

# Release Notes: Junos<sup>®</sup> OS Release 17.2R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

3 September 2020

<b>Contents</b>	<b>Introduction   11</b>
	<b>Junos OS Release Notes for ACX Series   11</b>
	<b>New and Changed Features   12</b>
	Release 17.2R3 New and Changed Features   13
	Release 17.2R2 New and Changed Features   13
	Release 17.2R1 New and Changed Features   13
	<b>Changes in Behavior and Syntax   15</b>
	General Routing   16
	Management   16
	Subscriber Management and Services   16
	<b>Known Behavior   17</b>
	High Availability (HA) and Resiliency   18
	<b>Known Issues   18</b>
	Layer 2 Features   19
	<b>Resolved Issues   19</b>
	Resolved Issues: 17.2R3   20
	Resolved Issues: 17.2R2   20
	Resolved Issues: 17.2R1   20
	<b>Documentation Updates   20</b>

## Migration, Upgrade, and Downgrade Instructions | 21

- Upgrade and Downgrade Support Policy for Junos OS Releases | 21

## Product Compatibility | 22

- Hardware Compatibility | 22

## Junos OS Release Notes for EX Series Switches | 23

### New and Changed Features | 23

- Release 17.2R3 New and Changed Features | 25

- Release 17.2R2 New and Changed Features | 25

- Release 17.2R1 New and Changed Features | 25

### Changes in Behavior and Syntax | 28

- General Routing | 29

- IP Tunneling | 29

- Management | 29

- Multicast | 30

- Network Management and Monitoring | 30

- Subscriber Management and Services | 31

### Known Behavior | 32

- General Routing | 32

- High Availability (HA) and Resiliency | 33

- Interfaces and Chassis | 33

### Known Issues | 34

- General Routing | 34

- Authentication and Access Control | 36

- EVPN | 36

- Infrastructure | 36

- Layer 2 Features | 36

- Network Management and Monitoring | 37

- Platform and Infrastructure | 37

- Virtual Chassis | 37

### Resolved Issues | 38

- Resolved Issues: 17.2R3 | 38

- Resolved Issues: 17.2R2 | 43

- Resolved Issues: 17.2R1 | 45

### Documentation Updates | 46

Migration, Upgrade, and Downgrade Instructions | 46

Upgrade and Downgrade Support Policy for Junos OS Releases | 47

Product Compatibility | 47

Hardware Compatibility | 48

Junos OS Release Notes for Junos Fusion Data Center | 48

New and Changed Features | 49

Release 17.2R3 New and Changed Features | 50

Release 17.2R2 New and Changed Features | 50

Release 17.2R1 New and Changes Features | 50

Changes in Behavior and Syntax | 64

Junos Fusion | 64

Known Behavior | 64

Junos Fusion Data Center | 65

Known Issues | 65

Resolved Issues | 66

Resolved Issues: 17.2R3 | 66

Resolved Issues: 17.2R2 | 67

Resolved Issues: 17.2R1 | 67

Documentation Updates | 67

Migration, Upgrade, and Downgrade Instructions | 68

Basic Procedure for Upgrading an Aggregation Device | 68

Preparing the Switch for Satellite Device Conversion | 70

Autoconverting a Switch into a Satellite Device | 72

Manually Converting a Switch into a Satellite Device | 75

Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion  
Topology | 78

Configuring Satellite Device Upgrade Groups | 79

Converting a Satellite Device to a Standalone Device | 80

Upgrade and Downgrade Support Policy for Junos OS Releases | 80

Downgrading from Release 17.2 | 81

Product Compatibility | 82

Hardware Compatibility | 82

## Junos OS Release Notes for Junos Fusion Enterprise | 83

### New and Changed Features | 84

Release 17.2R3 New and Changed Features | 84

Release 17.2R2 New and Changed Features | 84

Release 17.2R1 New and Changed Features | 84

### Changes in Behavior and Syntax | 86

#### Known Behavior | 86

Junos Fusion Enterprise | 87

#### Known Issues | 88

Junos Fusion Enterprise | 88

#### Resolved Issues | 89

Resolved Issues: 17.2R3 | 89

Resolved Issues: 17.2R2 | 90

Resolved Issues: 17.2R1 | 90

#### Documentation Updates | 90

#### Migration, Upgrade, and Downgrade Instructions | 91

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 91

Upgrading an Aggregation Device with Redundant Routing Engines | 93

Preparing the Switch for Satellite Device Conversion | 94

Converting a Satellite Device to a Standalone Switch | 95

Upgrade and Downgrade Support Policy for Junos OS Releases | 95

Downgrading from Release 17.2 | 96

