

Release Notes: Junos[®] OS Release 18.1R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

26 August 2021

Contents	Introduction 12
	Junos OS Release Notes for ACX Series 12
	New and Changed Features 13
	Release 18.1R3 New and Changed Features 13
	Release 18.1R2 New and Changed Features 13
	Release 18.1R1 New and Changed Features 13
	Changes in Behavior and Syntax 14
	Interfaces and Chassis 14
	Known Behavior 15
	Known Issues 15
	Resolved Issues 16
	Resolved Issues: 18.1R3 16
	Resolved Issues: 18.1R2 16
	Resolved Issues: 18.1R1 17
	Documentation Updates 18
	New Simplified Documentation Architecture 18
	Migration, Upgrade, and Downgrade Instructions 19
	Upgrade and Downgrade Support Policy for Junos OS Releases 19

Product Compatibility | 20

Hardware Compatibility | 20

Junos OS Release Notes for EX Series Switches | 21

New and Changed Features | 21

Release 18.1R3-S3 New and Changed Features | 22

Release 18.1R3 New and Changed Features | 23

Release 18.1R2 New and Changed Features | 23

Release 18.1R1 New and Changed Features | 25

Changes in Behavior and Syntax | 30

Release 18.1R3-S10 Changes in Behavior and Syntax | 31

Release 18.1R3-S7 Changes in Behavior and Syntax | 31

Release 18.1R3 Changes in Behavior and Syntax | 32

Release 18.1R2 Changes in Behavior and Syntax | 34

Release 18.1R1 Changes in Behavior and Syntax | 34

Known Behavior | 35

Infrastructure | 36

Interfaces and Chassis | 36

Platform and Infrastructure | 36

Virtual Chassis | 37

Known Issues | 37

General Routing | 38

Infrastructure | 39

Layer 2 Features | 39

Platform and Infrastructure | 39

Resolved Issues | 40

Resolved Issues: 18.1R3 | 40

Resolved Issues: 18.1R2 | 42

Resolved Issues: 18.1R1 | 44

Documentation Updates | 48

New Simplified Documentation Architecture | 48

Migration, Upgrade, and Downgrade Instructions | 49

Upgrade and Downgrade Support Policy for Junos OS Releases | 49

Product Compatibility | 50

Hardware Compatibility | 50

Junos OS Release Notes for Junos Fusion Data Center | 51

New and Changed Features | 52

Release 18.1R3 New and Changed Features | 53

Release 18.1R2-S2 New and Changed Features | 53

Release 18.1R2 New and Changed Features | 57

Release 18.1R1 New and Changed Features | 57

Changes in Behavior and Syntax | 58

Known Behavior | 58

Junos Fusion Data Center | 59

Known Issues | 61

Junos Fusion Data Center | 61

Resolved Issues | 62

Resolved Issues: 18.1R3 | 62

Resolved Issues: 18.1R2 | 62

Resolved Issues: 18.1R1 | 63

Documentation Updates | 63

New Simplified Documentation Architecture | 63

Migration, Upgrade, and Downgrade Instructions | 64

Basic Procedure for Upgrading an Aggregation Device | 64

Preparing the Switch for Satellite Device Conversion | 66

Configuring Satellite Device Upgrade Groups | 68

Converting a Satellite Device to a Standalone Device | 69

Upgrade and Downgrade Support Policy for Junos OS Releases | 70

Downgrading from Junos OS Release 18.1 | 70

Product Compatibility | 71

Hardware and Software Compatibility | 71

Hardware Compatibility Tool | 71

Junos OS Release Notes for Junos Fusion Enterprise | 72

New and Changed Features | 72

Release 18.1R3 New and Changed Features | 73

Release 18.1R2 New and Changed Features | 73

Release 18.1R1 New and Changed Features | 73

Changes in Behavior and Syntax | 74

Known Behavior | 74

Known Issues | 75**Junos Fusion Enterprise | 75****Resolved Issues | 76****Resolved Issues: 18.1R3 | 76****Resolved Issues: 18.1R2 | 76****Resolved Issues: 18.1R1 | 77****Documentation Updates | 77****New Simplified Documentation Architecture | 77****Migration, Upgrade, and Downgrade Instructions | 78****Basic Procedure for Upgrading Junos OS on an Aggregation Device | 79****Upgrading an Aggregation Device with Redundant Routing Engines | 81****Preparing the Switch for Satellite Device Conversion | 81****Converting a Satellite Device to a Standalone Switch | 83****Upgrade and Downgrade Support Policy for Junos OS Releases | 83****Downgrading Junos OS | 83****Product Compatibility | 84****Hardware and Software Compatibility | 84****Hardware Compatibility Tool | 84****Junos OS Release Notes for Junos Fusion Provider Edge | 86****New and Changed Features | 86****Release 18.1R3 New and Changed Features | 88****Release 18.1R2 New and Changed Features | 88****Release 18.1R1 New and Changed Features | 88****Changes in Behavior and Syntax | 89****Known Behavior | 90****Junos Fusion | 90****Known Issues | 91****Junos Fusion | 91****Resolved Issues | 92****Resolved Issues: 18.1R3 | 92****Resolved Issues: 18.1R2 | 92****Resolved Issues: 18.1R1 | 93****Documentation Updates | 93****New Simplified Documentation Architecture | 94**

Migration, Upgrade, and Downgrade Instructions | 94

Basic Procedure for Upgrading an Aggregation Device | 95

Upgrading an Aggregation Device with Redundant Routing Engines | 97

Preparing the Switch for Satellite Device Conversion | 98

Converting a Satellite Device to a Standalone Device | 99

Upgrading an Aggregation Device | 99

Upgrade and Downgrade Support Policy for Junos OS Releases | 99

Downgrading from Junos OS Release 18.1 | 100

Product Compatibility | 101

Hardware Compatibility | 101

Junos OS Release Notes for MX Series 5G Universal Routing Platforms | 102

New and Changed Features | 102

Release 18.1R3-S3 New and Changed Features | 103

Release 18.1R3 New and Changed Features | 104

Release 18.1R2 New and Changed Features | 104

Release 18.1R1 New and Changed Features | 104

Changes in Behavior and Syntax | 123

Release 18.1R3-S12 Changes in Behavior and Syntax | 123

Release 18.1R3-S7 Changes in Behavior and Syntax | 123

Release 18.1R3-S5 Changes in Behavior and Syntax | 124

Release 18.1R3 Changes in Behavior and Syntax | 124

Release 18.1R2 Changes in Behavior and Syntax | 127

Release 18.1R1 Changes in Behavior and Syntax | 128

Known Behavior | 132

EVPN | 133

General Routing | 133

Interfaces and Chassis | 134

MPLS | 135

Platform and Infrastructure | 135

Routing Protocols | 135

Services Applications | 135

Software Installation and Upgrade | 136

Subscriber Management and Services | 137

Known Issues | 137

- Class of Service(CoS) | 138**
- EVPN | 138**
- Forwarding and Sampling | 139**
- General Routing | 139**
- Infrastructure | 144**
- Interfaces and Chassis | 144**
- Layer 2 Ethernet Services | 145**
- Layer 2 Features | 146**
- MPLS | 146**
- Platform and Infrastructure | 147**
- Routing Protocols | 150**
- Services Applications | 151**
- Software Installation and Upgrade | 151**
- Subscriber Management and Services | 152**
- VPNs | 152**

Resolved Issues | 153

- Resolved Issues: 18.1R3 | 153**
- Resolved Issues: 18.1R2 | 164**
- Resolved Issues: 18.1R1 | 171**

Documentation Updates | 189

- New Simplified Documentation Architecture | 189**
- Subscriber Management Access Network Guide | 189**
- Subscriber Management Provisioning Guide | 190**

Migration, Upgrade, and Downgrade Instructions | 190

- Basic Procedure for Upgrading to Release 18.1 | 191**
- Procedure to Upgrade to FreeBSD 11.x based Junos OS | 191**
- Procedure to Upgrade to FreeBSD 6.x based Junos OS | 194**
- Upgrade and Downgrade Support Policy for Junos OS Releases | 196**
- Upgrading a Router with Redundant Routing Engines | 196**
- Downgrading from Release 18.1 | 196**

Product Compatibility | 197

- Hardware Compatibility | 197**

Junos OS Release Notes for NFX Series | 198

New and Changed Features | 199

Release 18.1R3 New and Changed Features | 199

Release 18.1R2 New and Changed Features | 199

Release 18.1R1 New and Changed Features | 199

Changes in Behavior and Syntax | 203

CLI | 204

Known Behavior | 204

Known Behavior: 18.1R3 | 205

Known Issues | 206

Known Issues: 18.1R3 | 206

Resolved Issues | 209

Resolved Issues: 18.1R3 | 209

Resolved Issues: 18.1R2 | 209

Resolved Issues: 18.1R1 | 210

Documentation Updates | 211

New Simplified Documentation Architecture | 211

Migration, Upgrade, and Downgrade Instructions | 212

Upgrade and Downgrade Support Policy for Junos OS Releases | 212

Basic Procedure for Upgrading to Release 18.1 | 212

Product Compatibility | 215

Hardware Compatibility | 215

Software Version Compatibility | 216

Junos OS Release Notes for PTX Series Packet Transport Routers | 218

New and Changed Features | 218

Release 18.1R3 New and Changed Features | 219

Release 18.1R2 New and Changed Features | 219

Release 18.1R1 New and Changed Features | 219

Changes in Behavior and Syntax | 227

Interfaces and Chassis | 227

Management | 229

Network Management and Monitoring | 229

Network Operations and Troubleshooting Automation | 229

Subscriber Management and Services | 230

Known Behavior | 231**General Routing | 231****Interfaces and Chassis | 232****Known Issues | 232****General Routing | 233****Infrastructure | 235****Interfaces and Chassis | 235****MPLS | 235****Resolved Issues | 235****Resolved Issues: 18.1R3 | 236****Resolved Issues: 18.1R2 | 237****Resolved Issues: 18.1R1 | 238****Documentation Updates | 241****New Simplified Documentation Architecture | 241****Migration, Upgrade, and Downgrade Instructions | 242****Upgrade and Downgrade Support Policy for Junos OS Releases | 242****Upgrading a Router with Redundant Routing Engines | 243****Basic Procedure for Upgrading to Junos OS Release 18.1 | 243****Product Compatibility | 247****Hardware Compatibility | 247****Junos OS Release Notes for the QFX Series | 248****New and Changed Features | 248****Release 18.1R3-S3 New and Changed Features | 249****Release 18.1R3 New and Changed Features | 252****Release 18.1R2 New and Changed Features | 254****Release 18.1R1 New and Changed Features | 255****Changes in Behavior and Syntax | 273****| 273****Known Behavior | 276****EVPN | 276****Interfaces and Chassis | 276****Layer 2 Features | 277****Multicast | 277****Platform and Infrastructure | 277**

Routing Protocols	280
Storage and Fibre Channel	281
Virtual Chassis	281
Services Applications	281
Known Issues	282
EVPN	282
Interfaces and Chassis	283
Layer 2 Features	283
MPLS	283
Platform and Infrastructure	284
Routing Protocols	286
Resolved Issues	287
Resolved Issues: 18.1R3	288
Resolved Issues: 18.1R2	291
Resolved Issues: 18.1R1	294
Documentation Updates	299
New Simplified Documentation Architecture	299
Migration, Upgrade, and Downgrade Instructions	300
Upgrading Software on QFX Series Switches	301
Installing the Software on QFX10002-60C Switches	303
Installing the Software on QFX10002 Switches	303
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	304
Installing the Software on QFX10008 and QFX10016 Switches	306
Performing a Unified ISSU	310
Preparing the Switch for Software Installation	311
Upgrading the Software Using Unified ISSU	311
Upgrade and Downgrade Support Policy for Junos OS Releases	313
Product Compatibility	314
Hardware Compatibility	314
Junos OS Release Notes for SRX Series	315
New and Changed Features	316
Release 18.1R3 New and Changed Features	316
Release 18.1R2 New and Changed Features	316

Release 18.1R1 New and Changed Features	317
Changes in Behavior and Syntax	323
Chassis Cluster	323
Juniper Sky ATP	323
VPN	323
Known Behavior	324
Chassis Clustering	324
J-Web	324
Platform and Infrastructure	325
Software Installation and Upgrade	325
User Interface and Configuration	325
VPNs	325
Known Issues	326
Application Layer Gateways (ALGs)	326
Chassis Clustering	327
Class of Service (CoS)	327
Flow-based and Packet-based Processing	327
Interfaces and Routing	328
Intrusion Detection and Prevention (IDP)	328
Platform and Infrastructure	328
Routing Policy and Firewall Filters	329
Routing Protocols	329
Software Installation and Upgrade	329
VPNs	329
Software Installation and Upgrade	330
VPNs	330
Resolved Issues	330
Resolved Issues: 18.1R3	331
Resolved Issues: 18.1R2	333
Resolved Issues: 18.1R1	335
Documentation Updates	341
New Simplified Documentation Architecture	342

Migration, Upgrade, and Downgrade Instructions | 343

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 343

Product Compatibility | 344

Hardware Compatibility | 344

Upgrading Using ISSU | 345

Compliance Advisor | 345

Finding More Information | 345

Documentation Feedback | 346

Requesting Technical Support | 347

Self-Help Online Tools and Resources | 347

Creating a Service Request with JTAC | 348

Revision History | 348

Introduction

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

These release notes accompany Junos OS Release 18.1R3 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

- New and Changed Features | 13
- Changes in Behavior and Syntax | 14
- Known Behavior | 15
- Known Issues | 15
- Resolved Issues | 16
- Documentation Updates | 18
- Migration, Upgrade, and Downgrade Instructions | 19
- Product Compatibility | 20

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- [Release 18.1R3 New and Changed Features | 13](#)
- [Release 18.1R2 New and Changed Features | 13](#)
- [Release 18.1R1 New and Changed Features | 13](#)

This section describes the features and enhancements in Junos OS Release 18.1R3 for ACX Series routers.

Release 18.1R3 New and Changed Features

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 18.1R3.

Release 18.1R2 New and Changed Features

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 18.1R2.

Release 18.1R1 New and Changed Features

Management

- **Support for NETCONF over SSH and custom YANG models (ACX Series)**—Starting in Junos OS Release 18.1R1, ACX Series routers support NETCONF OVER SSH and custom YANG models.

Client applications can access the NETCONF server using the SSH protocol and use the standard SSH authentication mechanism. After authentication, the NETCONF server uses the configured Junos OS login usernames and classes to determine whether a client application is authorized to make each request.

You can load custom YANG models on the router to add data models that are not natively supported by Junos OS but can be supported by translation. Doing this enables you to extend the configuration hierarchies and operational commands with data models that are customized for your operations. You can load custom YANG modules by using the **request system yang add** operational command.

[See [Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS.](#)]

SEE ALSO

Changes in Behavior and Syntax	14
Known Behavior	15
Documentation Updates	18
Known Issues	15
Resolved Issues	16
Migration, Upgrade, and Downgrade Instructions	19
Product Compatibility	20

Changes in Behavior and Syntax

IN THIS SECTION

- Interfaces and Chassis | 14

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.1R3 for the ACX Series routers.

Interfaces and Chassis

- **Modified output of show-ptp-clock command (QFX Series switches)**—Starting in Junos OS Release 18.1R1, the output of the **show-ptp-clock** command is modified to display the value of the **GMC Class** field as **248** for a PTP boundary clock when the lock state of the clock is **Acquiring**.

SEE ALSO

New and Changed Features	13
Known Behavior	15
Documentation Updates	18
Known Issues	15
Resolved Issues	16
Migration, Upgrade, and Downgrade Instructions	19
Product Compatibility	20

Known Behavior

There are no known limitations in Junos OS Release 18.1R3 for the ACX Series.

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

SEE ALSO

New and Changed Features 13
Changes in Behavior and Syntax 14
Documentation Updates 18
Known Issues 15
Resolved Issues 16
Migration, Upgrade, and Downgrade Instructions 19
Product Compatibility 20

Known Issues

There are no known issues in hardware and software in Junos OS Release 18.1R3 for the ACX Series.

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

SEE ALSO

New and Changed Features 13
Changes in Behavior and Syntax 14
Known Behavior 15
Documentation Updates 18
Resolved Issues 16
Migration, Upgrade, and Downgrade Instructions 19
Product Compatibility 20

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3 | 16](#)
- [Resolved Issues: 18.1R2 | 16](#)
- [Resolved Issues: 18.1R1 | 17](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

There are no resolved issues in Junos OS 18.1R3 Release for ACX Series routers.

Resolved Issues: 18.1R2

There are no resolved issues in Junos OS 18.1R2 Release for ACX Series routers.

Resolved Issues: 18.1R1

Alarms

- The major alarm about Fan & PSU Airflow direction mismatch was seen by removing management cable [PR1327561](#)

Dynamic Host Configuration Protocol

- ACX5000 line of routers did not forward DHCP-RELAY requests with IRB interface after upgrade. [PR1243687](#)

Firewall Filters

- On ACX Series routers, syslog error was seen on the output/egress firewall filter. [PR1316588](#)

Installation and Upgrade

- fxpc core was observed during ISSU upgrade. [PR1318771](#)

Layer 2 Features

- On ACX5000 line of routers, transit ARP packets were being punted to the RE. [PR1263012](#)

VPN

- On ACX5000 line of routers, memory leak was seen during Layer 3 VPN scaling test when committing Layer 3 VPN configuration. [PR1115686](#)

SEE ALSO

[New and Changed Features | 13](#)

[Changes in Behavior and Syntax | 14](#)

[Known Behavior | 15](#)

[Documentation Updates | 18](#)

[Known Issues | 15](#)

[Migration, Upgrade, and Downgrade Instructions | 19](#)

[Product Compatibility | 20](#)

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 18](#)

This section lists the errata and changes in Junos OS Release 18.1R3 for the ACX Series documentation.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 13](#)

[Changes in Behavior and Syntax | 14](#)

[Known Behavior | 15](#)

[Known Issues | 15](#)

[Resolved Issues | 16](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 19](#)

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

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

New and Changed Features 13
Changes in Behavior and Syntax 14
Known Behavior 15
Documentation Updates 18
Known Issues 15
Resolved Issues 16
Product Compatibility 20

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 20

Hardware Compatibility

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.

To determine the features supported on ACX Series routers in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at <https://pathfinder.juniper.net/feature-explorer/>.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

New and Changed Features 13
Changes in Behavior and Syntax 14
Known Behavior 15

[Documentation Updates | 18](#)

[Known Issues | 15](#)

[Resolved Issues | 16](#)

[Migration, Upgrade, and Downgrade Instructions | 19](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- [New and Changed Features | 21](#)
- [Changes in Behavior and Syntax | 30](#)
- [Known Behavior | 35](#)
- [Known Issues | 37](#)
- [Resolved Issues | 40](#)
- [Documentation Updates | 48](#)
- [Migration, Upgrade, and Downgrade Instructions | 49](#)
- [Product Compatibility | 50](#)

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- [Release 18.1R3-S3 New and Changed Features | 22](#)
- [Release 18.1R3 New and Changed Features | 23](#)

- [Release 18.1R2 New and Changed Features | 23](#)
- [Release 18.1R1 New and Changed Features | 25](#)

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for EX Series.

NOTE: The following EX Series switches are supported in Release 18.1R3: EX2300, EX3400, EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 18.1R3, J-Web is supported on the EX2300, EX3400, EX4300, and EX4600 switches in both standalone and Virtual Chassis setup.

The J-Web distribution model being used provides two packages:

- Platform package—Installed as part of Junos OS; provides basic functionalities of J-Web.
- Application package—Optionally installable package; provides complete functionalities of J-Web.

For details about the J-Web distribution model, see [J-Web Application Package Release 18.1A1 for EX2300, EX3400, EX4300, and EX4600 Switches](#).

Release 18.1R3-S3 New and Changed Features

EVPNs

- **Support for VMTO for ingress traffic (EX9200)**—Starting in Junos OS Release 18.1R3-S3, you can configure the PE device to support virtual machine traffic optimization (VMTO) for ingress traffic. VMTO eliminates the unnecessary ingress routing to default gateways when a virtual machine is moved from one data center to another.

To enable VMTO, configure **remote-ip-host** routes in the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. You can also filter out the unwanted routes by configuring an import policy under the **remote-ip-host routes** option.

[See [Configuring EVPN Routing Instances](#).]

- **Support for Multihomed Proxy Advertisement (EX9200)**—Junos now provides enhanced support to proxy advertise the Mac address and IP route entry from all PEs that are multi-homed to a CE device.

This can prevent traffic loss when one of the links to the PE fails. To support the multihomed proxy advertisement, all multi-homed PE devices should have the same multihomed proxy advertisement bit value. The multihomed proxy advertisement feature is enabled by default and Junos uses the default multihomed proxy advertisement bit value of 0x20.

[See [EVPN Multihoming Overview](#).]

- **Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks (EX9200)**—Starting in Junos OS Release 18.1R3-S3, you can configure OSPF, IS-IS, BGP, and static routing with bidirectional forwarding detection (BFD) on an IRB interface that is used as a routed interface in EVPN. This allows protocol adjacencies to be established between an IRB on a Layer 3 gateway and a CE device connected directly to a Layer 3 gateway or to a Layer 2 leaf device in an EVPN-VXLAN network.

[See [Supported Protocols on an IRB Interface in EVPN-VXLAN](#) .]

Release 18.1R3 New and Changed Features

- There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 18.1R3.

Release 18.1R2 New and Changed Features

Hardware

- **EX2300-24MP and EX2300-48MP switches**—Starting with Junos OS Release 18.1R2, two new models of EX2300 switches—EX2300-24MP and EX2300-48MP—are available. EX2300-24MP switch models have eight 100/1000/2500 BASE-T Gigabit Ethernet ports with PoE/PoE+ capability, 16 10/100/1000 BASE-T Gigabit Ethernet ports with PoE/PoE+ capability, and four built-in 10-Gigabit Ethernet uplink ports. EX2300-48MP switch models have 16 100/1000/2500 BASE-T Gigabit Ethernet ports with PoE/PoE+ capability, 32 10/100/1000 BASE-T Gigabit Ethernet ports with PoE/PoE+ capability, and six built-in 10-Gigabit Ethernet uplink ports.

[See [EX2300 Switch Hardware Guide](#).]

Interfaces and Chassis

- **Support for Multi-Gigabit Ethernet (EX2300)**—Starting in Junos OS Release 18.1R2, the Multi-Gigabit Ethernet feature is supported on EX2300-48MP and EX2300-24MP switches. This feature fulfills the high-speed requirements for a large and mid-size campus, and branch locations for the enterprise customers.

The **mge** interface is a rate-selectable (multirate) Gigabit Ethernet interface that can support speeds of 10 Gbps, 5 Gbps, and 2.5 Gbps over CAT5e/CAT6/CAT6a cables. In the EX2300, the **mge** interface supports 100 Mbps, 1 Gbps, and 2.5 Gbps speeds, which can be configured by using the **speed** configuration statement.

NOTE: Power over Ethernet (PoE) is supported on Multi-Gigabit Ethernet interfaces. PoE enables EX2300 switches to transfer electrical power through an Ethernet cable. PoE enables electric power, along with data, to be passed over a copper Ethernet LAN cable.

[See [Speed](#).]

- **Support for Power over Ethernet (EX2300-24MP and EX2300-48MP)**—Starting in Junos OS Release 18.1R2, Power over Ethernet (PoE) is supported on EX2300-24MP and EX2300-48MP switch models, including multigigabit interfaces. (PoE) permits electric power, along with data, to be passed over a copper Ethernet LAN cable.

EX2300 24MP switches support PoE (IEEE 802.3af) and PoE+ (IEEE 802.3at) and can simultaneously deliver up to 15.4 watts of standards-based 802.3af Class 3 PoE to a maximum of 24 ports or 30 watts of standards-based 802.3at PoE+ to a maximum of 12 ports, based on a total system budget of 380 watts.

EX2300 48MP switches support PoE (IEEE 802.3af) and PoE+ (IEEE 802.3at) and can simultaneously deliver up to 15.4 watts of standards-based 802.3af Class 3 PoE to a maximum of 48 ports or 30 watts of standards-based 802.3at PoE+ to a maximum of 24 ports, based on a total system budget of 740 watts.

[See [Understanding PoE on EX Series Switches](#).]

Restoration Procedures Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (EX Series)**—In Junos OS Release 18.1R2, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode. The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

[See [Saving a Rescue Configuration File](#).]

Virtual Chassis

- **Virtual Chassis support (EX2300-24MP and EX2300-48MP)**—Starting in Junos OS Release 18.1R2, multigigabit EX2300 switches can be interconnected into a Virtual Chassis and operate as one logical device managed as a single chassis, as follows:
 - Members can be any combination of up to four EX2300-24MP and EX2300-48MP switches.
 - Multigigabit EX2300 switches cannot be mixed with any other switch models (including any other EX2300 switches) in the same Virtual Chassis.

- Any 10-Gbps uplink ports installed with SFP+ transceivers can be configured as Virtual Chassis ports (VCPs) to interconnect the members. Multigigabit EX2300 switches do not have any dedicated or default-configured VCPs.

To configure a multigigabit EX2300 Virtual Chassis, use similar steps as for configuring other EX Series and QFX Series Virtual Chassis.

[See [Understanding EX2300 Virtual Chassis.](#)]

Release 18.1R1 New and Changed Features

Hardware

- **EX9251 switches**—Starting with Junos OS Release 18.1R1, EX9251 switches are available as a fixed configuration switch. It is an Ethernet-optimized switch that provides carrier-class Ethernet switching. It has a throughput of up to 400 gigabits per second (Gbps). The switch is available in two variants—with AC power supply and with DC power supply.

[See [EX9251 Switch Hardware Guide.](#)]

Authentication, Authorization, and Accounting (AAA) (RADIUS)

- **Access control and authentication (EX2300 and EX3400 switches)**—Starting with Junos OS Release 18.1R1, EX2300 and EX3400 switches support controlling access to your network using 802.1X authentication and MAC RADIUS authentication.
 - 802.1X authentication provides port-based network access control (PNAC) as defined in the IEEE 802.1X standard. QFX5100 switches support 802.1X features including guest VLAN, private VLAN, server fail fallback, dynamic changes to a user session, RADIUS accounting, and configuration of port-filtering attributes on the RADIUS server using VSAs. You configure 802.1X authentication at the `[edit protocols dot1x]` hierarchy level.
 - MAC RADIUS authentication is used to authenticate end devices independently of whether they are enabled for 802.1X authentication. You can permit end devices that are not 802.1X-enabled to access the LAN by configuring MAC RADIUS authentication on the switch interfaces to which the end devices are connected. You configure MAC RADIUS authentication at the `[edit protocols dot1x authenticator interface interface-name mac-radius]` hierarchy level.

This feature was introduced previously in an “X” release of Junos OS.

[See [Understanding Authentication on Switches.](#)]

- **TACACS+ authorization for operational commands using regular expressions (EX2300, EX3400, EX4300 switches and MX Series)**—Starting in Junos OS Release 18.1R1, you can configure authorizations for operational mode commands using regular expressions using the **allow-commands-regexps** and **deny-commands-regexps** statements. Authorizations can also be configured remotely by specifying Juniper Networks vendor-specific attributes (VSAs) in your TACACS+ authentication server's configuration.

[See [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies.](#)]

Class of Service (CoS)

- **Support for Class of service (EX2300 and EX3400 switches and EX3400 Virtual Chassis)**—Starting in Junos OS Release 18.1R1, when a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default class-of-service(CoS) settings or CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and any rewrite rules to re-mark packets.

This feature was previously supported in an “X” release of Junos OS.

[See [Junos OS CoS for EX Series Switches Overview.](#)]

High Availability (HA) and Resiliency

- **High availability features (EX3400 switches and EX3400 Virtual Chassis)**—Starting with Junos OS Release 18.1R1, high availability features are supported. High availability features refer to the hardware and software components that provide redundancy and reliability for network communications.

The following features are supported:

- Graceful Routing Engine switchover (GRES), nonstop active routing and nonstop bridging
- Virtual Router Redundancy Protocol (VRRP) support

VRRP enables you to provide alternative gateways for end hosts that are configured with static default routes. You can implement VRRP to provide a high availability default path to a gateway without the need to configure dynamic routing or router discovery protocols on end hosts.

[See [High Availability User Guide.](#)]

Layer 2 Features

- **Layer 2 features (EX3400 switches and EX3400 Virtual Chassis)**—Starting with Junos OS Release 18.1R1, the following Layer 2 features are supported:

- VLAN support

VLANs enable you to divide one physical broadcast domain into multiple virtual domains.

- Link Layer Discovery Protocol (LLDP) support

LLDP enables a switch to advertise its identity and capabilities on a LAN, as well as receive information about other network devices.

- Q-in-Q tunneling support

This feature enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites.

- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) support

These protocols enable a switch to advertise its identity and capabilities on a LAN and receive information about other network devices.

This feature was previously supported in an “X” release of Junos OS.

[See [Ethernet Switching User Guide](#).]

Layer 3 Features

- **Layer 3 feature support (EX2300 and EX3400 Switches)**—Starting with Junos OS Release 18.1R1, the Layer 3 features supported in Junos OS Release 15.1X53-D50 are now supported on EX2300 and EX4300 Switches.

Multicast

- **Layer 2 and Layer 3 multicast support (EX2300 switches and Virtual Chassis, EX3400 switches and Virtual Chassis)**—Starting in Junos OS Release 18.1R1, the following IPv4 and IPv6 multicast protocols are supported:
 - Internet Group Management Protocol (IGMP) v1, v2, and v3
 - IGMP snooping
 - Multicast Listener Discovery (MLD) protocol v1 and v2
 - MLD snooping
 - Multicast Source Discovery Protocol (MSDP)
 - Protocol Independent Multicast (PIM) sparse mode (SM), dense mode (DM), and source-specific multicast (SSM)

These features were previously supported in an “X” release of Junos OS.

[See [Multicast Protocols User Guide](#).]

Network Management and Monitoring

- **Pseudohardware RPM timestamps (EX4300 switches and EX4300 Virtual Chassis)**—Starting in Junos OS Release 18.1R1, you can configure a pseudo-hardware timestamp on the switch for real-time performance monitoring (RPM). RPM enables you to configure active probes to track and monitor traffic on the network. To achieve this, RPM exchanges a set of probes with other IP hosts in the network. These probes are sent from a source node to other destination devices in the network that requires tracking. To account for latency or jitter in the communication of probe messages, you can enable timestamping of the probe packets. On the EX4300 switch, RPM timestamping is performed in the software. The RPM probes at the requester and responder devices are timestamped in the Packet Forwarding Engine instead of the Junos OS process (rmpod) that runs on the Routing Engine. This

timestamping method is referred to as pseudo-hardware timestamping. You must configure the switch as both the RPM client (the requester) and the RPM server (the responder) to timestamp the RPM packet. You configure pseudohardware timestamps at the **[edit services rpm]** hierarchy level.

[See [Understanding Real-Time Performance Monitoring on EX Series Switches.](#)]

- **Port mirroring support (EX2300, EX2300-C, and EX3400 switches and EX3400 Virtual Chassis)**—Starting in Junos OS Release 18.1R1, port mirroring is supported on EX2300, EX2300-C, and EX3400 switches and EX3400 Virtual Chassis. Port mirroring copies packets entering or exiting a port or entering a VLAN 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.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Port Mirroring and Analyzers on EX2300, EX3400, and EX4300 Switches.](#)]

Port Security

- **IPv4/IPv6 source guard (EX4600 switches)**—Starting in Junos OS Release 18.1R1, you can configure the IP source guard access port security feature to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, it discards the packet. This feature is supported for IPv4 and IPv6 source addresses.

[See [Understanding IP Source Guard for Port Security on EX Series Switches.](#)]

- **MACsec license enforcement (EX3400, EX4300, EX4600, EX9200, QFX5100 switches and Junos Fusion Enterprise)**—Starting in Junos OS Release 18.1R1, Media Access Control Security (MACsec) requires the installation of a MACsec feature license. If the MACsec license is not installed, MACsec functionality cannot be activated. You add the MACsec license using the **request system license add** command.

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

Security

- **Distributed denial-of-service (DDoS) protection (EX2300 and EX3400 switches, EX2300 and EX3400 Virtual Chassis)**—Starting in Junos OS Release 18.1R1, you can configure DDoS protection that enables the switch to continue functioning while under attack. DDoS attacks use multiple sources to flood a network or switch with protocol control packets. This malicious traffic triggers a large number of exceptions in the network and tries to exhaust the system resources so that valid users are denied access to the network or server. DDoS protection identifies and suppresses malicious control packets while enabling legitimate control traffic to be processed.

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

- **Support for firewall filters (EX2300 and EX3400 switches, EX2300 and EX3400 Virtual Chassis)**—Starting in Junos OS Release 18.1R1, you can define firewall filters on the switch that define whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces.

This feature was previously supported in an “X” release of Junos OS.

[See [Firewall Filters for EX Series Switches Overview](#).]

- **Port security features (EX2300 and EX3400 switches, EX2300 and EX4300 Virtual Chassis)**—Starting in Junos OS Release 18.1R1, the following port security features are supported:
 - DHCP snooping (Pv4 and IPv6)—Filters and blocks ingress Dynamic Host Configuration Protocol (DHCP) server messages on untrusted ports, and builds and maintains a database of DHCP lease information, which is called the DHCP snooping database.
 - Dynamic ARP inspection (DAI)—Prevents Address Resolution Protocol (ARP) spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable DAI on a VLAN.
 - IPv6 neighbor discovery inspection—Prevents IPv6 address spoofing attacks. Neighbor discovery requests and replies are compared against entries in the DHCPv6 snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable neighbor discovery inspection on a VLAN.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity](#).]

- **Port mirroring to IP address (EX4600 switches and Virtual Chassis)**—Starting with Junos OS Release 18.1R1, you can send mirrored packets to an IP address over a Layer 3 network (for example, if there is no Layer 2 connectivity to the analyzer device).

[See [Understanding Port Mirroring](#).]

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (EX2300, EX3400, EX4300, EX4600, and EX9200 switches)**—Starting in Junos OS Release 18.1R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database. The ephemeral database provides a fast programmatic interface that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. Junos OS provides a default instance and up to eight user-defined instances of the ephemeral configuration database. The device’s active configuration is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database. Ephemeral configuration data is volatile and is deleted upon rebooting the device.

[See [Understanding the Ephemeral Configuration Database](#).]

Virtual Chassis

- **Virtual Chassis support (EX2300, EX3400)**—Starting in Junos OS Release 18.1R1, EX2300 or EX3400 switches can be interconnected into a Virtual Chassis and operate as one logical device managed as a single chassis, as follows:

- EX2300 Virtual Chassis: Up to four EX2300 and EX2300-C member switches, interconnected using any 10-Gbps SFP+ ports configured as Virtual Chassis ports (VCPs)
- EX3400 Virtual Chassis: Up to 10 EX3400 member switches, interconnected using the QSFP+ uplink ports (default-configured VCPs) or any SFP+ uplink ports configured as VCPs

To configure an EX2300 or EX3400 Virtual Chassis, use similar steps as for configuring other EX Series and QFX Series Virtual Chassis.

This feature was previously supported in an “X” release of Junos OS.

[See [Virtual Chassis User Guide for Switches](#).]

SEE ALSO

Changes in Behavior and Syntax	 30
Known Behavior	 35
Known Issues	 37
Resolved Issues	 40
Documentation Updates	 48
Migration, Upgrade, and Downgrade Instructions	 49
Product Compatibility	 50

Changes in Behavior and Syntax

IN THIS SECTION

- [Release 18.1R3-S10 Changes in Behavior and Syntax](#) | 31
- [Release 18.1R3-S7 Changes in Behavior and Syntax](#) | 31
- [Release 18.1R3 Changes in Behavior and Syntax](#) | 32
- [Release 18.1R2 Changes in Behavior and Syntax](#) | 34
- [Release 18.1R1 Changes in Behavior and Syntax](#) | 34

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.1R3 for the EX Series.

Release 18.1R3-S10 Changes in Behavior and Syntax

Routing Protocols

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the show interfaces mc-ae extensive command. The output now displays the following two additional fields:
 - Local Partner System ID-LACP partner system ID as seen by the local node.
 - Peer Partner System ID-LACP partner system ID as seen by the MC-AE peer node.

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

[See [show interfaces mc-ae](#).]

Release 18.1R3-S7 Changes in Behavior and Syntax

Routing Protocols

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the show interfaces mc-ae extensive command. The output now displays the following two additional fields:
 - Local Partner System ID-LACP partner system ID as seen by the local node.
 - Peer Partner System ID-LACP partner system ID as seen by the MC-AE peer node.

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

[See [show interfaces mc-ae](#).]

Release 18.1R3 Changes in Behavior and Syntax

Network Management and Monitoring

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a non-default routing instance and a non-default logical system (EX Series)**—Starting in Junos OS Release 18.1R3, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

Layer 2 Features

- **Configuration option for LLDP VLAN name type, length, and value (TLV) (EX3400, EX4300)**—Starting in Junos OS Release 18.1R3, you can configure the **vlan-name-tlv-option (name | vlan-id)** statement at the **[edit protocols lldp]** hierarchy level to select whether to transmit the VLAN name or simply the VLAN ID for the Link Layer Discovery Protocol (LLDP) VLAN name TLV when exchanging LLDP messages. By default, EX Series switches running Enhanced Layer 2 Software (ELS) transmit the VLAN ID for the LLDP VLAN name TLV, and the **show lldp detail** command displays the default string **vlan-vlan-id** for an interface's VLAN name in the **Vlan-name** output field. Switches that support the **vlan-name-tlv-option** statement behave the same as the default if you configure the **vlan-id** option with this statement. If you configure the **name** option, the switch transmits the VLAN name instead, and the **show lldp detail** command displays the VLAN name in the **Vlan-name** output field.

Security

- **Firewall warning message (EX2300 switches)**—Starting in 18.1R3, a warning message is displayed whenever a firewall term includes log or syslog with the accept filter action.

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (EX Series)**—Starting in Junos OS Release 18.1R3, the **jdhcpcd** process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

In earlier releases:

1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

[See [Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview](#).]

Virtual Chassis

- **New configuration option to disable automatic Virtual Chassis port conversion (EX4300 and EX4600 Virtual Chassis)**—Starting in Junos OS Release 18.1R3, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in an EX4300 or EX4600 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:
 - LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
 - The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
 - The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

Automatic VCP conversion is not needed when using default-configured VCPs on both sides of the link to interconnect two members. On both ends of the link, you can also manually configure network or uplink ports that are supported as VCPs, whether or not the automatic VCP conversion feature is enabled.

Deleting the **no-auto-conversion** statement from the configuration returns the Virtual Chassis to the default behavior, which reenables automatic VCP conversion.

[See [no-auto-conversion](#).]

Release 18.1R2 Changes in Behavior and Syntax

Interfaces and Chassis

- **EEE not supported on mge interfaces operating at 100-Mbps speed (EX2300-24MP and EX2300-48MP)**—In Junos OS Releases 18.1R2, if both Energy Efficient Ethernet (EEE) and 100-Mbps speed are configured on a rate-selectable (or multirate) Gigabit Ethernet (mge) port on EX2300-24MP and EX2300-48MP switches, the port operates only at 100-Mbps speed but EEE is not enabled on that port. EEE is supported only on mge interfaces that operate at 1-Gbps and 2.5-Gbps speeds.

Multicast

- **Support for per-source multicast traffic forwarding with IGMPv3 (EX2300 and EX3400)**—Starting in Junos OS Release 18.1R2, EX2300 and EX3400 switches forward multicast traffic on a per-source basis according to received IGMPv3 INCLUDE and EXCLUDE reports. In releases prior to this release, EX2300 and EX3400 switches process IGMPv3 reports, but instead of source-specific multicast (SSM) forwarding, they consolidate IGMPv3 INCLUDE and EXCLUDE mode reports for a group into one route for all sources sending to the group. As a result, with the prior behavior, receivers might get traffic from sources they didn't specify.

[See [IGMP Snooping Overview](#).]

Release 18.1R1 Changes in Behavior and Syntax

Management

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (EX9200 switches)**—Starting with Junos OS Release 18.1R1, the format of telemetry data exported through gRPC for NPU memory and memory utilization implements prefix compression. This change reduces the payload size of data exported. The following example shows the new format:

key: __prefix__

str_value: /components/component[name='FPC0:NPU0']/properties/property

key: [name='mem-util-edmem-size']/value

uint_value: 12345

Telemetry data is exported in key-value pairs. Previously, the data exported included the component and property names in a single key string.

[See [Guidelines for gRPC Sensors](#).]

- **Enhancement to LSP statistics sensor for Junos Telemetry Interface (EX9200 switches, QFX10000 switches, MX Series, and PTX Series)**—Starting with Junos OS 18.1R1, the telemetry data exported for the LSP statistics sensor no longer includes the phrase **and source 0.0.0.0** after the LSP name in the value string for the prefix key. This change reduces the payload size of data exported. The following is an example of the new format:

str_value: /mpls/lsp/constrained-path/tunnels/tunnel[name='LSP-4-3']/state/
counters[name='c-27810']/

Network Management and Monitoring

- **SNMP syslog messages changed (EX Series)**—Starting in Junos OS Release 18.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD—AgentX master agent failed to respond to ping. Attempting to re-register
 - NEW—AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD—NET-SNMP version %s AgentX subagent connected
 - NEW—NET-SNMP version %s AgentX subagent Open-Sent!

[See the [MIB Explorer](#).]

SEE ALSO

New and Changed Features 21
Known Behavior 35
Known Issues 37
Resolved Issues 40
Documentation Updates 48
Migration, Upgrade, and Downgrade Instructions 49
Product Compatibility 50

Known Behavior

IN THIS SECTION

- [Infrastructure | 36](#)
- [Interfaces and Chassis | 36](#)
- [Platform and Infrastructure | 36](#)
- [Virtual Chassis | 37](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 for the EX Series.

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

Infrastructure

- When the image is copied through FTP from a server to a switch, sometimes the ftpd WCPU might go high, causing the CLI to freeze for approximately 10 seconds. [PR1306286](#)
- On rare occasions, the EX2300-MP switch panics with fatal abort. This issue is seen when the rpd process is aborted and it occurs only when dtrace is enabled with continuous rpd process killing. [PR1329552](#)
- Issue is specific to downgrade (17.4T) and a core file is seen only once during downgrade due to timing issue in the SDK toolkit upgrade. After the upgrade, dcpfe recovers on its own and no issues will be seen after that. [PR1337008](#)

Interfaces and Chassis

- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including the master routing instance), but only one logical interface was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating incorrect configuration. This issue is fixed and it is now not allowed to configure the same IP address (the length of the mask does not matter) on different logical interfaces. [PR1221993](#)

Platform and Infrastructure

- On EX4300 switches, when 802.1X single-suplicant authentication is initiated, multiple **EAP Request Id Frame Sent** packets might be sent. [PR1163966](#)
- On EX4300 10G links, preexisting MACsec sessions might not come up after the following events (1) Process (pfex, dot1x) restarts or the system restarts. (2) The link flaps. [PR1294526](#)
- LAG interfaces flap during unified ISSU when fast LACP timers are configured. This might result in traffic loss during the unified ISSU. This issue occurs because Fast LACP timers are not supported on EX-92XX during unified ISSU. The fast LACP timer support needs to be added. [PR1316251](#)
- NSSU upgrade from Junos OS Release 15.1X53-D58 to Junos OS Release 18.1R1 will fail with ksyncd core in backup Routing Engine. [PR1344686](#)
- When upgrading from certain release to Junos OS Release 18.1R1 statistics daemon PFED might be seen generating core files. This issue is not service impacting. The issue can be cleared by rebooting the chassis or by deleting all files from /mfs. [PR1346925](#)

Virtual Chassis

- VC internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur and it is considered to be known behavior. Release note reference:

https://www.juniper.net/documentation/en_US/junos/information-products/topic-collections/release-notes/172/topic-118735.html#PR1347902

SEE ALSO

New and Changed Features	21
Changes in Behavior and Syntax	30
Known Issues	37
Resolved Issues	40
Documentation Updates	48
Migration, Upgrade, and Downgrade Instructions	49
Product Compatibility	50

Known Issues

IN THIS SECTION

- General Routing | 38
- Infrastructure | 39
- Layer 2 Features | 39
- Platform and Infrastructure | 39

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for the EX Series.

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

General Routing

- On an EX9200-12QS line card, interfaces with the default speed of 10 Gigabit Ethernet are not brought down even when the remote end of a connection is misconfigured as 40 Gigabit Ethernet. [PR1175918](#)
- On an EX9200-40XS line card, if you toggle the MACsec encryption option multiple times, encryption and protected MACsec statistics might be updated incorrectly. As a workaround, restart the line card. [PR1185659](#)
- When a configuration that offlines a Packet Forwarding Engine and another configuration that brings the Packet Forwarding Engine back online, is committed in quick succession, there could be Routing Engine and Packet Forwarding Engine out of sync errors logged in syslog. Most of the time these are benign errors, but sometimes they might result in Packet Forwarding Engine crashes. [PR1232178](#)
- Some configurations that are valid in Junos OS Release 12.3 are not valid for Junos OS Release 15.1. When you try to upgrade from Junos OS Release 12.3 to 15.1 with such configurations, after upgrade the device goes into amnesiac mode. [PR1313501](#)
- In a streaming telemetry scenario, if performing **commit full**, the na-grpd process might restart, causing disconnection of the streaming telemetry. [PR1326366](#)
- On EX2300 and EX3400 switches with SFP, when the actual receiver signal power exceeds 0.21 mW, the output of the command **show interfaces diagnostics optics** might display an incorrect value for the field **Receiver signal average optical power**. [PR1326642](#)
- On an EX3400 platform, when force-renew is initiated from a server, the renewing entry for the bounded client will not be displayed under **show dhcp-security binding**. [PR1328542](#)
- On EX4600 platforms, in some cases, the CoS (class of Service) configuration is not properly applied in the Packet Forwarding Engine, leading to unexpected egress traffic drop on some interfaces. [PR1329141](#)
- On an EX9251 switch, after you perform the **restart chassis-control** for the first time after the software image is upgraded or after the switch is rebooted, the MPC booting state changes from offline to online directly, without staying at present state during booting. This issue is seen occasionally. There is no functional impact because of this state change. [PR1332613](#)
- On an EX9251 switch, physical links might not come up if you perform frequent port profile changes while a line card reboot is in progress. [PR1340140](#)
- On an EX9251 switch, if there is a packet loop between aggregation devices because of a redundant link, one aggregation device in a dual aggregation device setup might reboot if you clear the DHCP relay bindings. You must remove the redundant link to prevent such a reboot. [PR1347507](#)
- On EX2300-48MP, EX2300-48T and EX2300-48P platforms, the **show virtual-chassis** command might not display the model name. [PR1362421](#)
- The **show interface ge-x/x/x** command indicates "Duplex: Half-duplex" when link-mode is set to automatic or is not set. This is a display issue and it works as auto. [PR1364659](#)
- The multicast router advertisement packets coming on a VLAN need to be flooded on the ports of all FPCs belonging to the same VLAN. Packets traversing through HighGig ports need to hit the h/w filter

to transmit packets to other FPCs. In issue state, filter is not applicable for highgig ports, so multicast RA packets are not traversing other FPCs. [PR1370329](#)

Infrastructure

- The **request system zeroize** command will result in the device going for a continuous reboot in non-FIPS mode. The restoration is to halt the boot sequence at the loader and install the media-net package through TFTP. [PR1337826](#)
- On EX2300, EX2300-C, and EX2300-MP platforms, if Junos OS is with FreeBSD kernel version 11 with the build date on or after 2019-02-12, the switch might stop forwarding traffic or responding to console. A reboot is required to restore the service. [PR1442376](#)

Layer 2 Features

- No error or warning is displayed when you commit association of IRB interfaces with VLAN using **set vlans VLAN150 I3-interface irb.X** without configuring the actual IRB interface using **set interfaces irb unit X**. [PR1359982](#)
- The **eswd[1200]: ESWD_MAC_SMAC_BRIDGE_MAC_IDENTICAL: Bridge Address Add: XX:XX:db:2b:26:81 SMAC is equal to bridge mac hence don't learn** message is seen in syslog every few minutes on the ERPS owner. Because the log is caused by ERPS PDU in ERPS setup, you can ignore the message. [PR1372422](#)

Platform and Infrastructure

- On EX4300, MACsec might not work properly on PHY84756 1G SFP ports, if AN is on and MACsec is configured on those ports. On the EX4300 device, all four uplink ports (PIC 2) are attached to PHY84756. On EX4300 fiber box, the last four ports of base board (PIC 0) and 8*1G/10G uplink ports (PIC 2) are attached to PHY84756. [PR1291724](#)
- On EX4300 switches, the filter-based forwarding (FBF) might not work properly after deactivating or activating. This issue occurs because stale entries are not being freed in ternary content addressable memory (TCAM), which leads to insufficient space in TCAM for processing filters. [PR1293581](#)

SEE ALSO

[New and Changed Features | 21](#)

[Changes in Behavior and Syntax | 30](#)

[Known Behavior | 35](#)

[Resolved Issues | 40](#)

[Documentation Updates | 48](#)

[Migration, Upgrade, and Downgrade Instructions | 49](#)

[Product Compatibility | 50](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3 | 40](#)
- [Resolved Issues: 18.1R2 | 42](#)
- [Resolved Issues: 18.1R1 | 44](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

General Routing

- On EX2300 and EX3400 switches, the bridge-id is assigned to "02:00:00:00:00:10" irrespective of the base-mac addresses. [PR1315633](#)
- The MAC movement within a VLAN is not working as expected when setting up 802.1X for multiple supplicant mode. [PR1329654](#)
- After the EX9200 FPC comes online, other FPCs might have 100 percent CPU usage and a traffic loss for up to 30 seconds. [PR1346949](#)
- A commit error is observed if the device is downgraded from Junos OS Release 18.2 to Junos OS Release 17.3R3. [PR1355542](#)
- On EX2300, EX3400, and EX4300MP platforms in a Virtual Chassis setup, dynamic ARP inspection might fail after a Virtual Chassis switchover when VSTP is enabled along with **no-mac-table-binding**. [PR1359753](#)
- On EX2300, EX3400, EX4300-MP and EX2300-MP platforms used as a transit switch, routed traffic sent out of IRB interface, uses old MAC address instead of the configured MAC address for the IRB. [PR1359816](#)

- On EX2300MP platforms, the fan count is wrong in **jnxFruName**, **jnxFilledDescr** and **jnxContainersCount**. [PR1361025](#)
- On EX4300-48MP, the dot1x protocol subsystem is taking long time to respond to management requests with the error **the dot1x-protocol subsystem is not responding to management requests**. [PR1361398](#)
- A nonexistent fan tray 1 is reported by chassisd on EX2300. [PR1361696](#)
- A unicast ARP packet loop might be observed in DAI scenario. [PR1370607](#)
- Port access list group is not properly reallocating TCAM slices. [PR1375022](#)
- On EX4300-48MP, the Syslog error **Error in bcm_port_sample_rate_set(ifl_cmd) : Reason Invalid port** is seen. [PR1376504](#)

EVPN

- The proxy ARP might not work as expected in an EVPN environment. [PR1368911](#)

Infrastructure

- EX4300 firewall rule **ip-options** used with commands other than "any" does not provide expected results. [PR1173347](#)
- Unable to provide management when em0 interface of FPC is connected to another FPC Layer 2 interface of the same Virtual Chassis. [PR1299385](#)
- The upgrade might fail if bad blocks occurs in the flash memory device or file system. [PR1317628](#)
- Need support for archiving dmesg file **/var/run/dmesg.boot***. [PR1327021](#)
- Core file is generated upon attempt to commit configuration. [PR1376362](#)

Layer 2 Features

- The dcpfe or fxpc process might crash on Packet Forwarding Engines with low memory when allocating huge memory. [PR1362332](#)

Network Management and Monitoring

- While toggling multiple times between baseline and CFM configs, all 30 CFM sessions are not up. [PR1360907](#)

Platform and Infrastructure

- The mismatch of VLAN ID between a logical interface and a VLAN configuration might result in traffic being silently discarded. [PR1259310](#)
- EX4300 crashes when it receives more than 120 KPPS ARPs on me0 interface. [PR1329430](#)
- The SNMP trap message are always sent out with log about **Fan/Blower OK** on EX4300-VC switch. [PR1329507](#)
- The **show spanning-tree statistics bridge** command output gives 0 for all VLAN instance IDs. [PR1337891](#)

- On MPC5, the inline-ka PPP echo requests are not transmitted when anchor-point is lt-x/2/x or lt-x/3/x in a pseudowire deployment. [PR1345727](#)
- Running RSI through the console port might cause the system to crash and reboot. [PR1349332](#)
- A high usage chassis alarm in /var does not clear from the EX4300 Virtual Chassis when a file is copied from fpc1 (master) to fpc0 (backup). [PR1354007](#)
- The ports using an SFP-T transceiver might be still up after a system halt. [PR1354857](#)
- The FPC would crash due to the memory leak caused by the VTEP traffic. [PR1356279](#)
- Some interfaces cannot be added under the STP configuration. [PR1363625](#)
- On EX4300 and EX4600 platforms, the l2ald process might crash in a dot1x scenario. [PR1363964](#)
- The Packet Forwarding Engine might crash on encountering frequent MAC move. [PR1367141](#)
- The **request system zeroize** command non-interactively might not erase the configuration on EX4300. [PR1368452](#)
- NTP broadcast packets are not forwarded out on VLAN Layer 2 ports. [PR1371035](#)
- On EX4300, lldp advertisement appears with incorrect auto-negotiation values. [PR1372966](#)

Resolved Issues: 18.1R2

General Routing

- The **hawkeye alarmd** transient error is observed on MX240, MX480, MX960, EX9200, and SRX5000 platforms. [PR1312336](#)
- On an EX3400 switch, MACsec is not supported on 10G uplink ports. [PR1325545](#)
- Traffic going through the aggregated Ethernet interface might be dropped if mastership changes. [PR1327578](#)
- The EX3400 switch floods unicast ARP replies in the VLAN when dynamic ARP inspection is enabled. [PR1331928](#)
- On an EX9200 switch, when an anchor FPC has no active child, bridge protocol data units (BPDUs) are not sent out to the other active child. [PR1333872](#)
- All the DHCP-Reply or DCHP-Offer packets might be discarded by the DHCP snooping if the DHCP snooping is not enabled in that VLAN. [PR1345426](#)
- On an EX2300 running Junos OS Release 15.1X53-D56 with the fxpc process, issuing the **accept-source-mac** command causes the CPU usage to spike up to 90 percent on an idle chassis. [PR1345978](#)
- The statistics PFED process might generate a core file on an upgrade between certain releases. [PR1346925](#)

- Starting in Junos OS Release 18.1R2, there is support for OPSFv3 authentication on EX Series switches. [PR1347630](#)
- Different behavior on the tagging of interfaces before and after reboot without any change in configuration. [PR1349712](#)
- On EX2300 and EX3400 switches, the **lACP mac re-write** protocol sends duplicate Link Aggregation Control Protocol (LACP), bridge protocol data unit (BPDU) with different destination MAC addresses. [PR1350329](#)
- After an EX2300 switch reboots, if you have ECMP next hop configured, the ECMP group might only be created on one Packet Forwarding Engine. [PR1351418](#)

Forwarding and Sampling

- After an EX9251 switch is set to factory default by zeroize, the DHCP service crashes. [PR1329682](#)

Infrastructure

- EX4300 firewall rule ip-options with knobs other than "any" doesn't provide expected results. [PR1173347](#)
- On an EX4600 switch, priority-based flow control (PFC) frames might not work. [PR1322439](#)
- The **interface LED** status might stay green even after disabling the interface and removing the cable. [PR1329903](#)

Interfaces and Chassis

- Some PoE devices might not receive PoE power from EX2300 or EX3400 switches due to a false report of Underload Latch. [PR1345234](#)
- On EX4600, the MC-lag after reboot of VRRP Master and Back up discards traffic to downstream switches. [PR1345316](#)

Platform and Infrastructure

- On the EX4300 Virtual Chassis switch , the FPC might crash and a PFEX core file might get generated. [PR1261852](#)
- Multicast receiver connected to the EX4300 switch might not be able to get the multicast streaming. [PR1308269](#)
- Autonegotiation is not working as expected between EX4300 and SRX5800. [PR1311458](#)
- IGMPv3 on EX4300 does not have the correct outgoing interfaces in the Packet Forwarding Engine that are listed in the kernel. [PR1317141](#)
- On an EX4300 platform, a MAC learning issue and new VLANs creation failure might occur for some VLANs. [PR1325816](#)
- On an EX4300 platform, when exhausting TCAM, the table filter is still programmed. [PR1330148](#)
- Internet Group Management Protocol (IGMP) packets are forwarded out of the redundant trunk group (RTG) backup interface. [PR1335733](#)

- MSTP might not work normally after permitting a commit. [PR1342900](#)
- On EX4300, the loopback filter is not blocking unauthorized BGP peers. [PR1343402](#)
- The firewall filter might not be programmed in the Packet Forwarding Engine even though TCAM entries are available. [PR1345296](#)
- The VLAN translation feature does not work for the control plane traffic. [PR1348094](#)
- Traffic drop might occur if LLC packets are sent with DSAP and SSAP as 0x88 and 0x8e. [PR1348618](#)

Routing Protocols

- Open Shortest Path First (OSFP) routes cannot be added to the routing table until the **lsa-refresh** timer expires. [PR1316348](#)
- The **igmp-snooping** protocol might be enabled unexpectedly. [PR1327048](#)

Resolved Issues: 18.1R1

Authentication and Access Control

- The LLDP-MED cannot forward the correct POE class. [PR1296547](#)
- The dot1x process might stop authenticating if continuous dot1x client reauthentication requests cannot get processed. [PR1300050](#)
- EX2300-C is missing the dot1xd_usr_authenticated help string. [PR1311465](#)

EVPN

- Split horizon label is not allocated after switching the configuration of ESI from 'single-active' to 'all-active'. [PR1307056](#)

Infrastructure

- Reboot logs are not shown on the mini-USB console even though **set system ports auxiliary port-type mini-usb** is configured. [PR1192388](#)
- The file system might be corrupted multiple times during image upgrade or commit operation. [PR1317250](#)
- PFC feature might not work on EX4600. [PR1322439](#)
- The ifinfo might generate core files on the EX4600 Virtual Chassis. [PR1324326](#)

Interfaces and Chassis

- On EX2300 and EX3400 IPV6 neighborship is not created on the IRB interface. [PR1198482](#)
- On the EX4300 Virtual Chassis: LACP flap is observed, after rebooting the master FPC with PDT configurations. [PR1301338](#)
- The interface might not work properly after the FPC restarts. [PR1329896](#)

MPLS

- QFX5100 and EX4600: Unified ISSU is not supported with MPLS configuration. [PR1264786](#)

Platform and Infrastructure

- After access is rejected, the dot1x process might crash due to memory leak. [PR1160059](#)
- On EX3400 and EX2300, LLDP, LACP, and MVRP protocols are not available under the **mac-rewrite** configuration. [PR1189353](#)
- The I2C log error message is printed. [PR1251604](#)
- EX3400 Virtual Chassis has tail drops on multicast queues due to incorrect shared buffer programming. [PR1269326](#)
- Traffic loss might be observed for about 10 seconds if the master member FPC reboots. [PR1283702](#)
- Doing load replace terminal and attempting to replace the interface stanza might terminate the current CLI session and leave the user session hanging. [PR1293587](#)
- Some packets might be dropped after GRE encapsulation on EX4300. [PR1293787](#)
- Syslogs contain messages with **%PFE-3: fpc0 ifd null, port 28 dc-pfe: %USER-3: ifd null, port 28 : %PFE-3: fpc0 ifd null, port 29 dc-pfe: %USER-3: ifd null, port 29**. [PR1295711](#)
- Eswd core file might be observed if **apply-groups** is configured under **interface-range**. [PR1300709](#)
- On EX4300 switches, when unknown unicast ICMP packets are received by an interface, packets are routed, so TTL is decremented. [PR1302070](#)
- Unknown IPv6 multicast traffic are dropped if mld-snooping is enabled. [PR1304345](#)
- The **show snmp mib walk** CLI command used for jnxMIMstMstiPortState does not display anything in Junos OS Release 17.1R2 on the EX4600 platform. [PR1305281](#)
- On EX2300 and EX3400 Virtual Chassis or standalone chassis, IP routing fails for destination routes (IPv4 or IPv6 routes) with prefix length of 32 or 128 when they point to ECMP nexthops. [PR1305462](#)
- Inconsistent IEEE P-bit marking occurs in 802.1Q header for OSPF packets. [PR1306750](#)
- The me0 link might stay up after the link is disabled. [PR1307085](#)
- Multicast receiver connected to EX4300 might not be able to get the multicast streaming. [PR1308269](#)
- Multicast receiver connected to EX4300 might not be able to receive the multicast streaming. [PR1308269](#)
- VLAN rewrite is not working on aggregated Ethernet interface for EX2300/3400. [PR1309998](#)
- Traceroute is not working in EX9200 device for routing instances running on Junos OS Release 17.1R3. [PR1310615](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- IGMP snooping might not learn multicast router interface dynamically. [PR1312128](#)
- The DHCP security binding table might not get updated. [PR1312670](#)

- The PoE-enabled port does not come up after reboot of the line card member in EX3400 Virtual Chassis. [PR1312983](#)
- A memory leak is seen for dot1xd. [PR1313578](#)
- The interface with 1G SFP might go down if **no-auto-negotiation** is configured. [PR1315668](#)
- Policer does not work for 224.0.0.X MC traffic to the kernel on EX4300s. [PR1313251](#)
- On EX2300 and EX3400 switches, access ports might incorrectly send VLAN-tagged traffic. [PR1315206](#)
- Need to replace the **show vlans evpn** command with the **show ethernet-switching evpn** command for EX92xx and QFX Series switches. [PR1316272](#)
- Image upgrade fails with the error message **ERROR: Failed to add Junos-**. [PR1317425](#)
- EX2300 interface statistics shows an incorrect bits-per-second (bps) value when the interface has line-rate traffic at 10 Gbps. [PR1318767](#)
- L2cpd core files might be seen if the interface is disabled under VSTP and enabled under RSTP. [PR1317908](#)
- A vmcore file might be seen, and the device might reboot after the ICL is changed from an aggregated Ethernet to a physical interface. [PR1318929](#)
- High latency might be observed between the Master Routing Engine and the other FPC. [PR1319795](#)
- EX3400 changes FAN speed frequently with Over Temperature alarm after a software upgrade. [PR1320687](#)
- VLAN might not be processed, which leads to improper STP convergence. [PR1320719](#)
- On the EX2300-48 platform, known unicast might be flooded if the source MAC address is on PFE1 and the destination MAC address is on PFE0. [PR1321612](#)
- Multicast traffic might not be forwarded to one of the receivers. [PR1323499](#)
- EX3400: MACsec not supported on 10G uplink ports. [PR1325545](#)
- L2cpd might create a core file. [PR1325917](#)
- EX Series switches do not send RADIUS request after modifying the interface-range configuration. [PR1326442](#)
- Packets with the DEI bit set in the L2 header are not forwarded on the EX3400 switches. [PR1326855](#)
- EX4600, QFX5100, and ACX5000: Major Alarm **Fan & PSU Airflow direction mismatch** is seen after removing the management cable. [PR1327561](#)
- DHCP packet duplication issue is seen on EX2300/EX3400. [PR1326857](#)
- New operational status detail command is added in **show poe** interface. [PR1330183](#)
- EX3400 CPU have hog when Continuous Telnet EC command are sent on more than 75 concurrent telnet session. [PR1331234](#)

- IP Directed broadcast traffic forwarding does not work on EX3400/EX2300 platform. Applications such as Wakeup-on-lan do not work without this support. [PR1331326](#)
- EX3400 floods unicast ARP replies when DAI is enabled. [PR1331928](#)
- EX2300-48T: "Base power reserved" value seen is higher than "Total power supplied" in **show chassis power-budget-statistics** command. [PR1333032](#)
- Group unknown is seen on **show filter hw 1 show_term_info** CLI after adding **tcam-group-optimization** CLI. [PR1333367](#)
- EX9200 -- Major Errors - MQSS Error code: 0x2203cb. [PR1334928](#)
- IGMP traffic going out of RTG backup link is causing a loop. [PR1335733](#)
- VLAN rewrite might not work properly on trunk ports. [PR1336174](#)

Routing Protocols

- An mcsnoopd core file is seen at **core @**
__raise,abort,__task_quit__,task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal
(enable_slip_detector=true, no_exit=true) at
../../../../src/junos/lib/libjtask/base/task_scheduler.c:275. [PR1305239](#)

User Interface and Configuration

- EX2300 Virtual Chassis committing from J-Web causes PHP process to spike high. [PR1328323](#)

SEE ALSO

[New and Changed Features | 21](#)

[Changes in Behavior and Syntax | 30](#)

[Known Behavior | 35](#)

[Known Issues | 37](#)

[Documentation Updates | 48](#)

[Migration, Upgrade, and Downgrade Instructions | 49](#)

[Product Compatibility | 50](#)

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 48](#)

This section lists the errata and changes in Junos OS Release 18.1R3 for the EX Series switches documentation.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 21](#)

[Changes in Behavior and Syntax | 30](#)

[Known Behavior | 35](#)

[Known Issues | 37](#)

[Resolved Issues | 40](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 49](#)

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](#).

