 19.3R1 supports 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 **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:** 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 are 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

<a href="#">What's Changed</a>	<a href="#">  272</a>
<a href="#">Known Limitations</a>	<a href="#">  278</a>
<a href="#">Open Issues</a>	<a href="#">  281</a>
<a href="#">Resolved Issues</a>	<a href="#">  286</a>
<a href="#">Documentation Updates</a>	<a href="#">  306</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  307</a>

# What's Changed

IN THIS SECTION

- [What's Changed in Release 19.3R3-S1](#) | [273](#)
- [What's Changed in Release 19.3R3](#) | [273](#)
- [What's Changed in Release 19.3R2-S6](#) | [275](#)



● [What's Changed in Release 19.3R2 | 276](#)

● [What's Changed in Release 19.3R1 | 276](#)

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

## What's Changed in Release 19.3R3-S1

### *Routing Protocols*

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.

## What's Changed in Release 19.3R3

### *General Routing*

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on `ge`, `et`, `xe` interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for `ge` interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (`ge-0/0/0`), was displayed. Now, the logical interface (`ge-0/0/0.16386`) is also displayed.
- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

### *Interfaces and Chassis*

- **Logical Interface is created along with physical Interface by default (QFX Series)**—In Junos OS Release 19.3R3 and later, by default, logical interfaces are created along with the physical interfaces on `ge`-, `et`-, `xe`- ports. In earlier Junos OS releases, by default, only physical interfaces are created.

For example, for **ge** interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface, for example, (`ge-0/0/0`), is displayed. Now, the logical interface, for example, (`ge-0/0/0.16386`) is also displayed.

- **Autonegotiation status displayed correctly (QFX5120-48Y)**—In Junos OS Release 19.3R3, the **show interfaces interface-name <media> <extensive>** command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed.

In the earlier Junos OS releases, incorrect autonegotiation status is displayed even when autonegotiation is disabled.

#### *Juniper Extension Toolkit (JET)*

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

[See [traceoptions \(Services\)](#).]

#### *Multicast*

- **IGMP snooping in EVPN-VXLAN multihoming environments (QFX5110)**—In an EVPN-VXLAN multihoming environment on QFX5110 switches, you can now selectively enable IGMP snooping only on those VLANs that might have interested listeners. In earlier releases, you must enable IGMP snooping on all VLANs associated with any configured VXLANs because all the VXLANs share VXLAN tunnel endpoints (VTEPs) between the same multihoming peers and require the same settings. This is no longer a configuration limitation.
- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 19.3R3, we provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... n packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). Before Junos OS Release 19.3R3, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** do not provide valid statistics for multicast traffic at Layer 2.

[See [show multicast snooping route](#).]

#### *Platform and Infrastructure*

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**—Starting in this release, the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group is renamed simply **arp**. This packet type option enables you to change default control plane DDoS protection policer parameters for ARP traffic. After this change, the **edit system ddos-protection protocols arp** protocol group includes **aggregate**, **arp**, and **unclassified** packet type options.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

- **Priority-based flow control (PFC) support (QFX5120-32C)**—Starting with JunosOS Release 19.3R3, QFX5120-32C switches support priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic.

## Routing Protocols

- **IGMP snooping in EVPN-VXLAN multihoming environments (QFX5110)**— In an EVPN-VXLAN multihoming environment on QFX5110 switches, you can now selectively enable IGMP snooping only on those VLANs that might have interested listeners. In earlier releases, you must enable IGMP snooping on all VLANs associated with any configured VXLANs because all the VXLANs share VXLAN tunnel endpoints (VTEPs) between the same multihoming peers and require the same settings. This is no longer a configuration limitation.
- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—In Junos OS Release, multiple loopback addresses export into `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier Junos OS releases, multiple secondary loopback addresses in TED were added into `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router-id.

## 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

### *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](#).]

## What's Changed in Release 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.
- **XML RPC equivalent included for the show bgp output-scheduler | display xml rpc CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.3R1, we've included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.3R1, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler](#).]

## 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 | 263](#)

---

[Known Limitations | 278](#)

---

[Open Issues | 281](#)

---

---

[Resolved Issues | 286](#)

---

[Documentation Updates | 306](#)

---

[Migration, Upgrade, and Downgrade Instructions | 307](#)

## Known Limitations

### IN THIS SECTION

- [Class of Service \(CoS\) | 279](#)
- [EVPN | 279](#)
- [Platform and Infrastructure | 279](#)
- [Infrastructure | 280](#)
- [Layer 2 Features | 280](#)
- [Routing Protocols | 281](#)

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](#)

## Platform and Infrastructure

- 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](#)
- Downgrade from a TVP image to a non-TVP image is not supported. However, upgrade from a non-TVP image to a TVP image is supported. [PR1345848](#)
- The 100-Gigabit Ethernet interface goes down after you configure and delete the Ethernet loopback configuration. [PR1353734](#)
- 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](#)
- With 288,000 MAC scale, the Routing Engine command **show ethernet-switching table summary** output displays the learned scale entries after a delay of around 60 seconds. [PR1367538](#)
- 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](#)
- On the QFX5120 switches, due to the additional hi-gig header, 100 percent throughput cannot be achieved when packets are forwarded through Virtual Chassis ports. [PR1453709](#)
- Observing 100 percent L2 MAC scaling traffic loss in QFX10002-60C platform after loading EVPN-VXLAN collapsed profile configurations. [PR1489753](#)

## 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](#)
- On QFX5000 platforms, you might see the `pools exhausted for Table:EGR_DVP_ATTRIBUTE` error message when statistics requests exceed the supported scale because of the limited pool resources used for statistics collection on the hardware. There is no functional impact except for statistics collection for some hardware counters for which flex counter allocation failed for the time, the limit is exceeded. The statistics counters start functioning normally without manual change when the pool comes back to normal limit. [PR1479826](#)



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, **I3 egress intf**, **egr I3 next hop**, and **ingress I3 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

<a href="#">What's New   263</a>
<a href="#">What's Changed   272</a>
<a href="#">Open Issues   281</a>
<a href="#">Resolved Issues   286</a>
<a href="#">Documentation Updates   306</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   307</a>

Open Issues

IN THIS SECTION

- [EVPN | 282](#)
- [High Availability \(HA\) and Resiliency | 282](#)
- [Infrastructure | 282](#)
- [Interfaces and Chassis | 282](#)
- [Junos Fusion Provider Edge | 283](#)
- [Layer 2 Features | 283](#)
- [Platform and Infrastructure | 283](#)
- [Routing Protocols | 285](#)
- [Virtual Chassis | 286](#)

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 switch is encapsulating VXLAN traffic and sending it to an incorrect destination MAC address when multiple remote VTEPs are in the same subnet and reachable by means of an IRB interface in a stretched VLAN. [PR1424698](#)
- In an EVPN-VXLAN core isolation scenario, the server is multihomed to the leaf devices through LACP interfaces. If graceful restart is enabled, when you reboot the system or restart routing on the leaf device, the core isolation does not work. If you reboot the system, the issue results in the leaf device silently dropping the traffic sent from the server during the time window between LACP coming up and BGP coming up. If you restart routing, there might be no traffic drop because of the graceful restart. [PR1461795](#)

## High Availability (HA) and Resiliency

- Syslog error message **kernel: GENCFG: op 51 (AE bias) failed; err 255 (Undefined)** is seen. These messages do not have any functionality impact. [PR1416004](#)

## Infrastructure

- The error message **ftpd[14105]: bl\_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)** is seen during FTP. [PR1315605](#)

## Interfaces and Chassis

- Multicast traffic can be flooded for 15 to 20 seconds to both MC-LAG peers, after the following sequence of steps:
  1. Disable or enable ICL.
  2. Reboot one of the MC-LAG peers.
  3. Disable or enable a member link of ICL.