#### Product Compatibility | 97

Hardware and Software Compatibility | 97

Hardware Compatibility Tool | 97

## Junos OS Release Notes for Junos Fusion Provider Edge | 98

### New and Changed Features | 98

Release 17.2R3 New and Changed Features | 99

Release 17.2R2 New and Changed Features | 99

Release 17.2R1 New and Changed Features | 99

### Changes in Behavior and Syntax | 100

#### Known Behavior | 100

#### Known Issues | 101

**Resolved Issues | 102****Resolved Issues: 17.2R3 | 102****Resolved Issues: 17.2R2 | 102****Resolved Issues: 17.2R1 | 102****Documentation Updates | 103****Migration, Upgrade, and Downgrade Instructions | 103****Basic Procedure for Upgrading an Aggregation Device | 104****Upgrading an Aggregation Device with Redundant Routing Engines | 106****Preparing the Switch for Satellite Device Conversion | 106****Converting a Satellite Device to a Standalone Device | 108****Upgrading an Aggregation Device | 110****Upgrade and Downgrade Support Policy for Junos OS Releases | 110****Downgrading from Release 17.2 | 111****Product Compatibility | 111****Hardware Compatibility | 111****Junos OS Release Notes for MX Series 5G Universal Routing Platforms | 112****New and Changed Features | 113****Release 17.2R3 New and Changed Features | 114****Release 17.2R2 New and Changed Features | 114****Release 17.2R1 New and Changed Features | 116****Changes in Behavior and Syntax | 144****Class of Service (CoS) | 145****EVPNs | 145****Forwarding and Sampling | 146****General Routing | 148****High Availability (HA) and Resiliency | 148****Interfaces and Chassis | 148****IP Tunneling | 151****Management | 151****MPLS | 152****Network Management and Monitoring | 153****Routing Protocols | 155****Services Applications | 157****Software-Defined Networking | 157**

Software Installation and Upgrade | 157  
Subscriber Management and Services | 157  
User Interface and Configuration | 160  
VPNs | 161

#### Known Behavior | 162

Flow-Based and Packet-Based Processing | 162  
General Routing | 162  
High Availability (HA) and Resiliency | 163  
Interfaces and Chassis | 163  
Network Management and Monitoring | 163  
Services Applications | 164  
Software-Defined Networking (SDN) | 165  
Software Installation and Upgrade | 165  
Subscriber Management and Services | 166  
User Interface and Configuration | 166

#### Known Issues | 167

Class of Service (CoS) | 167  
EVPN | 168  
Forwarding and Sampling | 169  
General Routing | 170  
High Availability (HA) and Resiliency | 175  
Infrastructure | 175  
Interfaces and Chassis | 176  
J-Web | 176  
Layer 2 Ethernet Services | 176  
Layer 2 Features | 177  
MPLS | 177  
Platform and Infrastructure | 179  
Routing Policy and Firewall Filters | 181  
Routing Protocols | 181  
Services Applications | 184  
Subscriber Access Management | 184  
User Interface and Configuration | 184  
VPNs | 184

**Resolved Issues | 185****Resolved Issues: 17.2R3 | 186****Resolved Issues: 17.2R2 | 232****Resolved Issues: 17.2R1 | 242****Documentation Updates | 249****Protocol Independent Routing Properties | 249****Subscriber Management Access Network Guide | 249****Subscriber Management Provisioning Guide | 250****Migration, Upgrade, and Downgrade Instructions | 250****Basic Procedure for Upgrading to Release 17.2 | 252****Procedure to Upgrade to FreeBSD 10.x based Junos OS | 252****Procedure to Upgrade to FreeBSD 6.x based Junos OS | 254****Upgrade and Downgrade Support Policy for Junos OS Releases | 256****Upgrading a Router with Redundant Routing Engines | 257****Downgrading from Release 17.2 | 257****Product Compatibility | 258****Hardware Compatibility | 258****Junos OS Release Notes for NFX Series | 259****New and Changed Features | 259****Release 17.2R3 New and Changed Features | 260****Release 17.2R2 New and Changed Features | 260****Release 17.2R1 New and Changed Features | 260****Changes in Behavior and Syntax | 265****Known Behavior | 265****Juniper Device Manager | 265****Known Issues | 266****Infrastructure | 266****IPSec | 266****Juniper Device Manager | 267****Junos Control Plane | 269****vSRX | 269****Resolved Issues | 270****Resolved Issues: 17.2R3 | 270****Resolved Issues: 17.2R2 | 270**