NOTE:

- : EX2300 or EX3400 switches running Junos OS Software Release 15.1X53-D57 or earlier revisions cannot be directly upgraded via CLI to Junos OS Software Release 18.1R1 because of configuration incompatibilities between the two releases related to the uplink port configurations. For example: Any configuration having interfaces on the uplink module (xe-0/2/*) will throw errors during the upgrade process. To work around this problem, please specify the validate option in the upgrade command to check for these errors, then remove the configuration that results in the errors, and use the no-validate option to do the upgrade.

Alternately, an intermediate upgrade to 15.1X53-D58 can be performed by keeping the configuration intact and then a subsequent upgrade to 18.1R1 is possible.
- NSSU is not supported on EX2300-VC/EX3400-VC from Junos OS Release 15.1X53 to Junos OS Release 18.1R1 or later releases. For example, NSSU is not supported from Junos OS Release 15.1X53-D58 to Junos OS Release 18.1R1 or Junos OS Release 15.1X53-D57 to Junos OS Release 18.2R1.

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

New and Changed Features 21
Changes in Behavior and Syntax 30
Known Behavior 35
Known Issues 37
Resolved Issues 40
Documentation Updates 48
Product Compatibility 50

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 50](#)

Hardware Compatibility

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.

To determine the features supported on EX Series switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature

information to find the right software release and hardware platform for your network. Find Feature Explorer at <https://pathfinder.juniper.net/feature-explorer/>.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

New and Changed Features 21
Changes in Behavior and Syntax 30
Known Behavior 35
Known Issues 37
Resolved Issues 40
Documentation Updates 48
Migration, Upgrade, and Downgrade Instructions 49

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

●	New and Changed Features 52
●	Changes in Behavior and Syntax 58
●	Known Behavior 58
●	Known Issues 61
●	Resolved Issues 62
●	Documentation Updates 63
●	Migration, Upgrade, and Downgrade Instructions 64
●	Product Compatibility 71

These release notes accompany Junos OS Release 18.1R3 for the Junos Fusion Data Center. 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

New and Changed Features

IN THIS SECTION

- [Release 18.1R3 New and Changed Features | 53](#)
- [Release 18.1R2-S2 New and Changed Features | 53](#)
- [Release 18.1R2 New and Changed Features | 57](#)
- [Release 18.1R1 New and Changed Features | 57](#)

This section describes the new features and enhancements to existing features in Junos OS Release 18.1R3 for Junos Fusion Data Center.

Release 18.1R3 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Data Center in Junos OS Release 18.1R3.

Release 18.1R2-S2 New and Changed Features

Class of Service

- **Class of service support (Junos Fusion Data Center)**—Starting in Junos OS Release 18.1R2-S2, Junos Fusion Data Center supports the standard Junos class of service (CoS) features and operational commands in a quad-aggregation device configuration. Each extended port on a satellite device is a logical extension to the aggregation device. Therefore, the default CoS policy on the aggregation device applies to each extended port. You can also create standard CoS policies for extended ports.

A cascade port is a physical port or interface on an aggregation device that provides a connection to a satellite device. Port scheduling is supported on cascade ports. Junos Fusion technology reserves a separate set of queues with minimum bandwidth guarantees for in-band management traffic to protect against congestion caused by data traffic.

[See [Understanding CoS in Junos Fusion Data Center](#).]

EVPN

- **Designated event forwarding of SNMP traps in an EVPN topology (Junos Fusion Data Center)**—Starting with Junos OS Release 18.1R2-S2, you can enable SNMP on the aggregation device and designate trap forwarding in an EVPN topology in Junos Fusion Data Center. In an EVPN topology, the satellite device generates an SNMP trap event when a change occurs on any of the associated satellite devices. This trap event information is sent to all connected aggregation devices who then sends the trap request to the SNMP server. Because each aggregation devices sends its own copy of the trap, the SNMP server receives multiple copies of the trap for the same event on the satellite device, thereby causing overhead to the SNMP server. To prevent the trap from being generated for each aggregation device, you can enable designated trap forwarding so that the trap request is only sent by the aggregation device selected as the designated router. You enable designate trap forwarding under the **[satellite-management]** hierarchy. Designated event forwarding is disabled by default.

[See [Understanding Designated Event Forwarding of SNMP Traps in an EVPN Junos Fusion Data Center](#).]

Interfaces and Chassis

- **Configuration synchronization for up to four aggregation devices (Junos Fusion Data Center)**—Starting in Junos OS Release 18.1R2-S2, configuration synchronization enables you to easily propagate, synchronize, and commit configurations from one aggregation device (AD) to another AD. Log in to either AD to manage the other three ADs, and use configuration groups to simplify the configuration process. You can create one configuration group each for the local ADs, and a global configuration common to all ADs.

Create conditional groups to specify when configurations are synchronized. Enable peers-synchronize at the **[edit system commit]** hierarchy to synchronize configurations and commits across ADs by default. NETCONF over SSH provides a secure connection between ADs. Secure Copy Protocol (SCP) copies configurations securely between them.

[See [Understanding Multichassis Link Aggregation Group Configuration Consistency Check.](#)]

- **Increased number of aggregated Ethernet interfaces (Junos Fusion Data Center)**—Starting in Junos OS Release 18.1R2-S2, you can configure up to 1750 aggregated Ethernet interfaces for a Junos Fusion Data Center system. To configure, include the **device-count** statement with a value of 1000 at the **[edit chassis aggregated-devices ethernet]** hierarchy level and add member links in each bundle.

[See [Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion.](#)]

Junos Fusion Data Center

- **Junos Fusion Data Center with four aggregation devices and EVPN infrastructure (Junos Fusion Data Center)**—Starting with Junos OS Release 18.1R2-S2, Junos Fusion Data Center supports four aggregation devices to which each satellite device can be multihomed in active-active mode. In this topology, the four aggregation devices comprise a core fabric in which Ethernet VPN (EVPN) is implemented as the control plane in which host and server MAC addresses, network reachability, and other states learned by an aggregation device are advertised to the other aggregation devices. For the data plane, the aggregation devices use Virtual Extensible LAN (VXLAN) encapsulation when forwarding a Layer 2 data packet to other aggregation devices. Namely, an aggregation device encapsulates a data packet in a VXLAN UDP header and sends the packet by means of the Layer 3 network to another aggregation device. Upon receipt of the packet, the aggregation device de-encapsulates the packet and forwards it as appropriate.

Junos Fusion Data Center with four aggregation devices and an EVPN architecture implements IEEE 802.1BR processing between the aggregation devices and satellite devices.

[See [Understanding EVPN in a Junos Fusion Data Center.](#)]

- **Layer 2 unicast forwarding on extended ports (Junos Fusion Data Center)**—Starting with Junos OS Release 18.1R2-S2, Junos Fusion Data Center supports Layer 2 unicast forwarding on extended ports. When a remote MAC address is learned from a Type-2 MAC route advertisement, the aggregation device determines the corresponding extended port next hop from the Ethernet Segment Identifier (ESI) carried in the MAC route advertisement. This extended port next hop is resolved in the set of local cascade interfaces that are used to reach that extended port. Traffic sent to a destination extended port only traverses the EVPN tunnel if the destination extended port cannot be resolved to a local cascade interface. For non-extended port destinations located on a remote aggregation device (or external Provider Edge (PE) device in the same EVPN), traffic is carried in the EVPN tunnel. When EVPN MAC aliasing is enabled, aggregation devices signal their reachability towards the destination extended port using the per-EVI Ethernet A-D route, so that a list of aggregation devices can be built for load-balancing even if those aggregation devices have not advertised that specific MAC route.

[See [Understanding EVPN in a Junos Fusion Data Center.](#)]

- **Satellite device support (QFX5110 and QFX5200)**—Starting with Junos OS Release 18.1R2-S2, you can configure QFX5110-48S and QFX5200-32C switches as satellite devices in a Junos Fusion Data Center topology. The satellite device in a Junos Fusion topology is managed and configured by the aggregation device. Junos Fusion Data Center uses QFX10002, QFX10008 and QFX10016 switches in the aggregation device role.

[See [Junos Fusion Data Center Software and Hardware Requirements](#).]

- **Flow-based uplink selection (Junos Fusion Data Center)**—Starting in Junos OS Release 18.1R2-S2, you can configure flow-based uplink selection for satellite devices by defining a chassis group to which the uplink traffic flows will be directed.

[See [Understanding Remapping Uplink Traffic Flows on a Junos Fusion Data Center](#).]

Multicast

- **Layer 2 multicast support with local replication in an EVPN topology (Junos Fusion Data Center)**—Starting with Junos OS Release 18.1R2-S2, Junos Fusion Data Center with EVPN combines elements of an EVPN multicast infrastructure with 802.1BR local replication to support Layer 2 multicast forwarding. In this environment, each extended port on a satellite device is multihomed to all aggregation devices and modeled as an EVPN Ethernet Segment (ES). One aggregation device is elected as the designated forwarder (DF) for each ES (based on the extended port's satellite device DF), and the IGMP snooping state is synchronized on all aggregation devices connected to that ES for faster convergence when DF re-election is required.

To forward multicast traffic, a source aggregation device employs local bias forwarding towards any locally reachable extended port multicast destinations, and uses ingress replication to the other aggregation devices in the EVPN/VxLAN tunnel acting as DFs for other ES destinations. Any forwarding aggregation device also uses 802.1BR local replication to destination satellite devices if you configure the **local-replication** statement at the **[edit forwarding-options satellite]** hierarchy level. Local replication, also referred to as egress replication at the satellite devices, helps distribute packet replication load and reduce traffic on cascade ports for multicast traffic by having the forwarding aggregation device send only one copy of a packet to each satellite device that has an extended port in the multicast group, and the satellite device then does the replication for its local extended ports.

[See [Multicast Forwarding at Layer 2 in a Junos Fusion Data Center with EVPN](#).]

- **Layer 3 multicast support in an EVPN topology (Junos Fusion Data Center)**— Starting with Junos OS Release 18.1R2-S2, Junos Fusion Data Center with EVPN includes support for sending Layer 3 multicast traffic between extended ports, or link aggregations of extended ports, that are located in different VLANs. Instead of running PIM on the aggregation devices, IGMP reports are sent to the gateways.

In the EVPN topology, two or more satellite devices are multihomed to four QFX10002 aggregation devices, and a multicast VLAN is provisioned between the aggregation devices and the external gateways. Both the source and receiver ports can be inside the fabric, or one can be external to the fabric while the other is internal. Likewise, participating servers can be connected to the fabric through the same or different tenants (in which case traffic must transit an external gateway so the gateway can handle the routing between tenants).

[See [Multicast Forwarding at Layer 3 in a Junos Fusion Data Center with EVPN](#).]

- **VLAN flooding support with local replication in an EVPN topology (Junos Fusion Data Center)**—Starting with Junos OS Release 18.1R2-S2, Junos Fusion Data Center with EVPN combines elements of an EVPN multicast infrastructure with 802.1BR local replication to support Layer 2 VLAN flooding. Local replication helps distribute packet replication load and reduce traffic on cascade ports for multicast and flooded VLAN traffic. In this environment, each extended port on a satellite device is multihomed to all aggregation devices and modeled as an EVPN Ethernet Segment (ES). One aggregation device is elected as the designated forwarder (DF) for each ES (based on the extended port's satellite device DF).

An aggregation device might initiate VLAN flooding (broadcasting or flooding the packet out to all interfaces in the VLAN) to learn the MAC address for a destination that is not already in its Ethernet switching tables. With local replication enabled, the aggregation device requests multicast ECIDs to represent the extended ports in the VLAN on each satellite device. You configure 802.1BR local replication to destination satellite devices by configuring the **local-replication** statement at the **[edit forwarding-options satellite]** hierarchy level.

[See [local-replication](#).]

Port Security

- **Storm control on extended ports (Junos Fusion Data Center)**—Starting in Junos OS Release 18.1R2-S2, storm control is supported on the extended ports of the satellite device in a Junos Fusion Data Center. You can configure storm control from the aggregation device to rate-limit broadcast traffic, multicast traffic, and unknown unicast traffic level so that the fabric drops packets when the specified traffic level is exceeded, thus preventing packets from proliferating and degrading the LAN. You specify the storm control level as the traffic rate in kilobits per second (Kbps) of the combined traffic streams, or as the percentage of available link bandwidth used by the combined traffic streams. If the storm control is exceeded, you can also configure the extended ports to shut down interfaces by using the **action-shutdown** command, or by temporarily disabling the interfaces by using the **port-error-disable** command. Additionally, you can disable storm control on registered or unregistered multicast traffic.

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

- **Firewall filter support on extended ports (Junos Fusion Data Center)**—Starting in Junos OS Release 18.1R2-S2, you can configure firewall filters on extended ports in a Junos Fusion Data Center. An extended port is a physical interface on the satellite device that is managed through the aggregation device. From the aggregation device, you can configure a firewall filter to accept or discard a packet before it enters or exits the port. If a packet is accepted, you can configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). To use a firewall filter, you must first configure the filter and then apply it to the port. Firewall filters are defined under the **[edit firewall]** hierarchy level. This feature was previously supported in an "X" release of Junos OS.

[See [Overview of Firewall Filters](#).]

Storage

- **DCBX Version 1.01 sequence number TLV mapping for multiple aggregation devices (Junos Fusion Data Center)**—Starting in Junos OS 18.1R2-S2, to ensure Junos Fusion Data Center with EVPN appears as a single device to a DCBX Version 1.01 peer, satellite devices maintain a sequence number TLV mapping table to coordinate DCBX message sequence numbers with those on multiple aggregation devices. The satellite device maintains a local sequence number that is exchanged with the DCBX peer, and maps that value to and from each aggregation device's corresponding local sequence number before relaying DCBX messages in either direction. The **show dcbx neighbors interface *interface-name*** command displays both the aggregation device local sequence number (in the **sequence-number** field) and the satellite device local sequence number (in the **satellite sequence-number** field) in the Local-Advertisement section of the output. In addition, the **peer-chassis sequence-number** field displayed with dual aggregation devices for Junos Fusion Data Center with MC-LAG is no longer included that section of the output.

[See [Understanding DCBX on Junos Fusion Data Center](#).]

Release 18.1R2 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Data Center in Junos OS Release 18.1R2.

Release 18.1R1 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Data Center in Junos OS Release 18.1R1.

SEE ALSO

[Changes in Behavior and Syntax | 58](#)

[Known Behavior | 58](#)

[Known Issues | 61](#)

[Resolved Issues | 62](#)

[Documentation Updates | 63](#)

[Migration, Upgrade, and Downgrade Instructions | 64](#)

[Product Compatibility | 71](#)

Changes in Behavior and Syntax

There are no changes in behavior and syntax for Junos Fusion Data Center in Junos OS Release 18.1R3.

SEE ALSO

New and Changed Features	 52
Known Behavior	 58
Known Issues	 61
Resolved Issues	 62
Documentation Updates	 63
Migration, Upgrade, and Downgrade Instructions	 64
Product Compatibility	 71

Known Behavior

IN THIS SECTION

- [Junos Fusion Data Center](#) | [59](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 for Junos Fusion Data Center.

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

Junos Fusion Data Center

- The license installed will not be deleted, unless it is explicitly deleted using the `request` command. After disabling the cascade port, the license count will be marked as zero only after the satellite information is purged from the neighbor database. Previously, this satellite neighbor information persisted only for 8 minutes; now, neighbor information is being held for 8 hours. This time delay is introduced to avoid repeating the initial recognition of the satellite device for interface-down events. As a workaround, delete the FPC instance for the satellite device to see the license removed for the corresponding satellite device. [PR1192886](#)
- In a Junos Fusion Data Center, auto-channelization is not supported on 100G interfaces. As a workaround, you can set the channelization using the **channel-speed** CLI statement at the `[edit policy-options satellite-policies extended-ports-template template-name pic pic-number port port-number]` hierarchy level.
- The license installed will not be deleted, unless it is explicitly deleted using the **request** command. After disabling the cascade port, the license count will be marked as zero only after the satellite information is purged from the neighbor database. Previously this satellite neighbor information persisted for only for 8 minutes; now neighbor information is being held for 8 hours. This time delay is introduced to avoid repeating the initial recognition of the satellite device for interface-down events. `user@host> show configuration | display set | grep et-0/0/30 set groups user-host-grp interfaces et-0/0/30 cascade-port set chassis satellite-management fpc 101 cascade-ports et-0/0/30 set interfaces et-0/0/30 disable {master:0}` `user@host> show chassis satellite terse`

Device	Extended Ports	Slot	State	Model	Total/Up
Version 100 Online	EX4300-48T	50/1	17.4-20170726_common_xxx.0	102 Online	QFX5200-32C-32Q
2/1 17.4-20170726_common_xxx.0	103 Online	QFX5110-48S-4C	3/2 17.4-20170726_common_xxx.0		

`{master:0}` `user@host> show chassis satellite neighbor`

Interface	State	Port	Info	System	Name	Model
SW Version et-0/0/30 Dn	et-0/0/18 Two-Way	et-0/0/18	sd102	QFX5200-32C-32Q		
17.4-20170726_common_xxx.0	et-0/0/12 Two-Way	et-0/0/50	sd103	QFX5110-48S-4C		
17.4-20170726_common_xxx.0	et-0/0/6 Two-Way	et-0/1/3	sd100	EX4300-48T		
17.4-20170726_common_xxx.0						

`{master:0}` `user@host> show system license`

License usage:	Licenses	Licenses	Licenses	Expiry	Feature	name	used	installed	needed	bgp	1	0	1	invalid	SD-QFX5100-48SH-48TH	
0	4	0	permanent	Licenses installed:	License identifier:	JUNOSxxxxxx	License version:	4	Software	Serial	Number:	99999B999999999	Customer ID:	USER-SWITCH	Features:	SD-QFX5100-48SH-48TH-4PK - SD 4 pack QFX5000-10-JFD permanent

`{master:0}` `user@host> show system license usage`

Licenses	Licenses	Licenses	Expiry	Feature	name	used	installed	needed	bgp	1	0	1	invalid	SD-QFX5100-48SH-48TH
0	4	0	permanent											

`{master:0}` `user@host> show system alarms`

Alarms	currently active	Alarm time	Class	Description
4		2017-08-29 13:14:27 UTC	Minor	BGP Routing Protocol usage requires a license
2017-08-28 17:25:27 UTC	Major	FPC0: PEM 1 Not Powered	2017-08-28 17:25:27 UTC	Major
FPC Management	1	Ethernet Link Down		

[PR1294951](#)
- In a Junos Fusion Data Center, configuration synchronization is not triggered when you issue the `rollback` command on the local aggregation device (AD). [PR1298747](#)
- During a satellite device (SD) upgrade in Junos Fusion, there is a condition that the Link Aggregation Control Protocol (LACP) defaulted PDU received from the peer device connected in a link aggregation

group (LAG) is incorrectly sent to an aggregation device (AD) with aggregated Ethernet e-Channel Identifier (ECID) instead of member ECID, which causes that LACP PDU to be received on another member in the LAG, resulting in the LAG interface flapping. [PR1321575](#)

- In a Junos Fusion Data Center with EVPN solution, when an Aggregation Device loses EVPN connectivity with rest of the aggregation devices, then LACP over extended ports on this core isolated aggregation device will be brought down until EVPN connectivity is restored. [PR1327784](#)
- On a Junos Fusion Data Center, traffic is dropped when the aggregation device (AD) goes down because the routes advertised by that AD are withdrawn. As a workaround, configure the hold-time on uplink interfaces using the **set interfaces interface hold-time up <300 * 1000>** command. [PR1331465](#)
- On the QFX10000, EVPN NSR Unicast is not currently supported. [PR1337645](#)
- This issue was seen with UBS scale: however, with PLM scale this issue is not seen in latest releases. [PR1338659](#)
- In the following Junos OS releases, the EVI-RT extended community as defined in draft-ietf-bess-evpn-igmp-mld-proxy-00 was not being attached to advertised Type 7 routes: 17.2, 17.3R1, 17.3R2, and 17.4R1. Starting with the following Junos OS releases, the EVI-RT is mandatory and Type 7 routes that do not carry this community will be ignored: 17.3R3, 17.4R2, 18.1R1 and later. An EVPN network enabled with IGMP snooping and having a mix of multi-homed peer PEs with some running the older releases (17.2, 17.3R1, 17.3R2, 17.4R1) and some running the newer releases (17.3R3, 17.4R2, 18.1R1 and later) will not be able to inter-operate since Type 7 routes advertised without the EVI-RT community by PEs running an older release will be ignored by PEs running the newer releases and will not result in creation of IGMP snooping state. [PR1341807](#)
- In Junos Fusion with EVPN solution when an Aggregation Device loses EVPN connectivity with rest of the aggregation devices, then LACP over extended ports on this core isolated Aggregation Device will be brought down until EVPN connectivity is restored. [PR1342045](#)
- When some of the EX4300 satellite devices are rebooted, they stay in offline state as packets get corrupted or dropped due to internal FIFO logic error. As a workaround, restart the satellite device PFE process. [PR1349508](#)
- On platforms with dual routing engines, GRES will result in considerable traffic loss/duplication for EVPN traffic. The traffic should restore eventually once convergence is complete. [PR1350744](#)
- Partial ingress traffic can be mirrored on Junos Fusion Data Center setup with a few triggers, such as aggregation device and satellite device reboot, configuration reboot of traffic generator connected to satellite devices etc. As a workaround, flap of mirror output interface resolves the issue. [PR1352827](#)
- In Junos OS, two different input filters cannot be configured on the same interface; if two filters are configured, only the second filter (the one that was configured most recently) takes effect. Ingress mirroring on extended ports in Junos Fusion Data Center (JFDC) can only be done by using firewall filters. Due to the Junos OS filter behavior, in JFDC, ingress mirroring on extended ports and other firewall filter configurations cannot be done on the same port. [PR1353065](#)

- Since EVPN graceful restart is not supported, restart of the rpd process will result in considerable traffic loss for EVPN traffic. The traffic should restore eventually once convergence is complete. [PR1353742](#)
- Two aggregation devices going down at the same time will result in considerable traffic loss/duplication for EVPN traffic. The traffic should restore eventually once convergence is complete. [PR1354443](#)
- On an aggregation device (AD), if the mirror destination interface is down then the mirror gets globally removed from all FPCs of the AD. The mirror is reprogrammed back when the mirror destination interface is up again. Until the mirror stays down, sampling won't be functional on the given AD. [PR1360003](#)

SEE ALSO

[New and Changed Features | 52](#)

[Changes in Behavior and Syntax | 58](#)

[Known Issues | 61](#)

[Resolved Issues | 62](#)

[Documentation Updates | 63](#)

[Migration, Upgrade, and Downgrade Instructions | 64](#)

[Product Compatibility | 71](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for Junos Fusion Data Center.

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

Junos Fusion Data Center

- Partial ingress traffic can be mirrored on Junos Fusion setup with some triggers, such as a reboot of the aggregation or satellite device or a configuration reboot of traffic generator connected to the satellite devices. As a workaround, flapping the mirror output interface resolves the issue. [PR1352827](#)

SEE ALSO

[New and Changed Features | 52](#)

Changes in Behavior and Syntax	58
Known Behavior	58
Resolved Issues	62
Documentation Updates	63
Migration, Upgrade, and Downgrade Instructions	64
Product Compatibility	71

Resolved Issues

IN THIS SECTION

- Resolved Issues: 18.1R3 | 62
- Resolved Issues: 18.1R2 | 62
- Resolved Issues: 18.1R1 | 63

This section lists the issues fixed in the Junos OS Release 18.1R3 for Junos Fusion Data Center.

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

Resolved Issues: 18.1R3

- In a Junos Fusion setup, a DCPFE core file is generated after executing the **sh shim bridge bd bd-index** command in the PFE. [PR1296738](#)

Resolved Issues: 18.1R2

Junos Fusion Data Center

- On a Junos Fusion topology with QFX10002 switches as aggregate devices having dual cascade links to each satellite devices for redundancy, duplicated multicast traffic might be seen on downstream devices and multicast receivers if the multicast traffic passes through the aggregate devices. As a workaround, deactivate and re-activate the VLAN in which duplicated multicast traffic is seen. [PR1316499](#)
- In a Junos Fusion setup, an aggregate device may show a plus sign on the ICL link for a satellite device. [PR1335373](#)

Resolved Issues: 18.1R1

Junos Fusion Data Center

- In a Junos Fusion topology with LAG on extended ports from satellite devices which are dual-homed to aggregation devices, the LAG interface might flap if rebooting one of the aggregation devices. [PR1315879](#)

SEE ALSO

[New and Changed Features | 52](#)

[Changes in Behavior and Syntax | 58](#)

[Known Behavior | 58](#)

[Known Issues | 61](#)

[Documentation Updates | 63](#)

[Migration, Upgrade, and Downgrade Instructions | 64](#)

[Product Compatibility | 71](#)

Documentation Updates

This section lists the errata or changes in Junos OS Release 18.1R3 for Junos Fusion Data Center documentation.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.

- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

New and Changed Features 52
Changes in Behavior and Syntax 58
Known Behavior 58
Known Issues 61
Resolved Issues 62
Migration, Upgrade, and Downgrade Instructions 64
Product Compatibility 71

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 64](#)
- [Preparing the Switch for Satellite Device Conversion | 66](#)
- [Configuring Satellite Device Upgrade Groups | 68](#)
- [Converting a Satellite Device to a Standalone Device | 69](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 70](#)
- [Downgrading from Junos OS Release 18.1 | 70](#)

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for Junos Fusion Data Center. Upgrading or downgrading Junos OS can 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 [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 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 command.

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

All other customers, use the following command.

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

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 can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

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 [Junos Fusion Hardware and Software Compatibility Matrices](#).

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 Junos switch can only be converted to SNOS 3.1 and higher.
- The Junos switch must be either set to factory default configuration 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, replacing *n* with the spin number:

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

Customers with QFX5100 switches, use the following command, replacing *n* with the spin number:

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

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 entering 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 or Expanding a Junos Fusion Data Center](#) for detailed configuration steps for each method.

Configuring Satellite Device Upgrade Groups

To simplify the upgrade process for multiple satellite devices, you can create a software upgrade group at the aggregation device, assign satellite devices to the group, and install the satellite software on a groupwide basis.

To create a software upgrade group and assign satellite devices to the group, include the **satellite** statement at the **[edit chassis satellite-management upgrade-groups *upgrade-group-name*]** hierarchy level.

To configure a software upgrade group and assign satellite devices to the group:

1. Log in to the aggregation device.
2. Create the software upgrade group, and add the satellite devices to the group.

```
[edit]
user@aggregation-device# set chassis satellite-management upgrade-groups
upgrade-group-name satellite satellite-member-number-or-range
```

upgrade-group-name is the name of the upgrade group, and the ***satellite-member-number-or-range*** is the member numbers of the satellite devices that are being added to the upgrade group. If you enter an existing upgrade group name as the ***upgrade-group-name***, you add new satellite devices to the existing software upgrade group.

For example, to create a software upgrade group named `group1` that includes all satellite devices numbered 101 through 120, configure the following:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management upgrade-groups group1 satellite
101-120
```

To install, remove, or roll back a satellite software version on an upgrade group, issue the following operational mode commands:

- **request system software add upgrade-group *group-name***—Install the satellite software on all members of the specified upgrade group.
- **request system software delete upgrade-group *group-name***—Remove the satellite software association from the specified upgrade group.
- **request system software rollback upgrade-group *group-name***—Associate an upgrade group with a previous version of satellite software.

Customers installing satellite software on EX4300 and QFX5100 switches referenced in a software upgrade group, use the following command:

```
user@aggregation-device> request system software add upgrade-group group-name
source/package-name
```

NOTE: Before issuing **request system software add upgrade-group *group-name***, you must issue a one-time command to expand the storage capacity. Use the **request system storage user-disk expand** command to increase the size of /user partition.

A copy of the satellite software is saved on the aggregation device. When you add a satellite device to an upgrade group that is not running the same satellite software version, the new satellite device is automatically updated to the version of satellite software that is associated with the upgrade group.

You can issue the **show chassis satellite software** command to see which software images are stored on the aggregation device and which upgrade groups are associated with the software images.

Converting a Satellite Device to a Standalone Device

In the event that you need to convert a satellite device to a standalone device, you will need to 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 even 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 18.1

To downgrade from Junos OS Release 18.1 to another supported release, follow the procedure for upgrading, but replace the 18.1 **jinstall** package with one that corresponds to the appropriate downgrade release.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

SEE ALSO

[New and Changed Features | 52](#)

[Changes in Behavior and Syntax | 58](#)

[Known Behavior | 58](#)

[Known Issues | 61](#)

[Resolved Issues | 62](#)

[Documentation Updates | 63](#)

[Product Compatibility | 71](#)

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 71](#)
- [Hardware Compatibility Tool | 71](#)

Hardware and Software Compatibility

For a complete list of all hardware and software requirements for a Junos Fusion Data Center, including which Juniper Networks devices function as satellite devices, see [Understanding Junos Fusion Data Center Software and Hardware Requirements](#) in the [Junos Fusion Data Center User Guide](#).

To obtain information about the components that are supported on the devices, and special compatibility guidelines with the release, see the Hardware Guides for the devices used in your Junos Fusion Data Center topology.

To determine the features supported in a Junos Fusion, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at <https://pathfinder.juniper.net/feature-explorer/>

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

- | |
|---|
| New and Changed Features 52 |
| Changes in Behavior and Syntax 58 |
| Known Behavior 58 |
| Known Issues 61 |
| Resolved Issues 62 |
| Documentation Updates 63 |
| Migration, Upgrade, and Downgrade Instructions 64 |

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- New and Changed Features | 72
- Changes in Behavior and Syntax | 74
- Known Behavior | 74
- Known Issues | 75
- Resolved Issues | 76
- Documentation Updates | 77
- Migration, Upgrade, and Downgrade Instructions | 78
- Product Compatibility | 84

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- Release 18.1R3 New and Changed Features | 73
- Release 18.1R2 New and Changed Features | 73
- Release 18.1R1 New and Changed Features | 73

This section describes the new features and enhancements to existing features in Junos OS Release 18.1R3 for Junos Fusion Enterprise.

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

Release 18.1R3 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 18.1R3.

Release 18.1R2 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 18.1R2.

Release 18.1R1 New and Changed Features

Junos Fusion Enterprise

- **Aggregation device support on EX9251 switches (Junos Fusion Enterprise)**—Starting with Junos OS Release 18.1R1, EX9251 switches are supported as aggregation devices in a Junos Fusion Enterprise. The aggregation device acts as the single point of management for all devices in the Junos Fusion Enterprise. Junos Fusion Enterprise supports the 802.1BR standard.

[See [Junos Fusion Enterprise Overview](#).]

SEE ALSO

[Changes in Behavior and Syntax | 74](#)

[Known Behavior | 74](#)

[Known Issues | 75](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

[Product Compatibility | 84](#)

Changes in Behavior and Syntax

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 18.1R3 for Junos Fusion Enterprise.

SEE ALSO

[New and Changed Features | 72](#)

[Known Behavior | 74](#)

[Known Issues | 75](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

[Product Compatibility | 84](#)

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 for Junos Fusion 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

[New and Changed Features | 72](#)

[Changes in Behavior and Syntax | 74](#)

[Known Issues | 75](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

[Product Compatibility | 84](#)

Known Issues

IN THIS SECTION

- [Junos Fusion Enterprise | 75](#)

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- On a Junos Fusion, when using LLDP, the **Power via MDI** and **Extended Power via MDI** TLVs are not transmitted. [PR1105217](#)
- On a Junos Fusion Enterprise, when the satellite devices of a cluster are rebooted, the output of the CLI command **show chassis satellite** shows the port state of the cascade ports as **Present**. [PR1175834](#)
- In a Junos Fusion Enterprise, it could take 6 to 30 seconds for the traffic to converge when on the aggregation device is powered **OFF** or powered **ON**. [PR1257057](#)
- In a Junos Fusion Enterprise, during RE switchover, the BUM traffic is duplicated to indirectly connected satellite devices. This is because there is no current support to notify the GRES event to indirectly connected satellite devices. [PR1298434](#)
- In a Junos Fusion Enterprise, after an automatic POE firmware upgrade, the satellite device reboots. [PR1359065](#)

SEE ALSO

[New and Changed Features | 72](#)

[Changes in Behavior and Syntax | 74](#)

[Known Behavior | 74](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

Resolved Issues

IN THIS SECTION

- Resolved Issues: 18.1R3 | [76](#)
- Resolved Issues: 18.1R2 | [76](#)
- Resolved Issues: 18.1R1 | [77](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

- A satellite device does not recover PoE after the device is offline for more than 10 minutes and rejoins the aggregation device. [PR1345365](#)
- The ppm-lite process might generate a core file on the Fusion satellite devices. [PR1364265](#)

Resolved Issues: 18.1R2

- In a Junos Fusion Enterprise in which port mirroring analyzers are configured, mirrored packets are dropped when the packets must traverse the interchassis link (ICL) link to reach destination extended ports. As a workaround, you can alternatively configure a remote switched port analyzer (RSPAN) VLAN with the extended ports and the ICL link as members and configure the RSPAN VLAN as the analyzer destination. [PR1211123](#)
- In a Junos Fusion setup, an aggregate device may show a plus sign on the ICL link for a satellite device. [PR1335373](#)
- Issue with 802.1X re-authentication in Junos Fusion Enterprise. [PR1345365](#)

Resolved Issues: 18.1R1

- Request chassis satellite beacon functionality to specific satellite device is not working, causing all the satellite devices to enable the beacon LED. [PR1272956](#)
- On a Junos Fusion Enterprise with dual aggregation devices (ADs), if you apply Routing Engine loopback filters and bring down the cascade port on one of the ADs, the satellite device (SD) on the AD where the cascade port is down goes to ProvSessDown due to a TCP session drop over the ICL interface. [PR1275290](#)
- Junos Fusion : SD EX4300 displaying U-Boot on LCD screen. [PR1304784](#)
- All 802.1X authentication sessions are removed when the AUTO ICCP link is disabled. [PR1307588](#)
- LACP aggregated Ethernet interfaces go to down state when performing **commit synchronize**. [PR1314561](#)
- Packets loss for 2-3 seconds is seen in every 5 minutes on Junos Fusion. [PR1320254](#)
- In a Junos Fusion Enterprise deployment, an SCPD core might be seen on an aggregation device when DACL on dot1x enabled port is installed on a single homed satellite device.. [PR1328247](#)
- When the ICCP and IFBDs are in transition--Down/Up--DHCP security binding entries might be missing from server database. [PR1332828](#)

SEE ALSO

[New and Changed Features | 72](#)

[Changes in Behavior and Syntax | 74](#)

[Known Behavior | 74](#)

[Known Issues | 75](#)

[Documentation Updates | 77](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

[Product Compatibility | 84](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 18.1R3 documentation for Junos Fusion.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we

organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 72](#)

[Changes in Behavior and Syntax | 74](#)

[Known Behavior | 74](#)

[Known Issues | 75](#)

[Resolved Issues | 76](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

[Product Compatibility | 84](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 79](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 81](#)
- [Preparing the Switch for Satellite Device Conversion | 81](#)
- [Converting a Satellite Device to a Standalone Switch | 83](#)

- Upgrade and Downgrade Support Policy for Junos OS Releases | 83
- Downgrading Junos OS | 83

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 Release 18.1R2:

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 command:

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

All other customers, use the following command:

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

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 Junos switch can only be converted to SNOS 3.1 and higher.
- The Junos switch must be either set to factory default configuration to factory default configuration 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 on EEOL releases and to review a list of EEOL releases, see <https://www.juniper.net/support/eol/junos.html>

Downgrading 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, follow the procedure for upgrading, but replace the **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

[Changes in Behavior and Syntax | 74](#)

[Known Behavior | 74](#)

[Known Issues | 75](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

[Product Compatibility | 84](#)

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 84](#)
- [Hardware Compatibility Tool | 84](#)

Hardware and Software Compatibility

For a complete list of all hardware and software requirements for a Junos Fusion Enterprise, including which Juniper Networks devices function as satellite devices, see [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#) in the [Junos Fusion Enterprise User Guide](#).

To determine the features supported in a Junos Fusion, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

[New and Changed Features | 72](#)

[Changes in Behavior and Syntax | 74](#)

Known Behavior 74
Known Issues 75
Resolved Issues 76
Documentation Updates 77
Migration, Upgrade, and Downgrade Instructions 78

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- New and Changed Features | 86
- Changes in Behavior and Syntax | 89
- Known Behavior | 90
- Known Issues | 91
- Resolved Issues | 92
- Documentation Updates | 93
- Migration, Upgrade, and Downgrade Instructions | 94
- Product Compatibility | 101

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- Release 18.1R3 New and Changed Features | 88
- Release 18.1R2 New and Changed Features | 88
- Release 18.1R1 New and Changed Features | 88

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for Junos Fusion Provider Edge.

Release 18.1R3 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.1R3.

Release 18.1R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.1R2.

Release 18.1R1 New and Changed Features

Hardware

- **Support for QFX5110 and QFX5200 as satellite devices in Junos Fusion Provider Edge**—Starting in Junos OS Release 18.1R1, you can use QFX5110-48S or QFX5200-32C switches as satellite devices in Junos Fusion Provider Edge.

[See [Satellite Device Hardware Models](#) and [Preparing the Satellite Device](#).]

Class of Service (CoS)

- **Support for dynamic mapping of extend ports to cascade ports for hierarchical CoS (Junos Fusion Provider Edge)**—Junos Fusion treats the cascade ports connecting the aggregation device to the satellite device as aggregated Ethernet ports with aggregation done automatically without configuration. By default the Junos Fusion implementation of hierarchical CoS applies the scheduler parameters across all cascade ports in **scale** mode. Because **scale** mode divides the configured shaper equally across the cascade ports, traffic drops can start before a customer reaches its committed rate for a particular flow. To avoid this problem, starting with Junos OS Release 18.1R1, you can set all cascade ports on an aggregation device to be in **replicate** mode and automatically target all of a customer's traffic to a specific cascade port. To do this, simply enable **target-mode** at the **[edit chassis satellite-management fpc fpc-number]** hierarchy level.

[See [Understanding CoS on an MX Series Aggregation Device in Junos Fusion](#).]

Junos Fusion

- **Junos Fusion Provider Edge support for Junos Node Slicing (MX960, MX2010, MX2020)**—Starting in Junos OS Release 18.1R1, you can configure an aggregation device on guest network functions (GNFs), or partitions created on a router, by using Junos Node Slicing. The Junos Fusion topology is composed of an aggregation device and multiple satellite devices. An MX Series router supports a maximum of 10 GNFs, with each GNF supporting a separate aggregation device. The aggregation device on a GNF supports a maximum of 10 satellite devices. The aggregation device acts as the single point of management for all devices in a Junos Fusion topology, while the satellite devices provide interfaces that send and receive network traffic.

For more information on Junos Node Slicing, see [Junos Node Slicing Overview](#).

NOTE:

- In a Junos Fusion Provider Edge topology that has a GNF configured as the aggregation device, you can only use EX4300 switches as satellite devices.
- Only the following line cards support the cascade port on the aggregation device: MPC7E-MRATE, MPC7E-10G, MPC8 and MPC9.

[See [Understanding Junos Fusion Provider Edge Components](#).]

SEE ALSO

[Changes in Behavior and Syntax | 89](#)

[Known Behavior | 90](#)

[Known Issues | 91](#)

[Resolved Issues | 92](#)

[Documentation Updates | 93](#)

[Migration, Upgrade, and Downgrade Instructions | 94](#)

[Product Compatibility | 101](#)

Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 18.1R3.

SEE ALSO

New and Changed Features	86
Known Behavior	90
Known Issues	91
Resolved Issues	92
Documentation Updates	93
Migration, Upgrade, and Downgrade Instructions	94
Product Compatibility	101

Known Behavior

IN THIS SECTION

- Junos Fusion | 90

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 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

- Configuration synchronization is not triggered when you issue the rollback command on the local aggregation device (AD). [PR1298747](#)
- In Junos OS, two different input filters cannot be configured on the same interface; if two filters are configured, only the second filter (the one that was configured most recently) takes effect. Ingress mirroring on extended ports in Junos Fusion Data Center (JFDC) can only be done by using firewall filters. Considering the Junos OS filter behavior described above, in JFDC, ingress mirroring on extended ports and other firewall filter configurations cannot be done on the same port. [PR1353065](#)
- Since EVPN GR is not supported, restart of rpd will result in considerable traffic loss for EVPN traffic. The traffic should restore eventually once convergence is complete. [PR1353742](#)
- Restart of rpd will result in considerable traffic loss/duplication for EVPN traffic. The traffic should restore eventually once convergence is complete. [PR1350040](#)

SEE ALSO

New and Changed Features	86
Changes in Behavior and Syntax	89
Known Issues	91
Resolved Issues	92
Documentation Updates	93
Migration, Upgrade, and Downgrade Instructions	94
Product Compatibility	101

Known Issues

IN THIS SECTION

- Junos Fusion | 91

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 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

- During Satellite Device (SD) upgrade in Junos Fusion, there is a race condition that the Link Aggregation Control Protocol (LACP) defaulted PDU received from the peer device connected in a link aggregation group (LAG) is incorrectly sent to an aggregation device (AD) with aggregated Ethernet E-channel Identifier (ECID) instead of member ECID, which causes that LACP PDU to be received on another member in the LAG, resulting in the LAG interface flapping. [PR1321575](#)

SEE ALSO

New and Changed Features	86
Changes in Behavior and Syntax	89

[Known Behavior | 90](#)[Resolved Issues | 92](#)[Documentation Updates | 93](#)[Migration, Upgrade, and Downgrade Instructions | 94](#)[Product Compatibility | 101](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3 | 92](#)
- [Resolved Issues: 18.1R2 | 92](#)
- [Resolved Issues: 18.1R1 | 93](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

There are no fixed issues in the Junos OS Release 18.1R3 for Junos Fusion Provider Edge.

Resolved Issues: 18.1R2

Junos Fusion

- Duplicated packets might be received on the multicast downstream devices and multicast receivers. [PR1316499](#)
- In Junos Fusion, the **show interfaces diagnostics optics satellite** command does not display any outputs. [PR1327876](#)
- In a Junos Fusion setup, an aggregate device might show a plus sign (+) on the ICL link for a satellite device. [PR1335373](#)

- SSH key-based authentication does not work with Junos Fusion. [PR1344392](#)
- AD failure (power off) in a DC fusion is causing complete or partial traffic loss for an extended period. [PR1352167](#)

Resolved Issues: 18.1R1

Junos Fusion

- Chassis alarms are not generated after the uplinks are made down from SD. [PR1275480](#)
- The LAG interface might flap if rebooting aggregation device. [PR1315879](#)
- Duplicated packets might be received on the multicast downstream devices and multicast receivers. [PR1316499](#)
- Not able to disable fpc-slot .Getting error Operation not supported for device assigned slot-id. [PR1321268](#)

SEE ALSO

New and Changed Features 86
Changes in Behavior and Syntax 89
Known Behavior 90
Known Issues 91
Documentation Updates 93
Migration, Upgrade, and Downgrade Instructions 94
Product Compatibility 101

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 94](#)

This section lists the errata and changes in Junos OS Release 18.1R3 for Junos Fusion Provider Edge.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 86](#)

[Changes in Behavior and Syntax | 89](#)

[Known Behavior | 90](#)

[Known Issues | 91](#)

[Resolved Issues | 92](#)

[Migration, Upgrade, and Downgrade Instructions | 94](#)

[Product Compatibility | 101](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 95](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 97](#)

- [Preparing the Switch for Satellite Device Conversion | 98](#)
- [Converting a Satellite Device to a Standalone Device | 99](#)
- [Upgrading an Aggregation Device | 99](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 99](#)
- [Downgrading from Junos OS Release 18.1 | 100](#)

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

The download and installation process for Junos OS Release 18.1R3 is different 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 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.

NOTE: We highly recommend that you see 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

For upgrades from Junos Release 14.2 and earlier:

```
user@host> request system software add no-validate reboot source/package-name
```

All other upgrades:

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


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 18.1R3 **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 only be converted to SNOS 3.1 and higher.
- The switch must be either set to factory-default configuration 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 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 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

In the event that you need to convert a satellite device to a standalone device, you will need to 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](#).

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 18.1R3, 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 18.1

To downgrade from Junos OS Release 18.1 to another supported release, follow the procedure for upgrading, but replace the 18.1 **jinstall** package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

SEE ALSO

[New and Changed Features | 86](#)

[Changes in Behavior and Syntax | 89](#)

[Known Behavior | 90](#)

[Known Issues | 91](#)

[Resolved Issues | 92](#)

[Documentation Updates | 93](#)

[Product Compatibility | 101](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 101](#)

Hardware Compatibility

Hardware Compatibility

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

To determine the features supported in a Junos Fusion, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. See the [Feature Explorer](#).

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

New and Changed Features 86
Changes in Behavior and Syntax 89
Known Behavior 90
Known Issues 91
Resolved Issues 92
Documentation Updates 93
Migration, Upgrade, and Downgrade Instructions 94

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

IN THIS SECTION

- New and Changed Features | 102
- Changes in Behavior and Syntax | 123
- Known Behavior | 132
- Known Issues | 137
- Resolved Issues | 153
- Documentation Updates | 189
- Migration, Upgrade, and Downgrade Instructions | 190
- Product Compatibility | 197

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- Release 18.1R3-S3 New and Changed Features | 103
- Release 18.1R3 New and Changed Features | 104
- Release 18.1R2 New and Changed Features | 104
- Release 18.1R1 New and Changed Features | 104

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for MX Series.

Release 18.1R3-S3 New and Changed Features

EVPNs

- **Support for VMTO for ingress traffic (MX Series and vMX)**—Starting in Junos OS Release 18.1R3-S3, you can configure the PE device to support virtual machine traffic optimization (VMTO) for ingress traffic. VMTO eliminates the unnecessary ingress routing to default gateways when a virtual machine is moved from one data center to another.

To enable VMTO, configure **remote-ip-host** routes in the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. You can also filter out the unwanted routes by configuring an import policy under the **remote-ip-host routes** option.

[See [Configuring EVPN Routing Instances](#).]

- **Support for Multihomed Proxy Advertisement (MX Series and vMX)**—Junos now provides enhanced support to proxy advertise the Mac address and IP route entry from all PEs that are multi-homed to a CE device. This can prevent traffic loss when one of the links to the PE fails. To support the multihomed proxy advertisement, all multi-homed PE devices should have the same multihomed proxy advertisement bit value. The multihomed proxy advertisement feature is enabled by default and Junos uses the default multihomed proxy advertisement bit value of 0x20.

[See [EVPN Multihoming Overview](#).]

- **Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks (MX Series and vMX)**—Starting in Junos OS Release 18.1R3-S3, you can configure OSPF, IS-IS, BGP, and static routing with bidirectional forwarding detection (BFD) on an IRB interface that is used as a routed interface in EVPN. This allows protocol adjacencies to be established between an IRB on a Layer 3 gateway and a CE device connected directly to a Layer 3 gateway or to a Layer 2 leaf device in an EVPN-VXLAN network.

[See [Supported Protocols on an IRB Interface in EVPN-VXLAN](#) .]

Release 18.1R3 New and Changed Features

Interfaces and Chassis

- **Enhanced fault management features**—Starting with Junos OS Release 18.1R3, MX Series routers support configuration of error thresholds and actions at the error scope and error category levels. Use the command **set chassis fpc fpc-slot error scope error-scope category category (fatal | major | minor) threshold error-threshold action (alarm | disable-pfe | get-state | offline | log | reset)** to configure a threshold and action for a particular error scope and category at the FPC level. You can also configure these features at the chassis level (at the **[edit chassis]** hierarchy). You can use the command **show chassis fpc errors** to view the error information at the error scope and category level.

Release 18.1R2 New and Changed Features

Class of Service (CoS)

- **Hierarchical CoS support for anchor point redundancy of pseudowire subscriber logical Interfaces (MX Series)**—Starting in Junos OS Release 18.1R2, full hierarchical CoS support is provided for stateful anchor point redundancy of pseudowire subscriber logical interfaces. Both transport and services logical interfaces created for the pseudowire subscriber logical interface are stacked on the underlying redundant logical tunnel control logical interface. This logical interface stacking model is used for both redundant and non-redundant pseudowire subscriber logical interfaces. Hierarchical CoS is supported and configured the same on both redundant and non-redundant pseudowire subscriber logical interfaces.

[See [CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces](#).]

Restoration Procedures Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (MX Series)**—In Junos OS Release 18.1R2, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode. The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

[See [Saving a Rescue Configuration File](#).]

Software Installation and Upgrade

- **ZTP support is added for MX VM host platforms (MX Series)**—In Junos OS Release 18.1R2, ZTP, which automates the provisioning of the device configuration and software image with minimal manual intervention, is supported on MX Series VM hosts. When you physically connect a supported device to the network and boot it with a factory configuration, the device attempts to upgrade the Junos OS software image automatically and autoinstall a configuration provided on the DHCP server.

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

Release 18.1R1 New and Changed Features

Authentication, Authorization, and Accounting (AAA) (RADIUS)

- **TACACS+ authorization for operational commands using regular expressions (MX Series)**—Starting in Junos OS Release 18.1R1, you can configure authorizations for operational mode commands using regular expressions using the **allow-commands-regexps** and **deny-commands-regexps** statements. Authorizations can also be configured remotely by specifying Juniper Networks vendor-specific attributes (VSAs) in your TACACS+ authentication server's configuration.

[See [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies](#).]

Class of Service (CoS)

- **Support for policer actions before ingress queuing (MX Series)**—Starting with Junos OS Release 18.1R1, on MPCs that support ingress queuing, you can implement policer actions on traffic before the traffic is assigned to ingress queues. To do this, create the desired policers, apply them to a standard firewall filter, and attach the filter as an ingress queuing policing filter [**iq-policing-filter filter-name**] to an interface at the [**edit interfaces interface-name unit logical-unit-number family family**] hierarchy level. The **iq-policing-filter** can only be attached to a static interface.

[See [Ingress Queuing Filter with Policing Functionality](#).]

- **Support for rewrite of the first three bits of IPv6 DSCP value (MX Series, vMX)**—Starting with Junos OS Release 18.1R1, MX Series routers with MPCs support rewrite rules that rewrite only the first three bits of the IPv6 DSCP value. Junos OS provides a new rewrite rule option, **inet6-precedence**, at the [**edit class-of-service rewrite-rules**] hierarchy level that allows you to set a 3-bit code point for a particular forwarding class and loss priority for IPv6 traffic. This new rewrite rule option can also be applied to packets entering an MPLS LSP.

[See [inet6-precedence \(CoS Rewrite Rules\)](#).]

Dynamic Host Configuration Protocol (DHCP)

- **DHCP support for management interface in non-default RI (MX Series)**—Starting in Junos OS Release 18.1R1, DHCPv4 and DHCPv6 clients are supported on management interfaces (fxp0 and em0) configured in the non-default management routing instance, **mgmt_junos**.

[See [Configuring a DHCP Client](#).]

EVPNs

- **Connectivity Fault Management Support in an EVPN network (MX Series)**—Starting with Junos OS Release 18.1R1, Junos OS supports connectivity fault management (CFM) Up maintenance association endpoints (MEPs) on the attachment circuits (ACs) that are connected to a provider edge (PE) router in an EVPN network. You can configure up MEPs to monitor multiple attachment circuits on the same PE router as part of the same maintenance domain or maintenance association.

To configure multiple Up MEPs, specify the **mepmep-id** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance association ma-name]** hierarchy level, with the MEP direction configured as **direction up**.

[See [Connectivity Fault Management Support for Layer 2 VPN](#).]

- **Support for ESI per logical interface in PBB-EVPN (MX Series Routers with MPC and MIC interfaces)**—Starting in Junos OS Release 18.1R1, in a Provider Backbone Bridging (PBB) and Ethernet VPN (EVPN) integrated network, the Ethernet segments of a multihomed Provider edge (PE) device (single-active and active-active redundancy modes) can be uniquely configured per logical interface by assigning the Ethernet segment identifier (ESI) value to the logical interface. In addition to this, the backbone MAC (B-MAC) address and redundancy mode can also be configured per ESI.

In earlier releases, the ESI of the multihomed devices were supported only on physical interfaces. With this feature, all the existing functionalities of ESI per physical interface is extended to ESI per logical interface.

[See [Provider Backbone Bridging \(PBB\) and EVPN Integration Overview](#).]

General Routing

- **Support for PTP over Ethernet encapsulation and G.8275.1 profile (MX10003 and MX204)**—Starting in Junos OS Release 18.1R1, MX10003 and MX204 routers support the following features:
 - **PTP over Ethernet**—PTP over Ethernet enables effective implementation of packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks. PTP over Ethernet uses multicast addresses for communication of PTP messages between the slave clock and the master clock.
 - **G.8275.1 profile**—G.8275.1 is a PTP profile for applications that require accurate phase and time synchronization. It supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support and is based on the second version of PTP defined in IEEE 1588. You can configure the G.8275.1 profile by including the **profile-type g.8275.1** statement at the **[edit protocols ptp]** hierarchy level.

[See [Configuring G.8275.1 Profile](#).]

- **VRF support for NTP (MX Series)**—Starting in Junos OS Release 18.1R1, NTP clients can send requests to servers that are reachable through VRF. The **set system ntp server address routing-instance routing-instance-name** and **set date ntp routing-instance routing-instance-name** commands let you specify the routing instance that the server can be reached through.

[See [Configuring the NTP Time Server and Time Services](#).]

- **Support for PTP, Synchronous Ethernet, and hybrid mode over link aggregation group (MX240, MX480, MX960, MX2010, MX2020)**—Starting in Junos OS Release 18.1R1, the MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E MPCs support Precision Time Protocol (PTP), Synchronous Ethernet, and hybrid mode over a link aggregation group (LAG).

Link aggregation is a mechanism of combining multiple physical links into a single virtual link to achieve linear increase in bandwidth and to provide redundancy in case a link fails. The virtual link is referred to as an aggregated Ethernet interface or a LAG.

[See [Precision Time Protocol Overview](#).]

High Availability and Resiliency

- **MX Series Virtual Chassis Unified ISSU support for MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E line cards (MX Series Virtual Chassis)**—Starting in Junos OS Release 18.1R1, MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E line cards support Unified ISSU in MX Series Virtual Chassis environments. Unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

[See [Unified ISSU in a Virtual Chassis](#).]

Interfaces and Chassis

- **New speed configuration option introduced to change 10-Gbps port to operate in 1-Gbps speed (MX204, MX10003)**—Starting in Junos OS Release 18.1R1, the 10-Gbps port can operate in 1-Gbps mode on MX204 and MX10003 routers. Currently, MX204 and MX10003 routers support different operation modes; that is, 10-Gbps, 40-Gbps, and 100-Gbps speed. When the port is operating in 10-Gbps speed, you can change the operating speed to 1Gbps using a new CLI option, [speed 1g/10g](#) at the **[edit interfaces intf-name together-options]** hierarchy level. Once you commit this configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed without any FPC, PIC, or interface bounce.

The MX10003 MPC has one fixed PIC and one MIC (non-MACsec MIC/MACsec MIC). The fixed PIC has 6 ports that can operate in 40-Gbps or 4X10-Gbps mode. The MIC has 12 ports that can operate in 100-Gbps, 40-Gbps, or 4X10-Gbps mode. With this new speed configuration option, you can configure the 4X10-Gbps port on the fixed PIC and the non-MACsec MIC to 1-Gbps mode. You can also configure one or all ports that operate in 10-Gbps mode to 1Gbps mode.

The MX204 contains two PICs—where one PIC contains 8 ports that can operate in 10-Gbps mode and the other PIC contains 4 ports that can operate in 4X10-Gbps, 40-Gbps, or 100-Gbps mode. Using this new speed configuration option, you can configure the 4X10-Gbps port on one of the fixed-port PICs to operate in 1-Gbps mode. And on the other fixed-port PIC, you can configure the 10-Gbps port to 1Gbps.

NOTE:

- On the MX10003 router, the MACsec MIC does not provide 1-Gbps speed. If you attempt to change the operating speed to 1-Gbps, syslog displays that this feature is not supported on the MACsec MIC.
- On MX204 and MX10003 routers, rate selectability at PIC level and port level does not support 1-Gbps speed.
- On MX204 and MX10003 routers, 1-Gbps operation mode is only supported in no-autonegotiation mode.

To view the speed configured for the interface, execute the **show interfaces extensive** command. The **Speed Configuration** output parameter in the command output indicates the current operation speed of the interface.

If the interface is configured with 1-Gbps speed, then **Speed Configuration** displays **1G**; if the interface is configured with 10-Gbps speed, **Speed Configuration** displays **AUTO**.

For example:

```
user@host>show interfaces xe-0/1/11:0 extensive
```

```
Physical interface: xe-0/1/11:0, Enabled, Physical link is Up
Interface index: 284, SNMP ifIndex: 609, Generation: 383
Link-level type: Ethernet, MTU: 9192, MRU: 9200, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
Loopback: None, Source filtering: Disabled, Flow control: Enabled,
Speed Configuration: 1G
...
```

In this example, the **Speed Configuration** output parameter displays **1G**, which means the operation speed of xe-0/1/11:0 interface is 1-Gbps speed.

NOTE:

- The interface name prefix must be xe.
- To set a port that is operating in 10-Gbps speed to 1-Gbps speed, use the new CLI option [speed 1g/10g](#) for the existing **set interfaces [intf-name] gigether-options** command.
- To view the speed configured for the interface, execute the **show interfaces extensive** command.

[See [MX10003 MPC Rate-Selectability Overview](#) and [MX204 Router Rate-Selectability Overview](#).]

- **Upgraded SSD size and RAM size (MX Series)**—Starting in Junos OS Release 18.1R1, the Routing Engines on the MX240, MX480, MX960, MX2010, MX2020 routers support Secure Boot BIOS. The SSD size and the RAM size of the following Routing Engines are upgraded to 2x200-GB and 128-GB respectively:

- RE-S-X6-128G-S on the MX240, MX480, and MX960 routers
- RE-MX2K-X8-128G-S on the MX2010 and MX2020 routers

[See [Salient Features of the Routing Engines with VM Host Support.](#)]

- **Limited encryption Junos OS image and boot restriction (MX10003)**—Starting with Junos OS Release 18.1R1, MX10003 router with LT-SKU supports only Junos Limited image. The Junos Limited image does not have data-plane encryption and is intended only for countries in the Eurasian Customs Union because these countries have import restrictions on software containing data plane encryption. Unlike the Junos Worldwide image, the Junos Limited image supports control plane encryption through Secure Shell (SSH) and Secure Sockets Layer (SSL), thus allowing secure management of the system. The MX10003 LT SKU boots only the encryption free Junos software and fails to boot if the fully encrypted Junos software is used for booting. The Junos upgrade and VMHost upgrade using non-limited version of Junos software fails on the MX10003 LT SKU. The command **show chassis hardware [models | clei-models | extensive]** displays the model number and helps identifying the different SKUs. An alarm, **Mixed Master and Backup RE types** is displayed when dissimilar Routing Engines are present on the chassis.

[See [Junos OS Editions.](#)]

- **Enhanced support for the non-default management instance mgmt_junos (MX Series)**—Starting in Junos OS Release 18.1R1, syslog IPv6 addresses, RADIUS packets, and Automation scripts support the non-default management instance mgmt_junos, when the **management-instance** statement is configured. For syslog, statements at the **[edit system syslog]** hierarchy level now support IPv6 addresses when connecting to a remote host or an archival site. RADIUS authentication, authorization, and accounting packets can be configured to use the mgmt_junos instance. Also, Automation (commit, event, JET, op, or SNMP) scripts now can be refreshed over the mgmt_junos instance. To enable the non-default VRF management instance, you must also configure the mgmt_junos routing instance at the **[edit system routing-instances]** hierarchy level.

[See [Management Interface in a Non-Default Instance.](#)]

- **Enhancement to increase the threshold of corrected single-bit errors (MPC7E, MPC8E, MPC9E on MX Series)**—In Junos OS Release 18.1R1, the threshold of corrected single-bit error is increased from 32 to 1024, and the alarm severity is changed from Major to Minor for those error messages. There is no operational impact upon corrected single bit errors. Also, a log message is added to display how many single-bit errors have been corrected between the reported events as follows:

EA[0:0]: HMCIF Rx: Link0: Corrected single bit error detected in HMC 0 - Total count 25

EA[0:0]: HMCIF Rx: Link0: Corrected single bit error detected in HMC 0 - Total count 26

[See [Alarm Overview.](#)]

IPv6

- **IPv6 packet (pps) and byte (bps) rates included in interface traffic statistics (MX series)**—Starting in Junos OS Release 18.1R1, the output of the following commands are modified:
 - The **show interfaces** command displays the input and output bytes (bps) and packets (pps) rates individually for IPv6 family in the IPv6 interface traffic statistics.
 - The **monitor interface** command displays the IPv6 interface traffic statistics along with input and output bytes (bps) and packets (pps) rates individually for IPv6 family.

[See [show interfaces](#) and [monitor interface](#).]

Junos OS XML, API and Scripting

- **Automation script library additions and upgrades (MX240, MX480, MX960, and vMX routers)**—Starting in Junos OS Release 18.1R1, devices running Junos OS that support Python automation scripts include new and upgraded Python modules. Python automation scripts can leverage new on-box Python modules, including **appdirs**, **asn1crypto**, **cffi**, **cryptography**, **idna**, **libffi**, **packaging**, **psutil**, **pyasn1**, **pyparser**, and **pyparsing**, as well as upgraded versions of existing modules. The **psutil** module is available only on devices running Junos OS with upgraded FreeBSD, and only a subset of functions is supported.

[See [Overview of Python Modules Available on Devices Running Junos OS](#).]

Management

- **Expanded support for chassis sensors for Junos Telemetry Interface (MX Series Transport Series Routers)**—Starting with Junos OS Release 18.1R1, Junos Telemetry Interface (JTI) provides new sensors that expand optics and power information.

To export telemetry data from Juniper equipment to an external collector requires both Junos Telemetry Interface (JTI) and gRPC to be configured.

Enhanced sensor information is also supported through operational mode commands **show chassis fpc detail**, **show chassis power detail**, and **show chassis pic fpc-slot id pic-slot id**.

Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

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

- **“ON CHANGE” sensor support through gRPC Network Management Interface (gNMI) for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 18.1R1, ON_CHANGE streaming of Address Resolution Protocol (ARP), Network Discovery Protocol (NDP), and IP sensor information associated with interfaces is supported on Junos Telemetry Interface (JTI).

Periodical streaming of OpenConfig operational states and counters has been supported since Junos OS Release 16.1, exporting telemetry data from Juniper equipment to an external collector. While useful in collecting all the needed information and creating a baseline “snapshot,” periodical streaming is less useful for time-critical missions. In such instances, you can configure ON_CHANGE streaming for an external collector to receive information only when operational states experience a change in state.

To support ON_CHANGE streaming, Google has developed a new specification called gRPC Network Management Interface (gNMI) for the modification and retrieval of configurations from a network element. Additionally, the gNMI specification can be used to generate and control telemetry streams from a network element to a data collection system. Using the new gNMI specification, one gRPC service definition can provide a single implementation on a network element for both configuration and telemetry as well as a single NMS element to interact with a device by means of telemetry and configuration RPCs.

Information about the RPCs supporting this feature can be found in the gNMI Proto file version 0.4.0 (the supported version) and the specification released by Google at:

- <https://github.com/openconfig/reference/blob/master/rpc/gnmi/gnmi-specification.md>
- <https://github.com/openconfig/gnmi/blob/master/proto/gnmi/gnmi.proto>

The telemetry RPC **subscribe** under gNMI service supports ON_CHANGE streaming. RPC **subscribe** allows a client to request the target to send it values of particular paths within the data tree. Values may be streamed (STREAM), sent one-off on a long-lived channel (POLL), or sent one-off as a retrieval (ONCE).

If a subscription is made for a top level container with a sample frequency of 0, leaves with ON_CHANGE support are streamed based on events. Other leaves will not be streamed.

NOTE: In order to permit a device to decide which nodes will be streamed as ON_CHANGE and which will SAMPLE, the collector should subscribe for TARGET_DEFINED with sample_interval.

Streaming telemetry data through gRPC requires you to download the OpenConfig for Junos OS module.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Junos Events Sensor for the Junos Telemetry Interface (JTI) (MX240, MX480, MX960, MX2010, MX2020 with MPC1, MPC2, MPC3, MPC4, MPC5, MPC6, MPC7, MPC8, or MPC9)**—Starting in Junos OS Release 18.1R1, the Junos events sensor is available for streaming system event data through JTI. Previously, only interval-based statistical sensors were available for use with JTI. With the Junos events sensor, system events that are available through system logging (syslog) can now be streamed to telemetry collection systems, allowing more data to be streamed and collected in one location. This helps to give a better picture of overall system health through one interface.

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

- **PIC services and IPSec sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 18.1R1, Junos Telemetry Interface (JTI) provides support for gRPC-based IKE and GPB UDP-based PIC sensors. These sensors provide visibility for IPSec services on different service complexes and nodes.

Exported data is defined using an IP address and a UDP port. When an export interval expires, the most recent statistics collected by the sensors are gathered, placed in the payload of a UDP packet, and forwarded to a collector. A timestamp indicating when counters are read is included with the exported

data to allow collectors to collate data. The timestamp also can determine if and when an event happened, such as a PIC hardware restart or if counters were cleared by means of the CLI.

The resource paths are:

- `/junos/services/spu/ipsec-vpn`
- `/junos/ike-security-associations/ike-security-association/`

To export telemetry data from Juniper equipment to an external collector requires both Junos Telemetry Interface (JTI) and gRPC to be configured.

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

- **Fabric statistics support on Junos Telemetry Interface (JTI) (MX Series)**—Starting with Junos OS Release 18.1R1, fabric statistics limited to streaming over GBP over UDP are now supported for export by means of gRPC. Statistics are exported whether encoded as native or as a third-party data model.

Fabric statistic data is collected and exported by the following two types of fabric sensors:

- Per Packet Forwarding Engine pair fabric sensor
- Summary Flexible Pic Concentrator (FPC) fabric sensor

Streaming telemetry data through gRPC requires you to download the OpenConfig for Junos OS module.

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

- **ON_CHANGE support for Junos Telemetry Interface (JTI) (MX Series)**—Starting with Junos OS Release 18.1R1, OpenConfig support through gRPC Remote Procedure Calls (gRPC) and JTI is extended to support client streaming and bidirectional streaming of telemetry sensor information.

APIs have been implemented in Junos based on Protobuf specifications released by Google for OpenConfig. These APIs perform configuration, operational state retrieval, and telemetry on Junos routers using gRPC as the transport mechanism.

Starting in Junos OS 18.1R1, client streaming and bidirectional streaming are supported. With client streaming, the client sends a stream of requests to the server instead of a single request. The server typically sends back a single response containing status details and optional trailing metadata. With bidirectional streaming, both client and server send a stream of requests and responses. The client starts the operation by invoking the RPC and the server receives the client metadata, method name, and deadline. The server can choose to send back its initial metadata or wait for the client to start sending requests. The client and server can read and write in any order. The streams operate completely independently.

Junos devices can be managed through API (RPC) prototypes:

- `rpc Capabilities (CapabilityRequest)`

Returns `(CapabilityResponse)`. Allows the client to retrieve the set of capabilities that is supported by the target.

- `rpc Get (GetRequest)`

Returns (GetResponse). Retrieves a snapshot of data from the target.

- **rpc Set (SetRequest)**

Returns (SetResponse). Allows the client to modify the state of data on the target.

- **rpc Subscribe (stream SubscribeRequest)**

Returns (stream SubscribeResponse). Allows a client to request the target to send it values for particular paths within the data tree. These values may be streamed (STREAM) or sent one-off on a long-lived channel (POLL), or sent as a one-off retrieval (ONCE). If a subscription is made for a top-level container with a sample frequency of 0, leaves with ON_CHANGE support are streamed based on events. Other leaves will not be streamed.

Juniper Extension Toolkit (JET) support provides insight to users regarding the status of clients connected to JSD. JET support for gRPC includes expanding the maximum number of clients that can connect to JSD from 8 to 30 (the default remains 5). To specify the maximum number of connections, include the **max-connections** statement at the **[edit system services extension-service request-response grpc]** hierarchy level.

To provide information regarding the status of clients connected to JSD, issue the enhanced **show extension-service client information** command and include the **clients** or **servers** options. The **clients** option displays request-response client information. The **servers** option displays request-response server information.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

MPLS

- **RSVP-TE pop-and-forward LSP tunnels (MX Series routers with MPCs and MICs)**—Pop-and-forward LSPs introduce the notion of pre-installed per traffic engineering link pop labels that are shared by RSVP-TE LSPs that traverse these links. A transit label-switching router (LSR) allocates a unique pop label per traffic engineering link with a forwarding action to pop the label and forward the packet over that traffic engineering link should the label appear at the top of the packet. Starting in Junos OS Release 18.1R1, you can configure pop-and-forward LSPs to significantly reduce the required forwarding plane state, enabling the pop-and-forward tunnels to couple the feature benefits of the RSVP-TE control plane with the simplicity of the shared MPLS forwarding plane.

All the existing RSVP-TE functionalities, such as bandwidth admission control, LSP priorities, preemption, auto-bandwidth, and MPLS fast reroute continue to work with pop-and-forward tunnels.

[See [RSVP-TE Pop-and-Forward LSP Tunnels Overview](#).]

- **Localization of next-hop-based dynamic tunnels**—(MX Series) Next-hop-based dynamic generic routing encapsulation (GRE) tunnels and MPLS-over-UDP tunnels distribute forwarding information to all line cards on a device. As a result, the origination and termination states of all tunnels are built on the Packet Forwarding Engines (PFEs) on every line card on the device, limiting the maximum number of tunnels supported on the device to the tunnel capacity of a single line card. Starting in Junos OS Release 18.1R1, you can configure next-hop-based dynamic tunnel localization to create the forwarding information only

on the PFE of a line card that is designated as the anchor PFE. The PFEs on the other line cards on the device have state forwarding information to steer the packets to the anchor PFE.

[See [Next-Hop-Based Dynamic Tunnel Localization Overview](#).]

- **Support for static segment routing label switched path (MX Series)**—Starting with Junos OS 18.1R1 release, a set of explicit segment routing paths are configured on the ingress router of a non-colored static segment routing label switched path (LSP) by configuring the **segment-list** statement at the **[edit protocols source-packet-routing]** hierarchy level. You can configure the segment routing LSP by configuring the **source-routing-path** statement at **[edit protocols source-packet-routing]** hierarchy level. The segment routing LSP has a destination address and one or more primary paths and optionally secondary paths that refer to the segment list. Each segment list consists of a sequence of hops. For non-colored static segment routing LSP, the first hop of the segment list specifies an immediate next hop IP address and the second to Nth hop specifies the segment identifies (SID) labels corresponding to the link or node which the path traverses. The route to the destination of the segment routing LSP is installed in inet.3 table. The adjacency segments, node segments, and prefix segments can be provisioned on transit routers by configuring static MPLS segment LSPs at the **[protocols mpls static-label-switched-path]** hierarchy level

[See [static segment routing lsp](#).]

Multicast

- **Translation for MVPN Type 5 routes to MSDP SA (MX Series)**—Starting in Junos OS Release 18.1R1, Junos supports MVPN-Type-5 route to MSDP-SA conversion as defined in RFC [draft-ietf-bess-mvpn-sa-to-msdp-00.txt](#). Previously, Junos only supported translation in the other direction, MSDP SA to MVPN Type 5.

The ability to convert next-generation multicast virtual private network (MVPN) Type 5 routes to Multicast Source Discovery Protocol (MSDP) source active (SA) makes it possible to reduce the number of MSDP sessions running between VPN customer rendezvous points (C-RPs). For example, instead of having MSDP running among all C-RPs in a deployment, the C-RPs could instead run their MSDP sessions with a single PE router configured for multiple MSDP peers. The PE router, now acting as a C-RP device, would receive MVPN SA Type 5 routes from the RP-PE or source PE router, convert those routes to MSDP, and then advertise the MSDP routes to its MSDP peers.

MVPN Type 5 SA routes are added to MVPN table and include a new Extended Community (EC), with the IPv4 address of the RP where the MVPN SA was generated. The Type 5 routes source and EC are additionally added to the MSDP table. Stale routes, including the EC, are removed via MSDP once the MVPN type 5 SA route is gone from the MVPN table.

Enable MVPN to MSDP conversion at the **[edit routing-instance *name* protocols mvpn mvpn-mode spt-only convert-sa-to-msdp]** hierarchy level.

You can verify whether MVPN type 5 routes are being correctly converted to MSDP SA by running the **[show msdp source-active instance *name*]** command.

[See [MVPN Concepts and Protocols](#).]

Network Management and Monitoring

- **sFlow support on MX Series devices**—Starting in Junos OS Release 18.1R1, you can configure sFlow technology (as a sFlow agent) on a MX Series device, to continuously monitor traffic at wire speed on all interfaces simultaneously. The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a remote monitoring station, which presents quantifiably accurate network traffic visibility information after collecting data for a reasonably long period. These remote monitoring stations are called collectors.

[See the [Understanding How to Use sFlow Technology for Network Monitoring on a MX Series Router](#).]

- **Resource monitor support for PS and RLT interfaces (MX Series)**—Starting in Junos OS Release 18.1R1, PS and RLT interfaces support resource monitoring throttling. If a configured resource limit is exceeded for any member of a PS or RLT interface, the resource monitor prevents subscriber login and increments a denied counter. The denied counters can be verified with the **show system resource-monitor summary** command. In addition, the **show subscribers** command displays subscribers per PIC/PFE/Slot for PS and RLT interfaces.

[See [resource-monitor](#).]

- **The bbe-mibd component is enhanced with additional MIB objects (MX Series)**—In the next-generation broadband edge architecture, subscribers are represented by flows instead of logical interfaces. The SNMP subagent bbe-mibd was implemented to handle SNMP requests for subscriber interfaces. As of Junos OS Release 18.1R1, the following MIB objects are made available for flow-based dynamic interfaces as part of bbe-mibd:

- ifChassisTable
- ipv6IfTable
- ipv6IfStatsTable
- jnxIpv6

[See the [SNMP MIB Explorer](#).]

- **Enhancement to Junos OS SNMP MIB PCC functionality (MX Series)**—Starting in Junos OS Release 18.1R1, Junos OS provides enhanced MIB support for Path Computation Clients. This enhancement enables the Path Computation Client (PCC) process to accept SNMP **get** and **getnext** commands for Path Computation Client Protocol (PCEP) peer and PCEP session tables and reply to them. This feature monitors PCEP interactions between a PCC and a Path Computation Element (PCE). Not all members of PCEP peer and PCEP session tables mentioned in the RFC (RFC 7420) are supported. For exceptions, see [Standard SNMP MIBs Supported by Junos OS](#).

[See [MIB Explorer](#). Name of MIB is **pcep.mib**.]

Operation, Administration, and Maintenance (OAM)

- **CFM Action Profile to Bring Down a Group of Logical Interfaces (MX Series Routers)**—Starting with Junos OS Release 18.1R1, you can create a CFM Action Profile and define an action to bring down a group of logical interfaces using CFM session configured on a single IFL. Following new configuration statements are introduced:
 - To mark the interface group down configure **interface-group-down** at the **[edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name*** hierarchy level.
 - To mark the interface group down for the action profile configured with the action **interface-group-down**, configure **interface-group(*interface-device-name* | *unit-list*)** at the **[edit protocols ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id* action-profile *profile-name*** hierarchy level. The ***interface-device-name*** represents an Ethernet interface device. The ***unit-list*** defines a string of logical unit numbers.

[See [Ethernet OAM Connectivity Fault Management](#).]

Routing Policy and Firewall Filters

- **Firewall filters and policers for abstracted fabric interface (MX Series)**—Starting with Junos OS Release 18.1R1, you can configure firewall filters and policers on an abstracted fabric (AF) interface, a pseudointerface that facilitates routing control and management traffic between guest network functions (GNFs) through the switch fabric. AF interfaces support single-rate two-color policer, single-rate three-color policer, two-rate three-color policer, and hierarchical policer. The AF interface firewall filters are supported on Inet, Inet6, MPLS, and CCC protocol families.

NOTE: The AF interface bandwidth is assigned to all FPCs linked to that AF interface. Therefore, a policer bandwidth limit configuration on an AF interface is applicable to all the PFEs associated with the AF interface.

[See [Understanding the Use of Policers in Firewall Filters](#).]

Routing Protocols

- **Support for BGP multipath at global level (MX Series)**—Starting with Junos OS Release 18.1R1, BGP multipath is available at the global level in addition to the group and neighbor level. In earlier Junos OS releases BGP multipath is supported only at the group and neighbor levels. A new configuration option **disable** is available at the **[edit protocols bgp multipath]** hierarchy level to disable BGP multipath for specific groups or neighbors. This allows you to configure BGP multipath globally and disable it for specific groups according to your network requirements.

[See [disable](#).]

- **Support for BGP Labeled Unicast traffic statistics collection (MX Series)** —Starting in Junos OS Release 18.1R1, you can enable traffic statistics collection for BGP labeled unicast traffic at the ingress router. In a network configured with segment routing, traffic statistics can be collected periodically based on the label stack received in the BGP route update and saved in a specified file. Traffic statistics collection is supported only for IPv4 and IPv6 address families.

[See [Enabling Traffic Statistics Collection for BGP Labeled Unicast](#).]

- **Multipath optimization to improve RIB learning rate**—Starting in Junos OS Release 18.1R1, you can defer multipath calculation until all BGP routes are received. When multipath is enabled, BGP inserts the route into the multipath queue each time a new route is added or whenever an existing route changes. When multiple paths are received through BGP add-path feature, BGP might calculate one multipath route multiple times. Multipath calculation slows down the RIB (also known as the routing table) learning rate. To speed up RIB learning, multipath calculation can be either deferred until the BGP routes are received or you can lower the priority of the multipath build job as per your requirements until the BGP routes are resolved.

To defer the multipath calculation configure **defer-initial-multipath-build** at **[edit protocols bgp]** hierarchy level. Alternatively, you can lower the BGP multipath build job priority using **multipath-build-priority** configuration statement at **[edit protocols bgp]** hierarchy level to speed up RIB learning.

[See [defer-initial-multipath-build](#).]

Security

- **Secure Boot (MX240, MX480, MX960, MX2010, and MX2020 that use Routing Engines RE-S-X6-128G or RE-MX2K-X8-128G)**—Starting in Junos OS Release 18.1R1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected and thus safeguarded from tampering or modification. Secure boot is enabled by default on supported platforms.

[See [Feature Explorer](#) and enter **Secure Boot**.]

Services Applications

- **Support for additional DS-Lite features on MS-MPCs and MS-MICs (MX Series routers)**—Starting in Junos OS Release 18.1R1, dual-stack lite (DS-Lite) running on MS-MPCs and MS-MICs adds support for the following features:
 - ALGs (TFTP, FTP, DNS, ICMP, RTSP, PPTP)
 - Configurable MTU per software concentrator
 - IPv6 fragmentation and reassembly
 - NAPT-44 port block allocation
 - Receiving and transmitting IPv4 fragments in IPv6
 - Traceroute through the software tunnel
 - Hairpinning with NAPT-44 EIF

Prior to Junos OS Release 18.1R1, DS-Lite did not support these feature on the MS-MPCs and MS-MICs.

[See [Tunneling Services for IPv4-to-IPv6 Transition Overview](#).]

- **Support for IPv6 version 9 templates for inline active flow monitoring (MX Series)**—Starting in Junos OS Release 18.1R1, you can apply version 9 flow templates to IPv6 traffic when using inline flow monitoring. In addition, fields have been added to several IPFIX and version 9 templates for inline flow monitoring to make the templates more uniform for each supported family.

[See [Understanding Inline Active Flow Monitoring](#).]

- **Support for additional filtering on show command output of RPM probes generated on an MS-MPC or MS-MIC (MX Series)**—Starting in Junos OS Release 18.1R1, you can use new filters to limit the output of the **show services rpm probe-results** and **show services rpm history-results** commands for real-time processing (RPM) probes that are generated on an MS-MPC or MS-MIC.

[See [show services rpm probe-results](#) and [show services rpm history-results](#).]

- **Support for generating IPv6 RPM probes on MS-MPCs and MS-MICs (MX Series)**—Starting in Junos OS Release 18.1R1, you can configure an MS-MPC or MS-MIC to generate **icmp6-ping** real-time performance monitoring (RPM) probes. Generating RPM probes on an MS-MPC or MS-MIC increases the number of probes that can run at the same time, compared to probes that are generated on the Packet Forwarding Engine.

[See [Configuring RPM Probes](#).]

Software Defined Networking

- **MS-MIC and MS-MPC support for Junos Node Slicing (MX480, MX960, MX2010, MX2020)**—Starting from Junos OS Release 18.1R1, Junos Node Slicing supports assignment of MS-MICs and MS-MPCs to guest network functions (GNFs). MS-MICs and MS-MPCs provide improved scaling and high performance, and possess enhanced memory (16 GB for MS-MIC; 32 GB per NPU of MS-MPC) and processing capabilities. The MS-MIC supports the Layer 3 services such as stateful firewall, NAT, IPsec, active flow monitoring, RPM, and graceful Routing Engine switchover (GRES).

[See [Multiservices MIC and Multiservices MPC \(MS-MIC and MS-MPC\) Overview](#).]

Subscriber Management and Services

- **Controlling search behavior for address allocation from linked pools (MX Series)**—Starting in Junos OS Release 18.1R1, you can use the **linked-pool-aggregation** statement at the **[edit access]** hierarchy level to change how addresses are allocated from linked IP address pools. When you configure the statement, addresses can be assigned from a later pool in the chain before an earlier pool is depleted. When the statement is not configured, IP addresses are assigned contiguously, so that all addresses are allocated from the matching pool and then the first pool in the chain before addresses are assigned from a linked pool.

[See [Configuring Address-Assignment Pool Linking](#).]

- **Support for Packet triggered subscriber functionality (MX Series)**—Starting with Junos OS 18.1R1, support for packet triggered subscriber functionality creates IP demux IFL on receiving a data packet from clients with pre-assigned IP address using a new demux configuration at the hierarchy level **[edit interfaces interface-name unit unit-number]**.

[See [IP demultiplexing interfaces on Packet-Triggered Subscribers Services Overview](#).]

- **BPCEF Gy Assume Positive CCR-T File Support (MX Series)**—Starting with Junos OS 18.1R1, broadband PCEF provides the file backup for OCS data when both primary and alternative paths to the OCS are not available. The CCR-GY-T frames are stored in the files on remote location. The backup is supported at the hierarchy **[edit access ocs partition partition-name]**.

[See [Gy File Backup Overview](#).]

- **Support for per-subscriber MTU for dynamic profiles of IP v4 or IPv6 protocol family (MX Series)**—Starting with Junos OS 18.1R1, maximum transmission unit (MTU) can be configured per subscriber for dynamic profiles. The value of MTU can be static or represented through

\$junos-interface-mtu variable. By default, the variable value is the MTU of the payload that is less than the MTU of the physical interface minus the family protocol overhead. A specific value is returned through RADIUS authentication through the framed MTU attribute. If the RADIUS fails to return framed MTU value for \$junos-interface-mtu variable then the default value from **interface-mtu** statement at **[edit dynamic-profiles profile-name predefined-variable-defaults]** hierarchy level. You can configure value for **mtu** statement at **[edit dynamic-profiles name interfaces name unit name family inet]** hierarchy level or at **[edit dynamic-profiles name interfaces name unit name family inet6]** hierarchy level.

[See [Per-subscriber support of maximum transmission unit for dynamic profiles.](#)]

- **Enhancements to dual-stack, single-session authentication and reauthentication (MX Series)**—Starting in Junos OS Release 18.1R1, reauthentication is supported in response to DHCP discover and solicit messages, in addition to the previously supported renew and rebind messages.

Reauthentication is supported for DHCP dual-stack, single-session subscribers when on-demand address allocation is configured. Both authentication and reauthentication for dual-stack, single-session supports are performed per family in the stack, using separate Access-Requests to the RADIUS server.

[See [RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers.](#)]

- **Excluding addresses or ranges to manage address allocation pools for DHCP local server (MX Series)**—Starting in Junos OS Release 18.1R1, you can exclude IPv4 or IPv6 individual addresses or ranges of consecutive addresses within an address pool from being allocated to subscribers. If you exclude an address that has already been allocated, the subscriber is logged out, the address is deallocated, and then marked for exclusion.

[See [Preventing Addresses from Being Allocated from an Address Pool.](#)]

- **Preventing validation of magic numbers in PPP peer-originated keepalive messages (MX Series)**—Starting in Junos OS Release 18.1R1, you can include the **ignore-magic-number-mismatch** statement to disable the Packet Forwarding Engine from validating PPP magic numbers received during PPP keepalive (Echo-Request/Echo-Reply) exchanges. Because validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends a magic number that does not match the number agreed upon during LCP negotiation. This prevents PPP from tearing down the session in the event of a mismatch. This capability is useful when the remote PPP peers include arbitrary magic numbers in the keepalive packets. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.

[See [Preventing the Validation of PPP Magic Number During PPP Keepalive Exchanges](#) and [Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile.](#)]

- **Local dynamic service profile activation on L2TP login (MX Series)**—Starting in Junos OS Release 18.1R1, you can use dynamic service profiles to apply services to all subscribers in a tunnel group or to all subscribers using a particular LAC without involving RADIUS. In multivendor environments, customers might use only standard RADIUS attributes to simplify management by avoiding the use of vendor-specific attributes (VSAs) from multiple vendors. However, VSAs are generally required to apply services. Local service profile activation enables you to avoid that problem. It can also be a way to provide default

services when RADIUS servers are down. You can also pass parameters to the services as they are applied, such as a downstream shaping rate for a CoS service.

[See [Applying Services to an L2TP Session Without Using RADIUS.](#)]

- **Changes to JSRC Provisioning for Dual-Stack Subscribers (MX Series)**—Starting in Junos OS Release 18.1R1, you can include the **dualstack-support** statement at the **[edit jsrc]** hierarchy level to configure JSRC provisioning for dual-stack subscribers so that JSRC reports information about the separate stacks for a given subscriber, using a single JSRC session. Accounting statistics are reported separately for each family. In earlier releases, the remote SRC peer is not informed about whether only on family or both families are active, and all statistics are aggregated across the families.

[See [JSRC Provisioning for Dual-Stack Subscribers.](#)]

- **Providing L2TP service rate limits at subscriber login (MX Series)**—Starting in Junos OS Release 18.1R1, you can specify the name of the dynamic service profile that provides values for the transmit (Tx) and receive (Rx) connect speeds for traffic between the LAC and the subscriber. When that profile name is returned in the Juniper Networks Activate-Service VSA (26-65) in the RADIUS Access-Accept message at subscriber login, the values are converted from Kbps to bps and stored in the session database. You can also modify the rates with parameters passed in the VSA or specify an adjustment for the values, up or down, in the CLI. The rates are sent in the ICCN message from the LAC to the LNS as AVP 24 and AVP 38.

[See [Specifying a Rate-Limiting Service Profile for L2TP Connection Speeds.](#)]

- **Flexible filtering of RADIUS attributes and VSAs (MX Series)**—Starting in Junos OS Release 18.1R1, a flexible configuration method is supported for filtering undesirable RADIUS standard attributes and vendor-specific attributes (VSAs). Attributes received from RADIUS take precedence over internally provisioned attributes; filtering ensures that the corresponding internally provisioned attribute values are used.

The flexible configuration enables you to specify RADIUS standard attributes with the attribute number, and to specify VSAs with the IANA-assigned vendor ID and the VSA number. Some attributes can be ignored when received in Access-Accept messages; other attributes can be excluded from Access-Request and accounting messages. Only those standard attributes and VSAs supported by your platform can be filtered. You can configure unsupported standard attributes, vendors, and VSAs, but the configuration has no effect.

In earlier releases, you must specify dedicated keywords (options) for attributes to filter. This method is still supported. If you configure filtering with both methods, attributes that are specified with either method are filtered.

[See [Enabling the Use of Local Values by Filtering RADIUS Attributes and VSAs.](#)]

- **Additional RADIUS attributes and VSAs supported for filtering (MX Series)**—Starting in Junos OS Release 18.1R1, you can filter the following attributes from RADIUS Access-Accept messages with the **ignore** statement:

Standard RADIUS attributes:

- Session-Timeout (27)
- Idle-Timeout (28)

Microsoft (IANA vendor-id 311) vendor specific attributes:

- MS-Primary-DNS-Server (26-28)
- MS-Secondary-DNS-Server (26-29)

[See [ignore](#).]

User Interface and Configuration

- **Ephemeral configuration database support for load replace and load override operations (MX Series)**—Starting in Junos OS Release 18.1R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database using **load replace** and **load override** operations, in addition to the previously supported **load merge** and **load set** operations. To perform a **load replace** or **load override** operation, set the **<load-configuration> action** attribute to **replace** or **override**, respectively.

[See [Configuring Ephemeral Database Instances](#).]

VPNs

- **Support for seamless migration from BGP-VPLS to EVPN (MX Series)**—Starting in Junos OS Release 18.1R1, a solution is introduced for enabling staged migration from BGP-VPLS toward EVPN on a site-by-site basis for every VPN routing instance. In this solution, the PE devices running EVPN and VPLS for the same VPN routing instance and single-homed segments can co-exist. The solution supports single-active redundancy of multi-homed networks and multi-homed devices for EVPN PEs. With single-active redundancy, the participant VPN instances may span across both EVPN PEs and VPLS PEs as long as single-active redundancy is employed by EVPN PEs.

[See [Migrating From BGP-VPLS to EVPN Overview](#).]

SEE ALSO

[Changes in Behavior and Syntax | 123](#)

[Known Behavior | 132](#)

[Known Issues | 137](#)

[Resolved Issues | 153](#)

[Documentation Updates | 189](#)

[Migration, Upgrade, and Downgrade Instructions | 190](#)

[Product Compatibility | 197](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Release 18.1R3-S12 Changes in Behavior and Syntax | 123](#)
- [Release 18.1R3-S7 Changes in Behavior and Syntax | 123](#)
- [Release 18.1R3-S5 Changes in Behavior and Syntax | 124](#)
- [Release 18.1R3 Changes in Behavior and Syntax | 124](#)
- [Release 18.1R2 Changes in Behavior and Syntax | 127](#)
- [Release 18.1R1 Changes in Behavior and Syntax | 128](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.1R3 for MX Series routers.

Release 18.1R3-S12 Changes in Behavior and Syntax

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.

Release 18.1R3-S7 Changes in Behavior and Syntax

Interfaces and Chassis

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the **show interfaces mc-ae extensive** command. The output now displays the following two additional fields:
 - Local Partner System ID—LACP partner system ID as seen by the local node.
 - Peer Partner System ID—LACP partner system ID as seen by the MC-AE peer node.

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

Release 18.1R3-S5 Changes in Behavior and Syntax

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 https://www.juniper.net/documentation/en_US/junos/topics/reference/configuration-statement/advertise-from-main-vpn-table-edit-protocols-bgp.html advertise-from-main-vpn-tables.

Release 18.1R3 Changes in Behavior and Syntax

MPLS

- **Bandwidth allocation**—For a label-switched path (LSP) that has both **bandwidth** and **minimum-bandwidth** for autobandwidth configured under the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level, the LSP bandwidth is adjusted differently.

The LSP is initiated with the bandwidth value configured under the **bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level. At the expiry of the **adjust-interval** timer, the LSP bandwidth gets adjusted based on the traffic flow.

If the bandwidth to be signaled is less than the value configured under the **minimum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the minimum bandwidth.

If the bandwidth to be signaled is greater than the value configured under the **maximum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the maximum bandwidth.

- **Loss of traffic over bypass MPLS LSPs**—If RSVP link or node protection is enabled along with global RSVP authentication, there is loss of traffic over bypass MPLS LSPs at the time of local repair, when the point of local repair (PLR) and the merge point devices have different versions of the Junos OS software installed on them. That is, one device is running a release prior to Junos OS Release 16.1, and the other device is running a release starting with Junos OS Release 16.1R4-S12.

Network Management and Monitoring

- **New context-oid option for trap-options configuration statement to distinguish the traps which come from a non-default routing instance and non-default logical system (MX Series)**—Starting in Junos OS Release 18.1R3, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

Services Applications

- **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. In case, 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 per NAT pool address.

Software Installation and Upgrade

- **ZTP is supported on MX Series PPC platforms (MX Series)**—Starting in Junos OS Release 18.1R3, zero touch provisioning (ZTP) is supported on MX Series PPC platforms (which are MX5, MX10, MX40, MX80, and MX104 routers). Before the fix, the ZTP process did not start to load image and configuration for MX PPC routers.

[See [Junos OS Installation Package Names](#).]

Subscriber Management and Services

- **Changed behavior for framed routes without a subnet mask (MX Series)**—Starting in Junos OS Release 18.1R3, the router connects the session but ignores a framed route when it is received from RADIUS in the Framed-Route attribute (22) without a subnet mask.

In earlier releases, the router installs the framed route with a Class A, B, or C subnet mask depending on the value of the first octet. When the octet < 128, the mask is /8; when 128 <= octet < 192, the mask is /16; and when the octet >= 192, the mask is 24.

- **DHCPv6 lease renewal for separate IA renew requests (MX Series)**—Starting in Junos OS Release 18.1R3, the jdhcpd process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

In earlier releases:

1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

[See [Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview](#).]

- **Bandwidth options match for inline services and tunnel services (MX Series)**—Starting in Junos OS Release 18.1R3, you can configure the same bandwidth options for inline services with the **bandwidth** statement at the **[edit chassis fpc slot-number pic number inline-services]** hierarchy level as you can configure for tunnel services with the **bandwidth** statement at the **[edit chassis fpc slot-number pic number tunnel-services]** hierarchy level.

[See [bandwidth \(Inline Services\)](#) and [bandwidth \(Tunnel Services\)](#).]

Release 18.1R2 Changes in Behavior and Syntax

High Availability (HA) and Resiliency

- **Command `show chassis in-service-upgrade` not available (MX10003)**—Starting in Junos OS Release 18.1R2, the command `show chassis in-service-upgrade` is not available for MX10003 routers. If you enter this command, the following output is shown: **error: command is not valid on the JNP10003 [MX10003]**. Earlier, the output shown for this command was **error: Unrecognized command (chassis-control)**.

Interfaces and Chassis

- On MX Series Routers with the RE-S-X6-64G and RE-MX2K-X8-64G Routing Engines, when the user changes the router configuration on a live system, or when the user deletes an interface that has active traffic, the message **select: protocol failure in circuit setup** is randomly displayed. However, there is no known functional impact.

MPLS

- **Support for `inet.0` and `inet.3` labeled unicast BGP route for protocol LDP (MX Series)**—Starting in Junos OS Release 18.1R2, LDP egress policy is supported on both `inet.0` and `inet.3` routing Information bases (RIBs), also known as routing tables, for labeled unicast BGP routes. If a routing policy is configured with a specific (`inet.0` and `inet.3`) RIB, the egress policy is applied on the specified RIB. If no RIB is specified and a prefix is present on both `inet.0` and `inet.3` RIBs for labeled unicast BGP routes, then `inet.3` RIB is preferred. However, prior to Junos OS Release 12.3R1 and starting with Junos OS Release 16.1R1, LDP egress policy is always preferred on `inet.0` RIB and support for `inet.3` RIB egress policy for labeled unicast BGP routes was disabled. In Junos OS Release 12.3R1 and later releases, up to Junos OS Release 16.1R1, LDP egress policy was supported in `inet.3` RIBs, in addition to `inet.0` RIBs, for labeled-unicast BGP routes.

Network Management and Monitoring

- A decrease in the MPLS label-switched path (LSP) statistics pauses the SNMP MIB `mplsLspInfoAggrOctets` count for one MPLS statistics gathering interval. In such cases, the `mplsLspInfoAggrOctets` value is updated only after completing one more interval of the MPLS statistics gathering.

Routing Protocols

- **IS-IS adjacency SID routes retained only when backup path is available**—Starting in Junos OS Release 18.1R2, when an IS-IS link flaps the adjacency SID routes are retained in the RIB, also known as the routing table, and the FIB, also known as the forwarding table, only if a backup path is available. In earlier Junos OS releases, adjacency SID routes were retained in the RIB and FIB even when a backup path was not available.

Software Defined Networking

- **The 32-bit `libstdc++6` package no longer required for Junos Node Slicing setup**—Starting in Junos OS Release 18.1R2, you need not install the additional 32-bit `libstdc++` package for Red Hat Enterprise Linux (RHEL) or Ubuntu to set up Junos Node Slicing.

Subscriber Management and Services

- **Wildcard supported for show subscribers agent-circuit-identifier command (MX Series)**—Starting in Junos OS Release 18.1R2, you can specify either the complete ACI string or a substring when you issue the **show subscribers agent-circuit-identifier** command. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

In earlier releases, starting with Junos OS Release 14.1, the command requires you to specify the complete ACI string to display the correct results. In Junos OS Release 13.3, you can successfully specify a substring of the ACI without a wildcard.

- **Changes in recommendations for maximum configuration database size (MX Series)**—Starting in Junos OS Release 18.1R2, we recommend that you allow the router to determine the appropriate size for the configuration database to optimize the amount of shared memory available for subscriber management. Do not configure a maximum size.

This recommendation applies to MX240, MX480, MX960, MX2008, MX2010, MX2020, and MX10003 routers when all the Routing Engines have at least 32GB of RAM each. When the Routing Engines have less RAM, we recommend that you configure the maximum size to no more than 300MB.

[See [Configuring Junos OS Enhanced Subscriber Management](#).]

Release 18.1R1 Changes in Behavior and Syntax

EVPNs

- **Change in the output for show evpn instance and show evpn database**—Starting in Junos OS Release 18.1R1, the output for **show evpn instance** and **show evpn database** displays a local interface with an interface name of `.local..number`, and no configuration. This interface is created to support configuration fault management (CFM). For example, **show evpn instance** displays the following sample output.

```
Number of local interfaces: 2 (2 up)
  Interface name  ESI                                     Mode           Status
  AC-Role
  .local..9      00:00:00:00:00:00:00:00:00:00  single-homed   Up
  Root
```

- **Support for LSP on EVPN-MPLS**—Starting in Junos OS Release 18.1R1, Junos OS supports the mapping of EVPN traffic to specific label-switched paths (LSPs). Prior to this release, the traffic policies mapping extended community to specific LSPs did not work properly.
- **Changes in the show route extensive output**—Starting in Junos OS Release 18.1R1, the output for **show route extensive** displays unknown evpn, opaque, and experimental extended communities as follows:

- EVPN: unknown iana evpn 0xtype:0xsubtype:0xvalue
- OPAQUE: unknown iana opaque 0xtype:0xsubtype:0xvalue
- EXP: unknown 0xtype:0xsub-type:0xvalue

where type, sub-type, and value are defined in RFC 4360 *BGP Extended Communities Attribute*, RFC7153 *IANA Registries for BGP Extended Communities*. Internet Assigned Numbers Authority (IANA) maintains a registry with information on the type and subtype field values at

<https://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xhtml>

Interfaces and Chassis

- **Modified output of the request vmhost zeroize command**—The command **request vmhost zeroize**, upon execution, prompts the user for confirmation to proceed. The following line is displayed:

```
user@host request vmhost zeroize
VMHost Zeroization : Erase all data, including configuration and log files ?
[yes,no] (no) yes
```

See [request vmhost zeroize](#).

- **Recovery of PICs that are stuck because of prolonged flow controls (MS-MIC, MS-MPC, MS-DPC, MS-PIC 100, MS-PIC 400, and MS-PIC 500)**—Starting in Junos OS Release 18.1R1, if interfaces on an MS-PIC, MS-MIC, MS-MPC, or MS-DPC are in stuck state because of prolonged flow control, Junos OS restarts the service PICs to recover them from this state. However, if you want the PICs to remain in stuck state until you manually restart the PICs, configure the new option **up-on-flow-control** for the **flow-control-options** statement at the **[edit interfaces mo-fpc/pic/port multiservice-options]** hierarchy level. In releases before Release 16.1R7, there is no action taken to recover service PICs from this state unless one of the options for the **flow-control-options** statement is configured, or service PIC is manually restarted.

[See [flow-control-options](#)]

Management

- **Enhancement to LSP statistics sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 18.1R1, the telemetry data exported for the LSP statistics sensor no longer includes the phrase **and source 0.0.0.0** after the LSP name in the value string for the prefix key. This change reduces the payload size of data exported. The following is an example of the new format:

```
str_value: /mpls/lsp/constrained-path/tunnels/tunnel[name='LSP-4-3']/state/
counters[name='c-27810']/
```

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 18.1R1, the format of telemetry data exported through gRPC for NPU memory and memory utilization implements prefix compression. This change reduces the payload size of data exported. The following example shows the new format:

key: __prefix__

str_value: /components/component[name='FPC0:NPU0']/properties/property

key: [name='mem-util-edmem-size']/value

uint_value: 12345

Telemetry data is exported in key-value pairs. Previously, the data exported included the component and property names in a single key string.

MPLS

- **New option for show mpls lsp autobandwidth command (MX Series)**—Starting in Junos OS Release 18.1R1, a new option – **name lsp-name** – is introduced in the **show mpls lsp autobandwidth** command to specify the name of the LSP for which the autobandwidth information is displayed. With the **name** option, the autobandwidth information specific to the LSP name that has been provided can be obtained in the command output.

[See [show mpls lsp autobandwidth](#).]

Network Management and Monitoring

- **SNMP syslog messages changed (MX Series)**—In Junos OS Release 18.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD—AgentX master agent failed to respond to ping. Attempting to re-register
 - NEW—AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD—NET-SNMP version %s AgentX subagent connected
 - NEW—NET-SNMP version %s AgentX subagent Open-Sent!

[See the [MIB Explorer](#).]

- **Customer-visible SNMP trap name changes (MX Series)**—Starting in Junos OS Release 18.1R1, on the source control board enhanced (SCBE), name changes include the CB slot when jnxTimingFaultLOSset and jnxTimingFaultLOSClear traps are generated in the case of BITS interfaces (T1 or E1). SNMP traps for the backup Routing Engine clock failure event have been added and the control board name is included in SNMP trap interface name (jnxClkSyncIntfName), for example, value: "external(cb-0)".

See [SNMP MIB Explorer](#).

Network Operations and Troubleshooting Automation

- **JET - Correction to escaped characters notification events (MX Series and vMX routers)**—Per RFC7159, certain characters must be escaped. Data returned from JET notification subscriptions contained escaped characters that were not required. This has been corrected to comply with RFC7159.
- **respawn-on-normal-exit option added to [edit system extensions extension-service application file <application-name>] hierarchy (MX Series routers and vMX)**—This option helps to ensure that daemonized Juniper Extension Toolkit (JET) applications that exit normally will restart without user intervention. Daemonized JET applications that exit unexpectedly will still restart without user intervention. This is the default behavior.

Routing Protocols

- **show isis database command output enhanced**— Starting in Junos OS Release 18.1R1, the output of **show isis database** command includes the *Extended IS Reachability* TLV type and length fields. The output also includes the SubTLV length of IS extended neighbors, which helps in understanding the order in which the IS-IS neighbors are packed in the *Extended IS Reachability* TLV.

[See [show isis database](#).]

Software Defined Networking

- **Revoking delegation of PCE-initiated LSPs**—Starting in Junos OS Release 18.1, for a PCC to revoke the delegation of PCE-initiated LSPs, the **lsp-cleanup-timer** must be greater than or equal to the **delegation-cleanup-timeout** at the **[edit protocol pcep pce pce-name]** hierarchy level. If not, the redelegation timeout interval for the PCC can be set to infinity, where the LSP delegations to that PCE remain intact until specific action is taken by the PCC to change the parameters set by the PCE.
- The output of the **show mpls lsp ingress locally-provisioned** command is expected to display only label-switched paths (LSPs) that have been provisioned locally by the Path Computation Client (PCC). However, the **locally-provisioned** option was displaying all the LSPs, instead.

Starting in Junos OS Release 18.1R1, the **locally-provisioned** option in the **show mpls lsp ingress** command is behaving as expected.

Subscriber Management and Services

- **Change to ICRQ message inclusion of the ANCP Access Line Type AVP (MX Series)**—Starting in Junos OS Release 18.1R1, the ICRQ message includes the ANCP Access Line Type AVP (145), when the received ANCP Port Up message includes a DSL-type of 0 (OTHER). In earlier releases, the AVP is not sent when the value is 0.

[See [Subscriber Access Line Information Handling by the LAC and LNS Overview](#).]

- **Support for IPv6 all-routers address in nondefault routing instance (MX Series)**—Starting in Junos OS Release 18.1R1, the well-known IPv6 all-routers multicast address, FF02::2, is supported in nondefault routing instances. In earlier releases it is supported only for the default routing instance; consequently IPv6 router solicitation packets are dropped in nondefault routing instances.
- **Correction to CLI for L2TP tunnel keepalives (MX Series)**—Starting in Junos OS Release 18.1R1, the CLI correctly limits to 3600 seconds the maximum duration that you can enter for the hello interval of an L2TP tunnel group. In earlier releases, the CLI allows you to enter a value up to 65,535, even though only 3600 is supported.

See [hello-interval \(L2TP\)](#).

User Interface and Configuration

- Junos OS prohibits configuring ephemeral configuration database instances that use the name **default** (MX Series)—Starting in Junos OS Release 18.1R1, user-defined instances of the ephemeral configuration database, which are configured using the **instance *instance-name*** statement at the **[edit system configuration-database ephemeral]** hierarchy level, do not support configuring the name **default**.

SEE ALSO

[New and Changed Features | 102](#)

[Known Behavior | 132](#)

[Known Issues | 137](#)

[Resolved Issues | 153](#)

[Documentation Updates | 189](#)

[Migration, Upgrade, and Downgrade Instructions | 190](#)

[Product Compatibility | 197](#)

Known Behavior

IN THIS SECTION

- [EVPN | 133](#)
- [General Routing | 133](#)
- [Interfaces and Chassis | 134](#)
- [MPLS | 135](#)
- [Platform and Infrastructure | 135](#)
- [Routing Protocols | 135](#)
- [Services Applications | 135](#)
- [Software Installation and Upgrade | 136](#)
- [Subscriber Management and Services | 137](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 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

- Routing instances of type "evpn" configured with a vlan-id will advertise MAC (type 2) routes with the VLAN value in the Ethernet tag field of the MAC route. Advertising MAC routes with a nonzero VLAN is incompatible with the EVPN VLAN-based service type. To enable interoperability between a Junos OS routing instance of type "evpn" and a remote EVPN device operating in VLAN-based mode, the Junos OS routing instance should be configured with **vlan-id none** so that the Ethernet tag in advertised MAC routes is set to zero. [PR945247](#)
- When the Routing Engine switchover in a scaled-up EVPN VPWS configurations (approximately 8000 EVPN VPWS), the rpd scheduler slip messages might be seen. [PR1225153](#)

General Routing

- Source-prefix filtering and protocol filtering of the CGNAT sessions are incorrect. For example, **show services sessions extensive protocol udp source-prefix <0:7000::2>** displays incorrect filtering of the sessions. [PR1179922](#)
- Chef for Junos OS supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev resources. The netdev resource developed for interface configuration has a limitation to configuring the XE interface. The netdev interface resource determines that speed is a configurable parameter that is supported on a GE interface but not on an XE interface. Hence, the netdev interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages chef-11.10.4_1.1.*.tgz chef-11.10.4_2.0.*.tgz in all platforms {i386/x86-32/powerpc}. [PR1181475](#)
- When some route/NH has been created by the application, the route can propagate to the rest of the system. KRT asynchronously picks up this state for propagation. There is no reverse indication to the application, if there was an error in propagating the state. The system is supposed to eventually reconcile. So, if SPRING-TE produces a <route> pair that looks legal from the application standpoint, but KRT is not able to download it to the kernel, because the kernel rejected the NH, the <route> gets stuck in the routing protocol process (rpd). In the meantime, the previous version of the route (L-ISIS in this case) that is downloaded still lingers in the kernel and Packet Forwarding Engine. [PR1253778](#)
- On rebooting, RHEL 7.3 servers report libvirtd[6282]: segfault at 10 ip 00007f87eab09bd0. No core file is left and no operational impact is known. [PR1287808](#)
- For CFP2-DCO-T-WDM-1 pluggable, Rx payload type shown incorrectly (shown 0 vs 7). [PR1300423](#)
- IPsec operations are optimized for smaller packet size (up to ~1900 bytes) on MS-MPC and MS-MIC platforms thus yielding higher throughput and lower latency for more common network deployments. A slightly higher latency might be seen if there are jumbo packets in the network. [PR1307867](#)

- Sometimes 1GE interface might remain down after certain events. The events that might cause 1GE interface to remain down are as follows:
 - The two interfaces are connected in loopback on 12 Port QSFP28 TIC on MX10003 or 4 Port QSPF28 fixed PIC on MX204 and configured as 1GE interfaces.
 - When two MX10003 boxes and configured in 1G are connected back to back and rebooted at same time with 1GE interfaces on 12 Port QSFP28 TIC. [PR1312403](#)
- With 1GE interfaces configured on either MX10003 or MX204, the available throughput per port will be in the order of approximately 990mbps instead of 1000mbps (1Gbps) [PR1318293](#)
- Starting in Junos OS Release 15.1, the enhanced subscriber management SNMP interface filters might not work for subscriber interfaces when "interface-mib" is part of subscriber dynamic-profile. Without "interface-mib" in subscriber dynamic profile, there is no change in behavior. [PR1324573](#)
- In MX Series routers, when Telemetry is enabled upon GRES, IKE sensor subscription from collector should restart new backup Routing Engine that still holds the old subscription state, which is functionally dormant, until it becomes an active Routing Engine again. When it becomes a master Routing Engine, no stale state is held. New IKE sensors are added with new subscription from remote collectors. [PR1340110](#)
- When a new instance of Virtual Route Reflector (VRR) is launched, the factory default configuration has DHCP client and auto image is turned on. Even after DHCP configuration is removed, access-internal default routes installed by DHCP client might persist and cause reachability problem. This typically happens during initial installation, and restart routing immediately might clear the problem. [PR1335925](#)

Interfaces and Chassis

- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including the master routing instance), but only one logical interface is assigned with the identical address after commit. There is no warning during the commit, only syslog messages indicating incorrect configuration. [PR1221993](#)
- In case of hw-assisted-pm mode of operation at responder, it takes few ms/seconds (based on the programmed scale) to program inline-responder entries once CCM comes up. So till inline-entry corresponding to a SLM session does not get programmed response will not be sent back to originator and originator will see loss. Once inline-responder entry gets programmed responses will be sent back to originator. [PR1311963](#)
- For MIC-3D-8OC3-2OC12-ATM on MX104 routers, ensure that the configured **cell-bundle-size** is less than 30 for an ATM interface that is configured with **atm-ccc-cell-relay** encapsulation. If the configured **cell-bundle-size** is greater than or equal to 30 and the traffic is passing through the interface at line rate, it might lead to AFEB crash.

[See [cell-bundle-size](#)]

- At JDM install time, each JDM instance generates pseudo random MAC addresses to be used for JDM's own management interface and for the associated GNFs' management interfaces. At GNF creation time, each GNF instance generates pseudo random MAC addresses to be used as the chassis MAC address pool for the forwarding interfaces of that GNF. Once generated, JDM and GNF MAC addresses are persistent, and will only be deleted when the JDM or GNF instance itself is deleted.

At a GNF, the Junos OS CLI command **show chassis mac-addresses** can be used to examine its chassis MAC address pool, and the Junos OS CLI command **show interfaces fxp0** can be used to examine the MAC address of its management interface.

At JDM, the CLI command **show interfaces jmgmt0** can be used to examine the MAC address of its management interface.

In case of MAC address duplication across JDM or GNF instances, you must delete and then reinstall the respective JDM or GNF instance and check again for duplication.

MPLS

- An SR-TE path with "0" explicit NULL as inner most label, SR-TE path does not get installed with label "0". [PR1287354](#)

Platform and Infrastructure

- The Junos OS does not launch **/usr/sbin/commitd -s** after every switchover. [PR1284271](#)

Routing Protocols

- When an Junos OS aggregation gateway uses a IPv6 address as a next hop for IPv4 aggregates announced to downstream, it might attract the traffic prematurely before Packet Forwarding Engines are programmed with more specific IPv4 routes. This happens when the IPv6 address is advertised in BGP inet6-labeled-unicast family. [PR1220235](#)
- BGP peer flap is seen when Routing Engine switchover is triggered from old backup Routing Engine. This issue is seen only with higher scales. The issue is related to slow draining out of new backup socket. [PR1325804](#)

Services Applications

- It is not recommended to configure **ms- interface** when AMS bundle in one-to-one mode has the same member interface. [PR1209660](#)

- Broadband-edge platforms do not support service-set integration with dynamic profiles when the service set is representing a carrier-grade NAT configuration. As a workaround, you can use next-hop service set configurations and routing options to steer traffic to a multiservices (ms) interface where NAT functionality can be exercised. The following configuration snippet shows the basics of statically configuring the multiservices interface next hop and a next-hop service set. Traffic on which the service is applied is forced to the interface inside the network by configuring that interface as the next hop. This configuration does not show other routing-options or NAT configurations relevant to your network.