This results in no traffic loss, and one of the MC-LAG nodes processes duplicate packets during this time period. [PR1422473](#)

- Flooding of ARP reply unicast packets is seen as a result of an ARP request sent for the device's VRRP MAC address. 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](#)

## Junos Fusion Provider Edge

- IGMP membership is not getting learned by the AD fully even when the IGMP queries are being sent out. [PR1419265](#)

## Layer 2 Features

- In case of QFX5000 Virtual Chassis and Virtual Chassis Fabric setups, when IGMP snooping is enabled, multicast traffic is forwarded based on IGMP joins/reports. But when the IGMP report times out, traffic should be dropped; instead it is flooded in the VLAN. This happens only in case of QFX5000 Virtual Chassis and Virtual Chassis Fabric; this issue is not seen on stand-alone QFX5000 devices. [PR1431893](#)

## Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as **nh\_ucast\_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- On QFX10002-60C switches, the **filter** operation with the **log** action is not supported for protocols other than Layer 2, IPv4, and IPv6. The following message is seen in firewall logs: **Protocol 0 not recognized**. [PR1325437](#)
- Backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- The QFX10000 platform drops the wireless access point (WAP) heartbeat packets; as a result, the WAP cannot work. [PR1352805](#)
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)
- On the QFX10000 line of switches with EVPN-VXLAN, the following error message is seen: **expr\_nh\_fwd\_get\_egress\_install\_mask:nh type Indirect of nh\_id: # is invalid**. [PR1367121](#)
- The user might not be able to stop the ZTP bootstrap process when a QFX10016 or QFX10008 switch with many line cards is powered on with the factory-default configuration. [PR1369959](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- 100-Gbps VCP links might go down (become unavailable) after the linecard 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 an FPC (DRAM and buffer utilization). [PR1394978](#)

- On QFX5000 platforms with a scaled setup of the aggregated Ethernet bundles and VLANs, if LACP is enabled, and there are scaled configuration changes, for example, deleting 4000 VLANs/VXLANs and reapplying them, some interfaces of the aggregated Ethernet bundle might go to the detached state. Due to this issue, the running routing protocols (for example, LACP and BGP) go down over the affected aggregated Ethernet bundles. [PR1406691](#)
- You might see multiple reconnect logs, JTASK\_IO\_CONNECT\_FAILED, during the device initialization. There is no functionality impact because of these messages. These messages can be ignored. [PR1408995](#)
- On QFX5110 and QFX5120 platforms, unicast RPF check in strict mode might not work properly. [PR1417546](#)
- On the QFX10000 line of devices, if an 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, then the analyzer might not provide sample packets that flow through the newly added child interface. [PR1417694](#)
- The **show ptp lock-status** command is not supported on the QFX5110-48S-4C device. [PR1426863](#)
- Power cycling while ISSU/ISSR is in progress does not seem 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 into a problematic state. As a workaround, use the following commands to clear ISSU/ISSR-specific sysctl and 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 does not damage anything beyond the ISSU. [PR1427563](#)

- The unified ISSU is not supported on QFX5200 switches and fails from Junos OS Release 17.2X75-D43.2 through some target versions. Also, dcpfe crash might be seen. [PR1438690](#)
- The unified ISSU fails and is not supported for QFX5200 from Junos OS Release 17.2X75-D4(x) through Junos OS Release 19.2R1. [PR1440288](#)
- On QFX5000 platforms, the port qualifier is supported. This installs two entries in the Packet Forwarding Engine, one with source-port and second one with destination-port with value as specified in the port configuration. [PR1440980](#)
- On the QFX10000 line of switches, removal of the EVPN-VXLAN Layer 3 gateway on the IRB interface from the spine switches might cause traffic to be silently discarded. As a workaround, configure all virtual gateways with unique IPv4 or IPv6 MAC addresses. [PR1446291](#)

- Whenever any member in an RSPAN VLAN is removed from that VLAN, you must reconfigure the analyzer session for that RSPAN VLAN. [PR1452459](#)
- On QFX5200-32C-32Q switches, VM core file is generated after upgrading from Junos OS Release 18.3 throttle image to Junos OS Release 19.3R1 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. [PR1455851](#)
- Fan display in the output of the **show chassis environment** command is not proper. [PR1457896](#)
- On the QFX5000 and QFX10000 lines of switches, 9.51 percent degradation with commit time and 12 percent degradation with VLAN commit convergence are observed. [PR1457939](#)
- On the QFX5110 switches, VXLAN VNI (multicast) scaling traffic issue is observed from the VXLAN tunnel to the Layer 2 interface. [PR1462548](#)
- On QFX Series platforms, if there are a lot of MAC moves, the system might stop new MAC learning and old MAC addresses might get stuck, aged, and deleted. Because of this issue, Layer 2 traffic forwarding and the customer service might be impacted. [PR1475005](#)
- If the archival function is enabled with the routing-instance statement and transfer mode is set as SFTP/SCP for archive-sites, it might not work and fail to transfer the files. [PR1507044](#)
- On QFX Series platforms, in rare case, HMC memory error might cause routes update failure and traffic impact on the Packet Forwarding Engine. [PR1515092](#)
- Disruptive switchover (no GRES or NSR configured) can lead to stale PPM entries programmed on the new master Routing Engine. If both GRES and NSR are activated after disruptive switchover and then a GRES is performed, the BFD sessions might flap continuously. [PR1518106](#)
- On all Junos platforms that support Open vSwitch database (OVSDb), the vgd core file might generated when the OVSDb server is restarted. The vgd daemon restart after the core file generation might cause traffic impact. This issue happens when OVSDb server is disconnected and the device sends some updates events to server. [PR1518807](#)
- SNMP MIB walk with jnxContentsDescr does not show power supply number on QFX platforms. [PR1529108](#)

## 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](#)
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, when the mini-PDT-base configuration is issued, the following error message is seen in the hardware: **BRCM\_NH-,brcm\_nh\_bdvlan\_ucast\_uninstall(), 128:I3 nh 6594 unintsall failed.** There is no functionality impact because of this error message. [PR1407175](#)

- On QFX5120 switches, transit traffic drop is seen with loop back filter configuration. [PR1426828](#)
- QFX5110 MC-LAG: L2\_L3\_INTF\_OPS\_ERROR messages are seen after node reboot. [PR1435314](#)
- On the QFX5100 Virtual Chassis or Virtual Chassis Fabric, traffic loss on multiple traffic streams after reboot and interface flapping of the Virtual Chassis node are observed. [PR1500508](#)
- On QFX10000 platforms, deleting the physical interface under an aggregated Ethernet interface might flap the BFD sessions formed on the remaining physical interfaces if the same IP address is configured on multiple units in different routing instances. [PR1516556](#)

Virtual Chassis

- QFX5110-48S reports false parity error messages such as `soc_mem_array_sbusdma_read`. The QFX5110-48S SDK can raise false alarms for parity error messages such as `soc_mem_array_sbusdma_read`. This is a false positive error message. [PR1276970](#)
- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. [PR1490552](#)