Resolved Issues: 17.2R1	270
Documentation Updates	271
Migration, Upgrade, and Downgrade Instructions	272
Upgrade and Downgrade Support Policy for Junos OS Releases	272
Basic Procedure for Upgrading to Release 17.2	272
Product Compatibility	276
Hardware Compatibility	276
Junos OS Release Notes for PTX Series Packet Transport Routers	277
New and Changed Features	277
Release 17.2R3 New and Changed Features	278
Release 17.2R2 New and Changed Features	278
Release 17.2R1 New and Changed Features	278
Changes in Behavior and Syntax	294
Forwarding and Sampling	295
General Routing	295
Interfaces and Chassis	295
Management	296
Network Management and Monitoring	297
Routing Protocols	298
Subscriber Management and Services	298
Known Behavior	299
Hardware	300
High Availability (HA) and Resiliency	300
Known Issues	301
General Routing	301
Interfaces and Chassis	304
Platform and Infrastructure	304
Routing Protocols	305
Resolved Issues	305
Resolved Issues: 17.2R3	306
Resolved Issues: 17.2R2	310
Resolved Issues: 17.2R1	311
Documentation Updates	312
Protocol-Independent Routing Properties	312



**Migration, Upgrade, and Downgrade Instructions | 312****Basic Procedure for Upgrading to Release 17.2 | 313****Upgrade and Downgrade Support Policy for Junos OS Releases | 315****Upgrading Using Unified ISSU | 316****Upgrading a Router with Redundant Routing Engines | 316****Product Compatibility | 317****Hardware Compatibility | 317****Junos OS Release Notes for the QFX Series | 318****New and Changed Features | 318****Release 17.2R3 New and Changed Features | 319****Release 17.2R2 New and Changed Features | 319****Release 17.2R1 New and Changed Features | 319****Changes in Behavior and Syntax | 338****Class of Service (CoS) | 339****General Routing | 340****EVPNs | 340****Interfaces and Chassis | 340****Network Management and Monitoring | 341****Management | 342****Routing Protocols | 343****Virtual Chassis | 343****Known Behavior | 344****EVPNs | 344****General Routing | 344****High Availability (HA) and Resiliency | 345****MPLS | 346****Layer 2 Features | 346****Routing Protocols | 346****Virtual Chassis | 346****Known Issues | 347****General Routing | 347****EVPN | 350****Interfaces and Chassis | 350****Layer 2 Features | 350**

Network Management and Monitoring	350
Platform and Infrastructure	350
Routing Protocols	350
Resolved Issues	351
Resolved Issues: 17.2R3	352
Resolved Issues: 17.2R2	359
Resolved Issues: 17.2R1	362
Documentation Updates	363
Migration, Upgrade, and Downgrade Instructions	364
Upgrading Software on QFX Series Switches	364
Installing the Software on QFX10002 Switches	366
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	366
Installing the Software on QFX10008 and QFX10016 Switches	368
Performing a Unified ISSU	372
Preparing the Switch for Software Installation	373
Upgrading the Software Using Unified ISSU	373
Product Compatibility	376
Hardware Compatibility	376
Upgrading Using ISSU	377
Compliance Advisor	377
Finding More Information	377
Requesting Technical Support	378
Self-Help Online Tools and Resources	378
Opening a Case with JTAC	379
Revision History	379

# Introduction

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

These release notes accompany Junos OS Release 17.2R3 for the ACX Series, EX Series, Junos Fusion Enterprise, Junos Fusion Data Center, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, and QFX Series. 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 | 12
- Changes in Behavior and Syntax | 15
- Known Behavior | 17
- Known Issues | 18
- Resolved Issues | 19
- Documentation Updates | 20
- Migration, Upgrade, and Downgrade Instructions | 21
- Product Compatibility | 22

These release notes accompany Junos OS Release 17.2R3 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

### IN THIS SECTION

- [Release 17.2R3 New and Changed Features | 13](#)
- [Release 17.2R2 New and Changed Features | 13](#)
- [Release 17.2R1 New and Changed Features | 13](#)

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

## Release 17.2R3 New and Changed Features

There are no new features or enhancements to existing features for ACX Series in Junos OS Release 17.2R3.

## Release 17.2R2 New and Changed Features

There are no new features or enhancements to existing features for ACX Series in Junos OS Release 17.2R2.