```

routing-options {
  static {
    route 0.0.0.0/0 {
      next-hop ms-3/0/0.1;
      preference 0;
    }
  }
  ...
}
services {
  service-set CGN {
    nat-rules CGN_SAMPLE;
    next-hop-service {
      inside-service-interface ms-3/0/0.1;
      outside-service-interface ms-3/0/0.2;
    }
  }
  nat {
    ...
  }
}

```

[See [Configuring Service Sets to be Applied to Services Interfaces.](#)]

Software Installation and Upgrade

- **Unified ISSU with active BBE subscribers using advanced services supported only to 18.1R2 and later 18.1 releases**—If you have active broadband edge subscribers that are using advanced services, you cannot perform a successful unified in-service software upgrade (ISSU) to a Junos OS 18.1 release earlier than 18.1R2. If you perform an ISSU to an 18.1 release earlier than 18.1R2, the advanced services PCC rules are not attached to subscribers.
- **Unified ISSU not supported with an active RPM configuration**—If you have an active real-time performance monitoring (RPM) configuration, you cannot perform a successful unified in-service software upgrade (ISSU) to a Junos OS 18.1 release. The warning **ISSU is not supported for RPM configuration** appears.

Subscriber Management and Services

- Before you make any changes to the underlying interface for a demux0 interface, you must ensure that no subscribers are currently present on that underlying interface. If any subscribers are present, you must remove them before you make changes.

SEE ALSO

New and Changed Features 102
Changes in Behavior and Syntax 123
Known Issues 137
Resolved Issues 153
Documentation Updates 189
Migration, Upgrade, and Downgrade Instructions 190
Product Compatibility 197

Known Issues

IN THIS SECTION

- [Class of Service\(CoS\) | 138](#)
- [EVPN | 138](#)
- [Forwarding and Sampling | 139](#)
- [General Routing | 139](#)
- [Infrastructure | 144](#)
- [Interfaces and Chassis | 144](#)
- [Layer 2 Ethernet Services | 145](#)
- [Layer 2 Features | 146](#)
- [MPLS | 146](#)
- [Platform and Infrastructure | 147](#)
- [Routing Protocols | 150](#)
- [Services Applications | 151](#)
- [Software Installation and Upgrade | 151](#)

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for MX Series routers.

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

Class of Service(CoS)

- CoS is not applied to Packet Forwarding Engine when VCP link is added. [PR1321184](#)

EVPN

- The Layer 2 address learning process (l2-ald) might generate a core file in a scaled L2 setup, including bridge domain, VPLS, EVPN, and so on. The l2-ald core file usually follows a kernel page fault that recovers on its own. In some cases, a manual restart of the process is needed to recover logs /kernel: %KERN-3-BAD_PAGE_FAULT: pid 69719 (l2ald), uid 0: pc 0x88beb5ce got a read fault at 0x6ca, x86 fault flags = 0x4 /kernel: %KERN-6: pid 69719 (l2ald), uid 0: exited on signal 11 (core dumped) init: %AUTH-3: l2-learning (PID 69719) terminated by signal number 11. A core file is generated. [PR1142719](#)
- When you run VPLS on MX Series routers (with Junos OS), you might have statements such as **mac-table-size** configured under the **[routing-instance protocols vpls]** hierarchy. You can also configure these statements under **[protocols evpn]** for an EVPN routing instance. When you migrate a VPLS instance to EVPN, you do not need to configure the attributes under **[protocols evpn]**; which means that the configuration under **[protocols vpls]** continue to be in effect. But when you eventually disable the VPLS protocol (so that it does not interfere with EVPN after all nodes are migrated), you must configure the statements under **[protocols evpn]**. Otherwise, the configuration reverts to default values for the routing instance and results in the removal of all dynamically learned MAC addresses, which is also known as MAC flush processing. [PR1312531](#)
- When there is a direct connection between leaf to leaf, there might be a scenario where MAC is learned on a VTEP tunnel from a remote switching gateway instead of on a local interface. The MAC in question is behind the CE connected to both leaves in active/active mode. There is a temporary loop during system startup. [PR1323182](#)
- Provider backbone bridging (PBB) EVPN is unable to flood traffic toward the core. To recover traffic, use the **restart l2-learning** command. In addition to this, there is a limitation in PBB EVPN active/active unicast traffic forwarding. If entropy in the traffic is not sufficient, then uneven load balancing causes a

problem on MH peer active/active routers. The two problems are applicable to MAC-in-MAC PBB-EVPN and do not affect any other scenario. [PR1323503](#)

- In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) deployment, the rpd process might crash on the new master after performing a GRES. [PR1333754](#)
- On MX Series devices running Junos OS, the l2ald daemon might crash during MAC address processing. The MAC learning process will be impacted during the period of l2ald crash. The l2ald recovers itself. [PR1347606](#)
- Bidirection L2 traffic floods for around 5 seconds for streams from SH to MH, when the **clear mac table** command is executed on MX Series routers because MAC takes time to develop in the system. The **clear mac table** command is a disruptive command that deletes all dynamic MACs in the system. [PR1360348](#)
- If EVPN is configured with CoS-based forwarding (CBF), traffic might be lost for the CBF services. [PR1374211](#)

Forwarding and Sampling

- When a policing filter is applied to an active LSP carrying traffic, the LSP resignals and drops traffic for approximately 2 seconds. It can take up to 30 seconds for the LSP to come up under the following conditions: (1) Creation of the policing filter and application of the same to the LSP through configuration occurs in the same commit sequence. (2) Load override of a configuration file that has a policing filter and policing filter application to the LSP is followed by a commit. [PR1160669](#)
- Heap memory leaks occur on DPC when the flow specification route is changed. [PR1305977](#)
- Firewall filter is not applied as input filter to extended port when used for L2VPN. [PR1311013](#)
- This issue affects unified ISSU only when filter lists are being used. Starting in Junos OS Releases 15.1F5, 15.1F6, 16.1R1, or later to Junos OS Releases 17.1R2, 17.1R3, 17.2R2, 17.2X75-D50, 17.3R1, or later, an error might occur that prevents firewall configuration changes from being properly applied. To avoid this issue, the configuration must explicitly set the **filter-list-template** or **no-filter-list-template** flag before the unified ISSU is done. [PR1345711](#)
- When a logical interface is made admin down, the accounting records for that logical interface will not be written. [PR1348249](#)

General Routing

- SIP session fails when the IPv4 SIP client in a public network initiates the SIP call with the IPv6 SIP client in the private network. [PR1139008](#)
- When PIC PB-4OC3-4OC12-SON-SF (4x OC-12-3 SFP) is replaced with PB-4OC3-1OC12-SON2-SFP (4x OC-3 1x OC-12 SFP) and a CLI commit is done, the replacement PIC type bounces. As a workaround, do a full commit immediately after bringing the replacement PIC online. Doing so will bounce the PIC one time, and the replacement PIC does not bounce with consecutive commits. [PR1190569](#)

- SMID daemon has stopped responding to the management requests after a L2TP daemon crash on an MX960 BNG. [PR1205546](#)
- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, fpc7 KERNEL/PFE APP=NH OUT OF SYNC: **error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562, type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none.** No service impact is seen in MPC2 and MPC3 type cards. [PR1205593](#)
- The `/etc/passwd` file is created in the process of the first commit when a pristine jinstall image is used to boot for the first time. If **event-options** is configured, the system tries to read the configuration from the available event scripts, which requires privileges obtained from the `/etc/passwd` file. This causes a circular dependency because the commit will not pass if the configuration includes **event-options** the first time a pristine image boots up, which is the case when an upgrade is performed. [PR1220671](#)
- When a configuration that moves Packet Forwarding Engine to offline and another configuration that brings the Packet Forwarding Engine back online, is committed in quick succession. There could be RE-PFE out of synchronization syslog errors. Most of the time these are benign errors, but sometimes these errors might result in Packet Forwarding Engine crash. [PR1232178](#)
- Sometimes, when PPPoE subscribers log in and log out from Junos OS Release 16.1 and later releases, the following messages are generated. `user@host> show log messages | match authd authd[5208]: sdb_app_access_line_entry_read_by_uifl: uifl key 'demux0.xxxxxxxx': snapshot failed (-7) authd[5208]: sdb_app_access_line_entry_read: uifl key 'demux0.xxxxxxxx': read failed.` These messages indicate that authd daemon for subscriber authentication is attempting to read private data for an underlying interface that no longer exists (-7 = SDB_DATA_NOT_FOUND). These messages have no impact and can be safely ignored. [PR1236211](#)
- Malicious LLDP crafted packet leads to privilege escalation, denial of service (CVE-2018-0007). Refer to <https://kb.juniper.net/JSA10830> for more information. [PR1252823](#)
- The following cosmetic error is observed as the output: `mshpmand[190]: msvcs_session_send: Plugin id 3 not present in the svc chain for session.` [PR1258970](#)
- When an interface comes online and both the OAM protocol and the MKA protocol try to establish their respective sessions, OAM takes the interface down and MKA fails to establish connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
- On MPC2E-NG, MPC3E-NG, MPC5E, MPC6E, MPC7E, MPC8E, and MPC9E line cards, a firewall performance feature fast-lookup-filter can be activated. Because of the transient parity error, the packet is dropped within the PPE with the **sync txn error** message. This issue might adversely impact traffic, which might eventually affect the service. [PR1266879](#)
- GNF configurations in a node-slicing setup currently do not support Junos OS snapshot and recovery mechanisms. [PR1268943](#)
- Dynamic endpoint (DEP) does not support the dh group group19, encryption algorithm aes-256-cbc and hash sha-384 in its list of default proposals. These must be configured explicitly in the configuration. [PR1269160](#)

- In an inline J-Flow, when the **template-refresh-rate** and **option-refresh-rate** are set with both **packets interval** and **seconds interval** configuration options, the **packets interval** configuration does not work. [PR1274206](#)
- Performance of X710 NIC is lower compared to that of 82599 NIC. 40G line rate can be achieved at 512 byte packet size for X710 NIC as compared to 256 bytes for 82599 NIC. [PR1281366](#)
- If a VM host snapshot is taken on an alternate disk and there is no further VM host software image upgrade, the expectation is that the currently corrupt VM host image can use the alternate disk to recover the primary disk state. However, if the host file system is corrupted, the node boots from the previous VM host software instead of booting from the alternate disk. [PR1281554](#)
- Because of the vendor code limitation, ungraceful removing of summit MACsec TIC from the chassis might crash or give unpredictable results. [PR1284040](#)
- This is an internal change because syslog usage is deprecated; however, there might be customer impact because of the syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful and relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. [PR1284643](#)
- Junos OS releases with a fix committed in Junos OS Releases 15.1R5-S4, 16.1R4-S3, 16.1R5, and 17.3R1 with XM-based line cards (MPC3E, MPC4E, MPC5E, MPC6E, MPC2E-NG, and MPC3E-NG) might report a **DDR3 TEMP ALARM** chassisd error log message. [PR1293543](#)
- The Routing Engine gets stuck and boots from the other solid-state drive (SSD) after a VM host reboot. [PR1295219](#)
- In some MX Series deployments running Junos OS, random syslog messages are observed for FPC cards. For example, you might see **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 are addressed as INFO level messages. On a Junos OS 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](#)
- When a GRES or NSR is performed on a base system (BSYS), the master Routing Engine on the guest network functions (virtual nodes or network slices) detects the BSYS chassisd restart and enters an NSR hold down delay. During this time, CLI commands evoke a switchover on the master Routing Engine, indicating that the system is not NSR ready. This situation is similar to that of a standalone MX Series router in which chassisd is restarted on the master Routing Engine. Note that the CLI command on the BU Routing Engine will succeed. This too is similar to standalone MX Series router behavior. [PR1298571](#)
- User-configured packet hashing options for inet family under enhanced-hash-key might not take effect for FPCs in MX Series routers. FPCs might keep using default behavior for hash calculation for IPv4 packets. [PR1302637](#)
- Internal latency is high during initial subscription of sensors when multiple sensors (in order of 15-20) are subscribed together. This is not observed with a lesser number of subscriptions. This occurs during a small period when sensors are being installed. [PR1303393](#)

- This type of crash indicates simultaneous operation on an ephemeral instance. When a process wants to open an ephemeral configuration in merge view, some other activity (like purging, deletion, or recreation) is being carried out on this ephemeral instance. The occurrence of this core file is rare. [PR1305424](#)
- Support for enterprise profile is only provided for 10-Gigabit Ethernet interfaces. Use of 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. [PR1310048](#)
- While upgrading Juniper Device Manager (JDM), there is a possibility that the jdmd process might not run after the upgrade. No errors are reported during the upgrade. [PR1313964](#)
- The **show dynamic-tunnels database summary** command might not show an accurate tunnels summary when the anchor Packet Forwarding Engine line card is not up. As a workaround, use the **show dynamic-tunnels database** and **show dynamic-tunnels database terse** commands instead. [PR1314763](#)
- Alarm is raised if mixed AC PEMs are present. This occurs because the criteria for checking whether mixed AC is present was changed. If the PEM is AC(HIGH) the first bit of pem_voltage is set and if it is AC(LOW) the second bit of pem_voltage is set. If both the first and second bit are set, mixed AC is present. [PR1315577](#)
- Making changes in Traffic Load Balancer services for one instance might lead to a refresh of other existing instances. [PR1318184](#)
- According to the MACsec extended package of Network Device Protection Profile (NDPP), there needs to be an option to specify the lifetime for connectivity association keys (CAKs) based on FCS_MACSEC_EXT.4.3 (EP - NDcPP Version 1.0). This requirement indicates that there must be a start time and an end time or a time span configured for the lifetime of the CAK. In Junos OS Release 18.1R1, the end time for CAKs cannot be specified, which is a limitation. Only the start time of the various CAKs can be specified through the CLI. [PR1318543](#)
- MPLS GRE dynamic tunnel localization does not work when chained composite next hop is enabled. [PR1318984](#)
- In JDM (running on secondary server), jdmd daemon might generate a core file when the guest network function (GNF) add-image is aborted by pressing Ctrl+C. [PR1321803](#)
- BGP signal tunnels are always next-hop-based tunnels. The GRE tunnels created dynamically by a BGP signal are always next-hop-based tunnels, even if the user has configured the static tunnels created by GRE to use the logical interface base. [PR1322941](#)
- If **commit full** is configured, the na-grpd process might restart and the streaming telemetry might disconnect. [PR1326366](#)
- When a restart chassis-control is done for the first time after the software image is upgraded or after the switch is rebooted, the MPC booting state changes from offline to online directly, without staying at present state during booting. This issue is not seen consistently. There is no functional impact because of this state change. [PR1332613](#)
- Under some race conditions with failover and multiple core interface flapping on an EVPN Virtual Extensible LAN (VXLAN) network, the rpd process might use high CPU memory, causing some issues in

intercommunication with the l2ald process, then causing the l2ald process to generate a core file and restart. [PR1333823](#)

- If a filter is configured with the **scale-optimized** statement, then having the action pointing to traffic-class-count does not increment. [PR1334580](#)
- On a next-generation Routing Engine, after upgrading Junos vmhost, the AI-script gets uninstalled. You need to reinstall these scripts. This is not the case on K2-RE. [PR1337028](#)
- On MX204, MX10003, or MPC7E, MPC8E, and MPC9E, or a 100G, 40G, or 10G interface might keep flapping or stay down because of an interoperation issue between the Juniper device and the remote transport device connected. [PR1337327](#)
- On MX204 and MX10003 routers, physical links might not come up if you perform frequent port profile changes while a line card reboot is in progress. [PR1340140](#)
- On an MX Series platform with 100M SFP used on MIC-3D-20GE-SFP-E and MIC-3D-20GE-SFP-EH, SFP might not work if it is not from Fiberxon or Avago. [PR1344208](#)
- First packet pertaining to J-Flow Packet Forwarding Engine sensor in UDP mode is missing after a line card reboot on PORTER-R platform. [PR1344755](#)
- On a single Routing Engine system, when the GRES configuration is removed, the Routing Engine mastership keepalive timer is not resumed to the default value with GRES enabled. [PR1349049](#)
- In some cases, OIR (removal followed by reinsertion) of a MIC on a FPC can lead to traffic destined to the FPC being silently dropped or discarded. The only way to recover from this is to restart the FPC. [PR1350103](#)
- During stress conditions, error log messages regarding route add, change, and delete might be incorrect. [PR1350713](#)
- When ephemeral DB instance is configured, if committing changes which are unrelated to IGMP/MLD (such as **set interfaces ge-0/0/1.0 description**), and the number of ephemeral commits reaches to ephemeral DB maximum size, the ephemeral DB purge might happen. Then it would purge all the commits and rollover. On this purge the mgd gives all the applications a FULL COMMIT view. And on this FULL COMMIT view IGMP/MLD deletes all configuration and adds it back again. This might cause PIM to prune the groups on those interfaces and send join messages again. Finally, the multicast traffic flapping and drop might be seen. [PR1352499](#)
- The **ipv4-flow-table-size** is used to configure the size of the IPv4 flow table in units of 256,000 entries. However, in an inline J-Flow scenario, if the statement **ipv6-extended-attrib** is configured, changing the flow table configuration or clearing the flow entries might lead to the condition in which even if the **ipv4-flow-table-size** has been changed to a number larger than 149, the maximum number of IPv4 flows still remains at 37,372,900. [PR1355095](#)
- VM crashes when the L2circuit pseudowires are terminated into EVPN over a pseudowire subscriber (PS) interface. This occurs when an abstract fabric (AF) is used in Junos Node Slicing if the anchored Packet Forwarding Engine for PS is different from the ingress Packet Forwarding Engine from which the traffic coming in is seen. [PR1355530](#)

- Because of the lack of support for reflecting the correct link state for the copper cable SFP transceiver on the MX204 platform, the link stays up unexpectedly when the copper cable is removed. [PR1356507](#)
- When an FPC is powered off by the command **show chassis environment fpc**, the status of the FPC under **show chassis FPC** changes to ---Bad Voltage--- on **show chassis fpc**. [PR1358874](#)
- Fabric drops are seen when the MPC5E 3D 2CGE+4XGE 100G interface becomes overloaded due to scheduling functionality. [PR1360822](#)
- Default NIC driver coming as E1000 when vFPC is deployed on vmware using ova image. [PR1365337](#)
- When the FPC VTY command **show agent sensors verbose** is used on MPC7E, the FPC might crash. [PR1366249](#)
- When LAG-enhanced is disabled, one child next hop is created for each member link of a LAG interface. During the non-GRES switchover, the kernel memory might exhaust that leads to the creation of the child-next-hop fails. Hence, the Routing Engine crashes. [PR1373079](#)
- Input and output session used for communicating between threads is freed because of FSM state transition. After freeing the memory, the fields of the I/O session are used for tracing. As a result, the rpd process generates a core file. [PR1374759](#)
- When an MX Series BNG acts as DHCP relay and the destination DHCP server is reachable through abstract fabric interfaces (AFIs), the packets received by the DHCP server on AF interfaces are dropped as Junos DHCP daemon (jdhcpd) is not AFI aware. [PR1377358](#)

Infrastructure

- A file system corruption might create a kernel core file. The Routing Engine reboots with the message **ffs_blkfree: freeing free block**. [PR1028972](#)
- System with high uptime with Unigen SSD might cause watchdog panic during upgrade from Junos OS Release 14.X to Junos OS 15.X and later releases. A reboot of the system is required prior to upgrade to avoid dirty file system which might trigger the panic. [PR1309483](#)
- To test features like NSR, the junos-panic package can be installed, and /usr/libexec/panic can be run by root. [PR1352217](#)

Interfaces and Chassis

- Identical IP addresses can be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance). [PR1221993](#)
- Out of sequence packets are seen with LSQ interface. [PR1258258](#)
- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might crash the cfmd process. This is because of using the older version of /var/db/cfm.db. [PR1281073](#)

- LAG member links running LACP in slow mode might get disassociated from the LAG bundle with a combination of restart interface-control and FPC offline or online trigger. The issue is seen with scale configuration on DUT. The scale details are 2800 CFM sessions 2800 BFD sessions 2043 BGP peers 3400 VRF instances. [PR1298985](#)
- Y.1731 Dwlly measurement is not supported on MPC6. [PR1303672](#)
- CFM session does not come up if configured on logical interface with vlan-id matching that of configured native vlan-id under the physical interface. [PR1325190](#)
- When the link speed of the aggregated Ethernet bundle is configured to oc192, certain sequence operation might lead to the aggregated Ethernet interface flap which affects the traffic. First, configure the member links. And then, remove a member link from the bundle. At last, add a member link back. [PR1355270](#)
- Post GRES 1 GE changes to 10 GE. [PR1326316](#)
- On MX Series platform, while configuring duplicate IP between sonet (so-) interface between other type interface, other type interface might not get the IP address. [PR1377690](#)

Layer 2 Ethernet Services

- After changing an outer vlan-tags, the logical interface is getting programmed with incorrect STP state (discarding), so the traffic is getting dropped. [PR1121564](#)
- This is in an internal change as syslog usage is deprecated. However, there might be customer impact because of the syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful or relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. [PR1284592](#)
- MX Series routers might display false positive CB alarm **PMBus Device Fail**. [PR1298612](#)

Layer 2 Features

- For router equipped with following line cards: T4000-FPC5-3D MX-MPC3E-3D, MPC5E-40G10G, MPC5EQ-40G10G, and MPC6E MX2K-MPC6E, if the router is working as VPLS PE, because of MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)

MPLS

- When using "mpls traffic-engineering bgp-igp-both-ribs" with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit area border router (ABR) when there are two or more such area border routers (ABRs). This causes interdomain RSVP LSPs to break. RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. [PR1048560](#)
- The issue occurs when a GRES is performed between the master and backup Routing Engines of different memory capabilities. For example, one Routing Engine has only enough memory to run routing protocol process (rpd) in 32 bit mode while the other is capable of 64 bit mode. The situation might be caused by using Junos OS Release 13.3 or later with the configuration statement **auto-64-bit** configured, or, by using Junos OS Release 15.1 or later even without the configuration statement. Under these conditions, the rpd might crash on the new master Routing Engine. As a workaround, this issue can be avoided by using the CLI command **set system processes routing force-32-bit**. [PR1141728](#)
- In CE-CE setup, traffic loss might be observed over secondary LSP on primary failover. [PR1240892](#)
- With non-stop-routing (NSR), when the routing protocol process (rpd) restarts on the master Routing Engine, the rpd on the backup Routing Engine might restart. [PR1282369](#)
- In case of CSPF disabled LSPs, if the primary path ERO is changed to unreachable strict hop, sometimes the primary path stays UP with the old ERO. The LSP does not switch to standby secondary. [PR1284138](#)
- Swapping the binding SID between colored and non colored static SR LSPs might cause rpd to generate a core file. [PR1310018](#)
- If there are some LSPs for which a router has made link protection available and when primary link failure is caused by FPC restart, this might generate a core file. [PR1317536](#)
- If inet address is not configured for the gr-interface, the gr-interface borrows address from loopback interface. From Junos OS Release 16.1R1, the RSVP creates a node-neighbor by default. There are duplicate neighbors with the same IP address because the gr- interface borrows address from loopback interface. The RSVP path lookup will fail because it gets confused with the node neighbor presence. So, the RSVP LSP will not come up when it goes through the gr-interface which is borrowing address from the loopback interface. [PR1340950](#)
- Executing a **restart chassisd** in a MX Series Virtual Chassis router might result in rpd to crash generating a core file.

- IGP OSPF/OSPF3 (area 0, LFA) IS-IS (level 2, LFA) LDP synchronization IPv4 and IPv6
- IBGP dual, redundant route reflection IPv4 and IPv6
- MPLS LDP (IGP synchronization, track IGP metric) RSVP (node link protection, adaptive, auto bandwidth, refresh reduction)
- L3VPN OSPF OSPF3 BGPv4 BGPv6 RIPv2 static MBGP NGEN-MVPN I3vpn cnh with ext space any to any hub and spoke MPLS access Ethernet access multicast extranet per VPN and per prefix labels SRX Series based network address translation SRX Series based firewall
- Direct Internet Access EBGP
- CoS BA/MF classification policing/shaping queuing/scheduling hierarchical queuing/shaping/scheduling 8 traffic classes
- BFD/OAM/CFM liveness detection
- Load balancing L2 aggregate Ethernet IP equal-cost multipath MPLS equal cost multi path
- High availability GRES/NSR ISSU fabric redundancy tail-end protection BGP Prefix Independent Convergence
- Security loopback filter arp policers control plane traffic policers urpf check with all feasible paths ttl filtering J-Flow/IPfix export only SRX Series based DDoS [PR1352227](#)
- The traceroute MPLS from Juniper Networks to Huawei routers does not work as expected due to unsupported TLV. [PR1363641](#)
- This issue occurs when, on optimizing the timer expiry, the traffic engineering database version number match indicates that a CSPF is already run for the path. If an optimization is not done with that version, the CSPF will run despite the version number match. (This occurs due to a per-path optimize-seq-no that is updated with a traffic engineering database sequence number only on optimization.) When path-cc-updated is false and CSPF fails for optimization, the path is disabled to avoid invalid ERO, making sure this does not interfere with global repair/local reversion. [PR1365653](#)
- With static label-switched path (LSP) for MPLS configured with next hop, the next hop might get stuck in dead state when changing the network mask and keeping the IP address unchanged for the outgoing interface through which the LSP next hop is reachable. [PR1372630](#)

Platform and Infrastructure

- Starting with Junos OS Release 13.1R1, if **no-fast-sync** is used with configure-private mode, the commit operation might throw errors after the configuration statements under choice (such as **protocol [ospf pim tcp]**) are deleted and added. The whole hierarchy is shown as changed when using the command of **show configuration | compare**. [PR1042512](#)
- The **login_getclass: unknown class 'j-idle-timeout'** error is getting displayed when the user has not configured timeout value for the root user. If the user has not configured timeout value, j-idle-timeout entry is not present in the **login.conf** file and an error message is displayed because j-idle-timeout class

is not found. Steps to reproduce: 1) Login to router as root. 2) Clear log messages. 3) Exit and go to CLI mode and give "show log messages". The login error is logged in the following messages: `user@host> start shell user root Password: user@host:/var/home/lab # cli user@host> clear log messages all user@host> exit user@host:/var/home/lab # cli user@host> show log messages Jan 5 14:55:06.132 MX-re0 mgd[96513]: %INTERACT-6-UI_CHILD_STATUS: Cleanup child '/usr/libexec/ui/clear-log', PID 96517, status 0 Jan 5 14:55:06.132 MX-re0 mgd[96513]: %INTERACT-6-UI_FILE_CLEARED: 'messages' logfile cleared by user 'lab' Jan 5 14:55:08.047 MX-re0 mgd[96513]: %INTERACT-6-UI_CMDLINE_READ_LINE: User 'lab', command 'exit ' Jan 5 14:55:08.048 MX-re0 mgd[96513]: %INTERACT-6-UI_LOGOUT_EVENT: User 'lab' logout Jan 5 14:55:10.310 MX-re0 cli: %USER-3: login_getclass: unknown class 'j-idle-timeout' <<<<<<<<<< Login error Jan 5 14:55:10.318 MX-re0 mgd[96527]: %DAEMON-7: check_regex_add: 1059 regex_add = 0 Jan 5 14:55:10.319 MX-re0 mgd[96527]: %INTERACT-6-UI_AUTH_EVENT: Authenticated user 'root' at permission level 'super-user' Jan 5 14:55:10.320 MX-re0 mgd[96527]: %INTERACT-6-UI_LOGIN_EVENT: User 'lab' login, class 'super-user' [96527], ssh-connection "", client-mode 'cli' Jan 5 14:55:15.496 MX-re0 mgd[96527]: %INTERACT-6-UI_CMDLINE_READ_LINE: User 'lab', command 'show log messages ' user@host> exit user@host:/var/home/lab # cat /var/etc/csh.login.inc | grep autologout user@host:/var/home/lab # cat /var/etc/login.conf | grep j-idle. No idle timeout values are seen in /var/etc/csh.login.inc and /var/etc/login.conf files. PR1097799`

- The error message **LUCHIP(5) GUMEM1[77a0] mismatch** might be seen after an MX Series router with an MPC card with LU chipset goes offline or online. [PR1221195](#)
- The issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have XE/GE links toward the downstream switch and LAG bundles as uplinks toward upstream routers. The XE/GE link is part of the physical loop in the topology. Spanning-tree protocols such as VSTP, RSTP, and MSTP are used for loop avoidance. Some MAC addresses are not learned on DUT when LAG bundles that are part of such bridge domains are flapped and other events such as spanning-tree root bridge change occur. [PR1275544](#)
- Due to a transient hardware error condition, the **CPQ Sram parity error** and **CPQ RLDRAM double bit ECC error** syslog errors on MQCHIP raise a major CM alarm. [PR1276132](#)
- This is a corner case and the issue is seen only if the open-config package is installed. Analytics configuration goes to the default ephemeral database, while all the daemons read the configuration through merge view. [PR1296702](#)
- An accuracy issue occurs with three color policers of both type single rate and two rate in which the policer rate and burst-size combination of the policer accuracy vary. This issue is present starting in Junos OS Release 11.4 on all platforms that use MX Series with MPCs/MICs ASIC. [PR1307882](#)
- When chassis control restarted with aggregated Ethernet and CoS rewrite configuration, **Platform failed to bind rewrite** error message might be seen in syslog. Issue is specific to aggregated Ethernet interfaces when chassis control is restarted. It is a timing issue that might occur when a logical interface deletion is delayed due to high scale and when logical interfaces come up again post restart they have different indices. Issue is only applicable when aggregated Ethernet interfaces are present. [PR1315437](#)

- On MX Series routers with MPC1E, MPC2E, MPC3D 16x 10GE, MPC3E, and MP4E or T4000 with type 5 card, if the interface is configured with **input-vlan-map** option, then the traffic with more than two VLAN tags might be incorrectly rewritten and sent out. As a result, the traffic might be dropped. [PR1321122](#)
- On MX104 platforms, the error **sdk-vmmd: %USER-3: is_platform_ XX: Platform could not be detected** is seen. [PR1321622](#)
- On all platforms, with dual Routing Engines and GRES enabled, if executing switchover, the firewall filter's state might be incorrect and NPC core file might be seen. [PR1324819](#)
- Traffic statistics might not match on PS after clearing interface statistics. [PR1328252](#)
- The facility is one of the following keywords: auth, authpriv, cron, daemon, kern, lpr, mail, mark, news, security (same as auth), syslog, user, uucp, and local0 through local7. The keyword security should not be used anymore, and mark is only for internal use and therefore should not be used in applications. However, you might want to specify and redirect these messages here. The facility specifies the subsystem that produced the message; for example, all mail programs log with the mail facility (LOG_MAIL) if they log using syslog. The priority is one of the following keywords, in ascending order: debug, info, notice, warning, warn (same as warning), err, error (same as err), crit, alert, emerg, panic (same as emerg). The keywords error, warn, and panic are deprecated and should not be used anymore. The priority defines the severity of the message. **Guest side: =====**
https://www.juniper.net/documentation/en_US/junos/topics/reference/general/syslog-facilities-severity-levels.html remote : sync syslog server config from Junos to Linux & modify rsyslog.conf set vmhost/app-engine syslog host <ip> any any set vmhost/app-engine syslog host <ip> match xxx. [PR1341549](#)
- In a Layer 3 VPN topology, traceroute to a remote PE device for a CE-facing network results in an ICMP TTL expired reply with a source address of only one of the many CE-facing networks. In Junos OS Releases 15.1R5, 16.1R3, and 16.2R1 and later releases, there is a kernel sysctl value, icmp.traceroute_l3vpn. Setting this to 1 will change the behavior to select an address based on the destination specified in the traceroute command. [PR1358376](#)
- If a tunnel interface is anchored on MX Series routers with FPC, and the **class-of-service host-outbound-traffic ieee-802.1 rewrite-rules** statement is configured, the host outbound traffic might get dropped when the traffic goes through this tunnel interface. [PR1371304](#)
- On MX Series routers, after the unified ISSU from Junos OS Release 14.2 to Junos OS Release 16.1, traffic drops on new added interfaces because of the unified ISSU hardware synchronization phase issue. [PR1371373](#)

Routing Protocols

- Continuous soft core files might be generated because of bgp-path-selection code. The routing protocol process (rpd) forks a child, and the child asserts to produce a core file. The problem is with route-ordering and it is automatically corrected after collecting the soft-assert-core file, without any impact to the traffic or service. [PR815146](#)
- In rare cases, rpd might generate a core file with error **rt_notbest_sanity: Path selection failure on** The core is "soft", which means there should be no impact to traffic or routing protocols. [PR946415](#)
- JTASK_SCHED_SLIP for rpd might be seen on restarting routing or OSPF protocol disable with scaled BGP routes in MX104 router. [PR1203979](#)
- 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 categor. This situation results in some unwanted BGP debug messages being logged to the BGP traceoption file. [PR1252294](#)
- LDP OSPF are 'in sync' state and the reason observed for this is "IGP interface down" with ldp-synchronization enabled for OSPF. **user@host> show ospf interface ae100.0 extensive** Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. As per the current analysis, "IGP interface down" is observed as the reason because although LDP notified OSPF that LDP synchronization was achieved. OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up. [PR1256434](#)
- Performance degradation occurs during computation of LFA and RLFA. This issue does not impact functionality. [PR1264564](#)
- Two multicast tunnel (mt) interfaces are seen for each of the PIM neighbors after VPN-Tunnel-Source activation or deactivation. However, ideally, the same tunnel source should be used for both IPv4 and IPv6 address families, if both are using the same PIM tunnel. [PR1281481](#)
- The routing protocol process Packet Forwarding Engine is out-of-sync during MoFRR convergence. [PR1284463](#)
- This is in an internal change as syslog usage is deprecated. However, there might be customer impact because of syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful or relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. [PR1284621](#)
- In resource public key infrastructure (RPKI) scenario, the validation replication database might have many more entries than the validation database after restarting the RPKI cache server and the validation session is reestablished. [PR1325037](#)
- When route target filtering (RTF) is configured for VPN routes and multiple BGP session flap, there is a slight chance that some of the peers might not receive the VPN routes after the flapped sessions come up. [PR1325481](#)

- In a largescale OSPF network (for example, there are more than 500 devices in an area), OSPF remote loop free alternate (rLFA) default PQ node selection algorithm does not provide proper protection paths. [PR1335570](#)
- When **clear validation database** is issued back-to-back multiple times, some validation entries were missing. This eventually recovers after up to 30 minutes (half of the Record Lifetime), when periodical full updates are made. [PR1326256](#)
- The issue is that when configuring anycast and prefix segments in SPRING for IS-IS, prefix-segment index 0 is not supported, even though user is allowed to configure 0 as an index. [PR1340091](#)
- There are scenarios where the application allocates and caches next-hop templates. This causes the next-hop template cache to grow continuously. But when the application clears its local cache, then memory is freed to the next-hop template cache. But the next-hop template cache does not have the code to shrink the cache and free the memory back. So the next-hop template memory is trapped in the cache and cannot be used for other purposes. But if the same BGP routes and next hop come up again, they will reuse the templates from the cache and not consume additional memory. [PR1346984](#)
- Starting with Junos OS Release 16.1, **show bgp neighbor** does not show the correct "Last traffic (seconds)" correctly anymore. [PR1361899](#)
- On devices running Junos OS, when OpenConfig is running with sensor for **/network-instances/network-instance/protocols/protocol/bgp**, changing the BGP import or export policy might cause rpd to generate a core file. [PR1366696](#)

Services Applications

- One of the internal HA queues gets corrupted, which results in mspmand generating a core file on the backup SDG. This issue occurs because sometimes different threads of mspmand might have different timestamps. [PR1291664](#)

Software Installation and Upgrade

- With unified ISSU, momentary traffic loss is expected. In EVPN E-Tree, in addition to traffic loss, the known unicast frames can be flooded for around 30 seconds during unified ISSU before all forwarding states are restored. This issue does not affect BUM traffic. As a workaround, nonstop bridging (NSB) can be configured at **[set protocols layer2-control nonstop-bridging]**. This reduces traffic flood to around 10 seconds in a moderate setup. [PR1275621](#)
- During unified ISSU at MX Series Virtual Chassis, the MX Series Virtual Chassis side might clear the TCP connection, causing BGP peerings to flap. [PR1368805](#)
- Unified ISSU could be aborted at "Timed out Waiting for protocol backup chassis master switch to complete" with MX Series Virtual Chassis configuration. [PR1371297](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3 | 153](#)
- [Resolved Issues: 18.1R2 | 164](#)
- [Resolved Issues: 18.1R1 | 171](#)

This section lists the issues fixed in the Junos OS 18.1R3 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.

Resolved Issues: 18.1R3

Authentication and Access Control

- The client moves back to connecting state when VSTP is enabled along with dynamic VLAN assigned once the port gets authenticated by dot1x. [PR1304397](#)

Class of Service (CoS)

- CoS traffic control profiles fail to apply on an aggregated Ethernet interface in a specific condition. [PR1355498](#)

EVPN

- On EVPN-VXLAN, MAC entry is incorrectly programmed in Packet Forwarding Engine, leading to traffic being silently dropped and discarded. [PR1231402](#)
- MPLS label leak leads to label exhaustion and the rpd process crashes. [PR1333944](#)
- Traffic might be lost on Layer 2 and Layer 3 spine node in a multihome EVPN scenario. [PR1355165](#)
- EVPN IRB configured with **no-gratuitous-arp-request** is still sending gratuitous ARP. [PR1356360](#)
- The rpd might crash if the EVPN instance refers to a vrf-export policy that does not have 'then community'. [PR1360437](#)
- Proxy ARP might not work as expected in an EVPN environment. [PR1368911](#)

Forwarding and Sampling

- Junos OS allows firewall filters with the same name under `[edit firewall]` and `[edit firewall family inet]` hierarchy levels. [PR1344506](#)

- The remote MAC might not be added in the forwarding table, which might cause traffic to be dropped in an EVPN scenario with RSVP and CBF configured. [PR1353555](#)
- The backup Routing Engine might write dummy interface accounting records after GRES. [PR1361403](#)

General Routing

- In timing hybrid mode, MX Series routers MPC2 line cards are not working with ACX Series routers with VLAN (native-vlan-id). [PR1076666](#)
- The chassis alarm message **Bottom Fan Tray Pred Fail** needs to be renamed to something less obscure. [PR1202724](#)
- The incorrect TBB Packet Forwarding Engine component temperature might be reported on the MX80. [PR1259379](#)
- Flexible PIC Concentrator (FPC) crashes or reboots while bringing up about 12,000 Layer 2 bit stream access (L2BSA) subscribers simultaneously. [PR1273353](#)
- Error messages are observed on vty session while running a script for IGMP snooping over EVPN-VXLAN. [PR1276947](#)
- Migrate from syslog API to **Errmsg API;/src/junos/usr.sbin/mobiled**. [PR1284625](#)
- Chassisd core file is generated after insertion of REMX2K-X8-64 in MX2000 platform along with older RE-S-1800x4. [PR1300083](#)
- The total number of packets is less than the sum of unicast, multicast, and broadcast packets due to lag in read depending on packet rate. Now the total packets is set to the sum. [PR1306656](#)
- The error message about **pfeman_inline_ka_steering_gencfg_handler** might be seen when FPC is restarted with BFD configured. [PR1308884](#)
- On MX10000 suppress chassis alarm is observed for switched-off PEMs. [PR1311574](#)
- Sensors that belong to the same producer with identical reporting interval are not streamed in parallel. [PR1315517](#)
- CoS is not applied to Packet Forwarding Engine when VCP link is added. [PR1321184](#)
- The rpd might crash when two next hops are installed with the same next-hop index. [PR1322535](#)
- NCP Conf-Ack/Conf-Req packets might be dropped constantly from Cisco MLPPP client on Tomcat. [PR1323265](#)
- With **auto-installation usb** configured, interface-related commits might not take effect because of the dcd error. [PR1327384](#)
- JDID thrashes continuously and continuous log messages are observed in syslog. [PR1333632](#)
- Master LED glows on master and backup RCB, while performing the image upgrade on the master with GRES and NSR enabled. [PR1335514](#)
- The RIP route updates might be partially dropped when NSR is enabled. [PR1335646](#)

- Where intermediate CA profiles are not present on the device, in some corner cases, the pkid can become busy and stop responding. [PR1336733](#)
- The hash value generated for 256-bit key length of AES-GCM-256 algorithm is incorrect. [PR1336834](#)
- AI-script does not get auto re-install upon a Junos OS upgrade on next generation Routing Engine. [PR1337028](#)
- On MX204, MX10003, or MPC7E, MPC8E, and MPC9E, or EX9200 a 100G, 40G, or 10G interface might keep flapping or stay down due to an interoperation issue between Juniper device and the remote transport device connected. [PR1337327](#)
- MPC throughput degradation might be seen after taking SBF2 offline or online. [PR1338216](#)
- The MX10003 MPC offline button is not effective. [PR1340264](#)
- A very few of subscribers show incorrect accounting values in a large-scale subscribers scenario. [PR1340512](#)
- SNMP walk might fail for LLDP-related OIDs. [PR1342741](#)
- The RLT interface might not be able to route and forward traffic in Junos OS Release 17.3. [PR1344503](#)
- IPv6 MAC resolve fails if the DHCPv6 client uses a non-EUI64 link-local address. [PR1347173](#)
- The FPC might crash because of the MIC error interrupt hogging. [PR1348107](#)
- Routing Engine mastership keepalive timer is not updated after the GRES configuration is removed. [PR1349049](#)
- PS over rLT does not work on MPC7 and MPC9; PS over LT for the same scenario works. [PR1350115](#)
- Scope and Category fields are now available for configuring Error threshold and Action. [PR1350603](#)
- When DHCP subscribers are in BOUND (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state, if dhcp-service is restarted, then the subscribers in this state are logged out. [PR1350710](#)
- DHCP relay agent will discard DHCP request message silently if the requested IP address has been allocated to the other client. [PR1353471](#)
- DHCPv6 relay ignores replies from server when renewing. [PR1354212](#)
- **PPE errors async txrxn** error occurs when FPC is restarted or removed. [PR1350909](#)
- After GRES, the BGP neighbors at the master Routing Engine might reset, and the BGP neighbors at the backup Routing Engine take a long time to establish. [PR1351705](#)
- The DHCP relay-reply packets are dropped in the DHCPv6 relay scenario. [PR1352613](#)
- Offlining MIC6-100G-CFP2 MIC using the CLI command might trigger an FPC crash. [PR1352921](#)
- Rpd permanently hogs CPU resources because of the logical system configuration commit. [PR1353548](#)
- The 3D 40x 1GE(LAN) RJ45 MIC is not recognized on MX104. [PR1353632](#)
- Traffic interruption is observed after multiple Routing Engine switchovers. [PR1354002](#)

- The rpd generates a core file when adding an inter-region template in routing instances. [PR1354629](#)
- An aggregated Ethernet operational state goes up even though some of the member interfaces configured under the aggregated Ethernet are down. [PR1354686](#)
- There is a memory leak on agentd when Junos Telemetry Interface is configured. [PR1354922](#)
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100 Gbps). [PR1355168](#)
- Alarm LED is not working in MX204 to indicate minor and major faults. [PR1355225](#)
- Packets destined to the Routing Engine might be dropped in the kernel when LACP is configured. [PR1355299](#)
- Fabric chip failure alarms are observed in a GRES scenario. [PR1355463](#)
- VM crash might be seen when terminating the Layer 2 circuit pseudowires into EVPN. [PR1355530](#)
- Syslog messages **ui_client_connect_to_kmd_instance: KMD-SHOW connect to kmd-instance failed kmd-instance RE, fpc slot 0, pic slot 0** are seen. [PR1355547](#)
- The chassis alarm is not reflecting the right state when INP0 and INP1 has AC voltage is out-of-range. [PR1355803](#)
- The CLI **flex-flow-sizing** command is not working on MX204. [PR1356072](#)
- The mpls-IPv4 template does not have correct src AS and dst AS as 4294967295 src Mask and DstMask as 0 after adding mpls-flow table size on the fly. [PR1356118](#)
- The rpd process crashes when issuing the command **show dynamic-tunnels database terse** for RSVP automatic mesh tunnels. [PR1356254](#)
- The L2c messages from PEM or PSM are reported if SNMP is enabled. [PR1356259](#)
- DHCP subscribers fail after reconfiguration of port from tagged to untagged mode. [PR1356980](#)
- When MX Series router is used for external node slicing for enhanced subscriber management functionality, the both sets of links (master and backup) between the external x86 server and the BSYS go down, bbe-smgd process might stuck. [PR1357252](#)
- Upgrading from Junos OS Release 15.1F2-S20 to 15.1X12 using **validate** throws a fabric mixed mode error. [PR1357423](#)
- Routing Engine switchover that occurs before the backup Routing Engine is GRES ready might cause line card restart, Routing Engine kernel crash and multiple chassisd crashes. [PR1357427](#)
- Rpd memory leak occurs for RT_NEXTHOPS_TEMPLATE. [PR1357897](#)
- Traffic might be sent to an incorrect RLT member interface after RLT switchover. [PR1358320](#)
- Incorrect traffic load balance might be seen even if locality-bias is configured on MX Series Virtual Chassis. [PR1358635](#)

- The **show chassis fpc** might show **Bad Voltage** for FPC powered off by configuration or CLI command after the command **show chassis environment fpc** is executed. [PR1358874](#)
- FRU model number is not displayed for few FRUs in /component sensor for the platforms MX10008 and MX10003. [PR1359300](#)
- The scheduled boot of both the Routing Engines fail with a special time format. It only boots master Routing Engine. [PR1359602](#)
- The IGMP membership report packets might not be forwarded over an interface on QFX10,000. [PR1360137](#)
- The rpd crashes and generates a core file at Assertion failed rpd[10169]: file `../../../../src/junos/usr/sbin/rpd/lib/rt/rt_attrib.c`, line 3329: `rt_template_get_rtn_ngw(nhp) <= 1` on doing Routing Engine switchover with SRTE routes. [PR1360354](#)
- The rpd scheduler slip might be seen when frequently deleting, modifying, and adding groups that are applied on the top level. [PR1361304](#)
- The route might be stuck after BGP neighbor and route flapping. [PR1362560](#)
- Executing **show route prefix proto ip detail** during route churn in a route scale scenario might lead to FPC crash. [PR1362578](#)
- Unexpected **DCD_PARSE_ERROR_SCHEDULER** messages are logged when MS-MPC and MS-MIC are brought offline or online. [PR1362734](#)
- The non-default routing instance is not supported correctly for NTP packet in subscriber scenario. [PR1363034](#)
- Traffic destined to the MAC or IP address of VRRP VIP gets dropped on platforms that have common TFEB terminals, such as MX5, MX10, MX40, MX80, and MX104. [PR1363492](#)
- **Suppress the pq3 iic** errors related to voltage devices that are not pmbus compliant are seen. [PR1363587](#)
- Maximum value for internal interfaces through xmlproxyd is max uint32, where as it reports correctly through uint64. [PR1363766](#)
- The L2 circuit on MPC7E, MPC8E, and MPC9E with asynchronous-notification and ccc configured might keep flapping when the circuit goes up. [PR1363773](#)
- A traffic loop might occur even though the port is blocked by RSTP in a ring topology. [PR1364406](#)
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with "Link-Layer-Down" flag set. [PR1365263](#)
- MS-MPC/MS-PIC might be crash if two or more service sets configured with the same prefix lists and SIP ALG is configured in NAT scenario. [PR1366259](#)
- On MX150, upgrading to Junos OS Release 18.1R1.9 fails - installing package `nfx-2-routing-data-plane-1.0-0.x86_64` needs 76MB on the /filesystem. [PR1366324](#)
- The next hop of MPLS path might be stuck in hold state which could cause traffic loss. [PR1366562](#)

- SNMP MIB walk for UDP flood gives different output statistics than CLI. [PR1366768](#)
- The DHCP lease query message is replied to with the incorrect source address. [PR1367485](#)
- The message **system resource-monitor fpc show none-existing pfe** is seen when swapping DPC to MPC3E. [PR1367534](#)
- Incorrect RTG interface status is shown with **show interface**. [PR1368006](#)
- DHCP relay binding state - rebinding state counter added to dhcpv4 and dhcpv6 binding sensors. [PR1368392](#)
- The **commit commit check** might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)
- Kernel crash might be seen after committing a DEMUX related configuration. [PR1370015](#)
- The packet whose size exceeds 8000 might be dropped by MS-MPC in ALG scenario. [PR1370582](#)
- VRRP get stuck on MX150 XE interface. [PR1371838](#)
- FPC high CPU utilization or crash is seen during hot-banking condition. [PR1372193](#)
- The Routing Engine might crash after non-GRES switchover. [PR1373079](#)
- On MX Series platform, when dcd process restart, the BGP session might flap because of the aggregated Ethernet interface flap after the physical interface in it is detached or attached. [PR1373188](#)
- The bbe-smgd process crashes continuously and generates a core file while deleting multicast group node from the tree. [PR1374530](#)
- PCE initiated LSPs remain "Control status became local" after removing PCE configuration. [PR1374596](#)

High Availability (HA) and Resiliency

- The error **error: not enough space in /var on re1** is observed while performing a unified ISSU upgrade from Junos OS Release 17.4-20180328.0 to Junos OS Release 18.2-20180416.0. [PR1354069](#)
- Virtual Chassis-Bm cannot synchronize with Virtual Chassis-Mm when the virtual chassis splits then reforms. [PR1361617](#)

Infrastructure

- The fxp0 interface does not accept IP address with "master-only" applied. [PR1341325](#)
- The kernel might crash and the system might reboot in an SNMP query reply scenario. [PR1351568](#)

Interfaces and Chassis

- MX Series routers might occasionally drop the first LCP configure request packet when operating in a PPPoE subscriber management configuration. [PR1338516](#)
- VRRP VIP becomes unreachable after deleting one of the logical interfaces. [PR1352741](#)
- For ps ifd encapsulation must be IFLE_ETHERNET_VPLS or IFLE_ETHERNET_VLAN_VPLS, which is taken care in this PR. [PR1352933](#)

- The FPC might be stuck at 100 percent for a long time when MC-AE with enhanced-convergence is configured with large-scale logical interfaces. [PR1353397](#)
- The aggregated Ethernet interface might flap when the link speed of the aggregated Ethernet bundle is configured to oc192. [PR1355270](#)
- FPC core file related to cfmman is seen. [PR1358192](#)
- On all Junos OS products, the CLI allows you to configure more than 2048 subinterfaces on a LAG interface starting with Junos OS Release 17.2R1. [PR1361689](#)
- Error messages like **ifname [ds-5/0/2:4:1] is chan ci candidate** are seen during a commit operation. [PR1363536](#)
- The EOAM LTM messages might not get forwarded after system reboot in a CFM scenario configured with the ccc interface. [PR1369085](#)
- The dcd process might go down when **vlan-id none** is configured for interface. [PR1374933](#)

Layer 2 Ethernet Services

- The MAC address might not be learned, because of spanning-tree state "discarding" in the kernel table after Routing Engine switchover. [PR1205373](#)
- The snmpget for OID: dot3adInterfaceName might not work. [PR1329725](#)
- ZTP infrastructure scripts are not included for MX Series PPC routers. [PR1349249](#)

Layer 2 Features

- VPLS instance stays in NP state after LDP session flaps. [PR1354784](#)
- Routing Engine kernel might crash when OSPFv3 is configured with IPsec key authentication over the IRB interface. [PR1357430](#)
- The dcpfe and fxpc process might crash on Packet Forwarding Engines with low memory when allocating huge memory. [PR1362332](#)
- On MX Series routers with MPC cards, if there are large scale numbers of LSI interfaces or pseudowire paths (for example, more than 64000) in VPLS scenario, the traffic will not be transmitted correctly. [PR1371994](#)

MPLS

- When minimum bandwidth and bandwidth commands are present in the configuration, the bandwidth selection of the LSP is inconsistent. [PR1142443](#)
- Fate-sharing group cost is not set back to the default value after a CLI change to remove explicit cost configuration. [PR1330161](#)
- After an MPLS LSP link flap and local repair, a new LSP instance tries to be signaled but it might get stuck. [PR1338559](#)
- The MPLS LSP does not come up after changing admin-group mapping. [PR1348208](#)

- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. [PR1359956](#)
- L2 circuit might flap after an interface goes down even if the LDP session stays up when l2-smart-policy is configured. [PR1360255](#)
- The rpd process might crash during P2MP LSPs churn. [PR1363408](#)
- The rpd process might crash after RSVP is deactivated and then re-activated multiple times globally. [PR1366243](#)
- The rpd might crash in BGP LU and LDP scenario. [PR1366920](#)
- The traceroute **mpls ldp** to Huawei fails until TTL expired. [PR1372924](#)

Multicast

- Some IGMP groups might have incorrect upstream interface because discard route is installed in PIM. [PR1337591](#)

Network Management and Monitoring

- The jnxDcuStatsEntry and jnxScuStatsEntry OIDs are missing after interface configuration changes. [PR1354060](#)
- SNMP process crashes during polling CFM statistics. [PR1364001](#)

Platform and Infrastructure

- The apply-path prefix is not inherited under the policy after commit. [PR1286987](#)
- MPC5 - inline-ka PPP echo requests are not transmitted when anchor-point is lt-x/2/x or lt-x/3/x in a pseudowire deployment. [PR1345727](#)
- Running RSI through the console port might cause the system to crash and reboot. [PR1349332](#)
- When viewing IPv6 addresses, **display rfc5952** does not work when combined with **display set**. [PR1349949](#)
- The VCP port might not come back up after removing it and adding it again. [PR1350845](#)
- Kernel crash occurs because the initialization of logical interface MAC filter function is missing for Packet Forwarding Engine extended port devices. [PR1353498](#)
- The FPC might crash because of the memory leak caused by the VTEP traffic. [PR1356279](#)
- Traffic black hole is seen along with **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region** for next-hop messages. [PR1357707](#)
- When **forwarding-class-accounting** statement is enabled, on an interface, inside of a routing instance of instance-type vrf, aggregate input forwarding-class statistics do not increment (egress statistics work fine). [PR1357965](#)
- The CLI functions are not triggering properly with **set security ssh-known-hosts load-key-file**, **set system master-password**. [PR1363475](#)

- Same vlan-id is not allowed on multiple logical interfaces of the same GR interface. [PR1365640](#)
- Qmon sensors are not working with hyper-mode is enabled. [PR1365990](#)
- Subscribers over aggregated Ethernet interface have tail drops affecting fragmented packets because of the QXCHIP buffer getting filled up. [PR1368414](#)
- On MX Series routers, if CoS rewrite is enabled globally at the chassis level, and there are control packets generated locally by Routing Engine that should be sent through the logical tunnel (LT) interface, all the control packets will be dropped in kernel. Because of this issue, any control packets generated locally by Routing Engine cannot be sent from the LT interface, and this will affect control protocol handshake on the LT interface, which thereby affects traffic. The transit control packets are not impacted. [PR1372738](#)
- JNH memory leaks in multicast scenario with MoFRR enabled. [PR1373631](#)

Routing Policy and Firewall Filters

- The policy might not clean up after deleting the configuration and cause rpd to generate a core file. [PR1357724](#)

Routing Protocols

- Multihop eBGP peering session exchanging EVPN routes might result in generating a rpd core file when BGP updates are sent. [PR1304639](#)
- Mcsnoopd process memory leak is observed. [PR1326410](#)
- The rpd process might crash and generate a core file while running streaming telemetry. [PR1347431](#)
- The rpd process might crash and generate a core file is seen after executing the Routing Engine switchover. [PR1349167](#)
- The rpd process might crash when BGP route damping and the BGP multipath feature are configured. [PR1350941](#)
- Static route flaps on commit when configured with resolve configuration statement. [PR1366940](#)
- If IS-IS shortcut is enabled and IS-IS **topologies ipv6-unicast** is configured, when any link with no IPv6 address configured in the MPLS LSP path is flapping (or bring down and then up), the route entry go through this flapping link might be missing for about 10 minutes, which might lead to traffic loss. The issue is because when the flapping link is down and then up, the flash route update checks both IPv4 and IPv6 address family, since IPv6 is not configured for this link, the flash route update is not triggered, hence the route entry is missing. [PR1372937](#)

Services Applications

- The jl2tpd might crash if the RADIUS server returns 32 tunnel-server-endpoints. [PR1328792](#)
- Selectively start ZLB delay timer at the Packet Forwarding Engine for LAC tunnels. [PR1338450](#)
- While performing an SNMP walk on the IKE SA that is deleted, IPsec tunnels might go down and an infinite loop scenario might be seen. [PR1348797](#)

- IPsec tunnels go down when SNMP walk is executed in a scaled scenario which polls IPsec phase 2 statistics. [PR1353240](#)
- NAT64 does not translate ICMPv6 Type 2 packet (packet is too big) correctly when MS-DPC is used for NAT64. [PR1374255](#)

Software Installation and Upgrade

- When the device is booted into single-user mode (recovery mode), and any change in configuration is made, such as setting the root password, then the commit will fail. [PR1368986](#)

Subscriber Access Management

- In a dual-stack subscribers scenario with NDRA pool configured, the linked pools are not used when the first NDRA pool is exhausted. [PR1351765](#)
- When attempting to scale clients `sdbsts_lock_holder.bbe-smgd.pid10686.core` core files are seen. [PR1358339](#)

Subscriber Management and Services

- PPPoE cannot dial in because the PADI messages dropped as **unknown iif** when aggregated Ethernet configuration is deactivated or activated. [PR1291515](#)
- The pfd process crashes and generates a core file in `pfd_process_session_state_notification_msg`, `pfd_timer_manager_c::remove_serv_id`, `pfd_delete_timer_id_by_serv_sid` (`serv_sid=0`, `serv_info=0x0`) at `../../../../src/junos/usr.sbin/pfd/pfd_timer.cc:16`. [PR1296969](#)
- The framed-route "0.0.0.0/0" is not installed in MX Series routers with Junos OS enhanced subscriber management releases. [PR1344988](#)
- In certain scenarios on MX Series Virtual Chassis with L2TP LNS, the DHCPv6 solicit packet might be dropped. [PR1348846](#)
- High CPU usage by the bbe-smgd process might be seen when L2BSA subscribers get stuck. [PR1351696](#)
- JI2tpd process might crash shortly after one of the L2TP destinations becomes unavailable. [PR1352716](#)
- Syslog error `dfw_bbe_filter_bind:1125 BBE Filter bind type 0x84 index 167806251 returned 1` is observed. [PR1354435](#)
- The ifinfo process might crash in MX Series BNG running L2BSA service. [PR1354712](#)
- The static-subscribers do not properly update firewall information on the Packet Forwarding Engine when dynamic configuration changes are made to active subscribers. [PR1354774](#)
- L2TP tunnel-switch clients in subscriber session database reference the incorrect routing instance. [PR1355396](#)
- The CLI command `show pppoe underlying-interfaces` in a scaled environment might cause bbe-smgd memory leak. [PR1356428](#)

- MPC or FPC might be unable to reply with request messages to Routing Engine in a high subscriber scale scenario. [PR1358405](#)
- In a system with a scale in excess of 80,000 dual-stack L2TP (LNS) subscribers, bbe-smgd processes might crash after a GRES followed by a FPC panic/restart. [PR1359290](#)
- The IPv6 subscriber might not be authenticated successfully and fail to access the network because of missing attributes (agent-circuit-id & agent-remote-id) for radius-access-request packets if they are the needed parameters for radius authentication. [PR1359520](#)
- On PPC-based MX Series platforms (MX104), packet-triggered subscribers and policy control (PTSP) subscribers will not be able to re-login on the BNG device after an initial login attempt (either successful or failed). [PR1359574](#)
- When an on-demand address allocation is activated for dual-stack PPPoE subscribers, the clients might fail to get IPv4 address after reinitializing the IPv4 connection and keeping the IPv6 connection at the same time. [PR1360846](#)
- In subscriber management scenarios with PPPoE access models, during unified ISSU, it is possible to lose a small number of active subscribers after the unified ISSU is completed if certain timing conditions occur. These timing conditions might trigger session database related discrepancies between the jpppd daemon and the underlying statesync infrastructure causing the subscriber record loss. These subscribers, however, should be able reconnect right away minimizing any service outage. [PR1360870](#)
- In a subscriber management environment, because of the timing issue bbe-smgd process on the backup Routing Engine might crash during either login of a subscriber with a multicast service, or activation of multicast service for an existing subscriber. [PR1362188](#)
- L2TP Access Concentrator (LAC) tunnel connection request packets might be discarded on the LNS device. [PR1362542](#)
- In MX Series router running Junos OS enhanced subscriber management feature, if the service dynamic profile variable name and the associated default value is configured to be the same, the MX Series routers will suffer very rapid memory leak in bbe-smgd and exhausts the Routing Engine memory. [PR1362810](#)
- Some subscriber might be stuck in terminating state in an L2TP scenario. [PR1363194](#)
- The L2TP subscribers might not be able to log in successfully due to the jl2tpd memory leak. [PR1364774](#)
- Subscriber filter is not removed from Packet Forwarding Engine when routing-services are enabled in the dynamic profile on an L2TP LNS. [PR1369968](#)
- Subscriber cannot negotiate MLPPP session with MX Series LNS when dynamic-profile name contains more than 30 characters. [PR1370610](#)
- Actual data rate downstream value is not included in the L2TP ICRQ message from the LAC. [PR1370699](#)
- SMGD process crashes and generates a core file after essmd restart with reference to `mmf_ensure_mapped (mmf=0xe8f0200, offset=4294967295, len=108)` at `../src/junos/lib/libmmf/mmf.c:1972`. [PR1372223](#)

User Interface and Configuration

- The **max-db-size** configuration might not work on some MX Series platforms. [PR1363048](#)

VPNs

- The rpd process might crash and generate a core file during configuration change. [PR1351386](#)
- In dual-homed next-generation MVPN the receipt of Type 5 withdrawal removes the downstream join states for some routes. [PR1368788](#)

Resolved Issues: 18.1R2

Class of Service (CoS)

- CoS wildcard configuration is applied incorrectly when the router restarts. [PR1325708](#)
- Remove CoS IDL from the jet IDL package and update the documentation for the same. [PR1347175](#)
- The Routing Engine might get into amnesiac mode after restarting when **excess-bandwidth-share** is configured. [PR1348698](#)

EVPN

- On deactivated ESI for PS at physical interface level, routing protocol process crashes and generates a core file for EVPN VPWS PWHT. [PR1332652](#)
- In an EVPN and NSR environment, the routing protocol process (rpd) crash and generates a core file on backup Routing Engine for any configuration changes on master Routing Engine. [PR1336881](#)
- The rpd process might crash when executing CLI command **show route evpn-ethernet-tag-id**. [PR1337506](#)
- In an EVPN and VXLAN environment, Packet Forwarding Engine crashes when BFD and VTEP flap. [PR1339084](#)
- Bring the IRB logical interface UP, even if L2 interfaces are absent but IM next hop is present. [PR1340723](#)
- The rpd might crash if the IRB interface and routing instance are deleted together in the same commit. [PR1345519](#)

Forwarding and Sampling

- The error messages about dfw_gencfg_handler might be seen during unified ISSU. [PR1323795](#)
- DHCP service crashes after EX9251 switch is set to factory default by zeroize [PR1329682](#)
- The error logical interface under VPLS might be blocked after MAC moves if the logical interfaces are on the same physical interface. [PR1335880](#)
- The l2ald crashes when a duplicate MAC is learnt by two different interfaces. [PR1338688](#)
- Commit failed when attempting to delete any demux0 unit numbers which are greater or equal to 1000000000. [PR1348587](#)

General Routing

- Memory leak is causing the rpd crash. [PR1052614](#)
- Unexpected MobileNext Gateway Activation license alarms when TDF gateway is configured. [PR1162518](#)
- SNMP trap sent for "PEM Input failure" alarm is not generated when a single input feed fails on MX960. [PR1189641](#)
- High priority fabric drops from MPC7E towards MPC3E next-generation. [PR1207417](#)
- High priority fabric drops from MPC7E towards MPC3E next-generation. [PR1226804](#)
- The error log messages **cc_mic_irq_status: CC_MIC(5/2) irq_status(0x1d) does not match irq_mask(0x20), enable(0x20), latch(0x1d)** are seen continuously for "MIC-3D-4OC3OC12-1OC48". [PR1231084](#)
- BSYS logs GNF owned pics does not support power-off configuration at commit when no such configuration is present. [PR1281604](#)
- On EVPN or VXLAN, inter-vrf traffic blackhole is seen after the routing is restarted repeatedly on redundant gateways. [PR1289091](#)
- The log message about shutdown time is incorrect when system exceeds chassis over temperature limit. [PR1298414](#)
- Utilization of **commit check** just after setting master-password can trigger improper decoding of configuration secrets. [PR1310764](#)
- The incorrect error number might be reported for syslog messages with a prefix of **%DAEMON-3-RPD_KRT_Q_RETRIES**. [PR1310812](#)
- The transient error **hawkeye alarmd** are observed on MX240, MX480, and MX960 Series routers. [PR1312336](#)
- The MPC with specific failure hardware might impact other MPCs in same chassis. [PR1319560](#)
- The rpd might crash when two next hops are installed with the same next hop index. [PR1322535](#)
- MS-MIC interface logical interfaces remain down after many offline or online iterations. [PR1322854](#)
- CLI command **request vmhost halt routing-engine other** does not halt the backup Routing Engine. [PR1323546](#)
- The snmp interface filter does not work when "interface-mib" is part of dynamic-profile. [PR1324573](#)
- Constant logging of **fm_feacap_sys_feature_get:Attribute DB init not yet done, reading from pvid (id: 18)**. [PR1328868](#)
- When an AMS bundle has a single MAMs added to it, the subinterfaces do not recover after the subinterface has been disabled. [PR1329498](#)
- Host-Outbound traffic do not rewrite ieee-801.pbits for dynamic subscriber logical interface over PS interface. [PR1329555](#)
- SNMP walks of Interfaces related MIB objects are slower than expected in a scaled configuration. [PR1329931](#)

- **Too many supplies missing in lower and upper zone** alarm flaps (set/clear) every 20 seconds if a zone does not have minute required PSMs. [PR1330720](#)
- In rare cases, a highly scaled node-slicing Guest Network Function (GNF) might fail to complete NSR replication after a NSR switch over. This is true even if the GNF is confirmed NSR ready before the switch occurs. [PR1331145](#)
- Two subscribers cannot reach the online state at the same time if they have an identical frame-route attribute value. [PR1334311](#)
- Hitless key-chain rollover feature has limitations when used on MIC-MACSEC-MRATE. [PR1335644](#)
- The MAC_STUCK might be seen on MS-MPC or MS-MIC linecards. [PR1335956](#)
- On MX2000 with SFB card installed, high amount of traffic volume on MPC7E, MPC8E or MPC9E might cause traffic drops with cell underflow messages. [PR1336446](#)
- FPC temperature might mismatch for MPC6, MPC8, and MPC9 on MX2000 platform. [PR1339077](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- CLI shows CB states online after pressing RCB offline button for more than four seconds. [PR1340431](#)
- VRRP sticks in master on upgrade or cold boot. [PR1341044](#)
- When discard interfaces are configured with IGMPv3, KRT queue gets stuck while deleting multicast next hop (MCNH) with error **EPERM -- Jtree walk in progress**. [PR1342032](#)
- The FPC is marked as down and stay down on the MX150 platform causing service effecting loss of traffic. [PR1343170](#)
- In MPLS or RSVP environment, LSP might stick in Dn state with **Record route: <self> ...incomplete**. [PR1343289](#)
- Queue counters are not getting displayed in the interface details for PORTER-R once the system reboots. [PR1343306](#)
- Support required for **show system resource-monitor subscribers-limit chassis extensive in Summit**. [PR1343853](#)
- MPC8 and MPC9E crashes and generates a core file with DHCPv6 on static VLAN logout. [PR1343965](#)
- The l2cpd process might generate a core file on executing **l2cpd_ifbd_attach** command (where ifbd=0x98914c0, vlan_id=1, line_vid=1) after disabling mc-ae on qfx10002-60c {default vlan-scenario} which is getting hit where delete is missed by l2cpd because it uses Sync socket read when it starts. [PR1344983](#)
- The routing protocol process (rpd) crash might be seen if the **no-propagate-ttl** configuration statement is set in a routing instance which has a specific route. [PR1345477](#)
- MAC address of multiple interfaces are found to be duplicate. [PR1345882](#)
- Routing Engine model is changed from JNP10003-RE1 to RE-S-1600x8. [PR1346054](#)
- New PPPoE users might fail to login. [PR1346226](#)

- VCCP-ADJDOWN detection is delayed on VC-Bm when deleting one VCP link on VC-Mm. [PR1346328](#)
- The twice-napt-44 sessions are not synchronizing to backup SDG with **stateful sync** configured. [PR1347086](#)
- Remove libstdc++ dependency on hypervisor to install JDM rpm or deb package. [PR1347921](#)
- Packet loop is detected when VRF multipath is enabled with **equal-external-internal** configuration statement under L3VPN instance and install-nexthop is enabled in forwarding-table export policy regarding that l3vpn route. [PR1348175](#)
- The **get config** configuration statement for hidden choices is not working with ODL controller. [PR1348503](#)
- MACsec ACK validation is added for boundary condition check and invalid values (entering '0x0' or '0x0000000000000000' as an error). [PR1348642](#)
- Chassisd memory leak issue is seen on MX10003 and MX204 platform and it might eventually switchover Routing Engine and crash. [PR1348753](#)