SEE ALSO

<a href="#">What's New   263</a>
<a href="#">What's Changed   272</a>
<a href="#">Known Limitations   278</a>
<a href="#">Resolved Issues   286</a>
<a href="#">Documentation Updates   306</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   307</a>

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 19.3R3 | 287](#)
- [Resolved Issues: Release 19.3R2 | 295](#)
- [Resolved Issues: Release 19.3R1 | 299](#)

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: Release 19.3R3

### *Platform and Infrastructure*

- On QFX5100 Virtual Chassis, **MacDrainTimeOut** and **bcm\_port\_update failed: Internal error** is observed. [PR1284590](#)
- Port LEDs do not work on QFX5100 in QFX5110-QFX5100 mixed mode Virtual Chassis. [PR1317750](#)
- During bootup, the error message **CMQFX: Error requesting SET BOOLEAN, illegal setting 66** is generated. [PR1385954](#)
- The 10-Gigabit Ethernet fiber interfaces might flap frequently when they are connected to other vendor's switch. [PR1409448](#)
- The **show interface** command shows **Media type: Fiber** on QFX5100-48T switches running the QFX 5e Series image. [PR1419732](#)
- VM core files are generated on QFX Series Virtual Chassis. [PR1421250](#)
- SFP-LX10 stays down until autonegotiate is disabled. [PR1423201](#)
- CoS rewrite rules applied under an aggregated Ethernet interface might not take effect after NSSU. [PR1430173](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- The default logical interfaces on channelized physical interfaces might not be created after ISSU/ISSR. [PR1439358](#)
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB interface. [PR1442587](#)
- CRC error might be seen on the VCPs of the QFX5100 Virtual Chassis. [PR1449406](#)
- The em0 route might be rejected after the em0 interface is disabled and then enabled. [PR1449897](#)
- FPC does not restart immediately after rebooting the system. That might cause packet loss. [PR1449977](#)
- On QFX10000 platforms, CoS classification does not work. [PR1450265](#)
- On QFX5000 switches, no warning or error is shown when the dual VLAN tag feature is configured on the physical interface. [PR1450455](#)
- On QFX5110 Virtual Chassis, the master FPC might come up in master state again after reboot instead of backup. [PR1454343](#)

- On QFX5120 switches, untagged hosts ARP/NS requests connected on the **encapsulation ethernet-bridge** interface are not being resolved. [PR1454804](#)
- A firewall filter might not be able to be applied in a particular Virtual Chassis or Virtual Chassis Fabric member as TCAM is running out of space. [PR1455177](#)
- In a 16+ member QFX5100 Virtual Chassis Fabric, the **FROM** column in the output of **show system users** shows incorrect information. [PR1455201](#)
- The priority-based flow control (PFC) feature does not work on the QFX10000 line of switches. [PR1455309](#)
- The cosd crash might be observed if the **forwarding-class-set** configuration is directly applied on the child interface of an aggregated Ethernet interface. [PR1455357](#)
- Telemetry traffic might not be sent out when the telemetry server is reachable through a different routing instance. [PR1456282](#)
- Link-up delay and traffic drop might be seen on mixed SP Layer 2 or Layer 3 and EP Layer 2 type configurations. [PR1456336](#)
- QFX5110 QSFP-100GBASE-SR4 made by the third party cannot link up. [PR1457266](#)
- The Packet Forwarding Engine process might crash after Routing Engine switchover on QFX10000 platforms. [PR1457414](#)
- On QFX5110 switches, port 51 has one LED blinking amber continuously. [PR1457516](#)
- Intermittent LAG interface flap might be seen on QFX platforms. [PR1458363](#)
- On QFX5210 switches, the LED does not light on port 64 and 65 after an upgrade to Junos OS Release 19.2R1. [PR1458514](#)
- On QFX5200 switches, the lightweight DHCPv6 relay agent functionality might be broken. [PR1459499](#)
- The **accept-source-mac** feature with VXLAN is not working on QFX5000 platforms. [PR1460885](#)
- On the QFX10000 line of switches, **show forwarding-options enhanced-hash-key** does not work. [PR1462519](#)
- On the QFX5000 line of switches, the fxpc process might generate a core file when you change MTU in a VXLAN scenario with firewall filters applied. [PR1462594](#)
- On QFX5100 devices, the interface output counter is double counted for self-generated traffic. [PR1462748](#)
- On QFX5100 Virtual Chassis and Virtual Chassis Fabric, the error **BRCM-VIRTUAL,brcm\_vxlan\_walk\_svp(),6916:Failed to find L2-iff for ifl:** might appear during cleanup EVPN-VXLAN configurations with Mini-PDT base configurations. [PR1463939](#)
- On the QFX10000 line of switches, the FPC process might restart during runtime. [PR1464119](#)
- On QFX10000 platforms, the interface might not come up on FPC restart. [PR1464650](#)



- On QFX5100-24Q switches, unable to attach a filter to an IRB interface when that filter contains an action to remark the DSCP value to a nonzero value. [PR1464883](#)
- PEM is not present spontaneously on QFX5210 switches. [PR1465183](#)
- On QFX5100-48T switches, a 10-Gigabit Ethernet interface might not come up or negotiate at 1-Gbps speed when connected with the BCM 10G/GbE 2+2P 57800-t rNDC card. [PR1465196](#)
- The QSFP-100G-PSM4 could not be correctly identified on QFX5200 or QFX5110 platforms. [PR1465214](#)
- The physical interface of aggregated Ethernet might take time to come up after disabling or enabling it. [PR1465302](#)
- Junos OS exhibits consistent fan and power supply numbering on whiteboxes (-O and -OZ). [PR1465327](#)
- In Virtual Chassis scenario, the broadcast and multicast traffic might be dropped over an IRB or a LAG interface. [PR1466423](#)
- BGP open messages with specific types of BGP optional capabilities causes BMP messages not to be encoded correctly when sent to the BMP collector. [PR1466477](#)
- On QFX10000 platforms, EBUF parity interrupt is not seen. [PR1466532](#)
- IPv6 traffic over Layer 3 VPN might fail. [PR1466659](#)
- Slow packet drops might be seen on QFX5000 platforms. [PR1466770](#)
- EPR iCRC errors in QFX10000 platforms might cause protocols to go down. [PR1466810](#)
- A few of the DHCP INFORM packets specific to a particular VLAN might be taking the incorrect resolve queue. [PR1467182](#)
- Ingress drops to be included at the CLI from interface statistics and added to InDiscards. [PR1468033](#)
- QFX5000 switches might generate VM core file due to the use of an unsupported QSFP. [PR1468368](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- MAC address might not be learned on a new extended port after vMotion in a Junos fusion data center environment. [PR1468732](#)
- On the QFX5100 line of switches, traffic loss might be seen with framing errors or runs if MACsec is configured. [PR1469663](#)
- If continuous interface flap occurs at ingress or egress of the PE devices, the IP routed packets might get looped on the MPLS PHP node. [PR1469998](#)
- Incorrect counter value is observed for the arrival rate and the peak rate for the DDoS commands. [PR1470385](#)
- The speed 10 Mbps might not be configured on the ge- interface [PR1471216](#)
- When the VTEP source interface is configured in the multiple routing instances, there might be traffic loss. [PR1471465](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)