- MGD core files are generated because of the issue in nsindb infrastructure. [PR1349288](#)
- The error message **Major PEM 0 Input Failure** might be observed for DC PEM. [PR1349179](#)
- Access internal routes remain even after AIU is completed. [PR1350401](#)
- The MTU value for subscriber's interface might be programmed incorrectly if the statements **routing-services** or **protocol pim** is configured in dynamic-profile. [PR1350535](#)
- The subinfo process might crash when executing **show subscribers address <> extensive** for a DHCPv6 address. [PR1350883](#)
- Dynamic physical interface creation fails when the SFP optic is plugged in MX150. [PR1351387](#)
- Node virtualization MSE after reinstalling one JDM server complains pull configuration failed, fallback to push configuration method. [PR1352503](#)
- On MX Series routers show chassis fpc errors does not show errors on GNF systems. [PR1352705](#)

High Availability (HA) and Resiliency

- The ksyncd process might crash continuously on the new backup Routing Engine after performing GRES. [PR1329276](#)
- Insufficient available space on hard disk lead by the crash information files is generated by ksyncd when GRES is configured in large-scale configuration scenario. [PR1332791](#)

Interfaces and Chassis

- L2TP subscribers might not be cleared if the access-internal routes fail to install. [PR1298160](#)
- MPC CPU might reach 100 percent when **otn ufec** statement is configured. [PR1311154](#)
- No route to IP address from directly connected route. [PR1318282](#)

- Unexpected log messages might be seen if a BGP session flaps in a dynamic-tunnels GRE scenario. [PR1326983](#)
- Unexpected log messages might be seen on a router for subscriber management. [PR1328251](#)
- The cfmd process core files are generated. [PR1329779](#)
- The dcd process might crash because of the memory leak and causes commit failure. [PR1331185](#)

Layer 2 Ethernet Services

- The memory leak might occur in l2cpd if the l2-learning process is disabled. [PR1336720](#)
- DHCP client is not able to connect if VLAN was modified on aggregated Ethernet interface associated with the IRB. [PR1347115](#)
- Restart FPC which host micro-bfd link might cause lacp to generate a core file. [PR1353597](#)

Layer 2 Features

- The rpd process memory leak is observed upon any changes in VPLS configuration like deleting or re-adding VPLS interfaces. [PR1335914](#)

MPLS

- Whenever there is a decrease in the statistics value across an LSP, the mplsLspInfoAggrOctets value take two intervals to get updated. [PR1342486](#)
- An LDP label is generated for serial interface subnet route unexpectedly. [PR1346541](#)
- In a very rare scenario, rpd might crash when LDP fails to allocate self-id for the P2MP FEC. [PR1349224](#)
- The rpd crash might happen in RSVP setup-protection scenario. [PR1349036](#)

Network Management and Monitoring

- The eventd process fails to startup with syslog configuration. [PR1353364](#)
- SNMP stops or becomes very slow after a very long period of time. [PR1328455](#)

Platform and Infrastructure

- MX204 performance degrades when using firewall filter with sampling action. [PR1303529](#)
- VPLS instance fails to learn MAC addresses upon pseudowire switchover. [PR1316459](#)
- Directories and files under /var/db/scripts lost execution permission or directory 'jet' is missing under /var/db/scripts causing **error: Invalid directory: No such file or directory** error during commit. [PR1328570](#)
- The TCP dump filter might not work in egress direction on PS and It logical interfaces. [PR1329665](#)
- Commit might fail with error reading from commit script handler. **commit script failure.** [PR1335349](#)
- While downgrading MX Series router from a later release, the router goes into amnesiac state. [PR1341650](#)
- Configuring the same DHCP server in different routing-instances is not supported in DHCP relay scenario. [PR1342019](#)

- Commit error on configuring same vlan-id on different logical interface of the same physical interface when **ethernet-bridge encapsulation** is configured. [PR1342229](#)
- Route corruption in Packet Forwarding Engine with **connectivity-fault-management** enabled for Layer 2 circuit. [PR1342881](#)
- ZTP is not supported for vmhost images on next generation Routing Engines on the MX Series platforms. [PR1343338](#)
- IPv4 GRPS traffic over aggregated Ethernet interface might be affected if enhanced hash key **gtp-tunnel-endpoint-identifier** is configured. [PR1347435](#)
- On MX Series routers, in an EVPN-VXLAN output policing action does not work on IRB interfaces for VNIs. [PR1348089](#)
- FPC CPU utilization with LT interfaces is pegged continuously at 100 percent. [PR1348840](#)
- ICMP error messages are not generated if 'don't fragment' packets exceed the MTU of the multiservice interface. [PR1349503](#)
- The CLI commands **system ddos-protection protocols unclassified** are missing on MX2020. [PR1349782](#)
- Suspect memory leak in chassisd. [PR1353111](#)

Routing Policy and Firewall Filters

- Access-internal route might fail to be leaked between routing instances when "from instance" is configured in the policy. [PR1339689](#)

Routing Protocols

- BGP extended communities with sub-type 4 erroneously displayed at LINK_BANDWIDTH. [PR1216696](#)
- The routing protocol process (rpd) generates a core file in the ASBR when BGP is deactivated in the ASBR before all stale labels have been cleaned up. [PR1233893](#)
- BGP traceoption logs are still written when it is deactivated. [PR1307690](#)
- In IS-IS and IPv6 scenario, rpd might crash when the neighbor router is restarted and causes routes to churn. [PR1312325](#)
- The primary path of MPLS LSP might switch to other address. [PR1316861](#)
- The inactive route cannot be installed in multipath next hop after disabling and enabling the next-hop interface in L3VPN scenario. [PR1317623](#)
- Traffic might get silently dropped and discarded temporarily when BGP GR is triggered and the direct interface flap. [PR1319631](#)
- When tracing BGP routes that contain the DF election community, BGP communities after this community might not display properly. [PR1323596](#)
- Manual GRES with MX Series Virtual Chassis results in some packet loss on core facing interfaces. [PR1329986](#)

- The LDP route in inet.3 is missing when both OSPF rLFA and LFA protections are available and rejected by backup selection policy. [PR1333198](#)
- In next-generation subscriber management, IGMP joins are not processed with **passive allow-receive** statement configured on IGMP interface. [PR1334913](#)
- The routing protocol process (rpd) generates a core file during delete and restore of BGP configuration. [PR1338567](#)
- Changes to the displayed value of AIGP is seen when **show route ... extensive** command is executed. [PR1342139](#)
- Traffic black-hole might be seen if local DUT receives BFD-down. [PR1342328](#)
- The rpd might crash when BGP flaps. [PR1342481](#)
- The rpd might crash if a route for RPF uses a qualified-next-hop. [PR1348550](#)
- The routing protocol process (rpd) might crash while restarting routing or deactivating IS-IS. [PR1348607](#)
- **source-as community** is not appended to rendezvous point (RP). A display issue is observed in **show route detail** command output. [PR1353210](#)

Services Applications

- SNMP MIBs do not yield data related to sp- interfaces. [PR1318339](#)
- Crash at `../src/junos/lib/libjuniper/mgmt-sock/mgmt_sock_select_info.c:35`. [PR1337406](#)
- UDP checksum inserted by MS-DPC after NAT64 is not valid when incoming IPv4 packet has UDP checksum set to 0. [PR1350375](#)
- The **show services stateful-firewall flows counter** shows ridiculously high numbers. [PR1351295](#)

Software Installation and Upgrade

- New versions of Junos OS does not have the tool for accessing aux port - `/usr/libexec/interposer`. [PR1329843](#)

Subscriber Access Management

- The L2TP LAC might drop packets that have incorrect payload length while sending packets to the LNS. [PR1315009](#)
- The multiple RADIUS servers having different dynamic request port is not supported. [PR1330802](#)
- Traffic drops on the MX Series router LNS because of software error or unknown family exception when traffic is destined to or coming from MLPPP subscriber if **routing-services** configuration statement is present in the dynamic-profile used by this subscriber. [PR1335276](#)
- The subscriber might get stuck in terminated state when JSRC synchronize state get stuck in "FULL-SYNC in progress". [PR1337729](#)
- MX Series router is sending IPv6 RA and the DHCPv6 advertisements before IPCPv6 Ack from CPE. [PR1344472](#)

- The ancpd process might generate a core file when clearing ancp subscribers in a scaled scenario when **enhanced-ip** is configured. [PR1344805](#)
- The bbe-smgd process might crash if there are 65535 L2TP sessions in a single L2TP tunnel. [PR1346715](#)
- Subscriber might experience SDB DOWN event and drop the clients' connections when issuing **show subscribers** command. [PR1336388](#)
- LNS subscribers on aggregated-inline service scale impacted. [PR1341659](#)
- **AC system error** counter in **show pppoe statistics** does not work. [PR1346231](#)
- The pfd process consumes 80 to 90 percent CPU running subscriber management on PPC based routers. [PR1351203](#)

VPNs

- The multicast might be rejected when Junos OS PE devices received C-Mcast route from other vendors' PE devices. [PR1327439](#)
- MVPN sender-site configuration is not allowed with S-PMSI. [PR1328052](#)
- The routing protocol process (rpd) crashes after committing interface related parameters (for example, MTU change, VRF RD/RT, QOS) on PS interface with vlan-ccc encapsulation and no vlan-id. [PR1329880](#)
- The routing protocol process (rpd) might continuously crash on the backup Routing Engine and some protocols might flap on the master Routing Engine if **hot-standby** is configured for I2circuit or VPLS backup-neighbor. [PR1340474](#)
- The rpd might crash on backup Routing Engine while changing the I2circuit virtual-circuit-id in an NSR scenario. [PR1345949](#)

Resolved Issues: 18.1R1

Application Layer Gateways (ALGs)

- IKEv2 negotiation might fail with IKE ESP ALG enabled in IKEv2 redirection scenario. [PR1329611](#)

EVPN

- EVPN traffic does not map to a specific LSP in the core. [PR1281415](#)
- BGP route refresh request might not be sent when the route target is modified. [PR1300332](#)
- Split horizon label is not allocated when the configuration of ESI is switched from **single-active** to **all-active**. [PR1307056](#)
- Discard EVPN route is installed on local PE device when the connection flaps on a remote PE device in a multihome EVPN topology. [PR1321125](#)
- FPC might stop functioning properly while deleting the VPLS configuration having the **no-tunnel-service** command enabled from the routing instance. [PR1324830](#)

- The core link flapping might result in an inconsistent global MAC count. [PR1328956](#)
- On restarting the router using **restart routing**, the rpd process generates a core file in provider edge (PE) router that has EVPN-VXLAN configuration. [PR1333331](#)

Forwarding and Sampling

- When subscriber services that are enabled for interim volume accounting goes down, the pfd process rarely generates a core file with the backtrace pfd_timer_manager_remove_serv_id. [PR1296969](#)
- There is a memory leak on mib2d when firewall MIBs are polled. [PR1302553](#)
- The remote CE1 MAC address might take along time to meaning not clear. [PR1304866](#)
- In a subscriber management environment, when the **show firewall templates-in-use** command is executed, dfwd process might crash during the execution if a CLI session disconnects before the complete output of this command is received. [PR1305284](#)
- ACCT_FORK_LIMIT_EXCEEDED log level is an error even when backup-on-failure feature is enabled for accounting files. [PR1306846](#)
- Second archive site in the accounting file configuration is not used when the first one uses SFTP protocol and is not reachable. [PR1311749](#)
- Accounting files with no records might be unexpectedly uploaded to the archive site. [PR1313895](#)
- The log message **dfwinfo: tvptest:dfwi_counter_output policer_byte_count support 0** might be noticed when issuing **show firewall** related command. These logs are harmless and intended for debug purpose.. [PR1315730](#)
- The commit might fail if the **next-hop learning** configuration statement is enabled for J-Flow v9. [PR1316349](#)
- The FPC CPU might reach 100 percent constantly when the shared bandwidth policer is configured. [PR1320349](#)
- Some firewall filter counters might not be created in SNMP. [PR1335828](#)

General Routing

- Maintenance association end-points (MEPs) persist to generate continuity check message (CCM) frames, after they are deleted from protocols OAM Ethernet CFM stanza. [PR1107542](#)
- Memory leak is seen on Layer 3 VPN configuration commit for L3VPN scaling test. [PR1115686](#)
- No warning is raised when the bridge family is configured with an interface-mode trunk but without VLAN tagging or flexible VLAN tagging. [PR1154024](#)
- Ksyncd process might not respond because of transient replication errors between Routing Engines. [PR1161487](#)
- Stale VBF states occur without SDB sessions. [PR1204369](#)

- Unable to deregister sub error (131072) for error (0x1b0001) for module MIC error messages seen on MPC5E card. [PR1221337](#)
- Changing the virtual switch interface type from IRB to regular bridge interfaces under the OpenFlow protocol are removed. The OpenFlow process (daemon) fails to program any flows. [PR1234141](#)
- The **multicast-replication** setting cannot be reflected in the redundancy environment after both Routing Engines are rebooted. [PR1240524](#)
- In a BGP and 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](#)
- The **chassisd[9132]: LIBJSNMP_NS_LOG_NOTICE: NOTICE: netsnmp_ipc_client_connection: unix connection error: socket(-1) main_session(0x9812f80)** error messages are seen after chassis control restarts. [PR1243364](#)
- Prolonged flow control core file is observed for the TFTP ALG traffic (10K simulated users). [PR1255973](#)
- When you plug in an SFP or SFP+ transceiver or remove it from any of the supported ports on an MX150, the ge-0/0/0 interface goes down and cannot be used. [PR1259112](#)
- GNF sometimes resets its MPC type 9 at NSR at a high scale. [PR1259910](#)
- The virtual MX Series router with FPC generates a core file - panic (format_string=format_string@entry=0x9e509c4 "Thread %s attempted to %s with irq priority at %d\n"). [PR1263117](#)
- Monitoring FPC temperature is not applicable on MX1RU platform, because MX1RU is a single board design with logical FPC. [PR1263315](#)
- On MX Series, the **show chassis led** command should not be displayed in possible completions of the **show chassis** command. [PR1268848](#)
- A low-memory condition putting the service PIC into the red zone on the MS-MIC or MS-MPC might cause the SIP ALG to generate a core file. [PR1268891](#)
- The load-based throttling is not enabled. [PR1271739](#)
- Aggregated Ethernet incorrect counters for output packets on child links for ae0 interface when configured with new feature 'revertive'. [PR1273983](#)
- On an MX104 platform with GRES enabled, the chassis network-services might not get set as "Enhanced-IP". [PR1279339](#)
- The jfirmware upgrade support is not available for Routing Engine BIOS. [PR1281050](#)
- BSYS logs GNF owned pics does not support power-off configuration at commit when no such configuration is present. [PR1281604](#)
- The kernel crash might happen in a rare corner case. [PR1282573](#)
- In a specific CE device environment in which asynchronous-notification is used, when the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. [PR1282875](#)

- The total number of corrected single-bit errors from HMC [x] exceeds the threshold value of 32. [PR1285315](#)
- LC, PFH, and Packet Forwarding Engine interfaces do not come up on the RE1. [PR1285606](#)
- A missing statement “Shared bandwidth policer not supported for interface ge-x/x/x” is noticed when a commit is unsuccessful in Junos OS Release 16.1R3. [PR1286330](#)
- During unified ISSU (FRU upgrade) micro BFD flapping is observed. [PR1288433](#)
- The interfaces might go to a down state after performing GRES. [PR1289493](#)
- The **request system zeroize** command deletes the **/var/db/scripts** directory, which is not re-created until the next USB or netboot recovery. [PR1289692](#)
- **jnxContainersType** MIB is not displayed for MX Series MICs and PICs as correctly as it is displayed on other Juniper Network platforms. [PR1289778](#)
- Incorrect temperature is displayed for MPC5 and MPC7 in the **show chassis fpc** command output. [PR1290771](#)
- The traffic traversing a label-switched path (LSP) with entropy label might get dropped after the bypass path goes down. [PR1291036](#)
- The routing protocol process (rpd) might generate a core file while it is restarted from the CLI. [PR1291110](#)
- The L2TP ICCN fast retransmission occurs after tunnels go down. [PR1291557](#)
- When GRES is enabled, restarting the chassisd process results in FPC restarting multiple times. [PR1293314](#)
- On an MPC6E, with inline flow monitoring enabled, the flow export rate remains less as compared to the configured export rate. [PR1294296](#)
- During PPPoE subscriber login, errors such as [**vbf_flow_src_lookup_enabled**] and [**failed to find iff structure, ifl**] were seen on the FPC. [PR1294710](#)
- The KRT queue might get stuck with the error **RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0**. [PR1295756](#)
- A [First_Net] commit error is thrown when you try to commit a configuration with the applied groups. [PR1298649](#)
- MX Series BNG does not respond to PADI after GRES on some ports or VLANs. [PR1298890](#)
- In certain conditions, the maximum count is reached, if a limit configured by a subscriber does not allow a second family of DHCP dual stack. [PR1298924](#)
- Software enhancements are made to AC NON-HC PEM that suppresses the I2C bus errors for PEM. [PR1299284](#)
- The asynchronous-notification feature cannot be implemented properly in a circuit that has MIC-3D-20GE-SFP-E or Tri Rate Copper SFP(740-013111). [PR1299574](#)
- ICMP or ICMPv6 error messages might be discarded while getting forwarded through an AMS interface. [PR1301188](#)

- A configured logical interface might not be created correctly after the configuration is committed. [PR1301823](#)
- In Junos Telemetry Interface (JTI) setup, the payload MTU might be much less than 16 KB when subscribing to a component sensor. [PR1301835](#)
- Duplicate keys are no longer exported by the physical interface related to Packet Forwarding Engine; these are only exported by MIB2D. [PR1301858](#)
- The rpd process might crash when NSR is enabled and the routing-instance specific configurations are committed. [PR1301986](#)
- Continuous interface flapping might lead to unwanted resetting of the MIC. [PR1302246](#)
- The rpd process might crash when the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statements are toggled. [PR1302504](#)
- The **chassisd.core-tarball.0.tgz** file is found during unified ISSU is aborted during the upgrade of a FRU. [PR1303086](#)
- Incorrect MTU might be seen on PPP interfaces when PPP MTU is not defined in the dynamic profile. [PR1303175](#)
- The inline-ka PPP echo requests are no longer generated for aggregated Ethernet interfaces. [PR1303249](#)
- The **request auto-configuration reconnect-pending** command is no longer available. [PR1303336](#)
- Blocking PPPoE or DHCP to initiate VLAN autosensing when the VLAN-OOB is in pending state. [PR1303338](#)
- Fan speed changes frequently on MX Series chassis. [PR1303459](#)
- When MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic, generating the error message **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)
- MX Series MIB polling returns a value that has **sdg**. Polling result should include **svc** generic value. [PR1303848](#)
- Truncated output appears for the **show pppoe lockout** command. [PR1304016](#)
- When either a MPC6E or SFB2 restarts, you might see link errors and training failure between fabric planes. [PR1304095](#)
- Effective rate of E3 in framed mode is limited to 30 Mbps on certain channelized MICs. [PR1304344](#)
- The RPF check strict mode causes traffic drop in the next-generation subscriber management release. [PR1304696](#)
- On an MX2000 with MPC9E and SFB2 installed, a certain high amount of traffic volume might cause traffic drops and generate cell underflow messages. [PR1304801](#)
- Commit fails with error: **ffp_intf_ifd_hier_tagging_config_verify: Modified IFD "si-1/1/0" is in use by BBE subscriber**, active L2TP LNS client. [PR1304951](#)
- The MX Series router sends immediate-interim reports for the services pushed by SRC. [PR1305425](#)

- When traceoptions are enabled on 32-bit Junos OS, the rpd process might generate core files. [PR1305440](#)
- JET daemonize application gets respawned even on normal exit. [PR1305615](#)
- The LIBJNX_REPLICATE_RCP_ERROR message is seen in the syslog when a backup Routing Engine is not present. [PR1305660](#)
- L2BSA subscriber's connection attempts failed with VLAN **profile-request-error**. [PR1305962](#)
- The network FPC command **start shell Packet Forwarding Engine** is not working on MX960. [PR1306236](#)
- L2BSA subscribers are not able to connect, and no new ANCP session get established during the RADIUS disaster backup procedure. [PR1306872](#)
- The smihelperd process generates core files when SNMP is polling for JUNIPER-SUBSCRIBER-MIB::jnxSubscriberGeneral.7.0. [PR1306966](#)
- The kmd process error **UI_DBASE_OPEN_FAILED** is seen because of too many open files. [PR1308380](#)
- License is lost during Routing Engine switchover in scale-subscriber scenario. [PR1308620](#)
- CoS applied to a subscriber DEMUX logical interface does not work. [PR1308671](#)
- All the MICs on one MPC, with PWHT subscribers configured, might go offline during the restart of an MPC installed in another slot. [PR1308995](#)
- Error messages **%PFE-3: fpc0 vbf_var_iflset_add:633: vbf container 11 not found in the msg for ifl .demux.6514** are often seen after MPC restart. [PR1309013](#)
- Incorrect timestamp values are found when RADIUS accounting stops packets. [PR1309212](#)
- On MX2020 and MX2010, after a smooth upgrade from SFB to SFB2, if one plane or SFB is restarted, the link training fails between those planes and the MPC6 line-cards. [PR1309309](#)
- When the Routing Engine mastership is switched, the bbe-mibd process might generate a core file. [PR1309341](#)
- The first access-request fails for L2BSA subscribers when the MTU of LACP aggregated Ethernet A10NSP interface. [PR1309599](#)
- DHCP client is stuck in selecting state while verifying untagged DHCP subscribers after modifying router configuration. [PR1309730](#)
- 9000 out of 10,000 terminated subscribers go down during the unified ISSU from Junos OS Release 16.1-20170922_161_r4_s 6.0 to Junos OS Release 17.3-20170923.0. [PR1309983](#)
- In the next-generation subscriber management release, memory leak is seen for the bbe-smgd process after the address pool is deleted or added. [PR1310038](#)
- The MS-MIC or MS-MPC might experience a high memory utilization in the subscriber management scenario. [PR1310064](#)
- **SPD_CONN_OPEN_FAILURE** and **SPC_CONN_FAILURE** log messages are seen in the log for SI interfaces when SNMP walk is on service PIC NAT OIDs. [PR1310081](#)

- The `krt_junos_sanity_check_ctrl_resp`: rtsock request finally succeeded after error 16 syslog message in Junos OS Release 17.1R1.8. [PR1310678](#)
- The local IPv6 interface in the NDRA prefix is not removed from the service interface when the subscriber dual-stack session is removed. [PR1310752](#)
- Utilization of **commit check** while setting master-password might trigger improper decoding of configuration secrets. [PR1310764](#)
- After the base systems reboot, the rpd process is unresponsive on one or more GNFs. [PR1310765](#)
- Bad Junos Telemetry Interface (JTI) packaging for MPLS sensor. [PR1310932](#)
- The FPC memory might get exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- The routing protocol process generates a core file when multiple session flaps are observed on scale setup. [PR1312169](#)
- Incorrect incrementing of the counter at the PPPoE session logical interface (IFL) might lead to incorrect Acct-input-packets value and incorrect Acct-input-octets value in accounting packet. [PR1312998](#)
- False overtemperature SNMP trap could be seen when using MPC5, MPC6, MPC7, MPC8, and MPC9. [PR1313391](#)
- MX-VC: BNG: IPv6 router-solicit packets are dropped in non-default RI, but for the default RI the packets are not dropped. [PR1313722](#)
- The **show version detail** command output displays severe error log messages **traffic-dird[20126]: main: swversion pkg: 'traffic-dird' name: 'traffic-dird' ret: 0**. [PR1313866](#)
- The mspmand process generates a core file because of the flow control seen while clearing CGNAT and SFW sessions. [PR1314070](#)
- The JDM link is incorrectly shown to be up when the underlying physical link is down. [PR1314180](#)
- The **show version detail | no-more** CLI hangs for more than 120 seconds on master Routing Engine and more than 60 seconds on the backup Routing Engine. [PR1314242](#)
- The smgd process generates a core file with reference to `bbe_cos_ifl_publish()` `bbe_cos_if.c:6543`. [PR1314651](#)
- The rpd might stops responding in an multicast-only fast reroute (MoFRR) scenario. [PR1314711](#)
- In MPC7E, the **IR-mode** configuration statement fails to commit. [PR1314755](#)
- An RPC error is observed when you try to commit the **system services subscriber-management enable** statement through NETCONF. [PR1314968](#)
- MPC might crash after unified ISSU is performed multiple times. [PR1314982](#)
- The output of the **show version detail** command displays the severe error log message **mobiled: main Neither BNG LIC nor JMOBILE package is present, exit mobiled**. [PR1315430](#)
- The output of the **show version detail** command displays the severe error log message **main: name: SRD ret: 0**. [PR1315436](#)

- The output of the **show subscribers summary port** command does not display the correct output when subscribers are connected over a pseudowire. [PR1315659](#)
- An rpd core file is generated when the **show route inetcolor.0** command is executed. [PR1316078](#)
- On MX Series routers, the fan speed might frequently keep changing between normal and full. [PR1316192](#)
- The **show auto-configuration out-of-band** CLI command shows the same output for different statements. [PR1316661](#)
- The demux interface sends neighbor solicitation with the source link MAC address that comprises zeros: 00:00:00:00:00:00. [PR1316767](#)
- Traffic Load Balancer (TLB) traffic statistics counters do not get updated in Junos OS Release 18.1. [PR1317077](#)
- A few issues are seen in the output of the **show configuration display JSON** command for example, the alphanumeric values do not match the JSON output. [PR1317223](#)
- Linux-based microkernel might panic because of concurrent update of mutable objects. [PR1317961](#)
- CoA shaping rate is not applied successfully after the unified ISSU from Junos OS Release 15.1R6.7 to Junos OS Release 16.1R6.2. [PR1318319](#)
- The rpd process might stop responding after link flapping is experienced on an adjacent router. [PR1318476](#)
- The bbe-smgd process might stop responding after GRES is performed. [PR1318528](#)
- Changed text reported in **show chassis hardware** for CFP2-DCO optical transceivers. [PR1318901](#)
- MS-MPC or MS-MIC might crash after a new IPsec tunnel is added. [PR1318932](#)
- Kernel core is seen when more than 256 routing instances are created. [PR1319781](#)
- In some cases after multiple NSR switchovers replication might not complete for BGP, LDP, or RIP. [PR1319784](#)
- Loading an xmlproxyd YANG module file bounces the telemetry sessions. [PR1320211](#)
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after a unified ISSU of MX Series Virtual Chassis. [PR1320370](#)
- PPP inline keepalive does not work as expected when the CPE aborts the subscriber session. [PR1320880](#)
- If OpenConfig is used for telemetry and BGP data is being streamed, the rpd stop responding when the configuration that removes a BGP peer from group is committed. [PR1320900](#)
- The MX Series router sends IPv6 RA and the DHCPv6 advertisements before IPCPv6 Ack from the CPE. [PR1321064](#)
- The bbe-smgd process generates a core file after massive clients logs out and logs in a PPPoE dual stack subscriber scenario. [PR1321468](#)
- A CoA-NAK with **Error-Cause = Invalid-Request** is sent back to the RADIUS server when a drop policy is applied under **radius-flow-tap** in an L2TP subscriber scenario. [PR1321492](#)

- In Junos Node Slicing, the hierarchy of **show system schema module** is broken. [PR1321682](#)
- The commit operation might get stuck after commit check is performed. [PR1322431](#)
- The rpd process might crash when OpenConfig package is upgraded with JT1 streaming data in the background. [PR1322553](#)
- MS-MIC logical interfaces remain down after many iterations of taking them offline and bringing them back online. [PR1322854](#)
- When RPT BBE regression test is performed, an incorrect output is observed while verifying the **show subscribers client-type vlan subscriber-state active logical-system default routing-instance default** command. [PR1322907](#)
- The **show system subscriber-management route routing-instance <xxx>** command shows unexpected outputs. [PR1323279](#)
- The CLI command **request vmhost halt routing-engine other** does not halt the backup Routing Engine. [PR1323546](#)
- After successive flaps on core interfaces in AA Multihoming EVPN VXLAN, some race conditions might trigger constant high CPU on backup Routing Engine, where rpd shows very high CPU. [PR1334235](#)
- The subscriber might fail to log-in after the interface is deactivated and re-activated. [PR1324446](#)
- Memory leakage might be seen in the mosquito-nossl daemon in an MQTT scenario. [PR1324531](#)
- For payload prefix resolved through SRTE **color multi-path protocol-nexthop**, initially route resolution works correctly; thereafter because of some network change events, the SRTE multipath next hop updates might get stuck in the async-ket IO thread. To recover, flap the corresponding BGP session. [PR1324669](#)
- SNMP values do not increase monolithically. [PR1325128](#)
- Approximately 3 percent of Packet Forwarding Engine forwarding capacity might be seen on the XM chip when temperature of the chip is higher than 67 degrees Celsius. [PR1325271](#)
- MACsec session might fail to establish on the MX10003. [PR1325331](#)
- On a SIP ALG, core files are generated and its memory is exhausted. [PR1326394](#)
- MACSec MKA transmits the upper limit of the interval. [PR1326526](#)
- In MX Series, BNG CoS service object is not deleted properly for TCP and scheduler. [PR1326853](#)
- A minor alarm "LCM Peer Connection un-stable" is observed on MX150. [PR1328119](#)
- In JDI BBE, when regression test is performed for **show class-of-service interface demux0 <demux-interface>**, "Adjustment overhead-accounting mode" does not provide the expected output. [PR1329212](#)
- The CLI command **show services nat mappings address-pooling-paired** times out and fails. [PR1330207](#)
- All the packets might get dropped if one route is advertised by BGP when a session is established through the subscriber interface. [PR1330737](#)

- FPC wedge with fragmented packets on LSQ interface - PT1: Head and Tail out of synchronization. [PR1330998](#)
- The bbe-smgd process might crash after the **clear ancp access-loop circuit-id <circuit-id>** command is run. [PR1332096](#)
- Inaccurate J-Flow records might be seen in the output interface and next hop. [PR1332666](#)
- On MX150 platform, when **set chassis alarm management-ethernet link-down ignore** is set, FPC mgmt 0 interface alarm is not ignored. [PR1332799](#)
- The subinfo process might crash and it might cause the PPPoE subscribers to get disconnected. [PR1333265](#)
- MPC8E or MPC9E report high temperature alarms and fan speed moving continuous through full and normal speed iterations. [PR1334750](#)
- The UID limit is reached in a large-scale subscriber scenario. [PR1334886](#)
- When using **show subscribers**, and when the FPC number has two digits, the interface and IPv6 address that get connected together for DHCPv6 prefix delegation (PD). [PR1334904](#)
- The any-any option cannot be configured in a traffic selector for either IPv4 or IPv6 traffic. [PR1334966](#)
- JET application might not respawn after a normal exiting. [PR1336107](#)
- BBE-SMGD process might generate a core file while configuring CoS **ifl-set**. [PR1336852](#)
- Error log message **sdb_db_interface_remove: del ifl:si <index> with licnese cnt non zero on** can be seen on LTS during subscriber logout. [PR1337000](#)
- CM2.0 configuration is hidden in Junos OS 18.1 Release because of systest resources. However, the **show chassis fpc** error has changed the output. [PR1337467](#)
- IPsec VPN, session and service set sensor protocol files are being added to the Junos Telemetry Interface packaging. [PR1339883](#)
- MAC address of multiple interfaces are found to be duplicate. [PR1345882](#)
- The FPC temperature mismatch is seen between **show chassis fpc** and **show chassis fpc detail** for MPC6, MPC8, and MPC9 on MX2K platform. [PR1339077](#)

High Availability (HA) and Resiliency

- After server links flap, the GNFs associated with the ports on the Control Board shows this status message **Switchover Status: Not Ready**. [PR1306395](#)

Infrastructure

- The syscalltrace.sh file might create a huge output file that can cause the router to run out of storage space. [PR1306986](#)
- Cleanup at thread exit causes memory leak. [PR1328273](#)

Interfaces and Chassis

- On MX240, MX480, and MX960, IPv6 neighborship is not created on the IRB interface. [PR1198482](#)
- The output value is incorrect when you query the optical power of OTN interfaces in the router. [PR1216153](#)
- Rate-Limit -dropped packets are not displayed by [**show interfaces <ifl or-> detail/extensive**] commands. [PR1249164](#)
- The monitoring interface on aggregated Ethernet logical interfaces displays an incorrect BPS value compared to that displayed on **show interface** command output. [PR1283831](#)
- The **delay-buffer-rate** command with an absolute value is allowed on an inline LSQ interface. [PR1300281](#)
- Some CFM sessions do not come up after router with MPC8/9E line cards are rebooted with the scaled configuration. [PR1300515](#)
- IRB interface shows incorrect bandwidth value. [PR1302202](#)
- AFEB might not come up if LFM is deactivated. [PR1306707](#)
- After executing the **request system reboot both** CLI command, the PPP daemon might become unresponsive. [PR1310909](#)
- The PPPoE subscriber might not be able to login correctly after it fails to authenticate in a subscriber scenario. [PR1311113](#)
- The jpppd process might generate a core file at telemetry_start_timer, mosquito_handle_connack, and telemetry_mqtt_publisher. [PR1311396](#)
- The ifinfo process might crash and generate a core file when you execute the CLI command **show interfaces <Name>** command with the name greater than 128 characters. [PR1313827](#)
- Benign error messages are seen during an unified ISSU of MX Series Virtual Chassis if unsupported FRUs are present. [PR1316374](#)
- There is no route to IP address from the directly connected route on the static VLAN DEMUX interface. [PR1318282](#)
- **IPv6 Framed Interface Id** field (from **show subscribers extensive** output) does not match the negotiated value. [PR1321392](#)
- Interfaces might not work properly after FPC restarts. [PR1329896](#)
- The transportd process might crash when there is an SNMP query on jnxoptIfOChSinkCurrentExtTable with unsupported interface index. [PR1335438](#)
- Traffic loss might be seen after deleting aggregated Ethernet bundle unit 1. [PR1329294](#)

Layer 2 Ethernet Services

- DHCPv6 client bound to IA_PD prefix on reception of DHCPv6 request for IA_NA, MX Series deletes the existing binding. [PR1286359](#)
- PPPoE or DHCP clients cannot log in to PPPoE or DHCP dual-stack subscriber scenario. [PR1298976](#)
- Multiple jdhcpd core files are observed in jdhcpd_update_groups at `../..../src/junos/usr/sbin/jdhcpd/jdhcpd_config.c:2290`. [PR1311569](#)
- DHCPv6 traffic might be dropped in a subscriber scenario. [PR1316274](#)
- The jdhcpd process might generate a core file after making DHCP configuration changes. [PR1324800](#)
- The **on-demand-address-allocation** option of the **dual-stack-group** statement does not work for IPv6. [PR1327681](#)
- The jdhcpd process crashes and generates a core file. [PR1334230](#)

MPLS

- Minor difference is seen between mpls.statistics and the adjusted bandwidth. [PR1259500](#)
- An ingress RSVP LSP fails to come up when the **clear rsvp lsp** command is run on the egress router. [PR1275563](#)
- The rpd might crash in LDP L2 circuit scenario. [PR1275766](#)
- The traffic is dropped during NSR switchover for RSVP P2MP provider tunnels are used by MVPN. [PR1293014](#)
- The traffic in P2MP tunnels might be lost when the next-generation MVPN uses RSVP-TE. [PR1299580](#)
- The rpd process might crash in rare scenarios where traffic engineering is configured. [PR1303239](#)
- kysncd process might crash after the backup Routing Engine is removed/uninstalled and then reinserted/reinstalled, [PR1303491](#)
- BGP multipath might not work if the interface flaps. [PR1305228](#)
- The configuration of the **explicit-null** statement might block host-bound traffic incoming from LSP. [PR1305523](#)
- RSVP node-hello works incorrectly after the next hop for the remote destination is changed. [PR1306930](#)
- On a router with UHP-based LSP configuration, the rpd process might crash when interfaces are down. [PR1309397](#)
- The rpd process might crash when LDP updates the label for BGP route. [PR1312117](#)
- The rpd might crash when LDP sessions and RSVP LSPs are flapped in an LDP over RSVP setup. [PR1318480](#)
- The IPv4 or IPv6 multicast traffic might get dropped in MX Series Virtual Chassis when the traffic comes in through the Layer 2 circuit and goes out through aggregated Ethernet member interfaces across Virtual Chassis members. [PR1320742](#)

- The rpd crashes with **ldp p2mp** configuration. [PR1321626](#)
- The rpd process crashes and generates a core file in jemalloc_block_mallocx because of a memory leak. [PR1321952](#)
- SNMP OID counters for mplsLspInfoAggrOctets might show a constant value for RSVP LSPs for longer time in case of route withdrawn scenario. [PR1327350](#)
- Local repair took about 150 milliseconds greater than expected 100 milliseconds. [PR1327988](#)
- Packet loss might be observed when auto-bandwidth for CCC connections is enabled. [PR1328129](#)
- The rpd process crashes on backup Routing Engine because of memory exhaustion. [PR1328974](#)

Network Management and Monitoring

- On MX Series platform, the Routing Engine does not reply to SNMP request. [PR1240178](#)
- When the SNMP configuration gets activated, the snmpd process starts to consume a lot of CPU time. [PR1300016](#)
- The syslog might generate duplicate entries of hostname and timestamp. [PR1304160](#)
- The mib2d process generates a core file when an FPC is reset during asynchronous statistic collection through SNMP. [PR1318302](#)
- With **interface-mib** configuration in a dynamic-profile, when multiple OIDs are queried in a SNMPGET and SNMPWALK, the router might reply with **No Such Instance currently exists at this OID** for some of the OIDs. [PR1329749](#)

Platform and Infrastructure

- Adaptive load balancing (ALB) functionality is supported only for unicast traffic. If the aggregate bundle contains logical interfaces for a bridge or VPLS domains, flooded traffic might be dropped. [PR821237](#)
- The Packet Forwarding Engine on an MS-MPC might crash with a large scale routes for MX Series routers. [PR1277264](#)
- The FPC **resource-monitor % mem free** values for next-hop forwarding are incorrect. [PR1287592](#)
- There might be Packet Forwarding Engine memory leak when the next-hop address that is defined in the next-hop group is reachable through multiple interfaces. [PR1287870](#)
- Dynamic MAC learning might fail on GRE tunnel interfaces. [PR1291015](#)
- RMOPD_HW_TIMESTAMP_INVALID is reported two to four times a day, which raises an alarm when polled through the jnxRpmResSumPercentLost MIB. [PR1300049](#)
- Traffic getting dropped in egress Packet Forwarding Engine because of hashing mismatch. [PR1300789](#)
- Packet Forwarding Engine might crash after the MPC is reset in a firewall filter scenario. [PR1300990](#)
- Classifiers do not get applied on the aggregated Ethernet member links when CoS is configured on MX Series routers with DPCs. [PR1301723](#)

- MX Series MPC wedges (might cause fabric blackhole and finally reboot the line card) when creating more than 4000 logical tunnel interfaces per Packet Forwarding Engine. [PR1302075](#)
- The **interface-mac-limit** configuration might fail for aggregated Ethernet interfaces. [PR1303293](#)
- The TWAMP Request-TW-Session message's Type-P Descriptor format is not RFC-compliant. [PR1305752](#)
- On MX Series routers with MPCs or MICs, the resource monitor (RSMON) thread might be stuck in a loop consuming 100 percent of FPC line card CPU. [PR1305994](#)
- System reach process ceiling <low> watermark due to auditd. [PR1305964](#)
- The source MACs might be leaked or not learned between different VPLS instances at the received VPLS PE devices. [PR1306293](#)
- The RPM probe with probe interval of 1 second fails in MX Series routers. [PR1308952](#)
- The expected error message is not observed during a Telnet session when a username longer than acceptable limit is used. [PR1312265](#)
- ICMP error messages are observed in the Packet Forwarding Engine. [PR1313668](#)
- Rate-limit configured with small temporal buffer size might cause packet loss. . [PR1317385](#)
- Multicast traffic is not forwarded on the newly added P2MP branch or receiver. [PR1317542](#)
- Multicast traffic might get duplicated when MoFRR is configured. [PR1318129](#)
- The GNF-associated MPC hangs during reboot after a unified ISSU. [PR1318394](#)
- Change in default severity of correctable ECC errors on MX Series routers from fatal to major. [PR1320585](#)
- Errors might be observed when **fabric-header-crc-enable** feature is enabled. [PR1320874](#)
- The rate limit with a lock protected variable of netisr queue, the count of packets in netisr queue becomes wrong. This leads to kernel crash or debugger command prompt. [PR1332153](#)
- RPM probes delegated to MS-MIC get stuck when any change is made on BGP group configuration. [PR1322097](#)
- The **no-propagate-ttl** configuration might not take effect when **chained-composite-next-hop ingress l3vpn extended-space** is configured. [PR1323160](#)
- MX Series Virtual Chassis MAC learning does not occur on specific interfaces. [PR1327723](#)
- The packet might get dropped in LSR when MPLS pseudowire payload does not have a control word and the packet's destination MAC address starts with 4 or 6. [PR1327724](#)
- Traffic loss might be observed on a logical tunnel interface. [PR1328371](#)
- Junos OS automation folder lost execution permissions. [PR1328570](#)
- SNMP pingResultsMinRtt, pingResultsMaxRtt, pingResultsAverageRtt and pingProbeHistoryResponse are marked as "1" instead of "0" if the response is not received from the RPM server. [PR1333320](#)

- Traffic loss might be seen for some flows because of the churn in the network. [PR1335302](#)
- Route corruption in Packet Forwarding Engine with connectivity-fault-management enabled for Layer 2 circuit. [PR1338854](#)

Routing Policy and Firewall Filters

- The rpd process might crash when **vrf-target auto** is configured for a routing instance. [PR1301721](#)
- The policy configuration might not be evaluated if policy expression is changed. [PR1317132](#)

Routing Protocols

- The command **show bgp summary** results incorrect while assisting a graceful restart. [PR1045151](#)
- BGP MIBv2 enterprise MIB objects for InetAddress types are not properly generating OIDs. [PR1265504](#)
- The rpd process might crash when BGP is deactivated or activated. [PR1272202](#)
- When the bfdd process restarts, an issue with next-generation MVPN and L2VPN route exchange causes MVPN and VPLS traffic to be dropped. [PR1278153](#)
- BGP updates might not be advertised to peers completely in certain conditions. [PR1282531](#)
- Some BGP-related traceoptions flag settings are not effective immediately when the configuration is committed, until the BGP sessions are flapped. [PR1285890](#)
- In IS-IS service request (SR) LAN scenario, advertising adjacencysegment identifiers might be missed for a few neighbors if the TLV length gets exhausted. This is not a common scenario. [PR1288331](#)
- Multihop BFD sessions flap continuously. [PR1291340](#)
- The link management protocol daemon (lmpd) repeatedly crash when a logical system is configured on the same router. [PR1294166](#)
- The rpd process might crash because of the AS PATH check error when RIB groups are added first and the routing instance later. [PR1298262](#)
- MSDP sessions might flap because data replication might get stuck between backup and master Routing Engine with a huge SA burst between peers. [PR1298609](#)
- When the device is restarted with Junos OS Release 17.4R1, the benign error message **channel 0: chan_shutdown_read: shutdown() failed for fd 10 [i0 o3]: Socket is not connected** messages may show up with no functionality impact are seen. [PR1300409](#)
- IBGP route damping does not take effect on VPN address families. [PR1301519](#)
- The rpd process might crash, generating a core file, when a multipath route is deleted. [PR1302395](#)
- The mcsnoopd process generates a core file at `__raise, abort, __task_quit__, task_quit, task_terminate_timer_callback, task_timer_dispatch, and task_scheduler_internal` (enable_slip_detector=true, no_exit=true) at `../src/junos/lib/libjtask/base/task_scheduler.c:275`. [PR1305239](#)

- The BFD session might flap when querying interface statistics through SNMP or CLI show command in vMX. [PR1305308](#)
- Junos OS Release 16.2 and later releases might give the following error: **Request failed: OID not increasing: ospfIfIpAddress.0.0.0.0.0** . [PR1307753](#)
- Qualified next-hop resolution fails in some scenarios when there is a next-hop interface specified. [PR1308800](#)
- BGP labeled unicast protection might break multicast reverse path forwarding (RPF). [PR1310036](#)
- When NSR is configured, the BGP session might flap if the connection between the master Routing Engine and the backup Routing Engine keeps flapping. [PR1311224](#)
- The rpd process might crash and generate a core file in **bgp_rt_send_message at ../../../../src/junos/usr/sbin/rpd/bgp/bgp_io.c:1460** . [PR1310751](#)
- Dedicated BFD does not work on MX Series platforms. [PR1312298](#)
- IS-IS SPF gets triggered by LSP updates containing changes in reservable bandwidth in TE extensions. [PR1313147](#)
- The routing protocol process (rpd) might crash and generate a core file. [PR1314679](#)
- BGP prefixes with three levels of recursion for resolution gets stuck with a stale next hop at the first level after a link-down event. [PR1314882](#)
- The SUB-TLV values are assigned for segment routing TE policy SUB-TLVs. [PR1315486](#)
- On a router with BGP Monitoring Protocol (BMP) configured, the rpd process might crash when the rpd process is gracefully terminated. [PR1315798](#)
- The link-state database (LSDB) entry cleanup might cause the rpd process to crash, if loop-free alternate is configured. [PR1317023](#)
- When two Route-reflector (RR) routers use PIC (protect core) to protect each other's BGP-LU (labeled-unicast) LSP, endless label oscillation might be seen. [PR1318093](#)
- The routing protocol process (rpd) crash is seen when deactivating static route if the next-hop interface is point-to-point (P2P) type. [PR1323601](#)
- Multiple next hops might not be installed for IBGP multipath route after IGP route updates. [PR1327904](#)
- With BGP, LDP, and IS-IS configurations, deleted IS-IS routes might still be visible in the RIB. [PR1329013](#)
- The rpd might crash on backup Routing Engine after BGP peer is deleted. [PR1329932](#)
- When prefix limit is reached, increasing maximum-prefixes does not take effect immediately. [PR1323765](#)
- BGP session get stuck in active state after remote end router is upgraded. [PR1335319](#)
- When the primary interface is back online, the discarded next hop address is retained until the BGP LU neighbor is cleared. This impacts the cloned route (S=0) only. [PR1333570](#)

Services Applications

- When configuring a NAT pool that is shared between PCP and standard NAT, the PCP mappings cannot be cleared. [PR1284261](#)
- The jl2tpd process might stop responding shortly after GRES. [PR1295248](#)
- L2TP subscribers might get stuck in a terminating state during login. [PR1298175](#)
- L2TP tunnel switch clients experience packets drops for large packets because of fragmentation in a L2TP tunnel switch. [PR1312691](#)
- When an l2tp subscriber BNG receives ANCP port up with TLV DSL-type=0 ("other"), the BNG does not include AVP 145 in the ICRQ packet. [PR1313093](#)
- L2TP tunnel Tx and Rx bytes count sometimes decrease when subscriber sessions are reduced within the tunnel. [PR1318133](#)
- In an L2TP scenario, the MRU might be changed to 1492 instead of the default 1500. [PR1319252](#)
- IPCP active mode remains disabled for MLPPP on LNS. [PR1319580](#)
- Stale L2TP routes might be seen when L2TP peer uses any UDP port other than the default. [PR1322197](#)
- L2TP tunnel switch might drop the first **CHAP Success** packet from LNS because of the delay in programming of the /136 route on the Packet Forwarding Engine. [PR1325528](#)
- In case the number of sessions addressed in CSURQ is more than about 107, not all CSURQ messages receive a response. [PR1330150](#)

Subscriber Access Management

- Service interim for DHCP subscribers does not work in a JSRC scenario. [PR1303553](#)
- The **show network-access aaa accounting** command might display additional entries. [PR1304594](#)
- Incorrect **Acct-Delay-Time** in RADIUS **Accounting-On** message is seen after the MX Series router, acting as a BNG, is rebooted. [PR1308966](#)
- Service interim for random users is missing in a JSRC scenario. [PR1315207](#)
- The delegated prefix from RADIUS is incorrectly parsed when the prefix is fewer than 20 bytes long. [PR1315557](#)
- The PPPoE subscribers might encounter connection failure during login. [PR1317019](#)
- The unified ISSU is allowed to proceed when the account is suspended. [PR1320038](#)
- Incorrect address assignment sequence is seen from linked IP pools. [PR1323829](#)
- The general authentication service considers the RADIUS attributes Framed-IPv6-Prefix = ::/64 and Delegated-IPv6-Prefix = ::/56 as valid parameters. [PR1325576](#)
- The MX204 does not send the RADIUS Accounting-Off message. [PR1327822](#)
- Subscriber management experiences SDB DOWN event; dfcd[4707]: %DAEMON-3: attempting to close SDB while DOWN. [PR1336388](#)

User Interface and Configuration

- The commit time increases every time. [PR1029477](#)
- The CLI session might be terminated while the **show configuration | compare rollback 1** command is issued. [PR1331716](#)

VPNs

- Next-generation MVPN IPv6 RP bootstrap type 3 S-PMSI AD route prefix ff02::d persists after BSR data stop. [PR1269234](#)
- Layer 2 circuits stitched through It peer interfaces might get stuck in **LD** (local site signaled down) state. [PR1305873](#)
- A nonoptimal route to source might be selected for next-generation MVPN with **unicast-umh-election** enabled. [PR1315011](#)
- Un-hide the **set protocols pim mvpn family inet6 disable** configuration to allow the users to disable the inet6 configuration on MVPN. [PR1317767](#)
- The routing protocol process (rpd) might stop responding after a unified ISSU in a large-scale scenario with PIM configuration. [PR1322530](#)
- Moving MC-LAG from LDP-based pseudowire to BGP-based pseudowire might cause an rpd crash. [PR1325867](#)
- In an next-generation MVPN and NSR configuration, the rpd process might crash and generate a core file on the backup Routing Engine. [PR1328246](#)

SEE ALSO

[New and Changed Features | 102](#)

[Changes in Behavior and Syntax | 123](#)

[Known Behavior | 132](#)

[Known Issues | 137](#)

[Documentation Updates | 189](#)

[Migration, Upgrade, and Downgrade Instructions | 190](#)

[Product Compatibility | 197](#)

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 189](#)
- [Subscriber Management Access Network Guide | 189](#)
- [Subscriber Management Provisioning Guide | 190](#)

This section lists the errata and changes in Junos OS Release 18.1R3 documentation for MX Series.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

Subscriber Management Access Network Guide

- The guide failed to include a feature that enables you to override the information that the LAC sends to the LNS in L2TP Calling Number AVP 22 when the LAC is configured to use the Calling-Station-ID format. You can configure the access profile to override that value for AVP 22 with any combination of the agent circuit identifier and the agent remote identifier received by the LAC in the PADR packet.

[See [Override the Calling-Station-ID Format for the Calling Number AVP.](#)]

- The guide incorrectly stated that the **linked-pool-aggregation** statement is located at the **[edit access address-assignment pool *pool-name*]** hierarchy level. In fact, this statement is located at the **[edit access]** hierarchy level.

[See [Configuring Address-Assignment Pool Linking](#).]

Subscriber Management Provisioning Guide

- Starting in Junos OS Release 15.1, the *Broadband Subscriber Sessions User Guide* and the [CLI Explorer](#) incorrectly included information about the **show extensible-subscriber-services accounting** command. This command is not present in the CLI. Instead, you can use accounting profiles to collect statistics from the Packet Forwarding Engine for Extensible Subscriber Services Manager (ESSM) subscribers. See [Flat-File Accounting Overview](#) for information about accounting for ESSM subscribers.

SEE ALSO

[New and Changed Features | 102](#)

[Changes in Behavior and Syntax | 123](#)

[Known Behavior | 132](#)

[Known Issues | 137](#)

[Resolved Issues | 153](#)

[Migration, Upgrade, and Downgrade Instructions | 190](#)

[Product Compatibility | 197](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 18.1 | 191](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 191](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 194](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 196](#)
- [Upgrading a Router with Redundant Routing Engines | 196](#)
- [Downgrading from Release 18.1 | 196](#)

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 18.1R1 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
MX204, MX240, MX480, MX960, MX2010, MX2020 MX10003, vMX	NO	YES

Basic Procedure for Upgrading to Release 18.1

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 [Junos OS Administration Library](#).

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

Platforms impacted: MX204, MX240, MX480, MX960, MX2010, MX2020, MX10003, and vMX.

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-x86-32-18.1R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-18.1R3.9-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-x86-32-18.1R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-18.1R3.9-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](#).

NOTE: After you install a Junos OS Release 18.1 **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

Platforms 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-18.1R3.9-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-18.1R3.9-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 18.1 **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 18.1

To downgrade from Release 18.1 to another supported release, follow the procedure for upgrading, but replace the 18.1 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

SEE ALSO

[New and Changed Features | 102](#)

[Changes in Behavior and Syntax | 123](#)

[Known Behavior | 132](#)

[Known Issues | 137](#)

[Resolved Issues | 153](#)

[Documentation Updates | 189](#)

[Product Compatibility | 197](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 197](#)

Hardware Compatibility

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

To determine the features supported on MX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

New and Changed Features	102
Changes in Behavior and Syntax	123
Known Behavior	132
Known Issues	137
Resolved Issues	153
Documentation Updates	189
Migration, Upgrade, and Downgrade Instructions	190

Junos OS Release Notes for NFX Series

IN THIS SECTION

- New and Changed Features | 199
- Changes in Behavior and Syntax | 203
- Known Behavior | 204
- Known Issues | 206
- Resolved Issues | 209
- Documentation Updates | 211
- Migration, Upgrade, and Downgrade Instructions | 212
- Product Compatibility | 215

These release notes accompany Junos OS Release 18.1R3 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

New and Changed Features

IN THIS SECTION

- [Release 18.1R3 New and Changed Features | 199](#)
- [Release 18.1R2 New and Changed Features | 199](#)
- [Release 18.1R1 New and Changed Features | 199](#)

This section describes the new features or enhancements to existing features in Junos OS Release 18.1R3 for NFX Series devices.

Release 18.1R3 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 18.1R3.

Release 18.1R2 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 18.1R2.

Release 18.1R1 New and Changed Features

Hardware

- **NFX150 platform**—Starting with Junos OS Release 18.1R1, the NFX150 Network Services Platform is available as a single platform that integrates routing, switching, and security functions. The NFX150 is a secure, automated, software-driven customer premises equipment (CPE) platform that delivers virtualized network and security services on demand. It is suited for small to medium-sized enterprises and acts as a secure router, SD-WAN CPE, or uCPE. The architecture of the NFX150 platform enables unified management of all its components through the Junos Control Plane (JCP). It also offers effective management of the system resources and reduced system boot time.

The NFX150 portfolio is available in the following variants:

- NFX150-S1—Rack-mount model with 2.2-GHz 8-core Intel CPU, 200-GB SSD, 16-GB RAM, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports.
- NFX150-S1E—Rack-mount model with 2.2-GHz 8-core Intel CPU, 200-GB SSD, 32-GB RAM, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports.

NFX150-S1 and NFX150-S1E support the following expansion modules:

- NFX-EM-6T2SFP—Expansion module with six 1-Gigabit Ethernet RJ-45 ports and two 1-Gigabit Ethernet SFP ports
- NFX-LTE-AE—Expansion module with an LTE modem supporting the frequency bands in Europe and North America.
- NFX-LTE-AA—Expansion module with an LTE modem supporting the frequency bands in Asia and Australia.
- NFX150-C-S1—Compact desktop model with 2.2-GHz 4-core Intel CPU, 8-GB RAM, 100-GB SSD, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports.
- NFX150-C-S1-AE—Compact desktop model with 2.2-GHz 4-core Intel CPU, 8-GB RAM, 100-GB SSD, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports. This device provides integrated LTE modem for Europe and North America.
- NFX150-C-S1-AA—Compact desktop model with 2.2-GHz 4-core Intel CPU, 8-GB RAM, 100-GB SSD, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports. This device provides integrated LTE modem for Asia, Australia, and New Zealand.
- NFX150-C-S1E-AE—Compact desktop model with 2.2-GHz 4-core Intel CPU, 16-GB RAM, 100-GB SSD, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports. This device provides integrated LTE modem for Europe and North America.
- NFX150-C-S1E-AA—Compact desktop model with 2.2-GHz 4-core Intel CPU, 16 GB RAM, 100 GB SSD, four 10/100/1000BASE-T RJ-45 LAN ports, and two 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ WAN ports. This device provides integrated LTE modem for Asia, Australia, and New Zealand.
- **Transceivers** –NFX150 supports the following optics:
 - 10-gigabit SFP+ transceivers: EX-SFP-10GE-USR, EX-SFP-10GE-SR, EX-SFP-10GE-LR, EX-SFP-10GE-ER, EX-SFP-10GE-DAC-1M, EX-SFP-10GE-DAC-3M, EX-SFP-10GE-DAC-5M, EX-SFP-10GE-DAC-7M
 - 1-gigabit SFP transceivers: EX-SFP-1GE-SX, EX-SFP-1GE-SX-ET, EX-SFP-1GE-LX, EX-SFP-1GE-LH, EX-SFP-1GE-LX40K, EX-SFP-GE80KCW1470, EX-SFP-GE80KCW1490, EX-SFP-GE80KCW1510, EX-SFP-GE80KCW1530, EX-SFP-GE80KCW1550, EX-SFP-GE80KCW1570, EX-SFP-GE80KCW1590, EX-SFP-GE80KCW1610

NOTE: USB and ER optics are displayed as SFP+-10G-ER in the **show system inventory hardware optics** command output.

Amphenol DAC 1M and 3M cables are displayed as **unknown** in the **show system inventory hardware optics** command output.

[See [NFX150 Network Services Platform Hardware Guide](#).]

Service Chaining

- **VNF service chaining**—Starting with Junos OS Release 18.1R1, the NFX150 device supports deploying and service chaining of multiple, secure, high-performance virtualized network functions (VNFs) as a single device. The Junos Control Plane (JCP) runs on the Junos VM and functions as the single point of management for the chassis and VNFs.

[See [Service Chaining on NFX Devices](#).]

Security

- **Secure Boot**—Starting with Junos OS Release 18.1R1, the NFX150 devices support secure boot implementation, which is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, bootloader, and kernel are cryptographically protected. No action is required to implement secure boot.

[See [Feature Explorer](#) and enter Secure Boot.]

Layer 2 Features and Protocols

- **Layer 2 features**—Starting with Junos OS Release 18.1R1, the NFX150 supports Layer 2 features such as VLANs, IGMP snooping, MLDv1 snooping, MLDv2 snooping, port mirroring, port security, and the Link Layer Discovery Protocol (LLDP).

[See [Services](#) and [Ethernet Switching](#).]

Layer 3 Features and Protocols

- **Layer 3 features**—Starting with Junos OS Release 18.1R1, the NFX150 supports Layer 3 features such as IP Security (IPsec), firewall filters, port mirroring, BFD, and class of service (CoS). It also supports Layer 3 protocols such as BGP, RIP, OSPFv1, OSPFv2, and IS-IS.

[See [IPsec](#) and [Security](#).]

Fault Management

- **OAM link fault management and connectivity fault management**—Starting with Junos OS Release 18.1R1, NFX150 devices support configuration of IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). CFM monitors Ethernet networks that might comprise one or more service instances for network-compromising connectivity faults.

[See [Fault Management](#).]

Network Service Orchestrator

- **Network Service Orchestrator**—Starting with Junos OS Release 18.1R1, NFX150 devices support Network Service Orchestrator, which is a client included in the base software of the NFX150 device, and connects to the Network Activator deployed on a cloud or server. The Network Activator application intelligently automates service life cycle management of managed VPN networks, in-region secured Internet connections, and out-of-region IPsec connections on NFX150 devices. This application enables the booting and configuration of the NFX150 device when it is first powered on.

[See [Network Activator Overview](#).]

Wireless WAN

- **Wireless WAN**— Starting with Junos OS Release 18.1R1, the following NFX150 device models provide wireless WAN support through the LTE module:
 - NFX150-S1
 - NFX150-S1E
 - NFX150-C-S1-AE
 - NFX150-C-S1-AA
 - NFX150-C-S1E-AE
 - NFX150-C-S1E-AA

[See [NFX150 Network Services Platform Hardware Guide](#).]

SEE ALSO

Changes in Behavior and Syntax 203
Known Behavior 204
Known Issues 206
Resolved Issues 209
Documentation Updates 211
Product Compatibility 215

Changes in Behavior and Syntax

IN THIS SECTION

- [CLI | 204](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.1R3 for the NFX Series.

CLI

- Starting with Junos OS Release 18.1R1, the host-os hierarchy level is replaced with the vmhost hierarchy level for NFX150 devices.

SEE ALSO

New and Changed Features 199
Known Behavior 204
Known Issues 206
Resolved Issues 209
Documentation Updates 211
Product Compatibility 215

Known Behavior

IN THIS SECTION

- [Known Behavior: 18.1R3 | 205](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 for the NFX Series.

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

Known Behavior: 18.1R3

- The file transfer rate from an external media over the network to an NFX150 device is around 40–50 Mbps.

[PR1290263](#)

- On NFX150 devices running Junos OS Release 18.1, service chaining can be achieved through front panel ports by using SR-IOV. For the switching to work through SR-IOV enabled front panel port, the physical NIC port must be up and operational.

[PR1319294](#)

- On NFX150 devices running Junos OS Release 18.1, you cannot use the **request system software scripts** command to add script packages on the Junos OS.

[PR1333061](#)

- On NFX150 devices running Junos OS Release 18.1, traffic shaping on tunnel interfaces such as IP-IP and GRE is not supported.

[PR1335582](#)

- On NFX150 devices running Junos OS Release 18.1, Transcend does not support Linux based SSD firmware upgrade mechanism in field for its SSD. Hence, field upgrade of Transcend SSD firmware cannot be provided for NFX150 devices.

[PR1347562](#)

SEE ALSO

New and Changed Features 199
Changes in Behavior and Syntax 203
Known Issues 206
Resolved Issues 209
Documentation Updates 211
Product Compatibility 215

Known Issues

IN THIS SECTION

- [Known Issues: 18.1R3 | 206](#)

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for the NFX Series.

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

Known Issues: 18.1R3

- On NFX150 devices, you cannot generate an ISO configuration image. [PR1316900](#)
- On NFX150 devices, connectivity fault management (CFM) is not supported on circuit cross-connect (CCC) interfaces. [PR1311588](#)
- There is no commit check if the PCI address is reused for different interfaces in a VNF. As a workaround, we recommend that you stop the VNF and then add or delete interfaces. [PR1205497](#)
- The **show chassis routing-engine** command displays the last reboot reason as **power cycle/failure** even for a normal system reboot. In addition, the logs record an abnormal shutdown message. [PR1232501](#)
- Configuring more than the available number of virtual functions for an SR-IOV front panel port, might result in a state where the user MAC addresses for such interfaces are not released back to the System MAC Pool on deletion of the VNF. [PR1259975](#)
- On NFX150 devices with LTE support running Junos OS Release 18.1, the **show system visibility cpu** command does not display CPU pinning information for LTE. There is no known workaround. [PR1347609](#)
- While changing port mapping configuration across FPC0 and FPC1 on NFX150 devices with expansion module, forwarding path simulation process for FPC0 may crash when FPC0 restarts for port mapping configuration to take effect. This results in an additional reboot of FPC0. After the reboot, FPC0 recovers automatically and appears online. [PR1347259](#)
- LTE functions as a kernel driver for modem packet handling and should not be treated as a customized VNF. The **request** command does not provide console support. [PR1348196](#)
- On NFX150 devices running Junos OS Release 18.1, manually loading the factory-default configuration on the device might not set up the necessary configurations for Remote Activation to work. As a workaround, before loading the factory default configuration on the device, ensure that the configuration for phone-home is deleted and committed. [PR1347308](#)

- On NFX150 devices running Junos OS Release 18.1, Dev key revocation is not supported by BIOS. Dev key revocation is to prevent customers from installing Dev signed image by mistake on their setup. [PR1344738](#)
- On NFX150 devices running Junos OS Release 18.1, enabling hugepages for VNFs and pre-reserving of hugepages are not supported. Hence, the following commands are not supported:
 - **set system memory hugepages**
 - **set virtual-network-functions *vnf-name* memory features hugepages**[PR1360998](#)
- On NFX150 devices running Junos OS Release 18.1, traffic statistics for 10-Gigabit Ethernet host interfaces are not displayed correctly. [PR1348720](#)
- On NFX150 devices running Junos OS Release 18.1, syslog messages do not display xauth client authentication information such as assigned IP address and DNS. [PR1305078](#)
- On NFX150 devices running Junos OS Release 18.1, FTP displays an error message, **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory.** [PR1315605](#)
- On NFX150 devices running Junos OS Release 18.1, CLI output for the **show security ipsec inactive-tunnels** command is incomplete. [PR1325763](#)
- On NFX150 devices running Junos OS Release 18.1, error messages are seen while rebooting the FPC0 interface. [PR1326487](#)
- On NFX150 devices running Junos OS Release 18.1, commit is successful with any message on the console while creating a VNF using CLI. However, VNF may not be created due to some errors. Syslog will show error messages with reasons for not creating the VNF. [PR1333057](#)
- On NFX150 devices running Junos OS Release 18.1, **file put** operation by a user with no super-user permissions might fail. [PR1333991](#)
- On NFX150 devices running Junos OS Release 18.1, **file copy** operation by a user with no super-user permissions might fail. [PR1333995](#)
- On NFX150 devices running Junos OS Release 18.1, extracting contents of an archived file by using the **tar -xvzf** command might fail.

[PR1334485](#)

- On NFX150 devices running Junos OS Release 18.1, the **op** command, which is used to execute python scripts that are residing on the JCP might fail and result in an error. As a workaround, delete the configuration knob **system scripts op allow-url-for-python** and re-run the **op** command by using CLI.

[PR1360806](#)

- During BIOS upgrade process, it does not display the existing BIOS version or the new BIOS version to which it is being upgraded. Similarly, it does not display the BIOS version when a lower version of BIOS is getting upgraded to a higher version of BIOS.

[PR1342573](#)

- On NFX150 devices running Junos OS Release 18.1, after upgrading the image, the SYSHMD error messages are observed only once.

[PR1341005](#)

- On NFX150 devices running Junos OS Release 18.1, after upgrading the image, FPC0, FPC1 IFL error messages are observed only once.

[PR1341583](#)

- On NFX150 devices running Junos OS Release 18.1, **request ca-certificate** command fails. CA Trust certificates cannot be installed on the device.

[PR1343474](#)

- On NFX150 devices running Junos OS Release 18.1, the MTU of an heth interface cannot be set. The configuration knob of **set vmhost interfaces heth-X-Y mtu** is not supported.

[PR1346876](#)

- On the NFX150 running Junos OS Release 18.1, the rssi value for Wireless Model interface cl-1/1/0 shows negative value.

[PR1344377](#)

SEE ALSO

New and Changed Features 199
Changes in Behavior and Syntax 203
Known Behavior 204
Resolved Issues 209
Documentation Updates 211
Product Compatibility 215

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3 | 209](#)
- [Resolved Issues: 18.1R2 | 209](#)
- [Resolved Issues: 18.1R1 | 210](#)

This section lists the issues fixed in the Junos OS main release and the 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: 18.1R3

Junos Control Plane (NFX150)

- On NFX150 devices running Junos OS Release 18.1, the **file-copy** operation by a user with no super-user permissions might fail. [PR1333995](#)

Resolved Issues: 18.1R2

Junos Control Plane (NFX150)

- Under some circumstances, FPC0 ukern of NFX150 may crash and restart. The FPC recovers automatically and it does not crash again after the recovery. There is no known workaround. [PR1347629](#)
- On NFX150 devices running Junos OS Release 18.1, jdmd core is observed after configuration changes failed to commit.
[PR1348783](#)
- On NFX150 devices running Junos OS Release 18.1, while changing port mapping configuration on FPC0 and FPC1 interfaces by using expansion module, memory corruption is detected in low memory and DMA Write errors are observed.
[PR1325585](#)
- On NFX150 devices running Junos OS Release 18.1, there could be a core related to Key Management Daemon (kmd) during some configuration changes. The issue is very rare.
[PR1330280](#)

- On NFX150 devices running Junos OS Release 18.1, vm core is observed while downloading the image from PHS.

[PR1330487](#)

- On NFX150 devices running Junos OS Release 18.1, mac-table entries are not updated with topology change notification (TCN).

[PR1326593](#)

- On NFX150 devices running Junos OS Release 18.1, with default LTE configuration, the PHC on the device will not be able to communicate with Juniper redirect server with LTE as the only link on the device. The name resolution of Juniper redirect server will fail without fixing this issue.

[PR1342499](#)

Juniper Device Manager (NFX250)

- On NFX250 devices running Junos OS Release 18.1, if the same VLAN ID is used in two different cross-connect configurations, the commit will not fail.

[PR1346698](#)

Resolved Issues: 18.1R1

Juniper Device Manager (NFX250)

- If a VNF requests for more memory than the available system memory, commit might go through without any errors resulting in VNF going into a shut off state. As a workaround, use the show system visibility memory command to check the available free memory before spawning a VNF. Alternatively, check the log files and the VNF shut off reason will be captured in /var/log/syslog file. [PR1221647](#)
- While spawning a VNF, there might not be a commit check for the valid image type supported. [PR1221642](#)

SEE ALSO

New and Changed Features 199
Changes in Behavior and Syntax 203
Known Behavior 204
Known Issues 206
Documentation Updates 211
Product Compatibility 215

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 211](#)

This section lists the errata and changes in Junos OS Release 18.1R3 for the NFX Series documentation.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 199](#)

[Changes in Behavior and Syntax | 203](#)

[Known Behavior | 204](#)

[Known Issues | 206](#)

[Resolved Issues | 209](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 212
- Basic Procedure for Upgrading to Release 18.1 | 212

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 18.1

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 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 [Junos OS Administration Library](#).

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 18.1R3 on NFX250 devices:

1. Using a Web browser, navigate to the NFX250 software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/?p=nfx250#sw>
2. 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.
3. In the **Install Package** section of the **Software** tab, select the software package for the release.
4. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.
7. Copy the software to the device or to your internal software distribution site.
8. Install the new package on the device. Use the following command to install the package:

```
root@jdm>request system software add
source/jinstall-host-nfx-2-flex-x86-64-18.1R3-secure-signed.tgz reboot
```

Replace **source** with the path name of the local directory on the device, for example, `/var/tmp`.

Adding the reboot command reboots the device after the upgrade is validated and installed. When the reboot is complete, the device 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 18.1R3 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command.

To download and install Junos OS Release 18.1R3 on NFX150 devices:

1. Using a Web browser, navigate to the NFX150 software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/?p=nfx150#sw>
2. 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.
3. In the **Install Package** section of the **Software** tab, select the software package for the release.
4. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.
7. Copy the software to the device or to your internal software distribution site.
8. Install the new package on the device. Use the following command to install the package:

```
root@nfx150>request vmhost software add  
source/jinstall-host-nfx-3-x86-64-18.1R3-secure-signed.tgz reboot
```

Replace **source** with the path name of the local directory on the device, for example, /var/public.

Adding the reboot command reboots the device after the upgrade is validated and installed. When the reboot is complete, the device 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 18.1R3 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command.

SEE ALSO

[New and Changed Features | 199](#)

[Changes in Behavior and Syntax | 203](#)

[Known Behavior | 204](#)

[Known Issues | 206](#)

[Resolved Issues | 209](#)

[Documentation Updates | 211](#)

[Product Compatibility | 215](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 215](#)
- [Software Version Compatibility | 216](#)

Hardware Compatibility

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

To determine the features supported on NFX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

Software Version Compatibility

This section lists the vSRX and Cloud CPE Solution software releases that are compatible with the Junos OS releases on the NFX150 and NFX250 platforms:

NFX150 Software Version Compatibility

This section lists the vSRX software releases that are compatible with the Junos OS releases on the NFX150 platform:

Table 1: Software Compatibility Details with only vSRX Installed

NFX150 Junos OS Release	vSRX
18.1R1	18.1R1
18.1R2	18.1R2
18.1R3	18.1R3

NFX250 Software Version Compatibility

This section lists the vSRX and Cloud CPE Solution software releases that are compatible with the Junos OS releases on the NFX250 platform:

Table 2: Software Compatibility Details with vSRX and Cloud CPE Solution

NFX250 Junos OS Release	vSRX	Cloud CPE Solution
15.1X53-D40.3	15.1X49-D40.6	Cloud CPE Solution 2.0
15.1X53-D41.6	15.1X49-D61	Cloud CPE Solution 2.1
15.1X53-D102.2	15.1X49-D61	Cloud CPE Solution 3.0
15.1X53-D47.4	15.1X49-D100.6	Cloud CPE Solution 3.0.1
15.1X53-D490	15.1X49-D143	Cloud CPE Solution 4.0
15.1X53-D495	15.1X49-D160	Cloud CPE Solution 4.1

Table 3: Software Compatibility Details with only vSRX Installed

NFX250 Junos OS Release	vSRX
15.1X53-D40.3	15.1X49-D40.6
15.1X53-D41.6	15.1X49-D40.6
15.1X53-D45.3	15.1X49-D61
15.1X53-D47.4	15.1X49-D78.3
17.2R1	15.1X49-D75
17.3R1	15.1X49-D100
15.1X53-D471	15.1X49-D143
18.1R1	18.1R1
18.1R2	18.1R2
18.1R3	18.1R3

SEE ALSO

New and Changed Features 199
Changes in Behavior and Syntax 203
Known Behavior 204
Known Issues 206
Resolved Issues 209
Documentation Updates 211

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- New and Changed Features | 218
- Changes in Behavior and Syntax | 227
- Known Behavior | 231
- Known Issues | 232
- Resolved Issues | 235
- Documentation Updates | 241
- Migration, Upgrade, and Downgrade Instructions | 242
- Product Compatibility | 247

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- Release 18.1R3 New and Changed Features | 219
- Release 18.1R2 New and Changed Features | 219
- Release 18.1R1 New and Changed Features | 219

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for PTX Series.

Release 18.1R3 New and Changed Features

There are no new features or enhancements to existing features for PTX Series in Junos OS Release 18.1R3.

Release 18.1R2 New and Changed Features

There are no new features or enhancements to existing features for PTX Series in Junos OS Release 18.1R2.

Release 18.1R1 New and Changed Features

Hardware

- **New Routing Engine RE-PTX-X8-128G (PTX5000)**—Starting in Junos OS Release 18.1R1, the RE-PTX-X8-128G Routing Engine is supported on the PTX5000 packet transport router. The Routing Engine has increased memory and storage to support node virtualization in future releases. The Routing Engine is equipped with an 8-Core 2.3-GHz processor, 128-GB memory, and 200-GB SSDs and also supports Secure Boot for enhanced boot security.

Class of Service (CoS)

- **Support for explicit-null packet classification using the EXP value from MPLS explicit-null labels (PTX Series)**—The default classification for explicit-null packets is based on the payload (IPv4 or IPv6 DSCP bits). Starting with Junos OS 18.1R1, PTX Series routers with third-generation FPCs (FPC3) support a new 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](#).]

- **Support for enabling a queue's buffer space to be 100 percent of the interface's buffer space (PTX Series)**—Starting in Junos OS 18.1R1, PTX Series devices provide a new CLI option that enables you to set a queue's buffer to be up to 100 percent of the interface's buffer. This option allows the queue's buffer to grow as large as 100 percent of the interface's buffer if and only if it is the only active queue for the interface. This option can be enabled by setting `buffer-size shared` at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.

[See [buffer-size \(Schedulers\)](#).]

Interfaces and Chassis

- **Support for the removal of child next-hop usage for aggregated Ethernet Interfaces and clients (PTX Series routers with FPC3-PTX-U2 and FPC3-PTX-U3)**—Starting in Junos OS Release 18.1R1, Junos OS supports removal of child next-hop usage for aggregated Ethernet Interfaces and clients. Removal of child next-hop usage helps reduce the memory and CPU resources required to support aggregated Ethernet Interfaces and improves the overall system performance and scaling numbers. This feature is enabled by default if the network services mode on the router is configured to `enhanced-mode`. You

can disable this feature by using the **set chassis aggregated-devices disable-lag-enhanced**. You must reboot the router for the configuration to take effect.

Previously, each unicast next hop over aggregated Ethernet Interfaces resulted in creation of a number of children next hops as well. For an aggregated Ethernet Interface with 16 member links, addition of one unicast next hop over the aggregated Ethernet Interface results in installing total of 17 next hops. As a result, with aggregated Ethernet configuration, the number of next hops supported decreases in proportion to the number of aggregated Ethernet links.

NOTE: Child next-hop optimizations are supported for aggregated Ethernet interfaces, interfaces that make use of aggregated Ethernet interfaces, and for both unicast and multicast scenarios.

[See [Aggregated Ethernet Interfaces Overview](#).]

- **Upgraded SSD size and RAM size (PTX5000)**—Starting in Junos OS Release 18.1R1, PTX5000, routers with the RE-PTX-X8-128G-S Routing Engine support Secure Boot BIOS. The SSD size and the RAM size of the Routing Engine is upgraded to 2x200 GB and 128 GB.

[See [Salient Features of the Routing Engines with VM Host Support](#).]

Junos OS XML API and Scripting

- **SLAX and Python scripts now can be sourced over the non-default VRF management instance (PTX Series)**—Starting in Junos OS Release 18.1R1, configuration of commit, event, JET, op, and SNMP scripts is upgraded to support the non-default management routing instance **mgmt_junos** as an option when specifying the source URL for refreshing or downloading SLAX and Python scripts.

[See [Using an Alternate Source Location for a Script](#) or [Configuring and Using a Master Source Location for a Script](#).]

Management

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 18.1R1, the format of telemetry data exported through gRPC for NPU memory and memory utilization implements prefix compression. This change reduces the payload size of data exported. The following example shows the new format:

key: __prefix__

str_value: /components/component[name='FPC0:NPU0']/properties/property

key: [name='mem-util-edmem-size']/value

uint_value: 12345

Telemetry data is exported in key-value pairs. Previously, the data exported included the component and property names in a single key string.

[See [Guidelines for gRPC Sensors](#).]

- **Physical interface operational status sensor (int-exp) support on Junos Telemetry Interface (JTI) (PTX Series)**—Starting with Junos OS Release 18.1R1, sensor int-exp (interface express) is supported 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 paths:

- Subscription path

/junos/system/linecard/intf-exp/

- OpenConfig path

/interfaces_exp/interface_exp[name='et-x/y/z:ch']/state/oper-statusdetails

Streaming telemetry data through gRPC requires you to download the OpenConfig for Junos OS module.

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

- **Expanded support for chassis sensors for Junos Telemetry Interface (MX Series and PTX3000 and PTX5000 Transport Series Routers)**—Starting with Junos OS Release 18.1R1, Junos Telemetry Interface (JTI) provides new sensors that expand optics and power information.

To export telemetry data from Juniper equipment to an external collector requires both Junos Telemetry Interface (JTI) and gRPC to be configured.

Enhanced sensor information is also supported through operational mode commands **show chassis fpc detail**, **show chassis power detail**, and **show chassis pic fpc-slot id pic-slot id**.

Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

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

- **“ON CHANGE” sensor support through gRPC Network Management Interface (gNMI) for Junos Telemetry Interface (MX Series) (PTX Series)**—Starting with Junos OS Release 18.1R1, ON_CHANGE streaming of Address Resolution Protocol (ARP), Network Discovery Protocol (NDP), and IP sensor information associated with interfaces is supported on Junos Telemetry Interface (JTI).

Periodical streaming of OpenConfig operational states and counters has been supported since Junos OS Release 16.1, exporting telemetry data from Juniper equipment to an external collector. While useful in collecting all the needed information and creating a baseline “snapshot,” periodical streaming is less useful for time-critical missions. In such instances, you can configure ON_CHANGE streaming for an external collector to receive information only when operational states experience a change in state.

To support ON_CHANGE streaming, Google has developed a new specification called gRPC Network Management Interface (gNMI) for the modification and retrieval of configurations from a network element. Additionally, the gNMI specification can be used to generate and control telemetry streams from a network element to a data collection system. Using the new gNMI specification, one gRPC service definition can provide a single implementation on a network element for both configuration and telemetry as well as a single NMS element to interact with a device by means of telemetry and configuration RPCs.

Information about the RPCs supporting this feature can be found in the gNMI Proto file version 0.4.0 (the supported version) and the specification released by Google at:

- <https://github.com/openconfig/reference/blob/master/rpc/gnmi/gnmi-specification.md>
- <https://github.com/openconfig/gnmi/blob/master/proto/gnmi/gnmi.proto>

The telemetry RPC **subscribe** under gNMI service supports ON_CHANGE streaming. RPC **subscribe** allows a client to request the target to send it values of particular paths within the data tree. Values may be streamed (STREAM), sent one-off on a long-lived channel (POLL), or sent one-off as a retrieval (ONCE).

If a subscription is made for a top level container with a sample frequency of 0, leaves with ON_CHANGE support are streamed based on events. Other leaves will not be streamed.

NOTE: In order to permit a device to decide which nodes will be streamed as ON_CHANGE and which will SAMPLE, the collector should subscribe for TARGET_DEFINED with sample_interval.

Streaming telemetry data through gRPC requires you to download the OpenConfig for Junos OS module.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **ON_CHANGE support for Junos Telemetry Interface (JTI) (PTX Series)**—Starting with Junos OS Release 18.1R1, OpenConfig support through gRPC Remote Procedure Calls (gRPC) and JTI is extended to support client streaming and bidirectional streaming of telemetry sensor information.

APIs have been implemented in Junos based on Protobuf specifications released by Google for OpenConfig. These APIs perform configuration, operational state retrieval, and telemetry on Junos routers using gRPC as the transport mechanism.

Starting in Junos OS 18.1R1, client streaming and bidirectional streaming are supported. With client streaming, the client sends a stream of requests to the server instead of a single request. The server typically sends back a single response containing status details and optional trailing metadata. With bidirectional streaming, both client and server send a stream of requests and responses. The client starts the operation by invoking the RPC and the server receives the client metadata, method name, and deadline. The server can choose to send back its initial metadata or wait for the client to start sending requests. The client and server can read and write in any order. The streams operate completely independently.

Junos devices can be managed through API (RPC) prototypes:

- **rpc Capabilities (CapabilityRequest)**

Returns (CapabilityResponse). Allows the client to retrieve the set of capabilities that is supported by the target.

- **rpc Get (GetRequest)**

Returns (GetResponse). Retrieves a snapshot of data from the target.

- **rpc Set (SetRequest)**

Returns (SetResponse). Allows the client to modify the state of data on the target.

- **rpc Subscribe (stream SubscribeRequest)**

Returns (stream SubscribeResponse). Allows a client to request the target to send it values for particular paths within the data tree. These values may be streamed (STREAM) or sent one-off on a long-lived channel (POLL), or sent as a one-off retrieval (ONCE). If a subscription is made for a top-level container with a sample frequency of 0, leaves with ON_CHANGE support are streamed based on events. Other leaves will not be streamed.

Juniper Extension Toolkit (JET) support provides insight to users regarding the status of clients connected to JSD. JET support for gRPC includes expanding the maximum number of clients that can connect to JSD from 8 to 30 (the default remains 5). To specify the maximum number of connections, include the **max-connections** statement at the **[edit system services extension-service request-response grpc]** hierarchy level.

To provide information regarding the status of clients connected to JSD, issue the enhanced **show extension-service client information** command and include the **clients** or **servers** options. The **clients**

option displays request-response client information. The **servers** option displays request-response server information.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface.](#)]

MPLS

- **Support for static adjacency segment identifier for aggregated Ethernet member links (PTX Series)**—Starting with Junos OS Release 18.1R1, you can configure a transit single-hop static label switched path (LSP) for a specific member link of an aggregated Ethernet (AE) interface, without enabling the **enhanced-ip** network services mode. A static labeled route is added with next-hop pointing to the AE member link of an aggregated interface. Label for these routes is picked from the segment routing local block (SRLB) pool of the configured static label range. This feature is supported for AE interfaces on PTX routers with FPC1, FPC2, and FPC3. The member-interface CLI statement is added under transit configuration to configure the AE member interface name. The static LSP label is configured from defined static label range.

NOTE:

- In the previous release, this feature was supported only when the **enhanced-ip** network services mode was enabled.
- If the ingress port for the OAM traffic is on FPC1 or FPC2, and the egress port (member link) has the Link Aggregation Control Protocol (LACP) Mux state as 'Detached', and the corresponding physical port (on FPC1, FPC2, or FPC3) is up, the traffic is forwarded at ingress.

[See [Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-hop Static LSP.](#)]

- **Support for segment routing statistics at the ingress (PTX Series)**— Starting in Junos OS Release 18.1R1, 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 at the ingress 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).

To enable recording of SR statistics, include the **sensor-based-stats per-interface-per-member-link ingress** statement at the **[edit protocols isis source-packet-routing]** hierarchy level.

[See [per-interface-per-member-link.](#)]

Network Management and Monitoring

- **Enhancement to Junos OS SNMP MIB PCC functionality (PTX Series)**—Starting in Junos OS Release 18.1R1, Junos OS provides enhanced MIB support for Path Computation Clients. This enhancement enables the Path Computation Client (PCC) process to accept SNMP **get** and **getnext** commands for Path Computation Client Protocol (PCEP) peer and PCEP session tables and reply to them. This feature monitors PCEP interactions between a PCC and a Path Computation Element (PCE). Not all members of PCEP peer and PCEP session tables mentioned in the RFC (RFC 7420) are supported. For exceptions, see [Standard SNMP MIBs Supported by Junos OS](#).

[See [MIB Explorer](#). Name of MIB is **pcep.mib**.]

Routing Policy and Firewall Filters

- **Filter-based GRE encapsulation (PTX Series)**—Starting with Junos OS Release 18.1R1, for PTX Series routers running third-generation line cards, 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.

The following commands are introduced for this feature:

set firewall tunnel-end-point *tunnel-name* gre.

set firewall tunnel-end-point *tunnel-name* ipv4.

set firewall tunnel-end-point *tunnel-name* ipv6.

[See [tunnel-end-point](#) and [Filter-Based Tunneling Across IPv4 Networks](#).]

- **Firewall filter enhancement for better resource optimization (PTX Series)**—When an interface specific firewall filter is configured with multiple Interface bind point instances, the PTX Series Help software allocates resources for each interface instance separately, and the resources consumption is directly proportional to the number of bind points. For better resource optimization, a new **scale-optimized** configuration statement is introduced starting in Junos OS Release 18.1R1 that optimizes interface specific firewall filters in the Packet Forwarding Engine itself. The ingress and egress traffic cannot have the same scale-optimized filter configured.

For more information, see [Guidelines for Configuring Firewall Filters](#).

Routing Protocols

- **Support for BGP multipath at global level (PTX Series)**—Starting with Junos OS Release 18.1R1, BGP multipath is available at the global level in addition to the group and neighbor level. In earlier Junos OS releases BGP multipath is supported only at the group and neighbor levels. A new configuration option **disable** is available at the **[edit protocols bgp multipath]** hierarchy level to disable BGP multipath for specific groups or neighbors. This allows you to configure BGP multipath globally and disable it for specific groups according to your network requirements.

[See [disable](#).]

Security

- **Secure Boot (PTX5000 with Routing Engine RE-PTX-X8-128G)**—Starting in Junos OS Release 18.1R1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected and thus safeguarded from tampering or modification. By default, Secure Boot is enabled on supported routers.

[See [Feature Explorer](#) and enter **Secure Boot**.]

Services Applications

- **Support for MPLS-over-UDP inner payload flow monitoring with IPFIX and version 9 formats (PTX Series)** —Starting with Junos OS Release 18.1R1 on PTX Series routers with an FPC3, PTX10K-LC1101, PTX10K-LC1102, or PTX1000 card, you can perform flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. This feature supports MPLS IPv4 and IPv6 payloads and both IPFIX and version 9 templates. Only ingress sampling is supported.

[See [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers](#).]

SEE ALSO

[Changes in Behavior and Syntax | 227](#)

[Known Behavior | 231](#)

[Known Issues | 232](#)

[Resolved Issues | 235](#)

[Documentation Updates | 241](#)

[Migration, Upgrade, and Downgrade Instructions | 242](#)

[Product Compatibility | 247](#)

Changes in Behavior and Syntax

IN THIS SECTION

- Interfaces and Chassis | 227
- Management | 229
- Network Management and Monitoring | 229
- Network Operations and Troubleshooting Automation | 229
- Subscriber Management and Services | 230

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 18.1R3 for the PTX Series.

Interfaces and Chassis

- **Modified output of the request vmhost zeroize command**—The command **request vmhost zeroize**, upon execution, prompts the user for confirmation to proceed. The following line is displayed:

```
user@host request vmhost zeroize
VMHost Zeroization : Erase all data, including configuration and log files ?
[yes,no] (no) yes
```

See [request vmhost zeroize](#).

- **Power supply alarm is not raised when the input switch status is OFF or power not connected (PTX10008, PTX10016)**—Starting in Junos OS release 17.4R2 and 18.1R2, the power supply alarm **A power supply input has failed** will not be raised if INP1/INP2 switch status is OFF and the power is not connected. Earlier, an alarm is raised for the Power Entry Module (PEM) that are not powered on as **Not Powered** irrespective of the switch state. Now, to know the power supply status, execute **show chassis power** or **show chassis power detail** CLI command. The **DC input** is the new output parameter that provides information about the status of the input feed.

Previous Behavior:

user@host> show chassis power

```
PEM 0:
  State:      Online
```

```

Capacity: 2500 W (maximum 2500 W)
DC output: 864 W (zone 0, 72 A at 12 V, 34% of capacity)

PEM 1:
State:      Online
Capacity: 2500 W (maximum 2500 W)
DC output: 864 W (zone 0, 72 A at 12 V, 34% of capacity)

System:
Zone 0:
Capacity:      7500 W (maximum 7500 W)
Allocated power: 6525 W (975 W remaining)
Actual usage:   2616 W
Total system capacity: 7500 W (maximum 7500 W)
Total remaining power: 975 W

...

```

Current Behavior:

user@host> show chassis power

```

PEM 0:
State:      Online
Capacity: 2500 W (maximum 2500 W)
DC input:  OK (No feed expected, Both feed connected)
DC output: 576 W (zone 0, 48 A at 12 V, 23% of capacity)

PEM 1:
State:      Online
Capacity: 2500 W (maximum 2500 W)
DC input:  OK (No feed expected, Both feed connected)
DC output: 576 W (zone 0, 48 A at 12 V, 23% of capacity)

...

```

[See [show chassis power](#).]

Management

- **Enhancement to LSP statistics sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS 18.1R1, the telemetry data exported for the LSP statistics sensor no longer includes the phrase **and source 0.0.0.0** after the LSP name in the value string for the prefix key. This change reduces the payload size of data exported. The following is an example of the new format:

```
str_value: /mpls/lsp/constrained-path/tunnels/tunnel[name='LSP-4-3']/state/
counters[name='c-27810']/
```

Network Management and Monitoring

- **SNMP syslog messages changed (PTX Series)**—In Junos OS Release 18.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD—AgentX master agent failed to respond to ping. Attempting to re-register
NEW—AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD—NET-SNMP version %s AgentX subagent connected
NEW—NET-SNMP version %s AgentX subagent Open-Sent!

[See the [MIB Explorer](#).]

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a non-default routing instance and a non-default logical system (PTX Series)**—Starting in Junos OS Release 18.1R3, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

Network Operations and Troubleshooting Automation

- **JET - Correction to escaped characters notification events (PTX Series routers)**—Per RFC7159, certain characters must be escaped. Data returned from JET notification subscriptions contained escaped characters that were not required. This has been corrected to comply with RFC7159.
- **respawn-on-normal-exit option added to [edit system extensions extension-service application file <application-name>] hierarchy (PTX Series routers)**—This option helps to ensure that daemonized Juniper Extension Toolkit (JET) applications that exit normally will restart without user intervention. Daemonized JET applications that exit unexpectedly will still restart without user intervention. This is the default behavior.

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (PTX Series)**—Starting in Junos OS Release 18.1R3, the `jdhcpd` process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

In earlier releases:

1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

[See [Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview](#).]

SEE ALSO

New and Changed Features 218
Known Behavior 231
Known Issues 232
Resolved Issues 235
Documentation Updates 241
Migration, Upgrade, and Downgrade Instructions 242
Product Compatibility 247

Known Behavior

IN THIS SECTION

- [General Routing | 231](#)
- [Interfaces and Chassis | 232](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 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.

General Routing

- For CFP2-DCO-T-WDM-1 pluggable, Rx payload type is shown incorrectly (shown as 0 instead of 7). [PR1300423](#)
- When a CFP2-DCO-T-WDM-1 plugged in a PTX Series PIC, when backward frr is enabled on the far end the convergence time is higher because extra delay (average 500 msec) is incurred in the triggering FRR, because of SW-based polling. [PR1303820](#)
- EPR queue is not getting drained during FRR test, and the packets already in the queue timeout eventually, causing this interrupt. This is a minor alarm. [PR1319520](#)
- This is expected behavior for TQ-chip ASICs. It is primarily due to strict-high priority queue and the shared shaper. Credits that are unused by an Output Queue (that is, the queue actual rate is less than the tx-rates) will cause the queue's credit bucket to reach its maximum value. Once a queue hits its maximum credit value, the remaining credits will be distributed to other queues. Once the other queues get transmit credits, they can then transmit. Thus with TQ-chip and the shared shaper, it is virtually impossible to completely shut off a queue by means of a guaranteed rate mechanism. [PR1319923](#)
- Aggregated Ethernet MIX feature is not supported. [PR1330204](#)

Interfaces and Chassis

- On PTX10008 and PTX10016 routers, if you remove the redundant Switch Interface Board (SIB) after upgrading Junos OS from Release 17.4R1 or Release 17.2X75-D90 to a later release, then an alarm is not generated. This is a known behavior and has no impact on the performance of the router.

SEE ALSO

[New and Changed Features | 218](#)

[Changes in Behavior and Syntax | 227](#)

[Known Issues | 232](#)

[Resolved Issues | 235](#)

[Documentation Updates | 241](#)

[Migration, Upgrade, and Downgrade Instructions | 242](#)

[Product Compatibility | 247](#)

Known Issues

IN THIS SECTION

- [General Routing | 233](#)
- [Infrastructure | 235](#)
- [Interfaces and Chassis | 235](#)
- [MPLS | 235](#)

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for the PTX Series.

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 a CFP2-DCO-T-WDM-1 is plugged into a PTX Series PIC, after repeated configuration rollback, link can take a long time to come up. [PR1301462](#)
- When a CFP2-DCO-T-WDM-1 is plugged into a PTX Series PIC, after FPC restart sometimes carrier frequency offset TCA is raised even when TCA not enabled. [PR1301471](#)
- Internal latency is high during initial subscription of sensors when multiple sensors (in order of 15-20) are subscribed together. This is not observed with lesser number of subscriptions. This is for a small period when sensors are being installed. [PR1303393](#)
- This type of crash indicates simultaneous operation on an ephemeral instance. When a process wants to open an ephemeral configuration in merge view, some other activity (like purging, deletion/re-creation) is being carried out on this ephemeral instance. The occurrence of this core file is rare. [PR1305424](#)
- Memory leak in chassisd daemon is noticed while streaming telemetry subscriptions are active. [PR1315672](#)
- On PTX10000, 100G LR4 optics with Part Number 740-061409 will show as QSFP-100G-LR4-T2 instead of QSFP-100G-LR4 and optics that shows as QSFP-100G-LR4 is not supported on PTX10000. [PR1322082](#)
- On PTX platform with FPC type 3, the error message could be observed when FPC card goes online or off-line. [PR1322491](#)
- On PTX Series platforms with cards such as FPC1 and FPC2 and class of service (CoS) used, a high-priority queue might not get the entire configured bandwidth.. [PR1324853](#)
- In a streaming telemetry scenario, if performing "commit full", the na-grpd daemon might restart, causing disconnection of streaming telemetry. [PR1326366](#)
- If a filter is configured with the scale-optimized command, traffic-class-count will not increment. [PR1334580](#)
- MPLS ingress LSP statistics are not supported. [PR1337814](#)
- The issue is with interoperability of TQ- chip with PTX Series routers based line cards. [PR1339481](#)
- RHI interface on PTX Series router can carry a maximum of 4 Mbps of traffic on OQ1 and hence NDP, which is mapped to OQ1 and a bandwidth assigned earlier of 500 pps is reduced to 100 pps due to this limitation. [PR1345938](#)
- PTX3000 reports chip to chip link (CCL) CRC errors while FPC3-SFF-PTX-1X is offlined through the CLI command or by pressing the offline button. The syslog error is generated by an FPC just before it goes offline, so there is no detectable traffic loss. ***** messages ***** Apr 2 08:43:00 fpc4 CMSNGFM: cmsngfpc_fm_send_spry_ctrl_ack: ev_id:11 fm_st:ALL fm_type:FPC_OFF fm_op:DEL Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 0 detected link error for S00F0_0(11,0,11)->FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: Logging statistics for FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: SOT:0x0000037649c2c43e Apr 2 08:43:00 fpc2 CCL: FrameCnt:0x000000000000419dc Apr 2 08:43:00 fpc2 CCL: LastCRCErrCnt:0x000000003 Apr 2 08:43:00 fpc2 CCL: AggrCRCErrCnt:0x0000000000000003 Apr 2 08:43:00 fpc2 CCL: AggrBERCnt:0x0000000000000001 Apr 2 08:43:00 fpc2 CCL: pe0-Avg-28nm-link-10-18 CRC error

```

history (last 5 polls): Apr 2 08:43:00 fpc2 CCL: 0x0 0x0 0x0 0x0 0x3 Apr 2 08:43:00 fpc2 CCL: FEC
Uncorrectable FEC Correctable Apr 2 08:43:00 fpc2 CCL: 00000004, 00000000 Apr 2 08:43:00 fpc2
CCL: 00000000, 00000000 Apr 2 08:43:00 fpc2 BEGIN Rx serdes info for asic pe0-0 serdes 18 Apr 2
08:43:00 fpc2 Signal & port condition for serdes_num 18 Apr 2 08:43:00 fpc2 Rx Signal : Signal Not
OK Apr 2 08:43:00 fpc2 Rx Electrical Idle : High Apr 2 08:43:00 fpc2 Rx Frequency Lock: Set Apr 2
08:43:00 fpc2 Rx Port : Ready Apr 2 08:43:00 fpc2 DFE TAPs : -- snip -- Apr 2 08:43:00 fpc2 CCL:
FrameCnt:0x00000000000041a0d Apr 2 08:43:00 fpc2 CCL: LastCRCErrCnt:0x00000003 Apr 2 08:43:00
fpc2 CCL: AggrCRCErrCnt:0x0000000000000003 Apr 2 08:43:00 fpc2 CCL:
AggrBERCnt:0x0000000000000001 Apr 2 08:43:00 fpc2 CCL: pe0-Avg-28nm-link-14-22 CRC error
history (last 5 polls): Apr 2 08:43:00 fpc2 CCL: 0x0 0x0 0x0 0x0 0x3 Apr 2 08:43:00 fpc2 CCL: FEC
Uncorrectable FEC Correctable Apr 2 08:43:00 fpc2 CCL: 00000004, 00000000 Apr 2 08:43:00 fpc2
CCL: 00000000, 00000000 Apr 2 08:43:00 fpc2 BEGIN Rx serdes info for asic pe0-0 serdes 22 Apr 2
08:43:00 fpc2 Signal & port condition for serdes_num 22 Apr 2 08:43:00 fpc2 Rx Signal : Signal Not
OK Apr 2 08:43:00 fpc2 Rx Electrical Idle : High Apr 2 08:43:00 fpc2 Rx Frequency Lock: Set Apr 2
08:43:00 fpc2 Rx Port : Ready Apr 2 08:43:00 fpc2 DFE TAPs : -- snip -- Apr 2 08:43:00 fpc2 CCL:
Logging errors for FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: BER Err Apr 2 08:43:00 fpc2 CCL: Frame
Lock Loss Apr 2 08:43:00 fpc2 CCL: Align Loss Apr 2 08:43:00 fpc2 CCL: Header Comparison Error Apr
2 08:43:00 fpc2 CCL: Header Preamble Error Apr 2 08:43:00 fpc2 CMSNGFM:
cmsngfpc_platform_fm_periodic: PFE 0 detected link error for S00F1_0(14,0,14)->FPC02FE0(1,00) Apr
2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 1 detected link error for
S00F0_0(11,0,11)->FPC02FE1(0,00) Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic:
PFE 1 detected link error for S00F1_0(14,0,14)->FPC02FE1(1,00) user@router> show chassis hardware
detail Hardware inventory: FPC 0 REV 43 750-057064 ACPV7514 FPC3-SFF-PTX-1X CPU BUILTIN
BUILTIN SMPC PMB FPC 2 REV 40 750-057064 ACPJ9145 FPC3-SFF-PTX-1X CPU BUILTIN BUILTIN
SMPC PMB FPC 4 REV 43 750-057064 ACPR8506 FPC3-SFF-PTX-1X CPU BUILTIN BUILTIN SMPC
PMB SIB 0 REV 10 750-057067 ACPJ8829 SIB3-SFF-PTX SIB 1 REV 10 750-057067 ACPJ8683
SIB3-SFF-PTX SIB 2 REV 10 750-057067 ACPJ8843 SIB3-SFF-PTX SIB 3 REV 10 750-057067 ACPJ8920
SIB3-SFF-PTX PR1348733

```

- If the output firewall filter is configured with the "syslog" option, the host interface might be wedged on a PTX1000 or a PTX Series platform with FPC type 3. [PR1354580](#)

Infrastructure

- A file system corruption might create a kernel core file. The Routing Engine reboots with the message `ffs_blkfree: freeing free block`. [PR1028972](#)

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 crash the `cfmd` process. This is because of using the older version of `/var/db/cfm.db`. [PR1281073](#)

MPLS

- When the `rpd` daemon is terminating, the process of signaling the deletion of all RSVP LSPs might take so long that a watchdog timer begins firing, resulting in an `rpd` core file. [PR1257367](#)

SEE ALSO

New and Changed Features	 218
Changes in Behavior and Syntax	 227
Known Behavior	 231
Resolved Issues	 235
Documentation Updates	 241
Migration, Upgrade, and Downgrade Instructions	 242
Product Compatibility	 247

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3](#) | [236](#)
- [Resolved Issues: 18.1R2](#) | [237](#)
- [Resolved Issues: 18.1R1](#) | [238](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

General Routing

- On a PTX1000, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3-20170721 fails frequently with sampling enabled. [PR1296533](#)
- FPC is being reported down in chassisd logs related to Streaming Telemetry, whereas the FPC is online. [PR1300795](#)
- PTX10000: Suppress chassis alarm for switched off PEMS. [PR1311574](#)
- PTX10000: Do not bounce FPC without warning or alarm for different port speed settings. [PR1311875](#)
- MPLS traceroute fails across PTX Series platforms. [PR1327609](#)
- Status LED on the chassis does not show UP on PTX10002-60c. [PR1332991](#)
- Telemetry software package is not running after system is returned to zero. [PR1336004](#)
- FPC got rebooted a few minutes after loading the configuration. [PR1346467](#)
- Threshold is not getting configured correctly in PTX Series router when threshold is configured using scope and category options. [PR1350841](#)
- BFD sessions do not come up on PTX3000. [PR1352112](#)
- The interface of 15 100G ports PIC might delay 60 seconds to come up. [PR1357410](#)
- P2MP LSP replication traffic loss occurs on aggregated Ethernet bundle after member link goes down on PTX Series router. [PR1359974](#)
- The route might be stuck after BGP neighbor and route flapping. [PR1362560](#)
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with the "Link-Layer-Down" flag set. [PR1365263](#)
- On PTX Series IPLC (OPT3-SFF-PTX FPC), a first J-UKERN crash triggers multiple secondary J-UKERN crashes. [PR1365791](#)
- The 'commit' or 'commit check' might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)

MPLS

- MPLS LSP statistics are not shown in the CLI command **show mpls lsp ingress statistics**. [PR1344039](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition. [PR1374102](#)

Platform and Infrastructure

- Running RSI through the console port might cause the system to crash and reboot. [PR1349332](#)
- Traffic black hole is seen along with JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH messages. [PR1357707](#)

Routing Protocols

- Protocol churn will create rpd crash. [PR1341466](#)
- Rpd core might be seen while running streaming telemetry. [PR1347431](#)

Resolved Issues: 18.1R2

General Routing

- Remove **show chassis spmb** command/response. [PR1244059](#)
- De-encapsulated GRE traffic is not passing to the egress port upon initial customer configuration. [PR1325104](#)
- On PTX5000 line packet transport routers with FPC3 linecards, PTX10000 line, and PTX1000 line output firewall filters that are configured with **syslog** and **discard** actions do not perform the **syslog** action. [PR1328426](#)
- PTX10000 line card might reboot continuously after upgrading to Junos OS Release 17.2R1 or later if HMC BIST fails. [PR1330618](#)
- Traffic stops flowing out of ae70 after some FPC restart iterations. [PR1335118](#)
- Member of IPv4 unicast next hops might be stuck in "Replaced" state after interface flaps. [PR1336201](#)
- FPC/FPC2/FPC E on PTX Series does not forward traffic. [PR1339524](#)
- PTX1008: 30-Port Coherent Line Card (DWDM-IC) does not come up. [PR1344732](#)

Interfaces and Chassis

- On PTX3000 line, when CPM is configured, the CFM filter is not getting programmed. The command **show oam ethernet connectivity-fault-management** is not checking interfaces ae0.0 extensive. [PR1335305](#)

Platform and Infrastructure

- On PTX1000 and QFX10002-60C: Python scripts/shell scripts cannot be executed during ZTP because verixec is enabled. [PR1334425](#)

Routing Protocols

- The rpd might constantly consume high CPU in BGP setup. [PR1315066](#)
- The primary path of MPLS LSP might switch to another address. [PR1316861](#)

- On PTX1000, rpd core file is generated at `ispfc_incrementally_mend_one_postf_sp_tree (postf_spf_res= <optimized out>, topo= <optimized out>)` at `../..../src/junos/usr.sbin/rpd/lib/igp-spf-compute/igp_spf_compute_ti_lfa.c:3364`. [PR1339296](#)

Resolved Issues: 18.1R1

General Routing

- `restart na-grpc-server` and `restart na-mqtt` for MTRE does not work. [PR1284121](#)
- Traffic passing LSP with entropy label might be dropped after the bypass path goes down. [PR1291036](#)
- The routing protocol process (rpd) might generate a core file while restarting the process from the CLI. [PR1291110](#)
- Incorrect SNMP OID values are sent in SNMP traps for removal or insertion of front panel display on PTX Series routers. [PR1294741](#)
- **LINK LED** is red when the port is disabled on PTX Series routers. [PR1294871](#)
- Classifier binding for logical interface (IFL) is lost when changing from trunk to access with L2 classifier configured. [PR1295043](#)
- After deleting and re-adding an interface that had a classifier bound, the binding to the default classifier is not restored. [PR1295477](#)
- FPC might crash after Routing Engine switchover. [PR1296282](#)
- An mgd core file is generated when downgrading from Junos OS Release 17.3-20170721 to Junos OS Release 16.1X65D40.2. The mgd core file is overwritten if downgrading is attempted multiple times. [PR1296504](#)
- On a PTX1000, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3-20170721 fails frequently with sampling enabled. [PR1296533](#)
- Upon configuring protect core, rpd keeps thrashing at `rt_attrib.c - nhp->rtn_n_gw > 1 && nhp->rtn_n_gw <= 64`. [PR1297044](#)
- Alarms and syslog errors are seen with priority strict-high on an AF4 queue, on the oversubscription cases (1X100G egress to 1X10G egress setup). [PR1297343](#)
- **Link errors** alarm messages might be seen after migrating to FPC3 on PTX3000. [PR1298841](#)
- The `disable-pfe` action upon Hybrid Memory Cube (HMC) fatal errors might have a system-wide impact on PTX Series platforms. [PR1300180](#)
- On PTX3000 platforms, powering on a FPC (PTX-IPLC-B-32) card might cause the other FPC cards to reboot. [PR1302304](#)
- The third-generation FPC (FPC3-SFF-PTX) might not boot on a PTX3000 with the Control Board or Routing Engine. [PR1303295](#)
- On PTX3000 and PTX5000 platforms, the 100G interfaces might not come up. [PR1303324](#)

- If MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic with the error message **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)
- Repeated log messages **%PFE-3 fpcX expr_nh_index_tree_ifl_get** and **expr_nh_index_tree_ipaddr_get** are observed when the sampling packet is discarded with log (or syslog) command under the firewall filter. [PR1304022](#)
- The **interface hold-time down** timer does not take effect on PTX5000 with optical interface. [PR1307302](#)
- Packet Forwarding Engine error messages are flooding as **expr_sensor_update_cntr_to_sid_tree** after deleting and rolling back **protocols isis source-packet-routing node-segment**. [PR1309288](#)
- After ZTP gets completed, the configuration on the backup Routing Engine is rolled back to the factory-default configuration in a few minutes. [PR1310117](#)
- The **SIB LED** on the FPD goes to green steady state even before an SIB comes online. [PR1311632](#)
- PTX Series routers might fail to grab a new region for the next hop during a link flap. [PR1311850](#)
- Rpd core observed after multiple session flaps on scale setup. [PR1312169](#)
- The 10g interface might flap if it is set to 100g speed. [PR1315079](#)
- Continuous logs from vhcclient are seen for all the commands executed. [PR1315128](#)
- ZTP config file fails to be loaded and committed, Cannot open configuration file **/config/auto_image_upgrade.conf**. [PR1315857](#)
- The physical interfaces (IFDs) might generate framing errors when ports are connected to an odd interface. [PR1317827](#)
- After jack-out/jack-in FPCs show "No-Power" for some time; however, FPCs eventually come up. [PR1319156](#)
- No traffic is flowing with IPV6 payload prefixes on PTX Series platform. [PR1319273](#)
- The rpd might crash when OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)
- The aggregated Ethernet interface might be in an error blocked state after enabling and disabling aggregated Ethernet member link in a large scale of routes scenario. [PR1323398](#)
- Fan tray 0(FAN-PTX-V-S) Failure alarm is set and cleared on multiple PTX Series routers. [PR1324905](#)
- On PTX1000, the local time on FPC might be different from the local time on Junos VM host. [PR1325048](#)
- Firewall filter is not supported on aggregated Ethernet. [PR1325237](#)
- On PTX5000 w/ FPC3 linecards, PTX10000, and PTX1000 platforms, output firewall filters that are configured with "syslog" and "discard" actions do not perform the "syslog" action. [PR1328426](#)
- Sensor configuration is not deleted from ephemeral database after collector disconnects. [PR1329134](#)

- Additional TX reset during link-down events can cause link instabilities. [PR1330708](#)
- PTX FPC might reboot in certain rare scenarios when an flapping interface configuration is committed. [PR1335161](#)

Infrastructure

- ixlv interface statistics not accounting properly. [PR1313364](#)
- PTX device may get to abnormal state due to the malfunction of the protection mechanism for F-Label. [PR1336207](#)

Interfaces and Chassis

- The transportd might crash when SNMP 1335438query on jnxoptIfOChSinkCurrentExtTable with unsupported interface index. [PR1335438](#)

MPLS

- Traffic drop occurs during NSR switchover for RSVP P2MP provider tunnels used by MVPN. [PR1293014](#)
- Traffic loss occurs for static LSP configured with the **stitch** command. [PR1307938](#)
- Rpd crashes on backup Routing Engine due to memory exhaustion. [PR1328974](#)

Platform and Infrastructure

- Continuous log messages occur: for example: **tftpd[23724]: Timeout #35593 on DATA block 85.** [PR1315682](#)
- PTX1000 & QFX10002-60C: Python scripts/shell scripts cannot be executed during ZTP as veriexec is enabled. [PR1334425](#)

Routing Protocols

- With BGP LU FRR in an inter-AS scenario, a very high FRR time is visible once the link is up. [PR1307258](#)
- Assignment of Sub-TLV values for segment routing TE policy Sub-TLVs occurs. [PR1315486](#)
- The rpd process might crash continuously on both Routing Engines when "backup-spf-options remote-backup-calculation" is configured in IS-IS protocol. [PR1326899](#)

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](#)

SEE ALSO

[New and Changed Features | 218](#)

[Changes in Behavior and Syntax | 227](#)

[Known Behavior | 231](#)[Known Issues | 232](#)[Documentation Updates | 241](#)[Migration, Upgrade, and Downgrade Instructions | 242](#)[Product Compatibility | 247](#)

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 241](#)

This section lists the errata and changes in Junos OS Release 18.1R3 documentation for the PTX Series.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.

- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 218](#)

[Changes in Behavior and Syntax | 227](#)

[Known Behavior | 231](#)

[Known Issues | 232](#)

[Resolved Issues | 235](#)

[Migration, Upgrade, and Downgrade Instructions | 242](#)

[Product Compatibility | 247](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 242](#)
- [Upgrading a Router with Redundant Routing Engines | 243](#)
- [Basic Procedure for Upgrading to Junos OS Release 18.1 | 243](#)

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.

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 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 acting as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Basic Procedure for Upgrading to Junos OS Release 18.1

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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 18.1R3:

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 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: After you install a Junos OS Release 18.1R3 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.

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 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.

Customers in the United States and Canada, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-18.1R3.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-18.1R3.SPIN-export-signed.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** (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.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 18.1 **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](#).

SEE ALSO

New and Changed Features 218
Changes in Behavior and Syntax 227
Known Behavior 231
Known Issues 232
Resolved Issues 235
Documentation Updates 241
Product Compatibility 247

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 247](#)

Hardware Compatibility

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

To determine the features supported on PTX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

[New and Changed Features | 218](#)

[Changes in Behavior and Syntax | 227](#)

[Known Behavior | 231](#)

[Known Issues | 232](#)

[Resolved Issues | 235](#)

[Documentation Updates | 241](#)

[Migration, Upgrade, and Downgrade Instructions | 242](#)

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- New and Changed Features | 248
- Changes in Behavior and Syntax | 273
- Known Behavior | 276
- Known Issues | 282
- Resolved Issues | 287
- Documentation Updates | 299
- Migration, Upgrade, and Downgrade Instructions | 300
- Product Compatibility | 314

These release notes accompany Junos OS Release 18.1R3 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Caveat: Juniper Networks does not recommend configuring and deploying EVPN-VXLAN on QFX Series platforms running Junos OS 18.1R1.

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.

New and Changed Features

IN THIS SECTION

- Release 18.1R3-S3 New and Changed Features | 249
- Release 18.1R3 New and Changed Features | 252
- Release 18.1R2 New and Changed Features | 254
- Release 18.1R1 New and Changed Features | 255

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for QFX Series.

NOTE: The following QFX Series platforms are supported in Release 18.1R3: QFX5100, QFX5110, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016.

Release 18.1R3-S3 New and Changed Features

EVPNs

- **Multicast support with IGMP snooping on spine and leaf devices in an EVPN-VXLAN centrally-routed bridging overlay network (QFX10002, QFX10008, QFX10016, and QFX5110 switches)**—Starting with Junos OS Release 18.1R3-S3, multicast forwarding with IGMP snooping is supported on spine and leaf devices in an EVPN-VXLAN centrally-routed bridging overlay network consisting of QFX10002, QFX10008, or QFX10016 switches as spine devices and QFX5110 switches as leaf devices. With IGMP snooping enabled in a centrally-routed bridging architecture, leaf devices forward multicast traffic at Layer 2 within a VLAN only, while spine devices perform forwarding within a VLAN and can also be configured with IRB interfaces to perform inter-VLAN routing. Spine or leaf devices forward multicast traffic on the access side only to interested listeners based on IGMP snooping state, but continue to use ingress replication to flood multicast traffic into the EVPN core to reach other spine or leaf devices that might serve interested listeners.

All of the following scenarios are supported for both intra-VLAN and inter-VLAN multicast traffic:

- Multicast source and multicast receivers within the EVPN-VXLAN network
- Multicast source external to the EVPN-VXLAN network and multicast receivers within the EVPN-VXLAN network
- Multicast receiver external to the EVPN-VXLAN network and multicast source within the EVPN-VXLAN network

To route multicast traffic from or to multicast sources and receivers external to the EVPN-VXLAN network, spine devices use PIM on a multicast VLAN through an external gateway (such as an MX Series router).

[See [Multicast Support in EVPN-VXLAN Overlay Networks](#).]

- **IPv6 data traffic support through an EVPN-VXLAN overlay network (QFX5110 switches)**—Starting with Junos OS Release 18.1R3-S3, QFX5110 switches that function as Layer 3 VXLAN gateways can route IPv6 data traffic through an EVPN-VXLAN overlay network. With this feature enabled, Layer 2 or 3 data packets from one IPv6 host to another IPv6 host are encapsulated with an IPv4 outer header and transported over the IPv4 underlay network. The Layer 3 VXLAN gateways in the EVPN-VXLAN overlay network learn the IPv6 routes through the exchange of EVPN Type 2 and Type 5 routes.

[See [Routing IPv6 Data Traffic through an EVPN-VXLAN Network With an IPv4 Underlay](#).]

- **MAC filtering, storm control, and port mirroring support in EVPN-VXLAN networks (QFX5100 and QFX5110 switches)**—Starting with Junos OS Release 18.1R3-S3, QFX5100 and QFX5110 switches support the following features in an EVPN-VXLAN overlay network:

- MAC filtering
- Storm control
- Port mirroring and analyzers

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

- **MAC filtering and storm control support in EVPN-VXLAN networks (QFX10002 and QFX10008 switches)**—Starting with Junos OS Release 18.1R3-S3, QFX10002 and QFX10008 switches support the following features in an EVPN-VXLAN overlay network:

- MAC filtering
- Storm control

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

- **Support for firewall filtering and policing on EVPN-VXLAN traffic (QFX5100 and QFX5110)**—Starting with Junos OS Release 18.1R3-S3, you can configure firewall filters and policers on VXLAN traffic in an EVPN topology. Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. Policing, or rate limiting, lets you control the amount of traffic that enters the switch and determines the actions to take when the traffic exceeds the defined limit. You configure firewall filters at the **[edit firewall]** hierarchy level. For each firewall filter that you apply to a VXLAN, you can specify **family ethernet-switching** to filter Layer 2 (Ethernet) packets or **family inet** to filter on IRB interfaces. The IRB interface acts as a Layer 3 routing interface to connect the XVLANS in collapsed or non-collapsed IP fabric topologies. You can only apply firewall filters and policers on CE-facing interfaces in the ingress direction (traffic entering the XVLAN). For IRB interfaces, you can only apply filtering at the ingress point of a non-encapsulated frame routed through the IRB interface.

This feature was previously supported in an "X" release of Junos OS.

This feature is not supported on a QFX5100 Virtual Chassis in an EVPN-VXLAN topology.

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#) and [Overview of Firewall Filters.](#)]

- **Support for VMTO for ingress traffic (QFX5100, QFX5110, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 18.1R3-S3, you can configure the PE device to support virtual machine traffic optimization (VMTO) for ingress traffic. VMTO eliminates the unnecessary ingress routing to default gateways when a virtual machine is moved from one data center to another.

To enable VMTO, configure **remote-ip-host** routes in the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. You can also filter out the unwanted routes by configuring an import policy under the **remote-ip-host routes** option.

[See [Configuring EVPN Routing Instances.](#)]

- **Support for Multihomed Proxy Advertisement (QFX5100, QFX5110, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016 switches)**—Junos now provides enhanced support to proxy advertise the Mac address and IP route entry from all PEs that are multi-homed to a CE device. This can prevent traffic loss when one of the links to the PE fails. To support the multihomed proxy advertisement, all multi-homed PE devices should have the same multihomed proxy advertisement bit value. The multihomed proxy advertisement feature is enabled by default and Junos uses the default multihomed proxy advertisement bit value of 0x20.

[See [EVPN Multihoming Overview](#).]

- **Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks (QFX Series)**—Starting in Junos OS Release 18.1R3-S3, you can configure OSPF, IS-IS, BGP, and static routing with bidirectional forwarding detection (BFD) on an IRB interface that is used as a routed interface in EVPN. This allows protocol adjacencies to be established between an IRB on a Layer 3 gateway and a CE device connected directly to a Layer 3 gateway or to a Layer 2 leaf device in an EVPN-VXLAN network.

[See [Supported Protocols on an IRB Interface in EVPN-VXLAN](#) .]

Routing Policy and Firewall Filters

- **Support for IPv6 Filter-Based Forwarding (QFX5200 switches)**— Starting with Junos OS Release 18.1R3-S3, you can use stateless firewall filters in conjunction with filters and routing instances to control how IPv6 traffic travels in a network. This is called IPv6 filter-based forwarding. To setup 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.

This feature was previously supported in an "X" release of Junos OS.

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

Security

- **Support for firewall filtering and policing on EVPN-VXLAN traffic (QFX5100, QFX5100 Virtual Chassis, QFX5110 switches)**— Starting with Junos OS Release 18.1R3-S3, you can configure firewall filters and policers on VXLAN traffic in an EVPN topology. For each firewall filter that you apply to a VXLAN, you can specify **family ethernet-switching** to filter Layer 2 (Ethernet) packets or **family inet** to filter on IRB interfaces. The IRB interface acts as a Layer 3 routing interface to connect the VXLANs in collapsed or non-collapsed IP fabric topologies. You can only apply firewall filters and policers on CE-facing interfaces in the ingress direction (traffic entering the VXLAN). For IRB interfaces, you can only apply filtering at the ingress point of a non-encapsulated frame routed through the IRB interface.

This feature was previously supported in an "X" release of Junos OS.

[See [Understanding EVPN with VXLAN Data Plan Encapsulation](#) and [Overview of Firewall Filters](#).]

Release 18.1R3 New and Changed Features

EVPNs

- **Layer 2 and 3 families, encapsulation types, and VXLAN on same physical interface (QFX5100, QFX5110, and QFX5200 switches)**— Starting with Junos OS Release 18.1R3, you can configure and successfully commit the following on a physical interface of a QFX5100, QFX5110, or QFX5200 switch in an EVPN-VXLAN environment:
 - Layer 2 bridging (**family ethernet-switching**) on any logical interface unit number (unit 0 and any non-zero unit number).
 - VXLAN on any logical interface unit number (unit 0 and any non-zero unit number).
 - Layer 2 bridging (**family ethernet-switching** and **encapsulation vlan-bridge**) on different logical interfaces (unit 0 and any non-zero unit number).
 - Layer 3 IPv4 routing (**family inet**) and VXLAN on different logical interfaces (unit 0 and any non-zero unit number).

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

Interfaces and Chassis

- **Support for connectivity fault management (CFM) (QFX5210 switches)**—Starting in Junos OS 18.1R3, you can use the connectivity fault management (CFM) feature to monitor an Ethernet network that may comprise one or more service instances. A service instance could be a VLAN or a collection of VLANs. CFM creates a maintenance domain (MD) entity that is a network or part of the network for which faults in connectivity are managed. An MD is associated with a level. The allocation of levels to the various network entities are decided based on their needs from an OAM perspective. For example, network entities such as operators, providers, and customers can be part of different administrative domains. Each administrative domain is mapped into one OAM domain. The OAM domain provides enough information for management, avoiding security breaches, and performing end-to-end monitoring. Configure CFM at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Junos on White Box

- **Junos on White Box**—Starting with Junos OS Release 18.1R3, the Junos on White Box software provides a disaggregated Junos that decouples the Junos operating system from Juniper Networks switches and runs as independent software on Open Compute Project (OCP)-compliant network hardware, enabling you to use that hardware in your data center (DC) networks and providing a robust, feature-rich network operating system for enabling the DC Fabric buildout. Junos for White Box is standalone software providing standards-based network protocols such as ISIS and BGP, overlay technology such as VXLAN with EVPN control plane, and full automation capabilities and is similar to the reliable, high performance Junos OS that powers the Juniper Networks QFX Series Data Center portfolio.

Key Junos OS features that enhance the functionality and capabilities of the White Box switches include:

- Software modularity, with process modules running independently in their own protected memory space and with the ability to do process restarts.
- Uninterrupted routing and forwarding, with features such as nonstop active routing (NSR) and nonstop bridging (NSB).

- Commit and rollback functionality that ensures error-free network configurations.
- A powerful set of scripts for on-box problem detection, reporting, and resolution.

Release 18.1R2 New and Changed Features

EVPNs

- **IPv4 inter-VLAN multicast forwarding modes for EVPN (QFX10000 switches)**—Starting with Junos OS Release 18.1R2, QFX10000 switches can forward IPv4 multicast traffic between VLANs in EVPN-VXLAN networks with these IP fabric architectures:
 - Two-layer IP fabric in which QFX10000 switches function as Layer 3 gateways, and QFX5100 or QFX5200 switches function as Layer 2 gateways. From their central location in the IP fabric, the QFX10000 switches on which IRB interfaces are configured can route multicast traffic from one VLAN to another. This mode of multicast forwarding is known as *centrally-routed mode*.
 - One-layer IP fabric in which QFX10000 switches function as both Layer 2 and Layer 3 gateways. From their location at the edge of the IP fabric, the QFX10000 switches on which IRB interfaces are configured can route multicast traffic from one VLAN to another. This mode of multicast forwarding is known as *edge-routed mode*.

To configure the multicast forwarding mode, you can specify the **irb** configuration statement with the **local-remote** option (centrally-routed mode) or the **local-only** option (edge-routed mode) in the **[edit forwarding-options multicast-replication evpn]** hierarchy level.

NOTE: We do not recommend specifying the **local-remote** option on some QFX10000 switches and the **local-only** option on the other QFX10000 switches in either of the IP fabric architectures. Doing so might cause the QFX10000 switches to forward the inter-VLAN multicast traffic inconsistently.

[See [Multicast Support in EVPN-VXLAN Overlay Networks](#).]

Restoration Procedures and Failure Handling

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (QFX Series)**—In Junos OS Release 18.1R2, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode. The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

[See [Saving a Rescue Configuration File](#).]

Release 18.1R1 New and Changed Features

Hardware

- **QFX10002-60C switch**—Starting in Junos OS Release 18.1R1, Juniper Networks introduces the QFX10002-60C switch. The Juniper Networks QFX10000 line of Ethernet switches provides cloud builders and data center operators with scalable solutions for both core and spine data center deployments. The 2 U fixed-configuration switch has 60 flexible configuration speed ports that can be set for 40-Gbps or 100-Gbps speeds. The QFX10002-60C also supports 10-Gigabit Ethernet when the ports are configured for 40-Gigabit Ethernet and channelized into 4 independent 10-Gigabit Ethernet ports. The QFX10002-60C is available with either AC or DC power supplies. The airflow is airflow out, where air comes into the vents in the port panel and exhausts through the field-replaceable units (FRU) panel. [See [QFX10002 Hardware Overview](#).]
- **QFX5210-64C switch**—Starting in Junos OS Release 18.1R1, Juniper Networks introduces the QFX5210-64C Switch. The 1 U fixed configuration switch is designed for cloud customers who need either a top-of-rack switch or a lean spine switch with flexible port speeds and high-port density. The Routing Engine and control plane are driven by the 2.2 GHz quad-core Intel; Xeon; CPU with 16 GB of memory and a 128-GB solid-state drive (SSD) for storage. The QFX5210-64C can be configured for 10/25/40/50/100 Gigabit Ethernet speeds. The switch comes standard with redundant fans and redundant power supplies. The QFX5210-64C can be ordered with either ports-to-FRUs or FRUs-to-ports airflow. The model is available with either AC or DC power supplies. [See [QFX5210 System Overview](#).]
- **QFX5200-48Y switch**— The Juniper Networks QFX5200 line of fixed-configuration access switches are designed for cloud builders and data centers deploying next-generation IP fabric networks. The QFX5200-48Y offers 48 ports of native 25-Gbps speed for downlinks and 6 ports of 100-Gbps speeds for uplinks. The 1 U fixed chassis switch allows a flexible configuration of the ports. The 40 downlink ports can be configured either as 10-Gbps speeds or 25-Gbps speeds while the 6 uplink ports can be configured for either 40-Gbps speeds or 100-Gbps speeds. The QFX5200-48Y comes standard with redundant fans and redundant power supplies. The QFX5200-48Y can be ordered with either

ports-to-FRUs (AFO) or FRUs-to-ports (AFi) airflow. The model is available with either AC or DC power supplies.

[See [QFX5200 Switch Hardware Guide](#).]

Authentication, Authorization, and Accounting (AAA) (RADIUS)

- **Access control and authentication (QFX5100 switches)**—Starting with Junos OS Release 18.1 R1, QFX5110 and QFX5200 switches support controlling access to your network using 802.1X authentication and MAC RADIUS authentication.
 - 802.1X authentication provides port-based network access control (PNAC) as defined in the IEEE 802.1X standard. QFX5100 switches support 802.1X features including guest VLAN, private VLAN, server fail fallback, dynamic changes to a user session, RADIUS accounting, and configuration of port-filtering attributes on the RADIUS server using VSAs. You configure 802.1X authentication at the **[edit protocols dot1x]** hierarchy level.
 - MAC RADIUS authentication is used to authenticate end devices independently of whether they are enabled for 802.1X authentication. You can permit end devices that are not 802.1X-enabled to access the LAN by configuring MAC RADIUS authentication on the switch interfaces to which the end devices are connected. You configure MAC RADIUS authentication at the **[edit protocols dot1x authenticator interface *interface-name* mac-radius]** hierarchy level.

[See [Understanding Authentication on Switches](#).]

Class of Service (CoS)

- **Support for data center quantized congestion notification (DCQCN) (QFX5100, QFX5110, QFX5200, QFX5210 switches)**—Remote Direct Memory Access (RDMA) provides the high throughput and ultra-low latency, with low CPU overhead, necessary for modern datacenter applications. RDMA is deployed using the RoCEv2 protocol, which relies on priority-based flow control (PFC) to enable a drop-free network. DCQCN is an end-to-end congestion control scheme for RoCEv2. Starting in Junos OS Release 18.1R1, Junos OS supports DCQCN by combining explicit congestion notification (ECN) and PFC to overcome the limitations of PFC to support end-to-end lossless Ethernet.

[See [Data Center Quantized Congestion Notification \(DCQCN\)](#).]

EVPN

- **Support for IGMP snooping for EVPN-VXLAN in a multihomed environment (QFX5110 switches)**—Starting in Junos OS Release 18.1R1, QFX5110 switches support IGMP snooping with Ethernet EVPN (EVPN). This feature is useful in an EVPN-VXLAN environment with significant multicast traffic. IGMP snooping enables PE devices to send multicast traffic to CE devices only as needed, which preserves bandwidth. To configure IGMP snooping, include the **igmp-snooping (all | vlan-number)** set of statements at the **[edit protocols]** hierarchy level. You must also include the proxy statement in the IGMP snooping configuration. All multihomed interfaces must have the same configuration.

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

- **EVPN control plane and VXLAN data plane support (QFX5210 switches)**—By using a Layer 3 IP-based underlay network coupled with an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network, you can deploy larger networks than those possible with traditional Layer 2 Ethernet-based architectures. With overlay networks, endpoints (bare-metal servers [BMSs] and virtual machines [VMs]) can be placed anywhere in the network and can remain connected to the same logical Layer 2 network, enabling the virtual topology to be decoupled from the physical topology.

The physical underlay network over which EVPN-VXLAN is commonly deployed is a two-layer IP fabric, which includes spine and leaf devices. The spine devices provide connectivity between the leaf devices, and the leaf devices function as Layer 2 VXLAN gateways and provide connectivity to the attached endpoints. Starting with Junos OS Release 18.1R1, you can deploy QFX5210 switches as leaf nodes in the EVPN-VXLAN overlay network.

[See [Understanding EVPN with VXLAN Data Encapsulation](#).]

- **EVPN proxy ARP and ARP suppression, and NDP and NDP suppression with or without IRB interfaces (QFX5100, QFX5110, and QFX5200 switches)**—Starting with Junos OS Release 18.1R1, QFX5100 and QFX5200 switches that function as Layer 2 VXLAN gateways and QFX5110 switches that function as Layer 2 or Layer 3 VXLAN gateways in an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment support proxy Address Resolution Protocol (ARP) and ARP suppression, and Network Discovery Protocol (NDP) and NDP suppression. The proxy ARP and ARP suppression, and NDP and

NDP suppression capabilities are enabled by default. Any interface configured on a Layer 2 or Layer 3 VXLAN gateway can deliver ARP requests from both local and remote hosts.

In addition, you can control the following aspects of the media access control (MAC)-IP address bindings database on a Layer 2 or Layer 3 VXLAN gateway:

- The maximum number of MAC-IP address entries in the database
- The amount of time a locally learned MAC-IP address binding remains in the database

[See [EVPN Proxy ARP and ARP Suppression, and NDP and NDP Suppression](#).]

- **Support for duplicate MAC address detection and suppression (QFX5100, QFX5110, and QFX5200 switches)**— When a MAC address relocates, PE devices can converge on the latest location by using sequence numbers in the extended community field. Misconfigurations in the network can lead to duplicate MAC addresses. Starting in Junos OS Release 18.1R1, QFX5100, QFX5110, and QFX5200 switches support duplicate MAC address detection and suppression.

You can modify the duplicate MAC address detection settings on the switches by configuring the detection window for identifying duplicate MAC address and the number of MAC address moves detected within the detection window before duplicate MAC detection is triggered and the MAC address is suppressed. In addition, you can also configure an optional recovery time that the switches wait before the duplicate MAC address is automatically unsuppressed.

To configure duplicate MAC detection parameters, use the **detection-window**, **detection-threshold**, and **auto-recovery-time** statements at the [edit routing instance *routing-instance-name* protocols evpn duplicate-mac-detection] hierarchy level.

To clear duplicate MAC suppression manually, use the **clear evpn duplicate-mac-suppression** command.

[See [Overview of MAC Mobility](#).]

Interfaces and Chassis

- **Generic routing encapsulation (GRE) support (QFX10002-60C switches)**— Starting with Junos OS Release 18.1R1, you can use GRE tunneling services to encapsulate any network layer protocol over an IP network. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first adds a GRE header and then adds an outer IP header that is used to route the packet. When it receives the packet, a switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect noncontiguous networks and to provide options for networks that contain protocols with limited hop counts.

[See [Understanding Generic Routing Encapsulation](#).]

- **Support for private VLANs and support for IRB in P-VLAN (QFX5210 switches)**— Starting with Junos OS Release 18.1R1, QFX5210 switches support private VLANs. VLANs limit broadcasts to specified users. Private VLANs (P-VLANs) take this concept a step further by splitting the broadcast domain into multiple isolated broadcast subdomains and essentially putting secondary VLANs inside a primary VLAN. P-VLANs restrict traffic flows through their member switch ports (called “private ports”) so that these

ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port is usually connected to a router, firewall, server, or provider network. Each P-VLAN typically contains many private ports that communicate only with a single uplink, thereby preventing the ports from communicating with each other.

Just like regular VLANs, P-VLANs are isolated on Layer 2 and require that a Layer 3 device be used to route traffic among them. P-VLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts. Service providers use P-VLANs to keep their customers isolated from one another.

[See [Understanding Private VLANs](#).]

Also starting with Junos OS Release 18.1R1, you can configure an integrated routing and bridging (IRB) interface in a private VLAN (P-VLAN) so that devices within community VLANs and isolated VLANs can communicate with each other and with devices outside the P-VLAN at Layer 3 without requiring you to install a router.

[See [Example: Configuring a Private VLAN Spanning Multiple Switches with an IRB Interface](#).]

- **FEC support for 25-gigabit and 50-gigabit channel speeds (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, you can configure forward error correction (FEC) clauses CL74 and CL91 on QFX5210 switches. FEC CL91 can be configured on 100-gigabit interfaces and FEC CL74 can be configured on 25-gigabit and 50-gigabit interfaces. Because the FEC clauses are applied by default on these interfaces, you must disable the FEC clauses if you do not want to apply them.

- To disable the FEC mode:

```
[edit]
set interfaces interface-name gigether-options fec none
```

- To reenable the FEC mode:

```
[edit]
set interfaces interface-name gigether-options fec (fec74|fec91)
```

or

```
[edit]
delete interfaces interface-name gigether-options fec none
```

- To check FEC status:

```
show interfaces interface-name
```

The output for the show command will list FEC statistics for a particular *interface-name*, including the FEC corrected errors count, the FEC uncorrected errors count, and the type of FEC that was disabled or enabled.

[See [FEC](#).]

- **Resilient hashing support for equal cost multipath routes (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, resilient hashing is now supported by equal cost multipath (ECMP) sets.

NOTE: Resilient hashing is not supported on link aggregations groups (LAGs).

[See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk Groups](#).]

- **Multichassis link aggregation groups (MC-LAG) (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG is an MC-LAG client that has one or more physical links in a LAG. This client does not need to detect the MC-LAG. On the other side of the MC-LAG are two MC-LAG QFX10008 switches. Each of these switches has one or more physical links connected to a single client. The switches coordinate with each other to ensure that data traffic is forwarded properly.

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

- **Auto-channelization of interfaces (QFX5210 switch)**—Starting in Junos OS Release 18.1R1, you can use the auto-channelization feature to divide and channelize data automatically by detecting the cable type. The mode and number of channels are decided based on the channel link status. On QFX5210, auto-channelization supports three modes of operation with unique port settings:
 - When 4x10G split cables are connected, the 40G port auto-channelizes to four 10G channels.
 - When 2x50G split cables are connected, the 100G port auto-channelizes to two 50G channels.
 - When 4x25G split cables are connected, the 100G port auto-channelizes to four 25G channels.
- **Channelization support (QFX10002-60C switches)**—Starting with Junos OS Release 18.1R1, you can use channelization functionality to subdivide a larger flexible optical interface into sub-interfaces or channels. The QFX10002-60C switch has 12 ASIC circuits (PE) as a part of a Packet Forwarding Engine, and each PE switch has 5 ports (one standalone MAC port and 4 channelized MAC ports). The standalone MAC ports cannot be channelized. The QFX10002-60C switch allows you to channelize 48 ports out of available 60 ports.

By default, the ports come up in a mode that does not support channelization. If you channelize a port in a PE switch for the first time, it would result in FPC reboot. But if you channelize another port in the same PE switch, the FPC will not be rebooted. If you channelize a port in a different PE switch, the FPC will be rebooted.

To enable channelization on an interface:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port port-number channel-speed speed
```

[See [Channelizing Interfaces](#).]