- DSCP marking might not work as expected if the fixed classifiers are applied to interfaces on QFX5000 platforms. [PR1472771](#)
- The detached interface in a LAG might process the xSTP BPDUs. [PR1473313](#)
- On QFX5000 switches, the **global-mac-table-aging-time** statement shows unexpected behavior with multihomed EVPN VXLAN ESI. [PR1473464](#)
- The l2ald crash might be observed when around 16,000 VLAN IDs share the same VXLAN tunnel and the Packet Forwarding Engine is rebooted. [PR1473521](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)
- On the QFX5000 line of switches in an EVPN-VXLAN scenario, continuous error log messages might be raised. [PR1474545](#)
- On the QFX5000 line of switches, the Layer 2 circuit might fail to communicate through VLAN 2. [PR1474935](#)
- sFlow does not work correctly if the received traffic goes out of more than one interface. [PR1475082](#)
- On the QFX5200 switches, the DAC cables are not being properly detected in the Packet Forwarding Engine. [PR1475249](#)
- There might be traffic drop on QFX5110 or QFX5120 switches acting as leaf switches in a multicast environment with VXLAN. [PR1475430](#)
- FPC major error is observed after the system boots up or the FPC restarts. [PR1475851](#)
- The QFX Series platforms exhibit invalid Packet Forwarding Engine PG counter pairs to copy src 0xfffff80, dst 0. [PR1476829](#)
- On the QFX10002-36Q/72Q switches, the following continuous error messages are logged on the device on getting adoption valid bit[8] asserted: **prds\_ptc\_wait\_adoption\_status: PECHIP[1] PTC[1]: timeout**. [PR1477192](#)
- Egress port mirroring might not work when the analyzer port and the mirrored port belong to a different FPC. [PR1477956](#)
- The default Virtual Chassis MAC persistence timer is incorrectly set to 20 seconds instead of 20 minutes. [PR1478905](#)
- The SLAX script might be lost after upgrading software. [PR1479803](#)
- The remaining interface might be still in **Down** state even though the number of channelized interfaces is no more than 5. [PR1480480](#)
- The ARP request packets for an unknown host might get dropped in a remote PE device in an EVPN-VXLAN scenario. [PR1480776](#)
- VLAN creation failure might be observed with the scaled VLAN and Layer 3 configuration. [PR1484964](#)
- The queue statistics are not as expected after configuring the IFD and logical interface shaping with the transmit rate and the scheduler-map. [PR1488935](#)

- After ISSU or ISSR, a port using SR4 or LR4 optics might not come up. [PR1490799](#)
- The BFD sessions start to flap when the firewall filter in loopback0 is changed. [PR1491575](#)
- Traceroute monitor with mtr version v.69 shows a false 10 percent loss. [PR1493824](#)
- On the QFX5120 switch in the MC-LAG scenario, traffic loss is observed. [PR1494507](#)
- ARP does not get refreshed after timeout on QFX10002-60C switches acting as an EVPN-VXLAN gateway. [PR1497209](#)
- Virtual Chassis is not stable with a channelization interface. [PR1497563](#)
- Outbound SSH connection flap or memory leak issue might be observed when pushing the configuration to an ephemeral DB with high rate. [PR1497575](#)
- An lcmd core file might be generated on QFX52100-64C switches. [PR1497947](#)
- Traffic might get dropped if an aggregated Ethernet member interface is deleted and then added or an SFP transceiver of the aggregated Ethernet member interface is unplugged/plugged. [PR1497993](#)
- On QFX5210 switches, unexpected behavior is seen for Port LED after upgrade. [PR1498175](#)
- Traffic loss might be seen in certain conditions under a multihomed PE scenario. [PR1498863](#)
- On QFX5100 and QFX5110 switches, firewall filter might not get applied. [PR1499647](#)
- BFD sessions flap after deactivating or activating the aggregated Ethernet interface or executing GRES. [PR1500798](#)
- On QFX5100 switches, ERPS might not work correctly. [PR1500825](#)
- Unable to set up LLDP neighborship on QFX5000 platforms. [PR1504354](#)
- Core files are generated if you add or delete ERP configuration multiple times and restart l2cpd. [PR1505710](#)
- On QFX10002, QFX10008, and QFX10016 platforms, PECHIP wedge caused by deactivating CoS ETS configuration affects traffic flow. [PR1509220](#)
- ARP replies might be flooded through the EVPN-VXLAN network as unknown unicast ARP reply. [PR1510329](#)
- The MAC learning might not work properly after multiple MTU changes on the access port in a VXLAN scenario. [PR1516653](#)
- The sFlow adaptive-sampling with the rate limiter statement enabled crosses sample rate 65535. [PR1525589](#)

### ***Class of Service (CoS)***

- On QFX5120 switches, the priority-based flow control (PFC) feature is not supported on 2-member Virtual Chassis currently because of the hardware limitation. [PR1431895](#)
- Shaping does not work after the reboot if **shaping-rate** is configured. [PR1432078](#)

- The traffic is placed in network-control queue on an extended port even if it comes in with different DSCP marking. [PR1433252](#)
- On QFX5120-32C switches, when you move unicast traffic to a multicast queue through an MF classifier, the show interface queue command does not display any status. [PR1459281](#)
- Traffic might be forwarded to an incorrect queue when a fixed classifier is used. [PR1510365](#)

### **EVPN**

- The rpd might crash after changing EVPN-related configuration. [PR1467309](#)
- The ESI of IRB interface does not update after an autonomous-system number change if the interface is down. [PR1482790](#)
- An l2ald memory leak might be observed in an EVPN scenario. [PR1498023](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)
- Unable to create a new VTEP interface. [PR1520078](#)

### **Forwarding and Sampling**

- Type 1 ESI/AD route might not be generated locally on an EVPN PE device in the all-active mode. [PR1464778](#)

### **High Availability (HA) and Resiliency**

- Unified ISSU is not supported on QFX5000 platforms. [PR1472183](#)

### **Infrastructure**

- The l2ald and eventd processes are hogging 100 percent after the **clear ethernet-switching table** command is issued. [PR1452738](#)

### **Interfaces and Chassis**

- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. [PR1459201](#)
- On a QFX Series platform, VRRPv3 MIBs are not working to poll VRRPv6-related objects. [PR1467649](#)
- Executing commit might become unresponsive due to a stuck dcd process. [PR1470622](#)
- Commit error is not thrown when a member link is added to multiple aggregation groups with different interface-specific options. [PR1475634](#)
- MC-LAG consistency check fails if multiple IRB units are configured with the same VRRP group. [PR1488681](#)

### ***Junos Fusion Enterprise***

- Loop detection might not work on extended ports in Junos fusion scenarios. [PR1460209](#)

### ***Layer 2 Features***

- The LLDP function might fail when a Juniper Networks device connects to a non-Juniper one. [PR1462171](#)
- A few MAC addresses might be missing from the MAC table in software on QFX5000 platforms. [PR1467466](#)
- After rebooting, an fxpc core file might be seen when committing the configuration. [PR1467763](#)
- Ingress traffic might be silently dropped if the underlying interface flaps in an EVPN-VXLAN scenario. [PR1469596](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)
- The FPC goes down when a 100-Gigabit Ethernet link comes up on the network port of the QFX5110-48S switches. [PR1499422](#)

### ***Layer 2 Ethernet Services***

- In an EVPN-VXLAN ERB scenario, DHCP relay-source lo0.1 is not used when enabled with anycast legacy IRB. [PR1455076](#)
- Member links state might be asynchronized on a connection between PE and CE devices in an EVPN active/active scenario. [PR1463791](#)
- Issues with DHCPv6 relay processing confirm and reply packets. [PR1496220](#)
- The MC-LAG might be down after disabling and then enabling the **force-up** statement. [PR1500758](#)
- After the switch is rebooted, the aggregated Ethernet interface might sometimes not come up. [PR1505523](#)