- **Dynamic port swap from 40G to 100G without restarting the Packet Forwarding Engine (QFX5110 switches)**—Starting in Junos OS Release 18.1R1, you can configure different system modes to achieve varying levels of port density on QFX5110-32Q switches without restarting the Packet Forwarding Engine. The QFX5110-32Q switch has fixed 32 front panel network ports. Four 100G ports can either function as 32x40G or 20x40G – 4x100G. You can combine the port configurations supported into default mode or non-oversubscribed mode. The **dcufe** restart is triggered with the mode change.

[See [Configuring the System Mode](#).]

- **Support for 128k vmembers and 96k Address Resolution Protocol (ARP) and Neighbor Discovery (ND) entries when using enhanced convergence in multichassis link aggregation groups (MC-LAG) (QFX10000 switches)**—Starting with Junos OS Release 18.1R1, the number of vmembers has increased to 128k, and the number of ARP and ND entries has increased to 96k. This increased scale is supported only when you enable the **enhanced-convergence** statement. Enhanced convergence improves Layer 2 and Layer 3 convergence time during multichassis aggregated Ethernet (MC-AE) link failures and restoration scenarios.

If you have configured an IRB interface over an MC-AE interface that has enhanced convergence enabled, then you must configure enhanced convergence on the IRB interface as well. Enhanced convergence must be enabled for both Layer 2 and Layer 3 interfaces.

To configure enhanced convergence, enable the **enhanced-convergence** statement at the **[edit interfaces ae unit-number aggregated-ether-options mc-ae]** at the Junos OS CLI hierarchy.

To configure enhanced convergence on an IRB interface, enable the **enhanced-convergence** statement at the **[edit interfaces irb unit unit-number]** at the Junos OS CLI hierarchy.

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

- **Support for additional 10G data ports (QFX5210 switches)**—Starting in Junos OS Release 18.1R1, QFX5210 switches support two additional 10G data ports. You can use the two additional data ports as revenue ports.
- **FEC support for 100-gigabit port speeds (QFX10002, QFX10008, and QFX10016 Switches)**—Starting with Junos OS Release 18.1R1, you can configure forward error correction (FEC) clause CL91 on QFX10000 series switches. FEC CL91 can be configured on 100-gigabit interfaces. FEC CL91 clause is applied by default on these interfaces. If you do not want to apply the FEC CL91 clause, you can disable it.
 - To disable the FEC mode:

```
[edit]
set interfaces interface-name gigether-options fec none
```

- To reenable the FEC mode:

```
[edit]
set interfaces interface-name gigether-options fec (fec74|fec91)
```

or

```
[edit]
delete interfaces interface-name gigether-options fec none
```

- To check FEC status:

```
show interfaces interface-name
```

The output for the show command will list FEC statistics for a particular *interface-name*, including the FEC corrected errors count, the FEC uncorrected errors count, and the type of FEC that was disabled or enabled.

[See [FEC](#).]

- **Support for Protocol Independent Multicast (PIM) Dual Designated Router Mode (QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 18.1R1, you can enable PIM dual designated router mode for a pair of Multichassis Link Aggregation Group (MC-LAG) peers managing VLAN multicast traffic and Layer 3 multicast traffic over IRB interfaces.

PIM dual designated router mode sets up one device in a pair of MC-LAG peers as a primary designated router (DR), and the other device as a standby or backup DR for redundancy in managing multicast packet forwarding. Both devices join the multicast forwarding tree and receive multicast traffic. If the primary device fails, the standby device quickly takes over forwarding multicast packets with minimal traffic disruption.

- **Link Aggregation Control Protocol (LACP) force-up enhancements (QFX5210 switches)**—Starting in Junos OS Release 18.1R1, if an aggregated Ethernet interface (AE) on a switch has multiple member links and one member link in that AE is in the force-up state with its peer's LACP down, and then if LACP comes up partially—that is, if LACP is established with a non-force-up member link—force-up is disabled on the member link on which force-up has been set, and that member link is ready for connection establishment through LACP. Force-up is eligible only if the server-side interface has LACP issues.
- **Channelization support (QFX10002-60C switches)**—Starting with Junos OS Release 18.1R1, you can use channelization functionality to subdivide a larger flexible optical interface into sub-interfaces or channels. The QFX10002-60C switch has 12 ASIC circuits (PE) as a part of a Packet Forwarding Engine,

and each PE switch has 5 ports (one standalone MAC port and 4 channelized MAC ports). The standalone MAC ports cannot be channelized. The QFX10002-60C switch allows you to channelize 48 ports out of available 60 ports.

By default, the ports come up in a mode that does not support channelization. If you channelize a port in a PE switch for the first time, it would result in FPC reboot. But if you channelize another port in the same PE switch, the FPC will not be rebooted. If you channelize a port in a different PE switch, the FPC will be rebooted.

To enable channelization on an interface:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port port-number channel-speed speed
```

[See [Channelizing Interfaces](#).]

- **Channelizing Ethernet interfaces (QFX5200 switches)**—Starting with Junos OS Release 18.1R1, you can channelize the 100-Gigabit Ethernet interfaces to two independent 50-Gigabit Ethernet. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated 40-Gigabit Ethernet interfaces or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables.

There are a total of 54 physical ports on the QFX5200 switch. Ports 0 - 47 can be used as 25-Gigabit Ethernet interfaces. Ports 48 - 53 can be used as either 40-Gigabit Ethernet or 100-Gigabit Ethernet interfaces. You choose the speed by plugging in the appropriate transceiver. They can also be channelized to 10G, 40G, or 100G.

[See [Channelizing Interfaces on QFX Switches](#).]

- **Channelizing Ethernet Interfaces (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, you can channelize the 100-Gigabit Ethernet interfaces to two independent 50-Gigabit Ethernet or to four independent 25-Gigabit Ethernet interfaces. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated 40-Gigabit Ethernet interfaces or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables.

There are a total of 64 physical ports on the QFX5210 switch. Any port can be used as either 100-Gigabit Ethernet or 40-Gigabit Ethernet interfaces. You choose the speed by plugging in the appropriate transceiver. They can also be channelized to 50G, 25G or 10G.

[See [Channelizing Interfaces on QFX Switches](#).]

IPv4

- **Generic routing encapsulation (GRE) support (QFX5200 and QFX5210 switches)**—Starting in Junos OS Release 18.1R1, you can use GRE tunneling services to encapsulate any network layer protocol over an IP network. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first adds a GRE header and then

adds an outer IP header that is used to route the packet. When it receives the packet, a switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect noncontiguous networks and to provide options for networks that contain protocols with limited hop counts.

[See [Configuring Generic Routing Encapsulation Tunneling](#).]

- **Layer 2, Layer 3, multicast, IPv4, IPv6, and hierarchical ECMP support (QFX5210-64C switches)**—Starting in Junos OS Release 18.1R1, the feature set supporting the QFX5200 switch for Junos OS Release 17.3 DCB also supports the QFX5210-64C switch.

IPv6

- **Layer 2, Layer 3, multicast, IPv4, IPv6, and hierarchical ECMP support (QFX5210-64C switches)**—Starting in Junos OS Release 18.1R1, the feature set supporting the QFX5200 switch for Junos OS Release 17.3 DCB also supports the QFX5210-64C switch.

Junos OS XML API and Scripting

- **SLAX and Python scripts now can be sourced over the non-default VRF management instance (QFX Series)**—Starting in Junos OS Release 18.1R1, configuration of commit, event, JET, op, and SNMP scripts is upgraded to support the non-default management routing instance **mgmt_junos** as an option when specifying the source URL for refreshing or downloading SLAX and Python scripts.

[See [Using an Alternate Source Location for a Script](#) or [Configuring and Using a Master Source Location for a Script](#).]

Layer 2 Features

- **Layer 2 features (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, the following Layer 2 features are supported:
 - **VLAN support**
VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
 - **Link Layer Discovery Protocol (LLDP) support**
LLDP enables a switch to advertise its identity and capabilities on a LAN, as well as receive information about other network devices.
 - **Q-in-Q tunneling support**
This feature enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites.
 - **Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) support**
These protocols enable a switch to advertise its identity and capabilities on a LAN and receive information about other network devices.

[See [Ethernet Switching User Guide](#).]

- **Layer 2, Layer 3, multicast, IPv4, IPv6, and hierarchical ECMP support (QFX5210-64C switches)**—Starting in Junos OS Release 18.1R1, the feature set supporting the QFX5200 switch for Junos OS Release 17.3 DCB also supports the QFX5210-64C switch.

Layer 3 Features

- **Layer 2, Layer 3, multicast, IPv4, IPv6, and hierarchical ECMP support (QFX5210-64C switches)**—Starting in Junos OS Release 18.1R1, the feature set supporting the QFX5200 switch for Junos OS Release 17.3 DCB also supports the QFX5210-64C switch.

Management

- **Support for the Junos Telemetry Interface (QFX5100 switches)**—Starting with Junos OS Release 18.1R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various network elements without involving polling. On QFX5100 switches, only gRPC streaming of statistics is supported. UDP streaming is not supported.

The following sensors are supported:

- Chassis components
- Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol
- Network Discovery Protocol table state

For resource path names for these sensors, see *Guidelines for gRPC Sensors (Junos Telemetry Interface)*

To provision sensors to stream data through gRPC, create a subscription and specify parameters using the **telemetrySubscribe** RPC. You must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **ARP and NDP telemetry support for Junos Telemetry Interface (JTI) (QFX5110)**—Starting with Junos OS Release 18.1R1, you can export Address Resolution Protocol (ARP) and Neighbor Discovery Protocol (NDP) statistics through the Junos Telemetry Interface for QFX5110 switches. Sensor support for ARP and NDP statistics is at the same level of support as for QFX10000 and QFX5200 switches in Junos OS Release 17.2R1.

To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

To export telemetry data from Juniper equipment to an external collector, both Junos Telemetry Interface (JTI) and gRPC must be configured.

For resource names and OpenConfig paths for these sensors, see [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).

MPLS

- **Support for equal-cost multipath routing on MPLS label-switching routers (QFX5210 switches)**—Starting in Junos OS Release 18.1R1, you can configure equal cost multipath (ECMP) routing on MPLS label-switched routers (LSRs). ECMP is a Layer 3 mechanism for load-balancing traffic to a destination over multiple equal-cost next hops. When a link goes down, ECMP uses fast reroute protection to shift packet forwarding to use operational links, thereby decreasing packet loss.

This feature was previously supported in an "X" release of Junos OS.

[See [Understanding ECMP Flow-Based Forwarding](#).]

- **MPLS support (QFX5210 switches)**—Starting in Junos OS Release 18.1R1, MPLS is supported on the QFX5210 switch. MPLS provides both label edge routers (LER) and label switch routers (LSR) and provides the following capabilities:

- Support for both MPLS major protocols, LDP and RSVP
- IS-IS interior gateway protocol (IGP) traffic engineering
- Class of service (CoS)
- Object access method, including ping, traceroute, and Bidirectional Forwarding Detection (BFD)
- Fast reroute (FRR), a component of MPLS local protection. (Both one-to-one local protection and many-to-one local protection are supported.)
- Loop-free alternate (LFA)
- 6 PE devices
- Layer 3 VPNs for both IPv4 and IPv6
- LDP tunneling over RSVP

This feature was previously supported in an "X" release of Junos OS.

[See [MPLS Overview for Switches](#).]

Multicast

- **Multicast-only fast reroute (MoFRR) (QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 18.1R1, QFX10002, QFX10008, and QFX10016 switches support MoFRR, which minimizes multicast packet loss in PIM domains when there are link failures. With MoFRR enabled, the switch maintains both a primary and a backup multicast packet stream toward the multicast source, accepting traffic received on the primary path and dropping traffic received on the backup path. Upon primary path failure, the backup path becomes the primary path and quickly takes over forwarding the multicast traffic. If alternative paths are available, a new backup path is created. When enabling MoFRR, you can optionally configure a policy for the (S,G) entries to which MoFRR should apply; otherwise, MoFRR applies to all multicast (S,G) streams.

[See [Understanding Multicast-Only Fast Reroute on Switches](#).]

- **Layer 2, Layer 3, multicast, IPv4, IPv6, and hierarchical ECMP support (QFX5210-64C switches)**—Starting in Junos OS Release 18.1R1, the feature set supporting the QFX5200 switch for Junos OS Release 17.3 DCB also supports the QFX5210-64C switch.

Network Management and Monitoring

- **Support for sFlow, port mirroring, and port mirroring to an IP address (QFX5210 switches)**—Starting in Junos OS Release 18.1 R1 the QFX5210 switch supports sFlow technology. sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring randomly samples network packets and sends the samples to a monitoring station called a collector. You can configure sFlow monitoring on the switch to continuously monitor traffic at wire speed on all interfaces simultaneously. sFlow monitoring also collects samples of network packets, providing you with visibility into network traffic information. You configure sFlow monitoring at the **edit protocols sflow** hierarchy level. sFlow operational commands include **show sflow** and **clear sflow collector statistics**. This feature was previously supported in an "X" release of Junos OS.

[See [Understanding How to Use sFlow Technology for Network Monitoring on a Switch](#).]

Also starting in Junos OS Release 18.1R1, you can use port mirroring on QFX5210 switches to copy packets entering or exiting a port or entering a VLAN and send the copies to a local interface for local monitoring or to a VLAN for remote monitoring. 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. This feature was previously supported in an "X" release of Junos OS.

[See [Understanding Port Mirroring](#).]

Finally, also starting in Junos OS Release 18.1R1, you can send mirrored packets to an IP address over a Layer 3 network (for example, if there is no Layer 2 connectivity to the analyzer device). This feature also enables you to apply an IEEE-1588 timestamp to the mirrored packets. This feature was previously supported in an "X" release of Junos OS.

[See [Understanding Port Mirroring](#).]

Port Security

- **MACsec license enforcement (EX3400, EX4300, EX4600, EX9200, QFX5100 switches and Junos Fusion Enterprise)**—Starting in Junos OS Release 18.1R1, Media Access Control Security (MACsec) requires the installation of a MACsec feature license. If the MACsec license is not installed, MACsec functionality cannot be activated. You add the MACsec license using the **request system license add** command.

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

Routing Protocols

- **Support for BGP multipath at global level (QFX Series)**—Starting with Junos OS Release 18.1R1, BGP multipath is available at the global level in addition to the group and neighbor level. In earlier Junos OS releases BGP multipath is supported only at the group and neighbor levels. A new configuration option **disable** is available at the **[edit protocols bgp multipath]** hierarchy level to disable BGP multipath for specific groups or neighbors. This allows you to configure BGP multipath globally and disable it for specific groups according to your network requirements.

[See [disable](#).]

Security

- **Distributed denial-of-service (DDoS) protection (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, you can use DDoS protection to enable the switch to continue functioning while under a DDoS attack.

[See [Understanding Distributed Denial-of-Service Protection on QFX Series Switches.](#)]

- **Support for firewall filters (QFX5210)**—Starting in Junos OS Release 18.1R1, you can define firewall filters on the switch that defines whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces. You configure firewall filters at the `[edit firewall]` hierarchy level.

This feature was previously supported in an “X” release of Junos OS.

[See [Overview of Firewall Filters.](#)]

- **Storm control support (QFX5210 switches)**—Starting in Junos OS Release 18.1R1, you can monitor traffic levels and take a specified action when a defined traffic level (called the storm control level) is exceeded, preventing packets from proliferating and degrading service. You can configure the switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs.

This feature was previously supported in an “X” release of Junos OS.

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

- **Support for policers (QFX5210 switches)**—Starting in Junos OS Release 18.1R1, you can use policers to apply limits to traffic flow and to set consequences for packets that exceed those limits. A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.

This feature was previously supported in an “X” release of Junos OS.

[See [Overview of Policers.](#)]

Software Defined Networking (SDN)

- **Layer 2 VXLAN gateway (QFX5210 switches)**—Virtual Extensible LAN (VXLAN) is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. You can use VXLAN tunnels to enable migration of virtual machines between servers that exist in separate Layer 2 domains by tunneling the traffic through Layer 3 networks. This functionality allows you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains.

Starting with Junos OS Release 18.1R1, you can manually create VXLANs on QFX5210 switches instead of using a controller such as a VMware NSX for vSphere or Juniper Networks Contrail controller. If you use this approach, you must also configure Protocol Independent Multicast (PIM) on the VTEPs so that they can create VXLAN tunnels between themselves.

[See [Understanding VXLANs.](#)]

- **OVSDB-VXLAN support with VMware NSX for vSphere (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, the Open vSwitch Database (OVSDB) management protocol provides a means through which an NSX for vSphere controller can communicate with QFX5210 switches and provision them as Layer 2 Virtual Extensible LAN (VXLAN) gateways. In an environment in which NSX for vSphere 6.3.5 or later is deployed, an NSX for vSphere controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and vice versa.

[See [Understanding the OVSDB Protocol Running on Juniper Networks Devices.](#)]

- **OVSDB-VXLAN support with VMware NSX for vSphere (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 18.1R1, the Open vSwitch Database (OVSDB) management protocol provides a means through which an NSX for vSphere controller can communicate with QFX5110 and QFX5200 switches and provision them as Layer 2 Virtual Extensible LAN (VXLAN) gateways. In an environment in which NSX for vSphere 6.3.5 or later is deployed, an NSX for vSphere controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and vice versa.

[See [Understanding the OVSDB Protocol Running on Juniper Networks Devices.](#)]

Software Installation and Upgrade

- **ZTP support (QFX10002-60C switch)**—Starting with Junos OS Release 18.1R1, ZTP, automates the provisioning of the device configuration and software image with minimal manual intervention, and is supported on QFX10002-60C VM hosts. When you physically connect a supported device to the network and boot it with a factory configuration, the device attempts to upgrade the Junos OS software image automatically and autoinstall a configuration provided on the DHCP server.

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

Storage and Fibre Channel

- **Support for FIP snooping and DCBX (QFX5210)**—Starting in Junos OS Release 18.1R1, QFX5210 switches support FCoE Initialization Protocol (FIP) snooping and Data Center Bridging Capability Exchange protocol (DCBX), which are technologies that help enable transporting converged Ethernet traffic. FIP snooping filters prevent FCoE devices from gaining unauthorized access to a Fibre Channel (FC) storage device or another FCoE device. DCBX discovers the data center bridging (DCB) capabilities of connected peers, and advertises the capabilities of applications on interfaces by exchanging information in the form of application type, length, and value elements (TLVs).

[See [Storage User Guide](#) and [Traffic Management User Guide for the QFX Series and EX4600 Switches.](#)]

- **Support for Converged Enhanced Ethernet (CEE) features (QFX5210)**—Starting in Junos OS Release 18.1R1, QFX5210 switches support the following data center bridging (DCB) traffic management features for transporting CEE traffic:

- Priority-based flow control (PFC) for traffic prioritization and managing link bandwidth for lossless traffic
- Buffer space management to prevent dropped traffic with PFC
- Congestion notification for managing link bandwidth, including Explicit Congestion Notification (ECN) and Data Center Quantized Congestion Notification (DCQCN)
- Data Center Bridging Capabilities Exchange protocol (DCBX)

CEE enables traffic differentiation at the link layer and sharing of links for both Ethernet and FCoE traffic.

[See [Traffic Management User Guide for the QFX Series and EX4600 Switches](#).]

System Management

- **Integrated software feature licenses (QFX5210 switches)**—Starting with Junos OS Release 18.1R1, the standard QFX Series premium feature license for BGP, Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDb) software license and the standard QFX Series advanced feature license for BGP, Intermediate System-to-Intermediate System (IS-IS), MPLS, and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDb) license are supported.

[See [Software Features That Require Licenses on the QFX Series](#).]

- **Support for the Precision Time Protocol (PTP) G.8275.2 enhanced profile (QFX5110-48S-4C switches)**—Starting in Junos OS Release 18.1R1, you can enable the G.8275.2 enhanced profile to support telecom applications that require accurate phase and time synchronization for phase alignment and time of day synchronization over a wide area network. This profile supports PTP over IPv4 unicast, ordinary and boundary clocks, and unicast negotiation.

To configure the G.8275.2 enhanced profile, enable the **g.8275.2.enh** statement at the **[edit protocols ptp profile-type]** Junos OS CLI hierarchy.

[See [Understanding the PTP G.8275.2 Enhanced Profile \(Telecom Profile\)](#).]

- **Support for request vmhost and show vmhost commands (QFX10002-60C switches)**—Starting in Junos OS Release 18.1R1, many of the **request system** and **show system** commands have been replaced with **request vmhost** and **show vmhost** commands.

Here is a list of the vmhost commands that are now supported:

- request vmhost cleanup
- request vmhost file-copy
- request vmhost halt
- request vmhost hard-disk-test
- request vmhost power-off
- request vmhost power-on

- request vmhost reboot
- request vmhost snapshot
- request vmhost software add
- request vmhost software rollback
- request vmhost zeroize
- show vmhost bridge
- show vmhost crash
- show vmhost hard-disk-test
- show vmhost hardware
- show vmhost information
- show vmhost logs
- show vmhost management-if
- show vmhost netstat
- show vmhost processes
- show vmhost resource-usage
- show vmhost snapshot
- show vmhost status
- show vmhost uptime
- show vmhost version

[See [VM Host Operations and Management](#) for more information.]

SEE ALSO

[Changes in Behavior and Syntax](#) | 273

[Known Behavior](#) | 276

[Known Issues](#) | 282

[Resolved Issues](#) | 287

[Documentation Updates](#) | 299

[Migration, Upgrade, and Downgrade Instructions](#) | 300

[Product Compatibility](#) | 314

Changes in Behavior and Syntax

IN THIS SECTION

- [xref target has no title]

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.1R3 for the QFX Series.

Interfaces and Chassis

- **Modified output of show-ptp-clock command (QFX Series switches)**—Starting in Junos OS Release 18.1R1, the output of the **show-ptp-clock** command is modified to display the value of the **GMC Class** field as **248** for a PTP boundary clock when the lock state of the clock is **Acquiring**.

Management

- **Enhancement to LSP statistics sensor for Junos Telemetry Interface (MX Series, PTX Series, QFX10000 switches, and EX9200 switches)**—Starting with Junos OS 18.1R1, the telemetry data exported for the LSP statistics sensor no longer includes the phrase **and source 0.0.0.0** after the LSP name in the value string for the prefix key. This change reduces the payload size of data exported. The following is an example of the new format:

```
str_value: /mpls/lsp/constrained-path/tunnels/tunnel[name='LSP-4-3']/state/
counters[name='c-27810']/
```

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (QFX5110, QFX5200, and QFX10000 switches)**—Starting with Junos OS Release 18.1R1, the format of telemetry data exported through gRPC for NPU memory and memory utilization implements prefix compression. This change reduces the payload size of data exported. The following example shows the new format:

```
key: __prefix__
str_value: /components/component[name='FPC0:NPU0']/properties/property
key: [name='mem-util-edmem-size']/value
uint_value: 12345
```

Telemetry data is exported in key-value pairs. Previously, the data exported included the component and property names in a single key string.

[See [Guidelines for gRPC Sensors](#).]

Network Management and Monitoring

- **SNMP syslog messages changed (QFX Series)**—In Junos OS Release 18.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD—AgentX master agent failed to respond to ping. Attempting to re-register
NEW—AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD—NET-SNMP version %s AgentX subagent connected
NEW—NET-SNMP version %s AgentX subagent Open-Sent!

[See the [MIB Explorer](#).]

Network Operations and Troubleshooting Automation

- **JET - Correction to escaped characters notification events** (QFX Series data center switches)—Per RFC7159, certain characters must be escaped. Data returned from JET notification subscriptions contained escaped characters that were not required. This has been corrected to comply with RFC7159.
- **respawn-on-normal-exit** option added to [edit system extensions extension-service application file <application-name>] hierarchy (QFX Series Data Center Switches)—This option helps to ensure that daemonized Juniper Extension Toolkit (JET) applications that exit normally will restart without user intervention. Daemonized JET applications that exit unexpectedly will still restart without user intervention. This is the default behavior.

Routing Policy and Firewall Filters

- **Support for configuring the GTP-TEID field for GTP traffic (QFX5000 line of switches)**—Starting in Junos OS Release 17.3R3 and 18.1R2, the **gtp-tunnel-endpoint-identifier** statement is supported to configure the hash calculation of IPv4 or IPv6 packets that are included in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations. The **gtp-tunnel-endpoint-identifier** configuration statement is configured at the **[edit forwarding-options enhanced-hash-key family inet]** hierarchy level.

In most of the cases, configuring **gtp-tunnel-endpoint-identifier** statement is sufficient for enabling GTP hashing. After enabling, if GTP hashing does not work, it is recommended to capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In this cases, use **gtp-header-offset** statement to set a proper offset value. Once the header offset value is resolved, run **gtp-tunnel-endpoint-identifier** command for enabling GTP hashing successfully.

[See [gtp-tunnel-endpoint-identifier](#) and [gtp-header-offset](#).]

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.

SEE ALSO

[New and Changed Features | 248](#)

[Known Behavior | 276](#)

[Known Issues | 282](#)

[Resolved Issues | 287](#)

[Documentation Updates | 299](#)

[Migration, Upgrade, and Downgrade Instructions | 300](#)

[Product Compatibility | 314](#)

Known Behavior

IN THIS SECTION

- [EVPN | 276](#)
- [Interfaces and Chassis | 276](#)
- [Layer 2 Features | 277](#)
- [Multicast | 277](#)
- [Platform and Infrastructure | 277](#)
- [Routing Protocols | 280](#)
- [Storage and Fibre Channel | 281](#)
- [Virtual Chassis | 281](#)
- [Services Applications | 281](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 for the QFX Series.

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

EVPN

- On QFX10000 switches configured as type-5 route peers, when only peer 1 advertises routes, that peer might not install the de-encapsulated next-hop (NH) route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any type-5 route. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. [PR1191092](#)
- EVPN/VXLAN implementations support up to 100 EVPN VLAN-based routing instances. Above 100 instances, MAC learning might behave incorrectly. [PR1287644](#)

Interfaces and Chassis

- When you commit a configuration change for IRB from VRRP to non-VRRP and the IRB address also changed to VRRP VIP. Junos OS loses the direct route from the IRB. This is a limitation. This issue was also logged in [PR1191371](#). [PR1319124](#)
- Because the link speed command cannot be hidden, configuring or committing the same should result in the intended functionality. Otherwise MC-LAG peer states will get impacted. [PR1329030](#)

- Forcing the LAG/MC-LAG feature up is not supported on the QFX10000 platform. [PR1332475](#)
- Supported ARP scale is 48,000 over MC-LAG interfaces [PR1334321](#)

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging has been enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. [PR1075230](#)
- In EVPN-VXLAN deployment with QFX1000 switches, when VXLAN enabled IRB interface is configured in the same routing instance as that of the underlay VTEP tunnel and if the remote VTEP interface IP is resolved over the IRB interface using routing protocols or static route, dc-pfe cores would be generated and all the interfaces would go down. The dc-pfe core files would be continuously generated until the configuration is corrected. [PR1261824](#)
- On QFX10016, after deleting and re-adding 1000 LAG interfaces, traffic drops could be seen until ARP is refreshed even though all LAG interfaces come up. [PR1289546](#)
- LAG-based resilient hashing is not supported on QFX5200 and QFX5210 switches. ECMP-based resilient hashing is supported on those switches. [PR1321505](#)
- QFX5210-64C: Resilient hashing is not supported for LAG interfaces. [PR1325499](#)
- Packet statistics are not supported for logical child members of aggregated Ethernet (AE) interface. [PR1335454](#)
- Supported global Vmember scale is 64000 when created over AE interfaces [PR1337569](#)

Multicast

- To use IGMP snooping on QFX5110 switches in an EVPN-VXLAN multihoming environment, you must enable IGMP snooping on all VLANs associated with any configured VXLANs. You cannot selectively enable IGMP snooping only on those VLANs that might have interested listeners, because all the VXLANs share virtual tunnel endpoints (VTEPs) between the same multihoming peers and must have the same settings. [PR1407557](#)

Platform and Infrastructure

- While scaling beyond 2000 VLAN/IRBs , L3 multicast traffic does not converge to 100 percentage and continuous drops are observed after bringing down /up the downstream interface up or down or while an FPC comes online after FPC restart. [PR1161485](#)
- On the QFX10000-12C-DWDM coherent line card, it is possible that sometimes the link flaps when MACsec is enabled on Ethernet interfaces. [PR1253703](#)
- ERPS convergence takes time after GRES switchover and hence traffic loss is observed for a brief period. [PR1290161](#)

- On QFX Series, the logical interface (IFD) and the physical interface (IFL) go down when traffic exceeds the rate-limit. Storm control is supported only on interfaces configured in family Ethernet-switching. Moreover, in this family, we support only one IFL per IFD. Due to this, bringing down the IFD is acceptable. Flexible VLAN tagging is not supported on the interfaces enabled for storm control. [PR1295523](#)
- On QFX10000 line platforms, with a high scale of 4000 VNIs or 200K MACs, or both, if a large configuration change happens with traffic flowing, then forwarding descriptor memory corruption might occur, leading to complete traffic loss on certain ports. The qualification shows that a system with 400 VNIs has been stable. However, other configurations like global MAC count and underlying MPLS LSPs can increase system load. [PR1296089](#)
- Em1 does not show correct speed when its other end is connected to 10m/100m ports. [PR1303902](#)
- One main requirement with CoS on the FC interface is that the FC interface should be brought down before applying any CoS configuration. Thus you need to bring down the interface, apply the CoS configuration, and bring up the interface. This is required due to HW (BCM) limitation. [PR1320425](#)
- IRBs interface on VXLAN that has IGMP snooping configured on that VXLAN are currently supported. If IRB is configured, then a dcd restart could lead to multicast traffic loss. [PR1322057](#)
- This issue occurs in an MH EVPN-VXLAN scaled scenario, with IGMP snooping configured: 1) For 10000 s,g scale : the trigger is to disable DF link for convergence. Total convergence for 10000 s,g scale is 4.5 secs with traffic rate of 60kpps Per flow convergence loss ranges from 3.16 secs to 5.66 secs. 2) For 8000 s,g scale, the trigger is to disable DF link for convergence: Total convergence for 8000 s,g scale is 2.86 secs with traffic rate of 60 kpps. Per flow convergence loss ranges from 1.86 secs to 3.73 secs. [PR1323155](#)
- When you perform PIC offline followed by online on the rombauer QIC module, the entire FPC that houses the Rombauer PIC will reboot. [PR1324362](#)
- 100G DAC/Copper cable is connected between QFX5210-64C and QFX10000 devices, links might not come up reliably. The rest of the 100G Optics/AOC, 40G Optics/DAC/Copper work well when connected between QFX5210-64C and QFX10000 devices. [PR1324600](#)
- Configuration of **mac-table-size** under **vlan switch-options** is not supported for QFX10002-60C. [PR1325315](#)
- QFX5210-64C : Irrespective of the physical interface speed, the speed displayed for Gr-interface is always 800 Mbps. [PR1325695](#)
- The **mac-learning-limit** option is not supported under VLAN switch-options for QFX10002-60C platform [PR1325752](#)
- The Broadcom chip has VLAN-based logical interface (IFL) statistics. Because for a given IFL both IPv4 and IPv6 use the same VLAN, statistics will count both IPv4 and IPv6 together. 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](#)

- Need to increase **global-mac-table-aging-time** and **global-mac-ip-table-aging-time** settings on Junos Fusion Provider Edge ADs: **set protocols l2-learning global-mac-table-aging-time 900 set protocols l2-learning global-mac-ip-table-aging-time 720**. [PR1328929](#)
- Configuring an IRB physical interface (IFD) static MAC address will not take effect. Only in logical interface (IFL) level static configuration works. [PR1329032](#)
- Because the scaling numbers for flex counters in Broadcom is less than the number of maximum multicast routes that can be installed in hardware and also the flex counters are shared among different entities like VFI, VRF, VFP, L3IIF, SOURCE_VP, MPLS_ENTRY, VLAN_XLATE, PORT_TABLE, L3_ENTRY_IPV4_MULTICAST, L3_ENTRY_IPV6_MULTICAST, L3_DEFIP, creation of counter will fail after the scale limit(70,000). [PR1330473](#)
- The use of flexible-vlan-tagging with two VLAN tags is not supported on Layer 3 logical interfaces on QFX5110-48S and QFX5200 switches. [PR1330510](#)
- All the UFT profiles except l3-profile while doing the multicast s,g entries scale test noticed the PFE mcast table occupancy is not upto 95%. This is a product limitation, Broadcom informed that they cannot do much about optimizing table utilization for all group range. [PR1332170](#)
- Error messages related to rt_pfe_veto might be seen when a large number of routes are learned and downloaded to FIB. It indicates slowness of the Packet Forwarding Engine to install the routes in HW and will not have any functionality impact [PR1333553](#)
- A few error messages related to function rt_mesh_group_add_check() will be seen during reboot and are harmless. [PR1335363](#)
- Analyzer is not supported on QFX10002-60C. [PR1335970](#)
- Inline and distributed BFD is not supported for IRB interfaces. Configure BFD timers according to guidelines for centralized mode. This problem is more pronounced in IS-IS because it needs more packets (L1 and L2) to maintain the sessions. [PR1339127](#)
- On QFX5110-48S, PTP delay-req packets might be generated at less than 128 PPS when the delay-request interval is configured as -7. [PR1339775](#)
- On QFX5000 platforms multihop BFD sessions might flap after a disruptive trigger in topology with aggressive BFD timeout < 1s. Examples of disruptive triggers: (a) restart routing and (b) reboot of router. [PR1340469](#)
- In an IPCLOS topology, when a spine/leaf is rebooted, you may see around 100 secs of traffic loss. The reason for this is that, Junos will start advertising routes before PFE route programming is completed, which can cause traffic loss. This is mainly a design trade off. [PR1341398](#)
- In a scaled VRRP scenario with 1000 groups, it takes around 17 seconds for all traffic to converge onto the backup node. [PR1341811](#)
- On switching platforms the LACP AE minimum-link with sync-reset enabled feature is not supported on an aggregate interface where MicroBFD is enabled. [PR1342657](#)

- On upgrading QFX10002 from Junos 15.1X53-D66 to Release 18.1R1 release, some of the 100G ports are not created. [PR1343970](#)
- When a **request system reboot now** is triggered it is observed it takes 10 seconds for the interfaces to go down. This issue is not observed in 18.2 images. [PR1344831](#)
- When you deactivate or activate IRB with VRRP configuration in a scaled setup with 1000 VRRP groups, convergence time will be around 10 to 30 seconds. [PR1345272](#)
- On any platform that does not clear out /mfs when installing a new software release such as EX and QFX Series, when upgrading from certain releases to Junos OS Release 18.1R1 the statistics daemon PFED might generate a core file. This issue does not impact service. [PR1346925](#)
- QFX-60C: Scheduler slip of sflowd daemon "sflowd[24814]: JTASK_SCHED_SLIP" observed whenever sflow configured 40g interface got channelized to 4x10g interface or non-channelized from 10g interface to 40g interface or Devices reboot or Whenever FPC Restart. [PR1358045](#)
- Accton AS7816-64X systems are shipping with 14 characters but Junos limitation is 12 characters. Accton serial number contains 781664X as first 7 characters and 78 should be added from **show chassis hardware** output when serial number is required. [PR1371126](#)

Routing Protocols

- Configuring link aggregation group (LAG) hashing with the **[edit forwarding-options enhanced-hash-key] inet vlan-id** statement uses the VLAN ID in the hashing algorithm calculation. On some switching platforms, when this option is configured for a LAG that spans FPCs, such as in a Virtual Chassis or Virtual Chassis Fabric (VCF), packets are dropped due to an issue with using an incorrect VLAN ID in the hashing algorithm. As a result, the vlan-id hashing option is not supported in a Virtual Chassis or VCF containing any of the following switches as members: EX4300, EX4600, QFX5100, or QFX5110. Under these conditions, use any of the other supported enhanced-hash-key hashing configuration options instead. [PR1293920](#)
- The route unidimensional limit is 1.6 million routes in Junos OS Release 18.1R1. [PR1320865](#)
- If you configure GRE tunneling with the underlying ECMP next-hop instead of a unicast next hop, traffic might be dropped. This scenario is not supported. [PR1332309](#)

Storage and Fibre Channel

- If the configuration changes or any aggregation devices (AD) restart, you might see inconsistency in the output of **show ethernet-switching table** and **show fip snooping satellite** on different ADs for some time. It takes time for the ADs to completely restart and hence MAC addresses might be learned over EVPN (DRP flag). When AD restart is complete, MAC addresses should be learned locally and hence the DRP flag moves to the S flag. It can take up to 10 minutes to get consistent output for **show** commands. The output for **show ethernet-switching table** on all ADs will show all the MAC addresses. However, the flags against the MAC addresses might be different on the ADs because the MAC addresses might be learned statically on some ADs and dynamically on others. The flag against the dynamic MAC addresses will be changed from D to S once those MAC addresses are relayed from the satellite device (SD) to the AD, which can take up to 10 minutes. However, there should not be any traffic drop. Traffic drop is expected only initially, when the AD has just been restarted. [PR1304173](#)

Virtual Chassis

- VC internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur and it is considered to be known behavior. Release note reference: https://www.juniper.net/documentation/en_US/junos/information-products/topic-collections/release-notes/17.2/topic-118735.html [PR1347902](#)

Services Applications

- You cannot configure analyzers on QFX10002-60C switches. The CLI configuration command **set forwarding-options analyzer** and the CLI operational command **show forwarding-options analyzer** are not supported on the switch. [PR1340607](#)

SEE ALSO

[New and Changed Features | 248](#)

[Changes in Behavior and Syntax | 273](#)

[Known Issues | 282](#)

[Resolved Issues | 287](#)

[Documentation Updates | 299](#)

[Migration, Upgrade, and Downgrade Instructions | 300](#)

[Product Compatibility | 314](#)

Known Issues

IN THIS SECTION

- [EVPN | 282](#)
- [Interfaces and Chassis | 283](#)
- [Layer 2 Features | 283](#)
- [MPLS | 283](#)
- [Platform and Infrastructure | 284](#)
- [Routing Protocols | 286](#)

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 18.1R3.

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

EVPN

- In a scaled setup, if **mac-move** is triggered more than four times, the detection might not be reliable. [PR1284315](#)
- Chained-composite-next-hop (CNH) is a must for EVPN pure type 5 with VXLAN encapsulation. Without this, the Packet Forwarding Engine will not program the tunnel NH. You have to explicitly set it on QFX5110. **set routing-options forwarding-table chained-composite-next-hop ingress evpn**. On the QFX10000, it is applied as part of default configuration. **user@router> show configuration routing-options forwarding-table | display inheritance defaults**. [PR1303246](#)
- In an EVPN collapsed L2/L3 multihomed GWs topology, when traffic is sent from IP fabric toward EVPN, some traffic loss is seen. If the number of hosts behind EVPN gateways is increased, the traffic loss becomes higher. [PR1311773](#)
- On QFX5000 and QFX10000 platforms, VTEP's MAC address is not learned in the Ethernet switching table though it is present in the EVPN database. [PR1371995](#)

Interfaces and Chassis

- A difference in error message reporting is seen while trying to configure 100G and 40G in a LAG. The QFX10002-72Q error message is more meaningful than the QFX10002-60C error message. [PR1340974](#)

Layer 2 Features

- On the QFX5210, there are issues with the latency test. The Issues with the latency tests: 10G latency values of cut-through are higher than store and forward, in the 40G latency test for the frame size 1280, higher latency value are seen. [PR1343579](#)
- On random initialization of QFX5100, the programming of the storm control profile is missed within hardware on random interfaces. This is not visible over CLI and the configuration still shows as intact. This happens as a result of interface speed not properly getting detected within the hardware. [PR1354889](#)
- No error/warning shown during this configuration commit. [PR1359982](#)
- In case of the access side interfaces as SP style interfaces, when a new IFL is added and if there is already an IFL on the IFD, there is a 20-50 msec traffic drop on existing IFL. [PR1367488](#)

MPLS

- There could be some lingering RSVP state that would keep some labeled routes programmed in the Packet Forwarding Engine longer than they should be. This RSVP state will eventually expire and then delete the RSVP MPLS routes from FIB. However, traffic loss is not anticipated due to this lingering state or the corresponding label routes in the FIB. In the worst case, in a network where there is persistent link flapping going on, this lingering state could interfere with the LSP scale being achieved. [PR1331976](#)
- The traffic loss was more than 50 ms while performing FRR. The traffic loss was well within 50 ms during FRR. However, the ingress nodes re-signals tunnels on detection of primary path failure detection and switches traffic to new tunnels. This occurs when transit LSR is not fully completed with the tunnel installation. Hence, more drop is observed during the overall FRR event. [PR1345843](#)
- Statistics of transit traffic do not increment LSP statistics signaled by RSVP-TE. [PR1362936](#)
- The issue occurs when on optimize timer expiry, the traffic engineering database version number match indicates a CSPF has already run for the path, if an optimization has not yet been done with that version, it will be run despite the version number match. (This occurs due to per-path optimize-seq-no that is updated with a traffic engineering database seq no only on optimization.) When the path is disabled to avoid invalid ERO, making sure this does not interfere with global repair/local reversion. [PR1365653](#)

Platform and Infrastructure

- When per-packet load balancing is removed or deleted, next hop index might change. [PR1198092](#)
- Single-bit and multiple-bit ECC errors are not logged on QFX5110 switches. [PR1251917](#)
- On QFX10000 series switches, at initialization, the port group module comes up after some time and negative ACKs are seen until the port group module is up. Once the port group module is up, negative ACKs are no longer observed. This is an expected behavior due to an Aggressive Link Scan feature introduced in Junos OS Release 17.2. [PR1271579](#)
- On QFX5110 Series switches, Digital optical monitoring (DOM) status via CLI is not correct for Junos OS Releases 15.1X53 through 17.x. The light level statistics can be seen in the FPC shell level. There is no traffic impact. [PR1305506](#)
- Traffic drop occurs on sending traffic over "et" interfaces due to CRC errors. [PR1313977](#)
- Family Ethernet-switching cannot be used when **flexible-vlan-tagging** is configured. It is unsupported. The behavior is nondeterministic with this configuration and there is a possibility of seeing a dcpfe core file. [PR1316236](#)
- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED stays unlit. [PR1317750](#)
- There might be a traffic loss on the ingress PE device if the EVPN MPLS is configured later on remote PE device or from the working condition EVPN MPLS is disabled and enabled later. [PR1319770](#)
- On a QFX10016, permanent traffic loss is seen for some hosts after the initial ARP timer expiry caused by an ARP entry is not synchronized between the two PE devices. [PR1322288](#)
- On the QFX10016 EVPN-VXLAN scaled testbed, it takes up to 3 minutes for traffic to converge when configuration related to a tenant (5 IRBs/VLAN) is added. [PR1323042](#)
- Port 0 of Qfx5100-48t does not come up in mixed VCF. As a workaround, use the **phy diag xe0 dsc** command from the BCM shell on reboot. This brings up the port, which stays up continuously until the next reboot. [PR1323323](#)
- QFX5210: No prune to RP was sent from LHR after shifting to the GR Interface, when the RP is in transit node (multicast over GRE tunnel scenario). [PR1323620](#)
- Traffic statistics for multicast stream on GR interfaces do not work on QFX5000 line platform. [PR1323622](#)
- Interface uptime has increased by 8 seconds from Junos OS Release 17.4R1 to Release 18.1R1. Note that SDK upgrades across releases impacts parameters such as login prompt appear time, FPC up time, and interface up time after switch reboot. [PR1324374](#)
- Persistent MAC is not enabled. [PR1325313](#)
- QFX10002-60C filter operation with log action is not supported for protocols other than L2/IPv4/v6 and the following message **Protocol 0 not recognized** is seen in firewall logs. [PR1325437](#)

- The management process (mgd) might panic after modifying AE interface members under **ethernet-switching vlan** stanza. After mgd panic, your remote session is terminated as a result. [PR1325736](#)
- In a streaming telemetry scenario, if performing **commit full**, na-grpd daemon might restart, causing disconnection of streaming telemetry. [PR1326366](#)
- Analyzer is not supported in QFX10002-60C. [PR1327288](#)
- On QFX5100 series platforms, in some cases, class of Service (CoS) configuration is not properly applied in Packet Forwarding Engine (PFE), leading to unexpected egress traffic drop on some interfaces. [PR1329141](#)
- In an EVPN-VXLAN scenario, ARP table information is not synchronized on two spines after reconfiguring an end host on a multihomed CE interface from IP1/MAC1 to IP1/MAC2. [PR1330663](#)
- On QFX52xx standalone devices with Vxlan configured, user configured Ingress ACL scale limit is 256 terms. [PR1331730](#)
- BFD session over AE flaps when member link carrying the BFD Tx flaps. [PR1333307](#)
- Changing MTU for GRE and underlying interfaces in single commit will be a caveat for the RLI Xellent: QFX: PFE: IP GRE (RLI NO: 34078). Refrain from committing MTU changes for GRE and underlying interfaces in single commit. For any GRE interface MTU update follow the mentioned workaround. [PR1335739](#)
- QFX 5200 ISSU with GR only support BGP, No OSPF support. [PR1336442](#)
- Changing MTU for GRE and underlying interfaces in single commit requires a caveat for the IPv4 GRE feature. Refrain from committing MTU changes for GRE and underlying interfaces in a single commit. For any GRE interface MTU update follow the mentioned workaround. [PR1339601](#)
- With Junos OS Release 18.1R1 image, when QFX5000 and 10000 boxes are upgraded through ZTP, the configuration commit might fail if the configuration is fetched through a python script. [PR1349240](#)
- When ZTP script fails to copy the ZTP configuration file from the DHCP server to the current directory location because of the read-only file system, then you need to specify the destination path to download that has read-write permission and has sufficient space to download. [PR1354197](#)
- On QFX10002, QFX10008, and QFX1016 spine nodes, the Virtual Extensible LAN (VXLAN) traffic might be lost if the VLAN tagged underlay traffic is received on Ethernet VPN (EVPN) type 2 and needs to be routed on to EVPN type5 tunnel. [PR1355773](#)
- On QFX5110, the FEC for 100G optics is not being displayed when the expected behavior is for FEC to be shown as NONE. On QFX10002, the FEC for 40g optics is being displayed as NONE when expected behavior is for FEC not to be displayed. On QFX10008, the FEC for 40G optics is being displayed as NONE when the expected behavior is for FEC not to be displayed. [PR1360948](#)
- When MC-LAG is configured with force-up enabled on MC-LAG Nodes, the LACP admin key should not match with the access/CE device. [PR1362346](#)
- On QFX10000 platform with IRB enabled, traffic might not be forwarded on some of the child members when the member link of the AE is added or deleted. [PR1362653](#)

- QFX52100: Filter with routing-instance applied to family inet logical interface (IFL) causes traffic to be discarded on unrelated interfaces. [PR1364020](#)
- pm4x25_line_side_phymod_interfa ERROR: u=0 p=81 interface type 16 not supported by internal SERDES for this speed 50000. This error messages is seen when channelization is detected in the build Junos OS Release 18.1R3. [PR1366137](#)
- The issue is observed if both local and remote end are Auto-Channelised and the local port QSFP is removed. [PR1370887](#)
- 100G DAC is not used by customer. [PR1373028](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- When one 50g port is made down with ifconfig command, other one also goes down in Junos OS Release 18.1R3. [PR1376389](#)
- LOC and Diag System LEDs on the front panel are not defined yet. [PR1380459](#)
- ifOutMulticastPkts , ifInBroadcastPkts and ifOutBroadcastPkts shows incorrect value in Junos OS Release 18.1R3 build in AS7816-64X. [PR1384069](#)
- BGP session bounce might sometimes prevent BUM traffic from being flooded to all remote VTEPs. [PR1373093](#)
- On QFX5000 platforms, there is a matching chassis:fpc:pic:port between the sxe interface and the et interface. If dcd process restarts, the BGP session might flap due to aggregated Ethernet interface flap after the physical interface in it is detached or attached. Any other configuration change operation resulting in sending SIGHUP to dcd would cause the aggregated Ethernet interface to flap. [PR1373188](#)
- The Junos license (JUNOS-FP-C2) is not getting installed on Junos white boxes. [PR1383274](#)

Routing Protocols

- On EX4600 and QFX5100 switches with Q-in-Q, if the native VLAN is configured on a Q-in-Q interface connected to a customer device, the packets going out with the native VLAN ID (customer-VLAN) are still tagged. [PR1105247](#)
- On QFX10000 line platforms, during route next-hop churn or earliest deadline first (EDF) job priority changes, memory corruption might occur, leading to processing issues and constant packet drop. [PR1243724](#)
- For the QFX10002 and QFX10008 switches, you might observe an increase in the convergence time of OSPF routes when compared to Junos OS 17.3 releases. An average increase of 1.5 seconds is seen for 100,000 OSPFv3 routes. [PR1297541](#)
- Performing GRES on the EVPN-VXLAN topology with uRPF results in total packet loss. [PR1322217](#)
- In the PVLAN configuration, the isolated VLAN and Community VLAN should not use the same VLAN Id. [PR1323520](#)

- VLAN range shown in community VLAN is 1..4094. Hence, VLAN 0 should not be configured as community VLAN in PVLAN. [PR1323719](#)
- When MoFRR is enabled, traffic statistics on the multicast route show double the outgoing traffic. Accounting is done for both the primary and backup route, hence the issue. When one of the upstream interfaces goes down, this issue will not be seen. [PR1326338](#)
- Higher convergence time for LFA with BFD occurs in Junos OS Release 18.1. [PR1337412](#)
- On QFX5210, when ICCP/ICL link is disabled/enabled, data-driven ARP learning is taking 2-3 seconds longer than on QFX5200-32C, leading to ~10 seconds of IPv4 and IPv6 traffic loss. [PR1338444](#)
- If permanent traffic loop is created in IPCLOS topology, PFE CPU utilization can go high which can result in ping drops [PR1341107](#)
- On a scaled setup, when the host table is full and the host entries are installed in LPM table, OSPF sessions might take more time to come up. [PR1358289](#)
- Disabling a LAG member that is part of an L3 IRB interface sometime causes traffic loss. [PR1359841](#)
- L3-GW is not supported on QFX5110 with SP style of configuration in Junos OS Release 17.3R3. [PR1363708](#)

SEE ALSO

New and Changed Features	 248
Changes in Behavior and Syntax	 273
Known Behavior	 276
Resolved Issues	 287
Documentation Updates	 299
Migration, Upgrade, and Downgrade Instructions	 300
Product Compatibility	 314

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 18.1R3](#) | [288](#)
- [Resolved Issues: 18.1R2](#) | [291](#)
- [Resolved Issues: 18.1R1](#) | [294](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

Class of Service (CoS)

- DST IP 224/4 match condition is programmed in HW as 224/24 in loopback FF entry rep=0. [PR1354377](#)

EVPN

- EVPN-VXLAN QFX10000: `jprds_dlu_alpha_add : 222 JPRDS_DLU_ALPHA KHT` addition failed. [PR1258933](#)
- When a VLAN uses an IRB interface as the routing interface, the `vlan-id` parameter must be set to `none` to ensure proper routing. This issue is platform independent. [PR1287557](#)
- Rpd has unreproducible core file with scaling EVPN-VXLAN configuration on QFX10000 platform. [PR1339979](#)
- On a scaled EVPN-VXLAN setup, loading the scaled configuration and the base configuration alternately for a few times can result in losing adjacency and hence the protocols will be down. [PR1349659](#)
- Traffic might be lost on Layer2 and Layer3 spine nodes in multihomed EVPN scenario. [PR1355165](#)
- The QFX10000 might drop transited traffic coming from an MPLS network to EVPN-VXLAN. [PR1360159](#)
- Increased risk of routing crash with temporary impact on traffic on QFX10000 or QFX5100 nodes with certain configuration changes or clearing L2 or L3 learning information in a high-scale EVPN-VXLAN configuration environment. [PR1365257](#)
- Ospf sessions are not coming up between MX and QFX10000 as ARP entries get deleted and added. [PR1366860](#)
- Proxy ARP might not work as expected in an EVPN environment. [PR1368911](#)
- QFX10000 / Import default ipv6 route to VRF causes infinite entries to get created in 'evpn ip-prefix-database' and become unstable. [PR1369166](#)

Infrastructure

- QFX5100: Enabling **mac-move-limit** stops ping on **flexible-vlan-tagging** enabled interface. [PR1357742](#)

Interfaces and Chassis

- MC-LAG peer does not send ARP request to the host. [PR1360216](#)

Layer 2 Features

- LACP packets are getting dropped with **native-vlan-id** configured after reboot. [PR1361054](#)
- The dcpfe/fxpc process might crash on Packet Forwarding Engines with low memory when allocating a huge memory. [PR1362332](#)
- QFX5000 Virtual Chassis acting as EVPN-VxLAN ARP proxy might cause ARP resolution to fail. [PR1365699](#)
- Hashing is not working for IPV6 packet encapsulated in VXLAN scenario [PR1368258](#)
- When native-vlan-id is configured for AE LACP session to multihomed server goes down [PR1369424](#)
- A port might still work even if it is deleted from an AE interface. [PR1372577](#)

MPLS

- LSP is not received by QFX5110. [PR1351055](#)
- **NO-propagate-TTL** acts on MPLS swap operation. [PR1366804](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition. [PR1374102](#)

Platform and Infrastructure

- The etherStatsCRCAlignErrors counters might disappear in the SNMP tree. [PR1329713](#)
- AI-script does not **auto re-install** unless it is manually done after a Junos upgrade. [PR1337028](#)
- The DF of an EVPN instance might flood all the ARP request back to the Ethernet segment. [PR1337275](#)
- On QFX5100 platforms, LR4 QSFP can take up to 15 minutes to come up after Virtual Chassis reboot. [PR1337340](#)
- On QFX10000 platforms, VRRP function does not work well when it is configured on subinterfaces. [PR1338256](#)
- On QFX5100, QFX5200, QFX5110, and EX4600 platforms, BPDU packets might get dropped and **bpdu-block-on-edge** might not work. [PR1343330](#)
- QFX5100: Fan RPM fluctuates when temperature sensor reaches its threshold. [PR1345181](#)
- Backup Routing Engine might experience a crash, causing vmcore to be generated on master Routing Engine, master Routing Engine performance will not be affected. [PR1346218](#)
- On QFX10000 platforms, syslog error messages might be seen in syslog after configuring multiple LAG interfaces under sFlow protocol. [PR1346493](#)

- QFX5100-48T 10G interface might be auto-negotiated at 100M speed instead of 10G. [PR1347144](#)
- Traffic with destination MAC matching the virtual gateway MAC might be silently dropped or discarded. [PR1348659](#)
- The BGP session might flap after changing the extended-vni-list under the EVPN hierarchy. [PR1349600](#)
- After upgrading to Junos OS Release 17.2R2 , QFX5100 40G port port has interoperability issues with some other vendors. [PR1349664](#)
- Dcpfe process might crash on QFX10000 switches. [PR1351503](#)
- QFX10002: Telemetry traffic does not leave the local device when the telemetry server is reachable through a routing instance. [PR1352593](#)
- QFX5100 ARP fails after change interface MAC address is changed. [PR1353241](#)
- RPC output is not showing failure when running **request system software add** with software already staged. [PR1353466](#)
- On QFX5110 platforms, SFP-LX10 might stay in up or down state when connected. [PR1353677](#)
- The Alarm errors might be seen during the startup on QFX10000. [PR1354582](#)
- Untagged packets may not be forwarded through the trunk port [PR1355338](#)
- Commit error is observed if the device is downgraded from Junos OS Release 18.2 to Release 17.3R3. [PR1355542](#)
- On LX10 SFPs on QFX5110 platforms, autonegotiation is not in effect with a new configuration. [PR1355746](#)
- "Load averages" output under **show chassis routing-engine** shows "nan" periodically. [PR1356676](#)
- The IGMP membership report packets might not be forwarded over an interface on QFX10000. [PR1360137](#)
- On QFX10000, virtual-gateway-address should be only configured on an IRB interface associated with a VXLAN VLAN. [PR1360646](#)
- The GTP traffic might not be hashed correctly for the AE interface. [PR1361379](#)
- On QFX10K platforms, the "clear services accounting statistics inline-jflow fpc-slot" command does not work. [PR1362396](#)
- QFX5100VC: Unable to connect management address through vme interface. [PR1362437](#)
- Traffic might not be forwarded when the member link of the AE is added or deleted. [PR1362653](#)
- 1G interface might stop working when **no-auto-negotiation** is configured. [PR1362977](#)
- OSPF might remain in initialization status after firmware upgrade loading the Junos OS Release 14.1X53-D47.4 image. [PR1362996](#)
- On QFX10008 and QFX10016 platforms, MPLS exp rewrite might not work for IPv6 and IPv4 traffic. [PR1364391](#)

- Root password recovery process does not work. [PR1365740](#)
- On QFX10002-60C and QFX10000-30C platforms, some interfaces do not come up during initialization after a reboot. [PR1368203](#)
- On QFX5100, QFX5110, and QFX5200 platforms, IS-IS adjacency goes down when MTU 9192 is configured. [PR1368913](#)
- The **commit** or **commit check** might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)
- On QFX10000 platforms, before the Junos OS Release 17.3R3, the maximum number of ESI logical interface (IFLs) was 4000 in the Packet Forwarding Engine. [PR1371414](#)
- TPI-50840 BUM traffic received on 5110 is not flooded to all remote vteps. [PR1373093](#)
- LLDP might stop fully working between a QFX10000 and non-Juniper device. [PR1374321](#)

Routing Protocols

- On QFX5110 platforms, setting MTU on a L3 interface does not take effect. [PR1345495](#)
- On QFX10000 platforms, NETCONF SSH TCP port 830 traffic hits host path or unclassified queue. [PR1345744](#)
- On QFX5100 platforms, parity errors in the L3 IPv4 table in the Packet Forwarding Engine memory might cause traffic to be silently dropped or discarded. [PR1364657](#)

Software Installation and Upgrade

- Commit might fail in single-user mode. [PR1368986](#)

Resolved Issues: 18.1R2

EVPN

- Sub interface from the same physical port do not work if configured under same VXLAN VLAN. [PR1278761](#)
- VXLAN traffic loss is observed after deleting and adding VLANs. [PR1318045](#)
- QFX5100: EVPN-VXLAN: leaf device forwarding traffic to the incorrect VTEP after MAC move / vmotion. [PR1335431](#)
- Configuration of VXLANs with and without **encapsulate-inner-vlan** cannot co exist causing traffic issues on access interfaces. [PR1337953](#)
- In EVPN/VXLAN environment, BFD flaps cause VTEP flaps and cause the Packet Forwarding Engine to crash [PR1339084](#)
- The rpd generates a core file on QFX Series switches with multiple VLANs with vlan-id zero, unique VNID. [PR1342351](#)

Interfaces and Chassis

- CVLANs range is 16 might not pass traffic in a Q-in-Q scenario. [PR1345994](#)

Layer 2 Features

- QFX5100: With multiple logical units configured on an interface, **input-vlan-map POP** is not removing outer vlan-tag when QinQ and VXLAN are involved. [PR1331722](#)
- Push is not working for VXLAN local switching with the QinQ. [PR1332346](#)
- Interface with **flexible-vlan-tagging** and **family ethernet-switching** does not work on the QFX10000 line.. [PR1337311](#)

MPLS

- The hot standby for I2circuit does not work on QFX5100, QFX5110, and QFX5200. [PR1329720](#)

Platform and Infrastructure

- C0 fiber link does not come up. [PR1298876](#)
- Packets such as TDLS without IP headers are looped between virtual gateways. [PR1318382](#)
- Autonegotiation is not working as expected between EX4300 and SRX5800. [PR1318382](#)
- The openflow session cannot be established correctly with controller and interfaces options configured on QFX5100 series switches. [PR1323273](#)
- The GRE traffic is not decapsulated by the firewall filter. [PR1325104](#)
- VLAN or VLAN bridge might not be added or deleted if there is an IFBD hardware token limit exhaustion. [PR1325217](#)
- Deleting one VXLAN might cause traffic loop on another VXLAN in a multihoming EVPN-VXLAN scenario with service provider style interface. [PR1327978](#)
- Directories and files under **/var/db/scripts** lost execution permission or directory 'jet' is missing under **/var/db/scripts** causing **error: Invalid directory: No such file or directory** error during commit. [PR1328570](#)
- The PTX10000 line card might reboot continuously after upgrading to Junos OS Release 17.2R1 or later if HMC BIST fails. [PR1330618](#)
- DHCP relay/server is not working on GRE interface on QFX10002-36Q (Elit). [PR1331158](#)
- PTP BC with its PTP slave interface configured on a 100-Gigabit Ethernet interface might get stuck in FREERUN state. [PR1331752](#)
- EVPN-VXLAN: DF drops multicast traffic. [PR1333069](#)
- Chassis reboots continuously when USB drive is connected after image recovery through USB and after CLI image install. [PR1335269](#)
- PTX1000 and QFX10002-60C: Python scripts/shell scripts cannot be executed during ZTP because verixec is enabled. [PR1334425](#)

- Supported scale for logical interface (IFL) based GRE tunnel on QFX10002-60C is 512. [PR1335681](#)
- SNMP jnxBoxDescr oid returns different value when upgrading to Junos OS Release 17.2. [PR1337798](#)
- The traffic coming from the remote VTEP PE device might be dropped. [PR1338532](#)
- The analyzer status might show as down when port mirroring is configured to mirror packets from an AE member. [PR1338564](#)
- The VXLAN traffic might not be transmitted correctly with IRB interface as underlay interface of VTEP tunnel. [PR1338586](#)
- DDOS counters for OSPF might not increment. [PR1339364](#)
- Reduced multicast scale with downstream IRB interfaces with snooping enabled. [PR1340003](#)
- QFX5200: Inconsistent result occurs after using **deactivate xxx** command in pfc-priority and no-loss context. [PR1340012](#)
- JDI-RCT : QFX5210-64C : IPv4 traffic routed out through the incorrect interface after rpd restart in leaf of IPCLOS profile. [PR1341381](#)
- While downgrading PTX from a later release, the router goes into amnesiac state. [PR1341650](#)
- JDI-RCT: EVPN-VXLAN: L3 traffic is not getting converged properly upon disabling the ECMP link between the spine and leaf devices with EVPN-VXLAN configurations. [PR1343172](#)
- Broadcast frames might be modified with the ethertype 0x8850. [PR1343575](#)
- EVPN-VXLAN: VLAN with flexible-tag mode , the xe statistics do not get updated for ingress traffic. [PR1343746](#)
- Implement **edit interfaces interface-name ether-options] configured-flow-control** option for QFX Series switches. [PR1343917](#)
- EVPN-VXLAN: ARP reply packet has auto generated virtual gateway MAC in Ethernet header. [PR1344990](#)
- The fxpc process might generate core files when removing a VXLAN configuration. [PR1345231](#)
- EVPN Type5: QFX5110 dcpfe generates core files at `src/pfe/common/pfe-arch/brcm/applications/virtual/brcm_vxlan.c:2185`. [PR1346980](#)
- Part numbers and serial numbers are not displayed for any of the optics/DAC connected. [PR1347634](#)
- The ARP might not update and packets might get dropped at the Routing Engine. [PR1348029](#)
- On the QFX10002-60C VMHOST, a crash was observed at `@ prds_if_ifl_get_gre_stats (ifl=0x9288a608, expr_ifl_l2d_stats=0x2cd3790c)`, just after configuring the GR Interface on it. [PR1348932](#)
- The pfed process is consuming 80-90 percent CPU usage when running subscriber management on PPC-based routers. [PR1351203](#)
- The GTP traffic might not be hashed correctly for aggregated Ethernet interface. [PR1351518](#)

Routing Protocols

- Diffserv bits/ToS bits are not getting copied from the inner IP header to GRE header, Wireshark captured attached with PR. [PR1313311](#)
- Some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update. [PR1320723](#)
- The dcpfe crash is seen in route leak scenario on QFX10000. [PR1334714](#)
- The rpf-check-policy does not work as expected. [PR1336909](#)
- QFX loopback firewall filter is not able to catch packets with martian source address. [PR1343511](#)
- vrf-fallback on the QFX5100 switch, is not supported in ALPM mode. [PR1345501](#)
- IPv6 packets with hop-by-hop header are not matched by filters. [PR1346052](#)

Resolved Issues: 18.1R1

Class of Service (CoS)

- For some of the frame sizes, throughput is not 100 percent. [PR1256671](#)

EVPN

- NH installation error messages are seen on QFX10000 .[PR1258930](#)
- VXLAN-EVPN: IPv6 Packet loss after normal traffic run rate. [PR1267830](#)
- Normal VRRP MAC is triggering a MAC move, and logical interfaces on the BD are getting shut down. [PR1285749](#)
- QFX10002 VXLAN with MPLS underlay has traffic loss at RSVP egress.[PR1289666](#)
- The **df-election-type** preference statements at the [**show interfaces esi**] hierarchy level are not supported on QFX10000 running Junos OS Release 17.3R1. [PR1300093](#)
- QFX5110-48S: L3 VPN traffic is dropped for some instances when EVPN-VXLAN configuration is removed and reapplied. [PR1307590](#)
- Dcpfe might crash on EVPN-VXLAN setup. [PR1315531](#)
- Core file link flap might result in inconsistent global MAC count. [PR1328956](#)
- EVPN-VXLAN: EVPN Type7 route is not synced across ESI peers when virtual-switching or EVPN instance exist. [PR1334408](#)
- QFX5100 -- EVPN-VXLAN -- Leaf forwarding traffic to incorrect VTEP after MAC move / vmotion. [PR1335431](#)

Interfaces and Chassis

- Multicast data packets are looping in MC-LAG. [PR1281646](#)
- ARP reply drop occurs in MC-LAG scenario. [PR1282349](#)

- Upgrading to Junos OS Release 16.1R5 without the **redundancy-group-id-list** statement prior in ICCP leads to commit failure during bootup. [PR1311009](#)

Layer 2 Features

- To set up PTP BC forwarding on a QFX10002, configure routing on the interface or add a static ARP entry on the remote PTP device. [PR1275327](#)
- Device transmits packets that exceed interface MTU. [PR1306724](#)
- The **bpdu-block-on-edge** statement does not work correctly when **fast-tune** is enabled. [PR1307440](#)
- jdhcpd core files are observed after making DHCP configuration changes. [PR1324800](#)
- Commit error occurs while configuring **native-vlan-id**. [PR1318881](#)
- NLB heartbeat packets might be dropped on QFX10000 and PTX Series. [PR1322183](#)
- ARP entry might be learned on STP blocking ports. [PR1324245](#)
- Junos Fusion MAC Learning failure occurs for device on Extended Satellite Interface. [PR1324579](#)
- The DHCP discover packets might be looped in an MC-LAG and DHCP-relay scenario. [PR1325425](#)
- QFX5100 : With multiple logical units configured on an interface, " input-vlan-map POP " is not removing outer vlan-tag when QinQ and VXLAN are involved. [PR1331722](#)
- Interface with **flexible-vlan-tagging** and **family ethernet-switching** does not work on QFX10K. [PR1337311](#)

MPLS

- QFX5100: ISSU is not supported with MPLS configuration. [PR1264786](#)
- Traffic drop during NSR switchover for RSVP P2MP provider tunnels used by MVPN occurs. [PR1293014](#)
- DHCP clients cannot get IP address over BGP-L3VPN. [PR1303442](#)
- MPLS forwarding might not happen properly for some LSPs. [PR1319379](#)
- The rpd might crash on backup RE due to memory exhaustion. [PR1328974](#)
- Hot standby for I2circuit does not work on QFX5100. [PR1329720](#)

Multicast

- aggregated Ethernet interface and IRB configuration issue causes kernel crash and causes either chassis or FPC to reboot. [PR1335904](#)

Platform and Infrastructure

- UFT for non local member is not shown in the CLI. [PR1243758](#)
- QFX5100 TVP: Not able to load TVP image on top of a non-TVP 5100 image while adding a QFX5100 switch to the Virtual Chassis. [PR1248145](#)
- Copper ports flap on QFX5100-48T when short-reach-mode is enabled. [PR1248611](#)

- After upgrading the QFX5100/EX4600 to Junos OS Release 16.1 from 15.1, commit warning. `/boot/ffp.cookie+` might be seen. [PR1283917](#)
- On QFX5100 switches, an AE interface might flap upon commit if an explicit speed is configured on an AE member interface. [PR1284495](#)
- BFD sessions might flap when BFD is configured over IRB interfaces. [PR1284743](#)
- Protocols might flap when disabling the AE member link. [PR1289703](#)
- Storm-control flags are not set after a Routing Engine switchover. [PR1290246](#)
- On QFX5100, the fxpc process generates a core file. [PR1294033](#)
- ULC-60S-6Q LC on QFX10008: The port becomes unusable after inserting a third-party SFP-T optic. [PR1294394](#)
- Oinker and TCP connection drop might be seen during large file SCP/FTP to the system (high `intr{virtio_p}` seen). [PR1295774](#)
- The 40-Gigabit Ethernet interface might not come up if a specific vendor's DAC cable is used. [PR1296011](#)
- The disable-pfe action upon hybrid memory cube (HMC) fatal errors might have a system-wide impact on PTX Series platforms. [PR1300180](#)
- QFX10008/10016: commit error is seen when configured with mixed speed. [PR1301923](#)
- If MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic with an error message **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)
- Systems running 32-bit Junos OS might generate rpd core file when traceoptions are enabled. [PR1305440](#)
- QFX5110-48S: Digital optical monitoring statistics cannot be received through the CLI in Junos OS Releases 15.1X53 through 17.x. [PR1305506](#)
- QFX5200: New apply group is not applying to the Virtual Chassis after a reboot. [PR1305520](#)
- QFX5100 crashes and the fxcp process generates a core file. [PR1306768](#)
- Some error messages might be observed on EVPN-VXLAN setup. [PR1307014](#)
- QSFP+4x10G-IR channelized interface goes down between QFX5200 and PTX5000. [PR1307400](#)
- Traffic stopped passing LSP after MPLS route change. [PR1309058](#)
- QFX5110 VC/VCF: Virtual Chassis members reboot before all members have image installed. [PR1309103](#)
- Run time pps statistics value might show zero for a subinterface of AE interface. [PR1309485](#)
- Traffic loss might be seen if sending traffic through the 40G interface. [PR1309613](#)
- Some log messages are seen on QFX5110 platform when plugging in an SFP-SX. [PR1311279](#)
- One aggregated Ethernet member does not send out sFlow sample packets. [PR1311559](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)

- CPU utilization is around 50 percent without any configuration. [PR1312520](#)
- QFX5100:5100-24q: After loading TVP image, unable to offline/online the EX4600-EM-8F PIC; shows as unsupported. [PR1313392](#)
- QFX10002-60C will support **show vmhost crash** to display core files in the host OS. [PR1314451](#)
- Transit traffic over GRE tunnel might hit CPU and trigger a DDoS violation on L3NHOP. [PR1315773](#)
- On switch platforms running under Junos OS with Enhanced Layer 2 Software (ELS) (EX4300/EX4600/EX9200/QFX5100/QFX10000), I2cpd might generate core files repeatedly if an interface is connected to VoIP product with LLDP and LLDP-MED enabled. [PR1317114](#)
- The optic interface still transmits power after it has been administratively shut down. [PR1318997](#)
- The packet might be dropped between 4-60 seconds when the master Routing Engine is rebooted in a virtual chassis. [PR1319146](#)
- Port speed is still showing 100G instead of 50G as IFD has been channelized to 50G. [PR1319884](#)
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX-VC ISSU. [PR1320370](#)
- The MACac address is stuck with "DR" flag on the spine node even though packets are received on the interface from source MAC. [PR1320724](#)
- FPCs are gone offline due to **CHASSISD_IPC_CONNECTION_DROPPED: Dropped IPC connection for FPC**. [PR1321198](#)
- The openflow session cannot be established correctly with controller on QFX5100 Series switches. [PR1323273](#)
- Update new firmware versions for jfirmware package for 100G-PSM4 and 100G-AOC issues. [PR1323321](#)
- EVPN Type 5: Unicast traffic getting is dropped on backup forwarder [PR1323907](#)
- VLAN or VLAN bridge might not be added or deleted if there is an IFBD HW token limit exhaustion. [PR1325217](#)
- MAC move is not expected when disabled globally with **set protocols l2-learning global-mac-move disable-action** [PR1325524](#)
- ARP request packets might not be flooded on QFX5110. [PR1326022](#)
- QFX5210-64C When the physical interface is down, **show chassis LED** CLI still showing as "Green". [PR1326078](#)
- QFX5100/EX4600/ACX5k: Major Alarm **Fan & PSU Airflow direction mismatch** occurs when removing management cable. [PR1327561](#)
- Deleting one VXLAN might cause traffic loop on another VXLAN in multi-homing EVPN/VXLAN scenario with Service Provider style interface. [PR1327978](#)
- Major alarm should be cleared once the chassis has more PEM units installed than the "minimum PEM" configuration. [PR1327999](#)

- Junos automation folder lost execution permissions. [PR1328570](#)
- Fan tray removal/insertion trap is not generated for the backup FPC. [PR1329031](#)
- QFX10000-60C : Although the set chassis fpc 0 pic command has the option of PIC numbers 0 to 2 , the switch has only 1 PIC.[PR1329105](#)
- After commit, members of VC or VCF are split and some members may get disconnected. [PR1330132](#)
- When configure total of 500 tunnels and all are part of routing-instance (500 routing-instance) and 500 BGP session with 20k routes. Adding or deleting configurations might occasionally result in FPC crash. [PR1331983](#)
- The error messages **out of HMC range** and **HMC READ failed** are seen. [PR1332251](#)
- The SOLICIT message of DHCPv6 is dropped. [PR1334680](#)
- Supported scale for IFL based GRE tunnel on QFX10002-60C is 512. [PR1335681](#)
- PTX1000 & QFX10002-60C: Python scripts/shell scripts cannot be executed during ZTP as veriexec is enabled.[PR1334425](#)
- CLI for beacon port state is not supported on QFX10002-60C. [PR1337125](#)
- The traffic coming from the remote VTEP PE might be dropped. [PR1338532](#)
- QFX5200 : Inconsistent result after using 'deactivate xxx' command on 'pfc-priority' and 'no-loss' context. [PR1340012](#)
- Implement **edit interfaces interface-name ether-options] configured-flow-control** option for QFX. [PR1343917](#)
- When upgrading from certain release to 18.1R1 statistics daemon PFED may be seen to core. This issue is not service impacting. The issue can be cleared by rebooting the chassis or by deleting all files from /mfs. [PR1346925](#)

Routing Policy and Firewall Filters

- The rpd might crash if **vrf-target auto** is configured under routing-instance [PR1301721](#)

Routing Protocols

- Filter-based forwarding (FBF) with next-ip/next-ip6/next-interface is not working [PR1289642](#)
- Remotely received traffic is not flooded to AC on FPC 1 when FPC 0 is offline.[PR1290500](#)
- An mcsnoopd core file is observed at
`__raise,abort,__task_quit__,task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal
(enable_slip_detector=true,no_exit=true) at ../../../../src/junos/lib/libtask/base/task_scheduler.c:275`
[PR1305239](#)
- GRE tunneled packets might be dropped. [PR1308438](#)
- QFX5100: Consistent hashing is not getting programmed. [PR1322299](#)

- QFX10002-60C is not supported as FHR in multicast PIM SM based network. [PR1324116](#)
- IS-IS L2 Hello packets are dropped when they come from a Brocade device. [PR1325436](#)
- vrf-fallback on QFX5K is not supported in ALPM mode. [PR1345501](#)

Virtual Chassis

- Sometimes multicast packets are received two or three time faster. [PR1306239](#)

SEE ALSO

[New and Changed Features | 248](#)

[Changes in Behavior and Syntax | 273](#)

[Known Behavior | 276](#)

[Known Issues | 282](#)

[Documentation Updates | 299](#)

[Migration, Upgrade, and Downgrade Instructions | 300](#)

[Product Compatibility | 314](#)

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 299](#)

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 18.1R2.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 248](#)

[Changes in Behavior and Syntax | 273](#)

[Known Behavior | 276](#)

[Known Issues | 282](#)

[Resolved Issues | 287](#)

[Migration, Upgrade, and Downgrade Instructions | 300](#)

[Product Compatibility | 314](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 301](#)
- [Installing the Software on QFX10002-60C Switches | 303](#)
- [Installing the Software on QFX10002 Switches | 303](#)
- [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 | 304](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 306](#)
- [Performing a Unified ISSU | 310](#)
- [Preparing the Switch for Software Installation | 311](#)

- Upgrading the Software Using Unified ISSU | 311
- Upgrade and Downgrade Support Policy for Junos OS Releases | 313

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 **18.1** 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 18.1 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-x86-64-18.1  
-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 18.1 **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-x86-64-18.1R3.9.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-x86-64-18.1R3.9.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 18.1R1.

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-x86-64-18.1R3.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-x86-64-18.1R3.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-18.1R3.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-x86-64-18.1R3.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 311](#)
- [Upgrading the Software Using Unified ISSU on page 311](#)

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-x86-64-18.1R3.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-x86-64-18.1R3.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-18.1R3.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

[New and Changed Features | 248](#)

[Changes in Behavior and Syntax | 273](#)

[Known Behavior | 276](#)

[Known Issues | 282](#)

[Resolved Issues | 287](#)

[Documentation Updates | 299](#)

[Product Compatibility | 314](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 314](#)

Hardware Compatibility

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.

To determine the features supported on QFX Series switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at <https://pathfinder.juniper.net/feature-explorer/>.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the [Hardware Compatibility tool](#).

SEE ALSO

New and Changed Features		248
Changes in Behavior and Syntax		273
Known Behavior		276
Known Issues		282
Resolved Issues		287
Documentation Updates		299
Migration, Upgrade, and Downgrade Instructions		300

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [New and Changed Features](#) | [316](#)
- [Changes in Behavior and Syntax](#) | [323](#)
- [Known Behavior](#) | [324](#)
- [Known Issues](#) | [326](#)
- [Resolved Issues](#) | [330](#)
- [Documentation Updates](#) | [341](#)
- [Migration, Upgrade, and Downgrade Instructions](#) | [343](#)
- [Product Compatibility](#) | [344](#)

These release notes accompany Junos OS Release 18.1R3 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.

New and Changed Features

IN THIS SECTION

- [Release 18.1R3 New and Changed Features | 316](#)
- [Release 18.1R2 New and Changed Features | 316](#)
- [Release 18.1R1 New and Changed Features | 317](#)

This section describes the new features and enhancements to existing features in Junos OS Release 18.1R3 for the SRX Series devices.

Release 18.1R3 New and Changed Features

There are no new features in Junos OS Release 18.1R3 for the SRX Series devices.

Junos OS Release 18.1R3 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/4200, SRX4600, SRX5400, SRX5600, and SRX5800. Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 12.1X44 through 15.1X49-D120. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D120 are not available in 18.1 releases.

Release 18.1R2 New and Changed Features

There are no new features in Junos OS Release 18.1R2 for the SRX Series devices.

Junos OS Release 18.1R2 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/4200, SRX4600, SRX5400, SRX5600, and SRX5800. Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 12.1X44 through 15.1X49-D120. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D120 are not available in 18.1 releases.

Release 18.1R1 New and Changed Features

IN THIS SECTION

- [Application Security | 317](#)
- [Authentication and Access | 319](#)
- [Chassis Cluster | 319](#)
- [Class of Service \(CoS\) | 319](#)
- [Flow-Based and Packet-Based Processing | 319](#)
- [Interfaces and Chassis | 319](#)
- [Multicast | 321](#)
- [Network Management and Monitoring | 321](#)
- [User Interface and Configuration | 321](#)
- [VPN | 321](#)

Junos OS Release 18.1R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/4200, SRX4600, SRX5400, SRX5600, and SRX5800. Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 12.1X44 through 15.1X49-D120. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D120 are not available in 18.1 releases.

Application Security

- **Data Loss Prevention (SRX Series)** —Starting in Junos OS Release 18.1, SRX Series devices support Data Loss Prevention (DLP) to redirect HTTP or HTTPS traffic to any server through Internet Content Adaptation Protocol (ICAP).

ICAP is a lightweight protocol for executing a remote procedure call on HTTP messages using **REQMOD** which encapsulate HTTP request messages and **RESPMOD** which encapsulate HTTP response messages.

See [SSL Proxy](#).

- **Optimizing SSL/TLS performance for HTTPS traffic (SRX Series, vSRX)** —Starting from Junos OS Release 18.1R1, SSL/TLS performance is optimized by minimizing the time required for performing the decryption by using the following methods:
 - Using optimized cipher suites
 - Maintaining the certificate cache

Enhanced SSL/TLS performance for HTTPS traffic results in improved website performance without compromising security, and maximizes user experience.

[See [SSL Proxy](#)].

- **SSL proxy support (SRX300, SRX320)**—Starting in Junos OS Release 18.1R1, SSL proxy support is available on SRX300 and SRX320 devices. SSL proxy acts as an intermediary, performing SSL encryption and decryption between the client and the server. SSL relies on digital certificates and private-public key exchange pairs for client and server authentication to ensure secure communication.

[See [SSL Proxy](#)].

Authentication and Access

- **IPv6 support for network access control (NAC) (SRX Series, vSRX)**—Starting with Junos OS Release 18.1R1, SRX Series devices support IPv6 for the network access control (NAC) system. You can configure a Web API client address with an IPv6 address and Web API supports IPv6 user or device entries obtained from Juniper Identity Management Service (JIMS). An SRX Series device can query JIMS periodically for batches of newly generated IPv6 users or devices for identity information. The SRX Series can query JIMS for identity information for an individual user or device based on the IPv6 address when the IPv6 traffic hits the SRX Series device. The SRX Series device firewall authentication can push IPv6 IP-user mapping information to JIMS.

[See [Understanding the SRX Series Advanced Query Feature for Obtaining User Identity Information from JIMS](#) .]

Chassis Cluster

- **VRRP and VRRPv3 support on redundant Ethernet interface to provide redundancy (SRX Series, vSRX)**—Starting with Junos OS Release 18.1R1, SRX Series devices in a chassis cluster support the Virtual Router Redundancy Protocol (VRRP) and VRRPv3 on reth interfaces to provide redundancy, route advertising, and load sharing. Using VRRP, a secondary node can take over a failed primary node within a few seconds with minimum VRRP traffic and without any interaction with the hosts.

[See [Understanding VRRP on SRX Series Devices](#).]

Class of Service (CoS)

- **Support for rewrite rules for both inner and outer VLAN tags on IEEE802.1 packets (SRX Series)**—Starting with Junos OS Release 18.1R1, SRX Series devices support applying rewrite rules to both inner and outer VLAN tags on IEEE802.1 packets. To apply rewrite rules to both inner and outer VLAN tags, set the **vlan-tag outer-and-inner** option at the **[edit class-of-service interfaces interface-name unit unit-number rewrite-rules ieee-802.1 rewrite-name]** hierarchy level.

[See [rewrite-rules \(CoS Interfaces\)](#)]

Flow-Based and Packet-Based Processing

- **Enhancement for show security flow statistics operational command (SRX Series, vSRX instances)**—Starting in Junos OS Release 18.1R1, the output of the **show security flow statistics** command has been modified. The **Packets forwarded** field has been split into the **Packets received** and **Packets transmitted** fields. The **Packets received** field displays the actual number of packets received, including those dropped by the system. The **Packet transmitted** field displays the number of packets returned to jexec for transmission. The **Packets forwarded/queued** field displays the actual number of packets forwarded excluding the dropped packets.

Additionally, a new field, **Packets copied** has been created to provide information about packets copied by other modules including fragmentation and TCP proxy.

[See [show security flow statistics](#).]

Interfaces and Chassis

- **Support for 4x10-Gigabit Ethernet Optical Breakouts (SRX4600)**—Starting in Junos OS Release 18.1R1, you can use optical breakout cable to configure four 10-Gigabit Ethernet interfaces on each 40-Gigabit Ethernet port on an SRX4600. By default, FPC 1 PIC 0 comes up with the default setting of four 40-Gigabit Ethernet ports. This new feature allows the 40 Gigabit Ethernet port to be configured in 4X10-Gigabit Ethernet mode by plugging in QSFP-4X10-Gigabit Ethernet optics connecting with 4x10-Gigabit Ethernet breakout cables. You use QSFP+ transceivers to connect the 40-Gbps (default speed) port to the breakout cable, which connects to four SFP+ transceivers at the other end thus converting that port into four 10-Gbps interfaces).

For example, on FPC 1 PIC 0, to configure each 40-Gbps port as four 10-Gbps interfaces, execute the **set chassis fpc 1 pic 0 pic-mode 10G** command.

After you commit the configuration, for the new configuration to take effect, you must reboot the device or chassis cluster. [See [SRX4600 Gateway Rate-Selectability Overview](#).]

- **Support for default 10-Gbps ports to operate at 1-Gbps speed (SRX4600)**—Starting in Junos OS Release 18.1R1, SRX4600 supports 1-Gbps port speed on the default 10-Gbps ports on its 8-port PICs and on two dedicated chassis cluster control ports on the 4-port chassis cluster PICs. The SRX4600 supports three different PIC types—8-port 10-Gigabit Ethernet PIC, 4-port 40-Gigabit or 100-Gigabit Ethernet PIC, and 4-port 10-Gigabit Ethernet PIC (in a chassis cluster). Out of the four ports on the 10-Gigabit Ethernet PIC in a chassis cluster, two ports are fabric ports and the other two ports are chassis cluster control ports. The two fabric ports do not support 1-Gbps speed. Only the two control ports of the chassis cluster support a port speed of 1 Gbps.

NOTE:

- The interface name prefix must be xe.
- You can configure a combination of 1-Gbps and 10-Gbps speed only on the 8-port 10-Gigabit Ethernet PIC. The chassis cluster control interfaces (that is, on the 4-port 10-Gigabit Ethernet PIC) do not support multiple speeds.

[See [SRX4600 Gateway Rate-Selectability Overview](#).]

Multicast

- **Layer 2 IGMP and MLD Snooping feature support (SRX1500)**—Starting with Junos OS Release 18.1R1, the SRX1500 supports the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) snooping feature in Layer 2 switching mode.

The snooping feature snoops the IGMP or MLD packets received by the switch interfaces and builds a multicast database. The SRX Series device uses the multicast database and forwards the multicast traffic only to the downstream interfaces of interested receivers. Using the multicast database to forward multicast packets helps ensure efficient use of network bandwidth.

[See [IGMP Snooping Overview](#) and [Understanding MLD Snooping](#).]

Network Management and Monitoring

- **Two-Way Active Measurement Protocol (TWAMP) support (SRX4100, SRX4200 and vSRX)**—Starting in Junos OS Release 18.1R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on SRX4100 and SRX4200 devices and on vSRX instances in addition to the existing support on SRX Series devices such as SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500. TWAMP is a standard protocol framework that defines control and test session separation based on the client/server architecture. The TWAMP-Control protocol is used to set up performance measurement sessions between a TWAMP client and a TWAMP server, and the TWAMP-Test protocol is used to send and receive performance measurement probes.

[See [Two-Way Active Measurement Protocol \(TWAMP\) Overview](#).]

User Interface and Configuration

- **Ephemeral configuration database support for load replace and load override operations (SRX Series)**—Starting in Junos OS Release 18.1R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database using **load replace** and **load override** operations, in addition to the previously supported **load merge** and **load set** operations. To perform a **load replace** or **load override** operation, set the **<load-configuration> action** attribute to **replace** or **override**, respectively.

[See [Configuring Ephemeral Database Instances](#).]

VPN

- **Binding trusted CAs or trusted CA group to an IKE policy (SRX Series and vSRX instances)**—Starting in Junos OS Release 18.1R1, you can group CA profiles (trusted CAs) in a trusted CA group and or bind a specific CA profile to an IKE policy. When a remote peer establishing a connection that matches this IKE policy, the particular CA profile or trusted CA group is used to validate the remote peer.

A group of trusted CA servers can be created with the trusted CA group configuration statement at the **[edit security pki]** hierarchy level; one or multiple CA profiles can be specified. The trusted CA server is bound to the IKE policy configuration for the peer at **[edit security ike policy policy certificate]** hierarchy level.

[See [Understanding Certificates and PKI](#) and [Understanding Certificate Authority Profiles](#).]

- **IPv6 support for AutoVPN and ADVPN with dynamic routing protocol (SRX Series and vSRX instances)**—Starting with Junos OS Release 18.1R1, IPv6 is supported on AutoVPN and Auto Discovery VPN (ADVPN) with point-to-multipoint secure tunnel mode. ADVPN can run with OSPFv3 routing protocol and AutoVPN can run with OSPFv3 and iBGP (internal BGP) routing protocols.

The **ospf3** option is introduced at the **edit protocol** hierarchy level to support IPv6 for AutoVPN and ADVPN with point-to-multipoint secure tunnel mode. In addition, the **show security ipsec next-hop-tunnels** command, which displays the IPsec VPN tunnels bound to a specific tunnel interface, is updated to add **family** and **tunnel ID** filters.

[See [Understanding AutoVPN](#) and [Understanding Auto Discovery VPN](#).]

- **IPv6 support for PKI (SRX Series and vSRX instances)**—Starting in Junos OS Release 18.1, the public key infrastructure (PKI) supports IPv6 address format for the Certificate Authority (CA) server and source addresses in a CA profile. The PKI provides an infrastructure for digital certificate management. In PKI, a CA is a trusted third party agency responsible for issuing and revoking certificates. The certificates are used to create secure connections between two or more entities.

[See [Understanding Certificate Authority Profiles](#).]

- **SSL remote access VPN support by bypassing an application-based firewall (SRX Series and vSRX instances)**—Starting with Junos OS Release 18.1R1, remote access VPN uses SSL to pass through an application level firewall using the third-party NCP Exclusive Remote Access Client on Windows, MAC OS, Apple iOS, and Android devices.

Most intermediate Internet-facing devices allow users to establish a session over SSL (HTTPS) to any Internet-based device. This solution allows users to establish a secure communication using a full SSL session when an intermediate device blocks IPsec or UDP traffic.

[See [Understanding SSL Remote Access VPNs with NCP Exclusive Remote Access Client](#).]

SEE ALSO

[Changes in Behavior and Syntax | 323](#)

[Known Behavior | 324](#)

[Known Issues | 326](#)

[Resolved Issues | 330](#)

[Documentation Updates | 341](#)

[Migration, Upgrade, and Downgrade Instructions | 343](#)

[Product Compatibility | 344](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.1R3 for the SRX Series.

Chassis Cluster

- The SRX5400, SRX5600, and SRX5800 devices operating in a chassis cluster might encounter the em0 or em1 interface link failure on either of the nodes, which results in split-brain condition. That is, both devices are unable to detect each other. If the failure occurs on the secondary node, the secondary node is moved to the disabled state.

This solution does not cover the following cases:

- em0 or em1 failure on primary node
- HA process restart
- Preempt conditions
- Control link recovery

Juniper Sky ATP

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 18.1R3, 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).

VPN

- **Default encryption algorithm for PKI certificates (SRX Series and vSRX)**—Starting in Junos OS Release 18.1R3, the default encryption algorithm that is used for validating automatically and manually generated self-signed PKI certificates is Secure Hash Algorithm 256 (SHA-256).

Prior to Junos OS Release 18.1R3, the default encryption algorithm is SHA-1.

[See [Understanding Certificates and PKI](#) and [request security pki local-certificate generate-self-signed \(Security\)](#).]

SEE ALSO

[New and Changed Features | 316](#)

[Known Behavior | 324](#)[Known Issues | 326](#)[Resolved Issues | 330](#)[Documentation Updates | 341](#)[Migration, Upgrade, and Downgrade Instructions | 343](#)[Product Compatibility | 344](#)

Known Behavior

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.1R3 for the SRX Series.

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

Chassis Clustering

- If you enable IP monitoring on the redundancy group when the reth interface has more than one physical interface configured, then IP monitoring may not work properly on the secondary node. This issue occurs because the backup node may send traffic using the MAC address of the lowest port in the bundle. If the reply does not come back on the same physical port, then the internal switch may drop the traffic. [PR1344173](#)

J-Web

- On SRX4100 and SRX4200 devices, as part of JDHCP changes DHCP relay configuration under **Configure > Services > DHCP > DHCP Relay** page is removed from J-Web in Junos OS Release 15.1X49-D60. The same DHCP relay can be configured using the CLI. [PR1205911](#)
- On SRX4100 and SRX4200 devices, as part of JDHCP changes DHCP client bindings under **Monitor** is removed for Junos OS Release 15.1X49-D60. The same bindings can be seen in CLI using the **show dhcp client binding** command. [PR1205915](#)
- On SRX Series devices, adding of 2,000 or more global addresses at a time to the SSL proxy profile exempted addresses can cause the web page to be unresponsive. [PR1278087](#)
- On SRX Series devices, you cannot view the custom log files created for event logging in J-Web. [PR1280857](#)
- On SRX Series devices, validation is not checked when the UTM policy is detached from the firewall policy rule after an SSL proxy profile is selected. [PR1285543](#)

- On SRX Series devices, uploading certificate using the browse button stores the certificate in the device at `/jail/var/tmp/uploads/`, which will be deleted upon executing the **request system storage cleanup** command. [PR1312529](#)
- On SRX Series devices, the values of address and address-range are not displayed in the inline address-set creation pop-up window of JIMS. [PR1312900](#)
- Application signature install or uninstall status above the grid remains in loading state when the device connectivity to the cloud server. Application signature database is not present or not responding. This in turn affects the status that is displayed in the J-Web. [PR1332768](#)

Platform and Infrastructure

- On SRX4600 devices, the USB flash drive is not available to Junos OS. However, the USB flash drive is available for the host OS (Linux) with full access. The USB flash drive is still used in the booting process (install and recovery functions). [PR1283618](#)
- When a USB device is under initialization, removing the USB device may cause the USB to stop working. [PR1332360](#)

Software Installation and Upgrade

- When you upgrade from Junos OS Release 15.1X49, the signature version is automatically refreshed to version 534. Hence, you need to download and install a new signature version; if not, some features such as SKYATP IMAP may be missing. [PR1324848](#)

User Interface and Configuration

- On SRX1500 devices, committing a configuration with a huge number of logical systems will take more time. This issue occurs because taking backup of previous configurations may take a little longer to finish. [PR1339862](#)

VPNs

- On SRX5400, SRX5600, and SRX5800 devices, when CoS is enabled on the st0 interface and the incoming traffic rate destined for the st0 interface is higher than 300,000 packets per second (pps) per SPU, the device may drop some of the high-priority packets internally and shaping of outgoing traffic may be impacted. We recommended that you configure the appropriate policer on the ingress interface to limit the traffic below 300,000 pps per SPU. [PR1239021](#)
- On SRX Series devices, IPsec traffic statistics counters return 32-bit values, which may quickly overflow. [PR1301688](#)

SEE ALSO

[New and Changed Features | 316](#)[Changes in Behavior and Syntax | 323](#)[Known Issues | 326](#)[Resolved Issues | 330](#)[Documentation Updates | 341](#)[Migration, Upgrade, and Downgrade Instructions | 343](#)[Product Compatibility | 344](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 18.1R3 for SRX Series devices.

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

Application Layer Gateways (ALGs)

- On all SRX Series devices with NAT configured, a memory overwrite issue occurs when the scaling RAS or H.323 traffic passes through the device and the device fails to perform NAT for RAS or H.323 traffic. As a result, the flowd process may stop. [PR1084549](#)
- On SRX Series devices with chassis cluster enabled and logical systems configured, when any ALG (except DNS ALG) is enabled and NAT is configured for the ALG sessions, the flowd process on the secondary node may stop. [PR1343552](#)
- When using the IPsec ALG, the IPsec tunnel payload is dropped after the IKE or IPsec tunnel reestablishment due to a session conflict. [PR1372232](#)
- If the SIP ALG is disabled, the SIP active sessions are affected. [PR1373420](#)

Chassis Clustering

- On SRX550M device, the SFP transceiver does not work after the chassis reboot. [PR1347874](#)

Class of Service (CoS)

- On all SRX Series devices, if the action of **forwarding-class** is configured in the output direction on a firewall filter, the host outbound traffic matching the same term of this firewall filter is blocked. [PR1272286](#)
- When the **host-outbound-traffic** command is configured in the class of service (CoS), the device stops working when a corrupted packet is arrived on the Packet Forwarding Engine. [PR1359767](#)

Flow-based and Packet-based Processing

- SRX1500 devices may power-off unexpectedly due to incorrect device temperature readings which reportedly is a too high temperature, leading to an immediate proactive power-off of the device to protect the device from overheating. When this condition occurs, the following log message is shown in file `/var/log/hostlogs/lcmd.log`: `Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor`. [PR1241061](#)
- On SRX1500 devices, the message `/kernel: kern.maxfiles limit exceeded by uid 0, please see tuning(7)` is displayed when the `kdm_savecore` process consumes the maximum open files allowed. As a workaround, use the **savecore -C** command to stop the file processing and clear the kernel crash flag, and reboot the device. [PR1277664](#)
- On 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 all SRX Series devices, filter-based forwarding (FBF) does not work when applied on IPsec tunnel interface (`st0.*`). [PR1290834](#)
- On SRX320, SRX340, SRX340, and SRX550 devices, the RPD process stops when you configure the **auto-bandwidth** option under the label-switched path (LSP) in the multiprotocol label switching (MPLS). [PR1331164](#)
- On SRX Series devices, when you run the **clear nhdb statistics** command on an SPU PIC, the SPC may reset. [PR1346320](#)
- The IPsec replay error for Z-mode traffic is observed. [PR1349724](#)
- On SRX Series devices in a chassis cluster, if an IPv6 session is being closed and at the same time the related data-plane Redundancy Group (RG1+) failover occurs, this IPv6 session on the backup node may hang and cannot be cleared. [PR1354448](#)
- The application layer protocol negotiation (ALPN) fails because the SSL proxy removes the ALPN extensions in the TLS packets. [PR1360820](#)

- In chassis cluster mode with the IPsec tunnel configured, packet loss is observed when the clear text packets are processed. [PR1373161](#)
- The Windows security log can overwrite the username that contains null to **N/A**. This issue causes the access privileges granted to that IP address to be lost. [PR1375514](#)

Interfaces and Routing

- Incorrect ingress packet per second is observed on the MPLS enabled interface. [PR1328161](#)
- On the SRX1500, when the LACP is configured with interfaces ae0 and ae1, the mac address is displayed as 00:00:00:00:00:00 and 00:00:00:00:00:01 for interfaces ae0 and ae1 respectively. [PR1352908](#)

Intrusion Detection and Prevention (IDP)

- On SRX Series devices, the output of **show security idp status** command does not accurately reflect the number of decrypted SSL or TLS sessions being inspected by IDP. [PR1304666](#)
- After an IDP signature automatic update is scheduled, the secondary node may not update the signatures. [PR1358489](#)

Platform and Infrastructure

- On SRX5400, SRX5600, and SRX5600 devices, when the control link is down, the secondary node becomes ineligible and then goes to disabled state. But the FPCs restart continuously after going to disabled state when the FPCs should remain offline until rebooted. [PR1170024](#)
- On SRX5600 and SRX5800 devices in a chassis cluster, when a secondary Routing Engine is installed to enable dual control links, the **show chassis hardware** command may display the same serial number for both the Routing Engines on both the nodes. [PR1321502](#)
- On SRX Series devices, the forwarding plane may failover from node 0 to node 1 when an SPC stops unexpectedly. [PR1331809](#)
- On SRX5600 and SRX5800 devices in a chassis cluster, when a secondary Routing Engine is installed to enable dual control links, the **show chassis hardware** command may display the same serial number for both the Routing Engines on both the nodes. [PR1342362](#)
- SSH to the device fails if the **phone-home: kern.maxfiles** limit is exceeded. [PR1357076](#)
- On SRX4100 and SRX4200 devices, the SRX Network Time Protocol (NTP) client may not stay synchronized to the NTP server and as a result the device clock often switches from NTP to local time. [PR1357843](#)
- When the secure copy protocol (SCP) fails to transfer the active configuration to an archive site, the archive site also fails. [PR1359424](#)

Routing Policy and Firewall Filters

- On SRX Series devices, DNS name entries in policies may not be resolved if the routing instance is configured under a system name server. [PR1347006](#)

Routing Protocols

- On SRX Series devices, RIP is supported in packet-to-packet DC mode on st0 interfaces. [PR1141817](#)
- A new CLI command is required to prevent traffic loss during a disaster recovery failover scenario. [PR1352589](#)

Software Installation and Upgrade

- On SRX1500 devices, the fan speed often fluctuates. [PR1335523](#)

VPNs

- IPsec uses ESP as the default protocol, if the user does not explicitly configure the protocol. [PR1061838](#)
- When an SRX Series device acts as an initiator behind the NAT, disabling NAT on the router in between causes an immediate new negotiation failure because of an attempt to disable NAT using the port 4,500. The next attempt succeeds by using the port 500. Disabling NAT and bringing down all the existing tunnels and re-establishing the tunnels with port 500 is the expected behavior. [PR1273213](#)
- On SRX Series devices, in case multiple traffic-selectors are configured for a peer with IKEv2 reauthentication, only one traffic-selector will rekey at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors will be cleared without immediate rekey. New negotiation of those traffic-selectors may trigger through other mechanisms such as traffic or by peer. [PR1287168](#)
- IPsec traffic statistic counters return 32-bit values. [PR1301688](#)
- During an RGO failover in ISSU, when you use the rekeys, the iked core process file are generated. [PR1340973](#)
- When NCP profile is changed on an existing IKE gateway, the SSL session corresponding to the existing tunnel is not affected. [PR1323425](#)
- If a period . is present in the CA profile name then the PKID may face issues, if the PKID is restarted at any point. [PR1351727](#)
- On SRX Series devices in a chassis cluster, configuration commit may succeed even though the external logical interface configuration (reth) associated with the Internet Key Exchange (IKE) VPN gateway configuration is deleted. This may lead to configuration load failure during the next device boot-up. [PR1352559](#)

Software Installation and Upgrade

- On SRX1500 devices, the fan speed often fluctuates. [PR1335523](#)

VPNs

- When an SRX Series device acts as an initiator behind the NAT, disabling NAT on the router in between causes an immediate new negotiation failure because of an attempt to disable NAT using the port 4,500. The next attempt succeeds by using the port 500. Disabling NAT and bringing down all the existing tunnels and re-establishing the tunnels with port 500 is the expected behavior. [PR1273213](#)
- On SRX Series devices, in case multiple traffic-selectors are configured for a peer with IKEv2 reauthentication, only one traffic-selector rekeys at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic-selectors are cleared without immediate rekey. New negotiation of those traffic-selectors might be triggered through other mechanisms such as traffic or peer. [PR1287168](#)
- When NCP profile is changed on an existing IKE gateway, the SSL session corresponding to the existing tunnel is not affected. [PR1323425](#)
- If a period . is present in the CA profile name then the PKID might face issues, if the PKID is restarted at any point. [PR1351727](#)

SEE ALSO

[New and Changed Features | 316](#)

[Changes in Behavior and Syntax | 323](#)

[Known Behavior | 324](#)

[Resolved Issues | 330](#)

[Documentation Updates | 341](#)

[Migration, Upgrade, and Downgrade Instructions | 343](#)

[Product Compatibility | 344](#)

Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

Resolved Issues: 18.1R3

Chassis Clustering

- On SRX Series devices in chassis cluster, minor **Potential slow peers are: FWDD0 XDPC1 XDPC8 FWDD1** alarm is observed which can be ignored. [PR1371222](#)

Flow-based and Packet-based Processing

- When you use CFLOW, the source address for flow packets is not displayed. [PR1328565](#)
- SSH to the loopback interface of SRX Series devices does not work properly when AppTrack is configured. [PR1343736](#)
- SNMP MIB walk provides incorrect data counters for total current flow sessions. [PR1344352](#)
- On SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 devices in chassis cluster, when CoS is configured on a interface, LACP communication stops due to failure of the fabric port and the connections between the SRX device and other devices breaks. [PR1350731](#)
- The flowd process generates a core file when the SIP ALG is enabled. [PR1352416](#)
- When the routing instance is configured, the **UTM Anti-Spam:DUT** process do not send the DNS query. [PR1352906](#)
- The IPsec VPN traffic may be dropped on pass-through SRX Series device after an IKE rekey. [PR1353779](#)
- On SRX Series devices, when AppTrack is configured the flowd process stop. [PR1354671](#)
- On SRX Series devices, the error message **error: Policy is out of sync between RE (Routing Engine) and PFE (Packet Forwarding Engine) node0.fpc0. Please resync before commit** is displayed if too many policies and addresses are configured. [PR1355528](#)
- The PIM register may stop the message from the source First Hop Router (FHR). [PR1356241](#)
- On SRX5000 devices, when the IPsec performance acceleration feature is enabled, packets going in or out of a VPN tunnel are dropped. [PR1357616](#)
- On SRX Series devices, if you disable one of the four reth interfaces, the traffic flow stops. [PR1360399](#)
- On the secondary control plane, a multicast session leak is observed for the PIM register. [PR1360373](#)
- On SRX5400, SRX5600, and SRX5800 devices, the MIB walk tool is not working when screens are applied to the security zones. [PR1364210](#)
- When RG0 failover occurs, the flowd process core files are generated. [PR1366122](#)
- On SRX300, SRX320, SRX340, and SRX345 devices, with LTE mini-PIM the DHCP relay packets are not forwarded. [PR1357137](#)

General Routing

- The **Pred Fail** Fan Tray chassis alarm is renamed to **Predicted Fail**. [PR1202724](#)

- On SRX Series devices, if the memory buffer is accessed without checking the mbuf and the associated external storage, the flowd process may stop. [PR1353184](#)

Interfaces and Routing

- On SRX Series devices, when the VPLS interface receives a broadcast frame, the device sends this frame back to the sender. [PR1350857](#)
- The **set protocols rstp interface all** command does not enable RSTP on all interfaces. [PR1355586](#)

Intrusion Detection and Prevention (IDP)

- The file descriptor may leak during a security package auto update. [PR1318727](#)

J-Web

- In J-Web when you click the **SKIP TO JWEB OPTIONS**, the Google Chrome browser automatically redirects. [PR1284341](#)
- When the J-Web fails to get resource information, the Routing Engine CPU usage is displayed as 100 percent. [PR1351416](#)
- The J-Web setup does not propagate the DHCP attributes from ISP to LAN. [PR1370700](#)

Layer 2 Features

- The DCPFE/FXPC process might stop and generate a core file. [PR1362332](#)

Layer 2 Ethernet Services

- The subnet mask address is not sent as a reply to the **DHCPINFORM** request. [PR1357291](#)

Platform and Infrastructure

- When you perform commits with apply-groups, VPN may flap. [PR1242757](#)
- On SRX5400, SRX5600, and SRX5800 devices log messages are seen often when IOC card has the same identifier as the SPC PIC card. [PR1357913](#)
- On SRX4100 devices, interfaces are shown as half-duplex, but there is no impact on the traffic. [PR1358066](#)

Routing Policy and Firewall Filters

- The TCP protocol ports 5800 and 5900 are added to junos-defaults to support VNC application. [PR1333206](#)
- On SRX Series devices, a large scale commit, for example, 70,000 lines security policy may stop the NSD process on the Packet Forwarding Engine. [PR1354576](#)
- The timeout value of junos-http is not accurate. [PR1371041](#)

Routing Protocols

- On SRX Series devices, dedicated BFD does not work. [PR1347662](#)

Unified Threat Management (UTM)

- The default actions under Web filtering profile do not work as expected. [PR1365389](#)

VLAN Infrastructure

- On SRX Series devices in transparent mode, the flowd process may stop when matching the destination MAC. [PR1355381](#)

VPNs

- On SRX5400, SRX5600, and SRX5800 devices, the chassis cluster control link encryption does not work. [PR1347380](#)
- S2S tunnels are not redistributed after IKE or IPsec are reactivated in a configuration. [PR1354440](#)
- On SRX5600 and SRX 5800 devices, during VPN to AutoVPN configuration migration, traffic loss is observed. [PR1362317](#)

Resolved Issues: 18.1R2

API

- On SRX320-POE devices, the REST API does not work when the relevant configuration is added under the **system services rest** hierarchy. [PR1347539](#)

Application Layer Gateways (ALGs)

- On SRX5400, SRX5600, and SRX5800 devices, when you use the SIP ALG and have multiple local SIP servers with consecutive IP addresses, the SIP session distribution over the SPUs might not be optimal. [PR1337549](#)

Authentication and Access Control

- The **uacd** process is not stable after upgrading to Junos OS Release 12.3X48 release. [PR1336356](#)
- On SRX Series devices, **show version detail** command displays the following error message: **Unrecognized command (user-ad-authentication)** when configuring the USERIDD. [PR1337740](#)
- New configuration is available to configure the **web-authentication** timeout. [PR1339627](#)

Chassis Clustering

- The FPC module is offline at the secondary node, after the primary node or the secondary node is restarted. [PR1340116](#)
- On SRX5400, SRX5600, and SRX5800 devices with DC PEM installed on the device, the output of **show chassis environment pem** and **show chassis power** commands do not accurately reflect the actual value. [PR1323256](#)

- IP monitoring is not working as expected when one node is in secondary-hold and the primary node priority becomes 0. [PR1330821](#)
- On SRX Series devices, the integrated routing and bridging (IRB) interface on high availability does not send the ARP request after clearing ARP. [PR1338445](#)
- When a PPPoE interface is configured over an Aggregate Ethernet (AE) or redundant ethernet (RETH) interfaces, reboot of the cluster nodes might occur in some cases. [PR1341968](#)

Class of Service (CoS)

- Packets are out-of-order on the SRX5K-SPC-4-15-320 card (SPC2) cards with IOC1 or FIOC cards. [PR1339551](#)

Flow-Based and Packet-Based Processing

- The forwarding plane drops the packets, when J-Flow version 9 related configuration is removed. [PR1351102](#)
- On SRX Series devices, packet reorder might occur in traffic when using Point-to-Point protocol (PPP). [PR1340417](#)
- The flowd process might stop when the SYN-proxy function is configured. [PR1343920](#)
- File download halts over a period of time when the TCP proxy is activated through antivirus or Sky ATP. [PR1349351](#)
- On SRX1500, SRX4100, and SRX4200 devices, if the Sky ATP cloud feeds updates, the packet forwarding engine might stop causing intermittent traffic loss. [PR1315642](#)

Intrusion Detection and Prevention (IDP)

- On SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices, if IDP and SSL forward proxy whitelist are used together, the device might generate a core file. [PR1314282](#)
- Unable to load IDP policy because of less available heap memory. [PR1347821](#)

J-Web

- Unable to delete the dynamic VPN user configuration using J-Web. [PR1348705](#)

Platform and Infrastructure

- SRX5400, SRX5600, and SRX5800 devices, the message **No Port is enabled for FPC# on node0** is observed in the chassis process (**chassisd**) log for every 5 seconds. [PR1335486](#)
- SRX1500 devices might encounter a failure while accessing the SSD drive. [PR1345275](#)
- On SRX300 devices, the **show system firmware** command displays old firmware image. [PR1345314](#)
- On SRX Series devices, mandatory argument is missing for **show usp policy counters** command in RSI. [PR1341042](#)
- Simultaneous commit triggers the configuration integrity check failure and halts the SRX. [PR1332605](#)

Routing Policy and Firewall Filters

- On SRX Series devices, if you configure a huge number of custom applications in the policies, the flowd process might stop. [PR1347822](#)
- The log messages **L2ALM Trying peer/master connection, status 26** is displayed on all SRX Series devices. [PR1317011](#)
- The flowd process stops when AppQoS is configured on the device. [PR1319051](#)

Routing Protocols

- When BGP traceoptions are configured and enabled, the traces specific to the messages sent to the BGP peer (BGP SEND traces) are not logged, but the traces specific to the received messages (BGP RECV traces) are logged correctly. [PR1318830](#)
- OpenSSL Security Advisory, refer to <https://kb.juniper.net/JSA10851> for more information. [PR1328891](#)

Software Installation and Upgrade

- On SRX Series devices, if power loss occurs few seconds after commit and if the Trusted Platform Module is enabled, the configuration integrity fails. [PR1351256](#).

VPNs

- For FIPS: PKID, the syslog for key-pair deletion is required for conformance. [PR1308364](#)
- The kmd process might generate a core file when all the VPNs are down. [PR1336368](#)

Resolved Issues: 18.1R1

Application Layer Gateways (ALGs)

- On SRX Series devices, SIP packets might drop when SIP traffic performs destination NAT. [PR1268767](#)
- H323 ALG does not work correctly with static NAT and VR. [PR1303575](#)
- H323 ALG decode Q931 packet error is observed even after H323 ALG is disabled. [PR1305598](#)
- HTTP ALG is listed within **show security match-policies**, when the HTTP ALG does not exist. [PR1308717](#)
- On SRX Series devices with SIP ALG enabled, the SIP ALG might drop SIP packets which have a "referred-by" or "referred-to header" field containing multiple header parameters. [PR1328266](#)
- When SIP ALG is enabled and NAT is used, cores might be observed and then the device might reboot after the cores. [PR1330254](#)

Authentication and Access Control

- PFE might stop working, resulting in generation of huge number of core files in a short period of time. [PR1326677](#)
- JIMS server stops responding to requests from SRX Series devices. [PR1311446](#)

- On SRX Series devices, incomplete RSI might be seen. [PR1329967](#)
- On SRX Series devices, sessions might be closed because of **idle Timeout junos-fwauth-adapter**. [PR1330926](#)

Chassis Clustering

- The ISSU or ICU operation might fail if the upgrade is initiated from Junos Space on multiple SRX clusters. [PR1279916](#)
- Warning messages are tagged with error tag wrongly in the RPC response from an SRX Series device when you configure a change through netconf. [PR1286903](#)
- On SRX Series devices, if you are running the User Firewall feature, under some condition, core files are seen with the flow process or user identification process. The Packet Forwarding Engine is restarted, and RG1+ failover occurs. [PR1299494](#)
- Flowd process core files are generated after adding 65536 VPN tunnels using traffic selector with the same remote IP. [PR1301928](#)
- ISSU might be unsuccessful if the control link recovery is configured. [PR1303948](#)
- On SRX1500, SRX4100 and SRX4200 devices, ISSU might fail if LACP and interface monitoring are configured. [PR1305471](#)
- File Descriptor might leak on SRX Series chassis clusters with Sky ATP enabled. [PR1306218](#)
- After the device is rebooted, IP monitoring on secondary node shows unknown status. [PR1307749](#)
- In an active/active cluster, route change timeout does not work as expected. [PR1314162](#)
- When ISSU is performed from a Junos OS Release prior to 15.1X49-D60 to a Junos OS Release 15.1X49-D60 or later, flowd process generates core files. [PR1320030](#)
- When RG0 failover or primary node reboot happens, some of the logical interfaces might not be synchronized to the other node if the system has around 2000 logical interfaces and 40,000 security policies. [PR1331070](#)
- The default-gateway route received by DHCP when some interface in the chassis cluster has been configured as a DHCP client is lost in about 3 minutes after RG0 failover. [PR1334016](#)

Flow-Based and Packet-Based Processing

- On SRX4100 and SRX4200 devices, packet loss is observed when the value of packet per second (pps) through the device is very high. This occurs due to the update of the **application interval statistics** statement, which has a default timer value of 1 minute. You can avoid this issue by setting the interval to maximum using the **set services application-identification statistics interval 1440** command. [PR1290945](#)
- If SDNS proxy is configured on SRX Series devices, the naming process might stop. [PR1307435](#)
- When executing operations for creating rescue configuration, some errors are reported but the rescue configuration is created. [PR1280976](#)

- RPM packets not account through LT interface under certain configurations. [PR1303445](#)
- Packet capture does not work after the value of the **maximum-capture-size** option is modified. [PR1304723](#)
- The **show host server name-server host** CLI command fails when the source address is specified under the name-server configuration. [PR1307128](#)
- Clear session takes 9 minutes to clear 57 million sessions. [PR1308901](#)
- On SRX Series devices, if destination NAT and session affinity are configured with multiple traffic selectors in IPsec VPN, the traffic selector match might fail. [PR1309565](#)
- The flow process might stop and generate a core file during failover between node 0 and node 1. [PR1311412](#)
- On SRX Series devices, the IPsec tunnel might fail to be established if datapath debug configuration include the options **preserve-trace-order**, **record-pic-history**, or both. [PR1311454](#)
- The SRX Series device drops packets citing the reason "Drop pak on auth policy, not authed". [PR1312676](#)
- When you commit configuration changes involving deletion of routing-instance with application-tracking and session-close log enabled for the zone a PFE core file is generated. [PR1312757](#)
- The flow process might stop if the SSL-FP profile is configured with whitelist. [PR1313451](#)
- On SRX550M devices, phone-home.core is generated after the zeroization procedure. [PR1315367](#)
- On SRX Series devices, the PIM register stop comes before the PIM register packet. The out-of-sequence packet causes the flow session build error. [PR1316428](#)
- On SRX Series devices, the **fin-invalidate-sessio** command does not work when the Express Path feature is enabled on the device. [PR1316833](#)
- Return traffic through the routing instance might drop intermittently after changing the zone and routing-instance configuration on the st0.x interface. [PR1316839](#)
- SRX300 devices DHCP client cannot obtain IP addresses. [PR1317197](#)
- Default route is lost after system zero. [PR1317630](#)
- SSL firewall proxy does not work if root-ca has fewer than four characters. [PR1319755](#)
- Software next-hop table is full with log messages RT_PFE: NH IPC op 1 (ADD NEXTHOP) failed, err 6 (No Memory) peer_class 0, peer_index 0 peer_type 10. [PR1326475](#)
- The FPC is dropped or gets stuck in present state when intermittent control link heartbeats are seen. [PR1329745](#)
- The OSPF peers are unable to establish neighbors between the LT interfaces of the logical systems. [PR1319859](#)
- Flow process generates core files on both nodes causing an outage. [PR1324476](#)

- On the SRX5000 line of devices with an SRX5K-MPC3-40G10G (IOC3) or an SRX5K-MPC3-100G10G, the IPv6 traffic might be dropped if the IOC3 with the service-offload (npcache) feature is applied. [PR1331401](#)
- Inaccurate Jflow records might be seen for output interface and next hop. [PR1332666](#)
- The whitelist function in syn-flood does not work. [PR1332902](#)

Interfaces and Chassis

- LLDP protocol is not supported on a reth interface but it can be configured. [PR1127960](#)
- Traffic is looped with MSTP for untag traffic from IxNetwork ports. [PR1259099](#)
- Unable to add IRB and aggregated Ethernet interfaces. [PR1310791](#)
- On SRX1500 devices, pp0.0 interface link status is not up. [PR1315416](#)
- An error is not seen at each commit or commit check if autonegotiation is disabled but the speed and duplex configurations are not configured on the interface. [PR1316965](#)
- RSI uses incorrect **show vlans** syntax. [PR1336267](#)

Intrusion Detection and Prevention (IDP)

- On SRX4600 devices, the maximum SSLRP session count is observed to be approaching 100,000. In the CLI, configuring a maximum of 100,000 sessions are allowed, whereas in SSLFP, 600,000 sessions are allowed. Thus, the **set security idp sensor-configuration ssl-inspection sessions** command is now modified to allow a maximum of 600, 000 sessions. However, for other devices the original session limit value of 100,000 is retained. [PR1329827](#)
- IDP policy compilation can be triggered even if changes that are not related to IDP are performed. [PR1283379](#)
- IDP signatures might not get pushed to the Packet Forwarding Engine if there is a policy in logical systems. [PR1298530](#)
- On SRX Series devices, IDP PCAP feature underwent improvements such as:
 - The first valid packet-log-id will no longer be generated as '0' as this was not compatible with third party tools.
 - The algorithm for assigning packet-log-id's is improved to reduce the likelihood of duplicate entries and id-rollover events, particularly among devices with multiple SPU's.

[PR1297876](#)

J-Web

- J-Web system snapshot throws error. [PR1204587](#)
- J-Web does not display all global address book entries. [PR1302307](#)
- J-Web removes backslash character on source identity object when committing changes. [PR1304608](#)

- In J-Web, the zone drop-down does not list the available zones while creating the zone address book or sets with Internet Explorer IE 10 or 11. [PR1308684](#)
- J-Web authentication fails when a password includes the backslash. [PR1316915](#)
- J-Web dashboard displays wrong last updated time. [PR1318006](#)
- J-Web display problems for security policies are observed. [PR1318118](#)
- J-Web does not display wizards on the dashboard. [PR1330283](#)

Layer 2 Ethernet Services

- Duplicate hops or more than expected hop count is seen in Layer 2 traceroute. [PR1243213](#)
- Ping to VRRP(VIP) address failed when VRRP is on VLAN tagging. It only affects Trio-based IOC2 and IOC3 in SRX5000 line of devices. Other devices are not affected. [PR1293808](#)
- DHCPv6 prefix delegation does not start with the first available subnet. [PR1295178](#)
- In DHCP relay configuration, the option **VPN** has been renamed to **source-ip-change**. [PR1318487](#)
- DHCP rebind and renew packets is not calculated in BOOTREQUEST. [PR1325872](#)

Network Address Translation (NAT)

- SCTP packet has incorrect SCTP checksum after the SRX Series device implements NAT on the payload. [PR1310141](#)
- Active source NAT causes an NSD error and the session closes. [PR1313144](#)
- On SRX340 and SRX345 devices, configuring the source NAT pool larger than 1024 fails. [PR1321480](#)
- Arena utilization on a FPC spikes and then resumes to a normal value. [PR1336228](#)

Network Security

- On SRX Series devices, the Sky ATP connection leak causes the service plane to be disconnected from the Sky ATP cloud. [PR1329238](#)

Network Management and Monitoring

- DHCP packets are dropped by the dot1x module, if the port is a multiple-suppllicant port. [PR1296734](#)
- On SRX Series devices, the Routing Engine does not reply to an SNMP request. [PR1240178](#)
- SRX1500 devices might power-off unexpectedly because of incorrect device temperature readings, which reported very high temperature, leading to an immediate proactive powering -off of the device to protect the device from overheating. However, in these cases the temperature was not actually too high and a power-off would not be required. When this occurs, the following log message is shown in file `/var/log/hostlogs/lcmd.log`: `Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor. PR1241061`

- On SRX Series devices, when J-flow is enabled for multicast traffic, **extern nexthop** is installed during the multicast composite next hop. However, when you uninstall the composite next hop, it does not free the **extern nexthop**, which results in the jtree memory leak. [PR1276133](#)
- SRX300 device is unresponsive as a result of cf/var: filesystem full error. [PR1289489](#)
- CLI options are available to manage the packet forwarding engine handling the ARP throttling for NHDB resolutions. [PR1302384](#)

Platform and Infrastructure

- SRX Series devices do not process traffic because of an IPv6 NDP packets burst. [PR1293673](#)
- Inconsistent flow-control status on reth interfaces is observed. [PR1302293](#)
- On SRX5400, SRX5600, SRX5800 devices, SPC2 XLP stops processing packets in the ingress direction after repeated RSI collections. [PR1326584](#)
- On SRX5400, SRX5600, and SRX5800 devices, the packet captured by datapath-debug on an IOC2 card might be truncated. [PR1300351](#)
- When Security Intelligence (SecIntel) is configured, IPFD CPU utilization might be higher than expected. [PR1326644](#)

Routing Policy and Firewall Filters

- BGP traceoption logs are written even when it is deactivated. [PR1307690](#)
- The nsd process might stop responding when the name of a logical system is replaced. [PR1307876](#)
- The number of address objects per policy for SRX5400, SRX5600, SRX5800 devices is increased from 4096 to 16,000. [PR1315625](#)

Routing Protocols

- On SRX1500 devices, the IS-IS adjacency remains down when using an IRB interface. [PR1300743](#)
- Dedicated BFD does not work on SRX Series devices. [PR1312298](#)
- In a chassis cluster device with BMP configured, the rpd process might stop responding when the rpd process gracefully terminates. [PR1315798](#)

Software Installation and Upgrade

- The **request system reboot node in/at** command results in an immediate reboot instead of rebooting at the allotted time. [PR1303686](#)

Unified Threat Management (UTM)

- On SRX Series, if Sophos antispam or Sophos antivirus interfaces are in a routing-instance, the feature might not work as expected. [PR1311694](#)
- The ISSU upgrade might fail because of the generation of Packet Forwarding Engine core files. [PR1328665](#)

VPNs

- The IRB interface does not support VPN. [PR1166714](#)
- Output hangs while checking pki ca-certificate ca-profile-group details. [PR1276619](#)
- Next hop tunnel binding (NHTB) is not installed occasionally during rekey for VPN using IKEv1. [PR1281833](#)
- Traffic through tunnel fails without configuring th authentication algorithm under IPsec proposal on SRX1500 devices. SRX5600 it works correctly.[PR1285284](#)
- ADVPN tunnels flap with spoke error **no response ready yet**, this issue leads to IKEv2 timeout. [PR1305451](#)
- On SRX Series devices, core files are observed under certain conditions with VPN and when NAT-T is enabled. [PR1308072](#)
- SNMP for jnxIpSecTunMonVpnName does not work. [PR1330365](#)
- The kmd process core files might be seen when all the VPNs are down. [PR1336368](#)
- On SRX Series devices, ESP packet drops in IPsec VPN tunnels with NULL encryption algorithm configuration are observed. [PR1329368](#)

SEE ALSO

New and Changed Features 316
Changes in Behavior and Syntax 323
Known Behavior 324
Known Issues 326
Documentation Updates 341
Migration, Upgrade, and Downgrade Instructions 343
Product Compatibility 344

Documentation Updates

IN THIS SECTION

- [New Simplified Documentation Architecture | 342](#)

This section lists the errata and changes in Junos OS Release 18.1R1 for the SRX Series device documentation.

New Simplified Documentation Architecture

- With the release of Junos OS Release 18.1, Juniper is simplifying its technical documentation to make it easier for you to find information and know that you can rely on it when you find it. In the past, we organized documentation about Junos OS software features into platform-specific documents. In many cases, features are supported on multiple platforms, so you might not easily find the document you want for your platform.

With Junos OS Release 18.1, we have eliminated the platform-specific software feature documents. For example, if you want to find documentation on OSPF, there is only one document regardless of which platform you have. Here are some of the benefits of our new simplified architecture:

- Over time, you will see better search results when looking for Juniper documentation. You will be able to find what you want faster and be assured that is the right document.
- If a software feature is supported on multiple platforms, you can find information about all the platforms in one place.
- Because we have eliminated many documents that covered similar topics, you will now find one document with all the information.
- You can know that you are always getting the most current and accurate information.

SEE ALSO

[New and Changed Features | 316](#)

[Changes in Behavior and Syntax | 323](#)

[Known Behavior | 324](#)

[Known Issues | 326](#)

[Resolved Issues | 330](#)

[Migration, Upgrade, and Downgrade Instructions | 343](#)

[Product Compatibility | 344](#)

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 12.3X48, 15.1X49, 17.3, and 17.4 are EEOL releases. You can upgrade from Junos OS Release 15.1X49 to Release 17.3 or from Junos OS Release 15.1X49 to Release 17.4.

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

[New and Changed Features | 316](#)

[Changes in Behavior and Syntax | 323](#)

[Known Behavior | 324](#)

[Known Issues | 326](#)

[Resolved Issues | 330](#)

[Documentation Updates | 341](#)

[Product Compatibility | 344](#)

Product Compatibility

Hardware Compatibility

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

To determine the features supported on SRX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>

SEE ALSO

New and Changed Features	316
Changes in Behavior and Syntax	323
Known Behavior	324
Known Issues	326
Resolved Issues	330
Documentation Updates	341
Migration, Upgrade, and Downgrade Instructions	343

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.

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information

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


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

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. apps.juniper.net/compliance/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

26 August 2021—Revision 19, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 18, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 November 2020—Revision 17, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 October 2019—Revision 16, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 October 2019—Revision 15, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 June 2019—Revision 14, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 May 2019—Revision 13, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 April 2019—Revision 12, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 March 2019—Revision 11, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

19 February 2019—Revision 10, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 January 2019—Revision 9, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 January 2019—Revision 8, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 January 2019—Revision 7, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 December 2018—Revision 6, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 December 2018—Revision 5, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 November 2018—Revision 4, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 October 2018—Revision 3, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 September 2018—Revision 2, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 September 2018—Revision 1, Junos OS Release 18.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 August 2018—Revision 6, Junos OS Release 18.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 August 2018—Revision 5, Junos OS Release 18.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 July 2018—Revision 4, Junos OS Release 18.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 June 2018—Revision 3, Junos OS Release 18.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 June 2018—Revision 2, Junos OS Release 18.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

1 June 2018—Revision 1, Junos OS Release 18.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 April 2018—Revision 6, Junos OS Release 18.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

19 April 2018—Revision 5, Junos OS Release 18.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 April 2018—Revision 4, Junos OS Release 18.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 April 2018—Revision 3, Junos OS Release 18.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 April 2018—Revision 2, Junos OS Release 18.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 March 2018—Revision 1, Junos OS Release 18.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

Copyright © 2021 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.