### ***MPLS***

- On QFX10002 switches, the **show mpls static-lsp | display xml** command produces invalid XML. [PR1469378](#)
- Traffic might be silently dropped and discarded on a PE device when a CE device sends traffic to the PE device and the destination is resolved with two LSPs through one upstream interface. [PR1475395](#)
- Traffic might be lost over a QFX5100 switch acting as a transit PHP node in the MPLS network. [PR1477301](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. [PR1493431](#)

### ***Routing Protocols***

- OSPF VRF sessions take a long time to come up when the host table is full and host routes are in LPM table. [PR1358289](#)
- Invalid VRRP mastership election on QFX5110 Virtual Chassis peers. [PR1367439](#)

- Host-destined packets with the **filter log** action might not reach the Routing Engine if log/syslog is enabled. [PR1379718](#)
- On QFX5100, BGP IPv4 or IPv6 convergence and RIB install or delete time degraded. [PR1414121](#)
- PIM (S,G) joins can cause MSDP to incorrectly announce source active messages in some cases. [PR1443713](#)
- CRC errors might be seen on QFX5100 Virtual Chassis. [PR1444845](#)
- A core file might be generated when you add or remove an EVPN Type-5 routing instance. [PR1455547](#)
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, the **brcm\_ipmc\_route\_counter\_delete:3900Multicast stat destroy failed (-10:Operation still running)** error is observed after unified ISSU with Mini-PDT base configurations. [PR1460791](#)
- The **other querier present interval** timer cannot be changed in an IGMP/MLD snooping scenario. [PR1461590](#)
- An mcsnoopd crash might be seen if one BD/VLAN is configured as part of EVPN and it has any multicast router interfaces (static/dynamic). [PR1468737](#)
- Traffic might not be forwarded over an ECMP link in an EVPN-VXLAN scenario. [PR1475819](#)
- GRE transit traffic is not forwarded in a VRRP scenario. [PR1477073](#)
- LACP interface might not go to **collecting and distributing** state along with the **force-up** statement. [PR1484523](#)
- FPC might go to NotPrsnt state after upgrading with non-QFX5100-24Q image in a Virtual Chassis/Virtual Chassis Fabric setup. [PR1485612](#)
- CPU port queue gets full due to excessive pause frames being received on interfaces. This causes control packets from the CPU to all ports to be dropped. [PR1487707](#)
- The BGP route-target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- The rpd might crash on QFX10000 due to rpd resolver problem of indirect next hop. [PR1494005](#)
- Firewall filter might not work in certain conditions in a Virtual Chassis setup. [PR1497133](#)
- Traffic drop might be observed after modifying the FBF firewall filter. [PR1499918](#)
- The filter installation fails if the number of filter entries configured exceeds 1000. [PR1514570](#)

### *User Interface and Configuration*

- The version information under the configuration is changed from Junos OS Release 19.1. [PR1457602](#)

### **Resolved Issues: Release 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 sxr 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: Release 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 QFX5K 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 inner 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 OVSD. [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, 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

<a href="#">What's New   263</a>
<a href="#">What's Changed   272</a>
<a href="#">Known Limitations   278</a>
<a href="#">Open Issues   281</a>
<a href="#">Documentation Updates   306</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   307</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 documentation for the QFX Series.

SEE ALSO

<a href="#">What's New   263</a>
<a href="#">What's Changed   272</a>
<a href="#">Known Limitations   278</a>
<a href="#">Open Issues   281</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Upgrading Software on QFX Series Switches | 307
- Installing the Software on QFX10002-60C Switches | 310
- Installing the Software on QFX10002 Switches | 310
- 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 | 311
- Installing the Software on QFX10008 and QFX10016 Switches | 313
- Performing a Unified ISSU | 317
- Preparing the Switch for Software Installation | 318
- Upgrading the Software Using Unified ISSU | 318
- Upgrade and Downgrade Support Policy for Junos OS Releases | 320

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-R3.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-x86-64-19.3R3.9.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-x86-64-19.3R3.9.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-x86-64-19.3R3.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-x86-64-19.3R3.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-x86-64-19.3R3.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-x86-64-19.3R3.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)
Routing 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 318](#)
- [Upgrading the Software Using Unified ISSU on page 318](#)

## 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-x86-64-19.3R3.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 | 263](#)

[What's Changed | 272](#)

[Known Limitations | 278](#)

[Open Issues | 281](#)

[Resolved Issues | 286](#)

[Documentation Updates | 306](#)

## Junos OS Release Notes for SRX Series

#### IN THIS SECTION

- [What's New | 322](#)
- [What's Changed | 330](#)
- [Known Limitations | 335](#)
- [Open Issues | 338](#)
- [Resolved Issues | 340](#)
- [Documentation Updates | 357](#)
- [Migration, Upgrade, and Downgrade Instructions | 357](#)

These release notes accompany Junos OS Release 19.3R3 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- [What's New in Release 19.3R3 | 322](#)
- [What's New in Release 19.3R2 | 322](#)
- [What's New in Release 19.3R1 | 322](#)

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

### What's New in Release 19.3R3

There are no new features in Junos OS Release 19.3R3 for the SRX Series devices.

### What's New in Release 19.3R2

There are no new features in Junos OS Release 19.3R2 for the SRX Series devices.

### What's New in Release 19.3R1

#### *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 | 330](#)

[Known Limitations | 335](#)

[Open Issues | 338](#)

[Resolved Issues | 340](#)

[Documentation Updates | 357](#)

[Migration, Upgrade, and Downgrade Instructions | 357](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 19.3R3 | 331](#)
- [What's Changed in Release 19.3R2-S6 | 332](#)
- [What's Changed in Release 19.3R2 | 333](#)
- [What's Changed in Release 19.3R1 | 333](#)

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

## What's Changed in Release 19.3R3

### *Application Layer Gateways (ALGs)*

- **Disable the do not fragment flag from packet IP header (SRX Series and vSRX)**—Starting in Junos OS Release 19.3R3, we've introduced the **clear-dont-frag-bit** statement at the **[edit security alg alg-manager]** hierarchy level to disable the do not fragment flag from the packet IP header. This configuration allows the packet to be split after NAT is performed.

In Junos OS releases earlier than Release 19.3R3, when the ALG performs payload-NAT, sometimes the size of the packet becomes bigger than the maximum transmission unit (MTU) of the outgoing interface. If the packet IP header has the do not fragment flag, this packet cannot be sent out.

[See [alg-manager](#).]

### *Authentication and Access Control*

- **Enabling and disabling SSH login password or challenge-response authentication (SRX Series)**—Starting in Junos OS Release 19.3R3, you can disable either the SSH login password or the challenge-response authentication at the **[edit system services ssh]** hierarchy level.

In Junos OS releases earlier than Release 19.3R3, you can enable and disable both SSH login password and the challenge-response authentication simultaneously at the **[edit system services ssh]** hierarchy level.

[See [Configuring SSH Service for Remote Access to the Router or Switch](#).]

### *Flow-based and Packet-based Processing*

- **Self-generated IKE packets chooses outgoing interface matching source IP Address (SRX Series)** — A self-generated Internet Key Exchange (IKE) packet always select the ECMP outgoing interface that matches source IP address. Note that filter-based forwarding for self-generated traffic with rerouting is not supported.

### *J-Web*

- **Deactivated policy rules are not visible in the J-Web UI (SRX Series)**—Starting in Junos OS 19.3R3 Release, J-Web does not support disabling or enabling the security firewall or global policy rules. The policy rules that are deactivated through the CLI are also not visible in the J-Web UI. As a workaround, use the CLI to disable or enable the policy rules on the device.

### *Juniper Extension Toolkit (JET)*

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

See [traceoptions \(Services\)](#).

### *Juniper Sky ATP*

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 19.3R3, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Sky Advanced Threat Prevention (ATP).

#### *Network Management and Monitoring*

- **Packet Forwarding Engine data display command (SRX Series)**—Starting in Junos OS Release 19.3R3, you can now view the Packet Forwarding Engine data by using the **usp flow config** and **usp flow stats** options for the **show pfe data** command.

### **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

### *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](#).]

## What's Changed in Release 19.3R1

### *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

[What's New | 322](#)

[Known Limitations | 335](#)

[Open Issues | 338](#)

[Resolved Issues | 340](#)

[Documentation Updates | 357](#)

[Migration, Upgrade, and Downgrade Instructions | 357](#)

## Known Limitations

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.

## Authentication and Access Control

- **Enhanced user firewall support**—In Junos OS Release 19.3R3, for SRX300 devices with eUSB (SRX300, SRX320, SRX340, and SRX345), the SRX Series user firewall (UserFW) module tries to synchronize user entries from the domain controller or Juniper Identity Management Service (JIMS) after booting up. If the historical login events expired on the domain controller, then the SRX Series UserFW module is unable to retrieve those user entries after the UserFW module boots up.

[See [User Authentication Entries in the ClearPass Authentication Table](#).]

## 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 on the 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](#)
- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is automatically 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](#)



## Logical Systems and Tenant Systems

- In case of logical systems, secure wire cannot work with the user firewall AD integrated solution, 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](#)

## Platform and Infrastructure

- When an NTP server is newly added to the Junos configuration using a domain name, a DNS server IP address needs to have been already configured and committed in a previous commit. Otherwise the commit will fail due to the NTP server domain name failing to be resolved to an IP address. As a workaround, use an IP address for the NTP server configuration. [PR1411396](#)

## VPNs

- In the output of the show security ipsec inactive-tunnels command, Tunnel Down Reason is not displayed as this functionality is not supported in Junos OS Release 18.2R2 and later. [PR1383329](#)
- In the HA design for SRX Series devices, the anti-replay window is synchronized 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](#)
- The IPsec VPN tunnel does not come up, and the gateway lookup failed error is seen, after the configurations are changed from IPv4 to IPv6 tunnels in the script. [PR1431265](#)
- On the SRX5000 line of devices with SPC3 cards, sometimes IKE SA is not seen on the device when the st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)
- Per-tunnel debugging configuration is not synchronized to the backup node. It needs to be configured again after RGO failover. [PR1450393](#)

## SEE ALSO

[What's New | 322](#)

[What's Changed | 330](#)

[Open Issues | 338](#)

[Resolved Issues | 340](#)

[Documentation Updates | 357](#)

[Migration, Upgrade, and Downgrade Instructions | 357](#)

## Open Issues

### IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 338](#)
- [Intrusion Detection and Prevention \(IDP\) | 338](#)
- [J-Web | 339](#)
- [Routing Policy and Firewall Filters | 339](#)
- [VPNs | 339](#)

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

- On an SRX4600 device, when the next hop is set to the st0 interface, the output of the show route forwarding-table command displays the next-hop IP address twice. [PR1290725](#)
- 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](#)
- The dynamic applications that relied on SSL proxy for matching have been deprecated. These dynamic applications are HTTPS, SMTPS, POP3S, and IMAPS. Use SSL, SMTP, POP3, and IMAP in all policies moving forward as they will now match for both encrypted and decrypted streams. [PR1444767](#)
- TCP session cannot time out properly upon receiving the TCP RESET packet, and the session timeout does not change to 2 seconds. [PR1467654](#)

### Intrusion Detection and Prevention (IDP)

- On the SRX Series devices, the commit or show command for IDP might not work if you keep running SNMP queries when large-scale IDP is used. [PR1444043](#)
- When intelligent inspection status changes, syslog is not generated on SRX300 and SRX500 line of devices. [PR1448365](#)

## J-Web

- Due to set chassis auto-image-upgrade in the 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](#)
- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is automatically 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](#)

## Routing Policy and Firewall Filters

- The SSL reverse proxy feature must be used instead of the SSL inspection feature. SSL inspection on IDP level is being deprecated in favor of SSL reverse proxy. [PR1450900](#)
- If a huge number of policies are configured on SRX Series devices and some policies are changed, the traffic that matches the changed policies might be dropped. [PR1454907](#)

## VPNs

- On SRX Series devices, if multiple traffic selectors are configured for a peer with Internet Key Exchange version 2 (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. A new negotiation of these traffic selectors is triggered through other mechanisms—for example, by traffic or by a peer. [PR1287168](#)
- When multiple traffic selectors are configured on a particular VPN, the iked process checks for a maximum of one DPD probe that is sent to the peer for the configured DPD interval. The DPD probe is sent to the peer if traffic flows over even one of the tunnels for the given VPN object. [PR1366585](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, a new behavior has been introduced that differs from the behavior on the older SPC2 card. The SRX Series device with AutoVPN configuration can now accept multiple IPsec tunnels from a peer device (with the same source IP address and port number) using different IKE IDs. [PR1407356](#)
- On 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 SPC3 cards, sometimes IKE SA is not seen on the device when the st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)
- Per-tunnel debugging configuration is not synchronized to the backup node. It needs to be configured again after RGO failover. [PR1450393](#)
- In an IPsec VPN scenario on SRX5400, SRX5600, and SRX5800 platforms, the iked process treats retransmission of IKE\_INIT request packets as new connections when the SRX Series device acts as a responder of IKE negotiation. This causes IKE tunnel negotiation to fail, and IPsec VPN traffic might be impacted. [PR1460907](#)
- The SRX5000 line of devices with SPC3 does not support simultaneous IKE negotiation in Junos OS Releases 19.2, 19.3, 19.4, and 20.1. [PR1497297](#)

#### SEE ALSO

[What's New | 322](#)

[What's Changed | 330](#)

[Known Limitations | 335](#)

[Resolved Issues | 340](#)

[Documentation Updates | 357](#)

[Migration, Upgrade, and Downgrade Instructions | 357](#)

## Resolved Issues

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: Release 19.3R3

#### IN THIS SECTION

- [Application Layer Gateways \(ALGs\) | 341](#)
- [Authentication and Access Control | 341](#)
- [Chassis Clustering | 341](#)

- Flow-Based and Packet-Based Processing | 341
- Interfaces and Chassis | 344
- Intrusion Detection and Prevention (IDP) | 344
- J-Web | 344
- Layer 2 Ethernet Services | 344
- MPLS | 344
- Network Address Translation (NAT) | 344
- Network Management and Monitoring | 344
- Platform and Infrastructure | 344
- Routing Policy and Firewall Filters | 345
- Routing Protocols | 345
- Unified Threat Management (UTM) | 345
- VLAN Infrastructure | 346
- VPNs | 346

### ***Application Layer Gateways (ALGs)***

- On the SRX5000 line of devices, the H323 call with NAT64 could not be established. [PR1462984](#)
- The flowd or srpxfe process might stop when an ALG creates a gate with an incorrect protocol value. [PR1474942](#)
- SIP messages that need to be fragmented might be dropped by the SIP ALG. [PR1475031](#)
- FTPS traffic might be dropped on SRX Series or MX Series devices if FTP ALG is used. [PR1483834](#)

### ***Authentication and Access Control***

- SRX Series: Unified Access Control (UAC) bypass vulnerability (CVE-2020-1637). [PR1475435](#)

### ***Chassis Clustering***

- IP monitoring might fail on the secondary node. [PR1468441](#)
- An unhealthy node might become primary in SRX4600 devices in a chassis cluster scenario. [PR1474233](#)

### ***Flow-Based and Packet-Based Processing***

- The show security group-vpn server statistics |display xml is not in expected format. [PR1349959](#)
- The show security pki local-certificate logical-system all command is not showing any output. [PR1414628](#)
- The trusted-ca and root-ca names or IDs should not be the same within an SSL proxy configuration. [PR1420859](#)

- The control logical interface is not created by default for LLDP. [PR1436327](#)
- The SPC card might stop on the SRX5000 line of devices. [PR1439744](#)
- With NCP remote access solution, in a PathFinder case (for example, where IPsec traffic has to be encapsulated as TCP packets), TCP encapsulation for transit traffic is failing. [PR1442145](#)
- In the BERT test for E1 interface, bits count is not within the range. [PR1445041](#)
- Introduction of default inspection limits for application identification to optimize CPU usage and improve resistance to evasive applications. [PR1454180](#)
- The SRX Series devices stop and generate several core files. [PR1455169](#)
- The srpxfe or flowd process might stop if the sampling configuration is changed. [PR1462610](#)
- The tunnel packets might be dropped because the gr0.0 or st0.0 interface is wrongly calculated after a GRE or VPN route change. [PR1462825](#)
- The flowd or srpxfe process might crash if the tunnel information-related show command is issued during the tunnel deleting process. [PR1469123](#)
- On the SRX300 line of devices, you might encounter Authentication-Table loading slowly while using user identification. [PR1462922](#)
- A core file might be generated when you perform an ISSU on SRX Series devices. [PR1463159](#)
- The PKI daemon keeps leaking memory on SRX Series devices. [PR1465614](#)
- The jbuf process usage might increase up to 99 percent after Junos OS upgrade. [PR1467351](#)
- The rpd process might stop after several changes to the flowspec routes. [PR1467838](#)
- The Packet Forwarding Engine might generate core files because SSL proxy is enabled on NFX Series and SRX Series devices. [PR1467856](#)
- Server unreachable is detected; ensure that port 443 is reachable. [PR1468114](#)
- Tail drop on all ports is observed when any switch-side egress port gets congested. [PR1468430](#)
- FTP data connection might be dropped if SRX Series devices send the FTP connection traffic through the dl interface. [PR1468570](#)
- RPM test probe fails to show that round-trip time has been exceeded. [PR1471606](#)
- 19.3R2: pkid crash at bn\_i2c (pval=0x1d, cont=0x0, putype=0xffffc8e8, it=0xc8c848b8 < BIGNUM\_it>) at ../../../../src/crypto/openssl/crypto/asn1/x\_bignum.c:127. [PR1471878](#)
- Look-up failure for expected e-mail address in DUT. [PR1472748](#)
- Packet drop might be observed on the SRX300 line of devices when adding or removing an interface from MACsec. [PR1474674](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- ECMP load balancing does not happen when the RG1 node 0 is secondary. [PR1475853](#)

- The flowd or srpxfe process might stop when deleting user firewall local authentication table entry. [PR1477627](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- Recent changes to JDPI's classification mechanism caused a considerable performance regression (more than 30 percent). [PR1479684](#)
- The flowd or srpxfe process might stop when advanced anti-malware service is used. [PR1480005](#)
- On Web proxy, memory leak in association hash table and DNS hash table. [PR1480760](#)
- The jsqsyncd process synchronizes its databases every second even there is no change. [PR1482428](#)
- IMAP curl sessions get stuck in the active state if AAMW IMAP block mode is configured. [PR1484692](#)
- The show chassis temperature-thresholds command displays extensive FPC 0 output. [PR1485224](#)
- The flowd process might stop and impact services if J-Flow version 9 is configured. [PR1486528](#)
- Commit does not work after the installation through boot loader. [PR1487831](#)
- If a cluster ID of 16 or multiples of 16 is used, the chassis cluster might not come up. [PR1487951](#)
- CPU board inlet increases after OS upgrade from Junos OS Release 15.1X49 to Junos OS Release 18.x. [PR1488203](#)
- All interfaces remain in the down state after the SRX300 line of devices power up or reboot. [PR1488348](#)
- There is a risk of service interruption on all SRX Series devices with a dual-stacked CA server. [PR1489249](#)
- Continuous drops seen in control traffic, with high data queues in one SPC2 PIC. [PR1490216](#)
- Phone client stop seen while doing SRX345 device ZTP with CSO. [PR1496650](#)
- Outbound SSH connection flap or memory leak issue might be observed when pushing configuration to ephemeral database with high rate. [PR1497575](#)
- Traffic interruption happens due to MAC address duplication between two Junos OS devices. [PR1497956](#)
- Don't use capital characters for source-identity when using the show security match-policies command. [PR1499090](#)
- J-Flow version 9 does not display the correct outgoing interface for APBR traffic. [PR1502432](#)
- Fabric interface might be monitored down after chassis cluster reboot. [PR1503075](#)
- SOF asymmetric scenario not working phase-1 solution. [PR1507865](#)
- VRRP does not work on the reth interface with a VLAN ID greater than 1023. [PR1515046](#)
- A logic issue was corrected in SSL proxy that could lead to an srpxfe or flowd core file under load. [PR1516903](#)
- The PPPoE session does not come up after zeroization on SRX Series devices. [PR1518709](#)

### ***Interfaces and Chassis***

- Static route through the dI0.0 interface is not active. [PR1465199](#)
- MAC limiting on Layer 3 routing interfaces does not work. [PR1465366](#)

### ***Intrusion Detection and Prevention (IDP)***

- On SRX Series devices, updating the IDP security package offline might fail. [PR1466283](#)
- The IDP attack detection might not work in a specific situation. [PR1497340](#)
- IDP's custom-attack time-binding interval command was mistakenly hidden within the CLI. [PR1506765](#)

### ***J-Web***

- The Policy Rules grid will not populate the deactivated security firewall or global policy rules on the Rules page. [PR1460161](#)
- IPv6 address objects containing a,b,c,d,e cannot be configured through J-Web. [PR1464110](#)
- J-Web security resources dashboard widget was not being populated correctly. [PR1464769](#)
- The J-Web users might not be able to configure PPPoE using the PPPoE wizard. [PR1502657](#)

### ***Layer 2 Ethernet Services***

- Member links state might be asynchronized on a connection between PE and CE devices in EVPN active/active mode. [PR1463791](#)

### ***MPLS***

- BGP session might keep flapping between two directly connected BGP peers because of the incorrect use of the TCP-MSS. [PR1493431](#)

### ***Network Address Translation (NAT)***

- The flowd or srpxfe process might stop when traffic is processed by both ALGs and NAT. [PR1471932](#)
- Issuing the show security nat source paired-address command might return an error. [PR1479824](#)

### ***Network Management and Monitoring***

- The flowd or srpxfe process might stop immediately after committing the J-Flow version 9 configuration or after upgrading to affected releases. [PR1471524](#)
- SNMP trap coldStart agent-address becomes 0.0.0.0. [PR1473288](#)

### ***Platform and Infrastructure***

- On certain MPC line cards, cm errors need to be reclassified. [PR1449427](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)



- On SRX1500 and the SRX4000 line of devices, physically disconnecting the cable from the fxp0 interface causes hardware monitor failure and redundancy group failover, when the device is the primary node in a chassis cluster. [PR1467376](#)
- On SRX Series devices, Packet Forwarding Engine memory might be used up if the security intelligence feature is configured. [PR1472926](#)
- Support LLDP protocol on reth interface. [PR1473456](#)
- Certificate error while configuration validation during Junos OS upgrade. [PR1474225](#)
- The RGx might fail over after RG0 failover in a rare case. [PR1479255](#)
- Packets get dropped when the next hop is IRB over It interface. [PR1494594](#)
- On the SRX1500 device, the factory-default configuration for ge-0/0/0 and ge-0/0/15 should be set with family inet DHCP. [PR1503636](#)

### ***Routing Policy and Firewall Filters***

- Security policies cannot synchronize between Routing Engine and Packet Forwarding Engine on SRX Series devices. [PR1453852](#)
- Traffic log shows an incorrect custom-application name when the alg ignore option is used in application configuration. [PR1457029](#)
- Some domains are not resolved by the SRX Series devices when using DNS address book. [PR1471408](#)
- The count option in the security policy does not take effect even if the policy count is enabled. [PR1471621](#)
- Support for dynamic tunnels on SRX Series devices was mistakenly removed. [PR1476530](#)
- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. [PR1492436](#)
- The srxpfe or flowd process might crash due to memory corruption within JDPI. [PR1500938](#)
- Traffic fails to hit the policies with matching source-end-user-profiles. [PR1505002](#)

### ***Routing Protocols***

- SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. [PR1454177](#)
- The rpd might stop when both instance-import and instance-export policies contain the as-path-prepend action. [PR1471968](#)
- The BGP route-target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

### ***Unified Threat Management (UTM)***

- Increase the scale number of UTM profile or policy for the SRX1500 device, and the SRX4000 and SRX5000 lines of devices. [PR1455321](#)

- The utmd process might pause after deactivating a UTM configuration with predefined category upgrading used. [PR1478825](#)
- UTM websense redirect supports IPv6 messages. [PR1481290](#)

### **VLAN Infrastructure**

- ISSU failed from Junos OS Release 18.4R2.7 to Junos OS Release 19.4, with secondary node PICs in present state after upgrading to Junos OS Release 19.4. [PR1468609](#)

### **VPNs**

- After RG1 failover, IKE phase 1 SA is getting cleared. [PR1352457](#)
- IPsec VPN missing half of the IKE SA and IPsec SA showing incorrect port number when scaling to 1000 IKEv1 AutoVPN tunnels. [PR1399147](#)
- The established tunnels might remain unchanged when an IKE gateway is changed from AutoVPN to Site-to-Site VPN. [PR1413619](#)
- The show security ipsec statistics command output displays buffer overflow and wraps around 4,---,---,--- count. [PR1424558](#)
- IKE SA does not get cleared and is showing a very long lifetime. [PR1439338](#)
- Some IPsec tunnels flap after RGs failover on the SRX5000 line of devices. [PR1450217](#)
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. [PR1455951](#)
- IPsec tunnels might lose connectivity on SRX Series devices after chassis cluster failover when using AutoVPN point-to-multipoint mode. [PR1469172](#)
- The kmd process might crash continually after chassis cluster failover in an IPsec ADVPN scenario. [PR1479738](#)
- On SRX Series devices with SPC3, when overlapping traffic selectors are configured, multiple IPsec SAs get negotiated with the peer device. [PR1482446](#)

## **Resolved Issues: Release 19.3R2**

### **IN THIS SECTION**

- [Application Layer Gateways \(ALGs\) | 347](#)
- [Authentication and Access Control | 347](#)
- [Application Security | 347](#)
- [Chassis Clustering | 347](#)
- [Flow-Based and Packet-Based Processing | 347](#)
- [Interfaces and Chassis | 348](#)

- J-Web | 348
- Network Address Translation (NAT) | 348
- Network Management and Monitoring | 348
- Platform and Infrastructure | 348
- Routing Policy and Firewall Filters | 349
- Services Applications | 349
- Unified Threat Management (UTM) | 349
- VPNs | 349

### ***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: Release 19.3R1**

#### **IN THIS SECTION**

- [Application Layer Gateways \(ALGs\) | 350](#)
- [Application Security | 350](#)
- [Authentication and Access Control | 350](#)
- [Chassis Clustering | 351](#)

- Class of Service (CoS) | 351
- Flow-Based and Packet-Based Processing | 351
- Infrastructure | 353
- Interfaces and Routing | 353
- Installation and Upgrade | 353
- Interfaces and Chassis | 353
- Intrusion Detection and Prevention (IDP) | 353
- J-Web | 353
- Network Address Translation (NAT) | 354
- Network Management and Monitoring | 354
- Platform and Infrastructure | 354
- Routing Policy and Firewall Filters | 355
- Unified Threat Management (UTM) | 355
- User Interface and Configuration | 356
- VPNs | 356

### ***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 srpxfe 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 jsqsyncd process on SRX Series chassis clusters. [PR1424884](#)
- The flowd or srpxfe process might stop when SCCP or MGCP ALG works on SRX Series chassis clusters. [PR1426722](#)
- RGO 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 useful 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 committing 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

<a href="#">What's New   322</a>
<a href="#">What's Changed   330</a>
<a href="#">Known Limitations   335</a>
<a href="#">Open Issues   338</a>
<a href="#">Documentation Updates   357</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   357</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 19.3R3 documentation for the SRX Series.

#### SEE ALSO

<a href="#">What's New   322</a>
<a href="#">What's Changed   330</a>
<a href="#">Known Limitations   335</a>
<a href="#">Open Issues   338</a>
<a href="#">Resolved Issues   340</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   357</a>

## Migration, Upgrade, and Downgrade Instructions

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

## Upgrade and Downgrade Support Policy for Junos OS Releases 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

<a href="#">What's New</a>	<a href="#"> </a>	<a href="#">322</a>
<a href="#">What's Changed</a>	<a href="#"> </a>	<a href="#">330</a>
<a href="#">Known Limitations</a>	<a href="#"> </a>	<a href="#">335</a>
<a href="#">Open Issues</a>	<a href="#"> </a>	<a href="#">338</a>
<a href="#">Resolved Issues</a>	<a href="#"> </a>	<a href="#">340</a>
<a href="#">Documentation Updates</a>	<a href="#"> </a>	<a href="#">357</a>

# Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the [High Availability User Guide](#).

For additional information about using ISSU on security devices, see the [Chassis Cluster User Guide for SRX Series Devices](#).

For information about ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\)](#) Web application.

## Licensing

Starting in 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that have been developed at Juniper Networks over the past several years.

The major components of the framework are:

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

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

## Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. <https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. [prsearch.juniper.net](https://prsearch.juniper.net).
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](https://apps.juniper.net/hct/home)


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

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. [apps.juniper.net/compliance/](https://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](mailto: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](mailto: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 4, Junos OS Release 19.3R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 3, Junos OS Release 19.3R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 October 2020—Revision 2, Junos OS Release 19.3R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 1, Junos OS Release 19.3R3— 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.