## Release 17.2R1 New and Changed Features

### *Hardware*

- **Support for fixed and tunable DWDM Optics, 1GE and 10GE BIDI Optics (ACX Series)**—Starting in Junos OS Release 17.2R1, ACX Series Universal Metro Routers support fixed and tunable 1-Gigabit Ethernet and 10-Gigabit Ethernet BIDI DWDM optics.

### *Interfaces and Chassis*

- **Support for Ethernet ring protection switching (ACX Series, ACX500, ACX5000)**—Starting in Junos OS Release 17.2R1, ACX Universal Metro Routers support Ethernet ring protection switching (G.8032v2). With the G.8032v2 capability, the ACX Series routers support manual commands (force switch, manual switch, and clear commands) and interconnection of multiple Ethernet rings without virtual channels. ERPS on the ACX5000 line of routers supports Aggregated Ethernet (AE) interfaces.

[See [Ethernet Ring Protection Switching Overview](#)]

### *Management*

- **Support for device family and release in Junos OS YANG modules (ACX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

[See [Understanding Junos OS YANG Modules](#).]

### *Network Management and Monitoring*

- **Support for sFlow agent (ACX5000)**—Starting in Junos OS Release 17.2R1, ACX5000 line of routers supports sFlow agent. sFlow is a statistical sampling based network monitoring protocol for high speed switched or routed networks. The sFlow monitoring system consists of an sFlow agent (embedded in a switch or router or in a stand alone probe) and a central data collector, or sFlow analyzer.

sFlow technology uses the following two sampling mechanisms:

- **Packet-based sampling**—Samples one packet out of a specified number of packets from an interface enabled for sFlow technology.

- Time-based sampling—Samples interface statistics at a specified interval from an interface enabled for sFlow technology.
- Adaptive sampling—Monitors the overall incoming traffic rate on the device and provides feedback to the interfaces to dynamically adapt their sampling rate to traffic conditions.

[See [Overview of sFlow Technology](#) and [Configuring sFlow Technology](#).]

### ***Operation, Administration, and Maintenance (OAM)***

- **Support for ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on aggregated Ethernet interfaces (ACX Series, ACX5000)**—Starting in Junos OS release 17.2R1, you can configure ITU-T Y.1731 standard-compliant Ethernet loss measurement (ETH-LM), Ethernet synthetic loss measurement (ETH-SLM), and Ethernet delay measurement (ETH-DM) capabilities on aggregated Ethernet (AE) interfaces. These performance monitoring functionalities are supported on ACX Series and ACX5000 line of routers.

[See [Understanding Ethernet OAM Link Fault Management for ACX Series Routers](#)]

### ***Routing Protocols***

- **Support for IS-IS flooding groups (ACX5000)**—Starting with Junos OS Release 17.2R1, you can configure flooding groups with IS-IS on the ACX5000 line of routers. This feature limits link-state PDU flooding over IS-IS interfaces. An LSP that is not self-originated is flooded only through the interface belonging to the flood group that has the configured area ID in the LSP. This helps minimize the routes and topology information, thus ensuring optimal convergence. You can segregate both level 1 and level 2 networks into flood groups by using area IDs as tags to identify a flood group. Configure interfaces with specific area IDs to modify the flooding behavior as per your requirements.

To enable IS-IS flooding groups, include the flood-group flood-group-area-ID statement at the [edit protocols isis interface] hierarchy level.

[See [Understanding IS-IS Flood Group](#)]

### ***Software Installation and Upgrade***

- **Support for In-Service Software Upgrade (ACX5000)**—Starting with Junos OS Release 17.2R1, Junos OS for ACX5000 Universal Metro Routers supports ISSU, the ability to do software upgrades between two different software releases with minimal disruption to network traffic and no disruptions in the control plane. As a prerequisite, you need to have the graceful Routing Engine switchover (GRES), nonstop active routing (NSR), and nonstop bridging (NSB) enabled in the routing engine to support ISSU on ACX5000 line of routers.

[See [Understanding In-Service Software Upgrade \(ISSU\) in ACX5000 Series Routers](#)]

### ***Timing and Synchronization***

- **Support for PHY timestamping in boundary clock mode (ACX Series)**—Starting in Junos OS Release 17.2R1, ACX Series Universal Metro Routers supports timestamping at the physical layer, also known

as PHY timestamping, in boundary clock mode. To enable PHY timestamping on ACX Series routers, configure **clock-mode** as boundary clock at the [edit protocols ptp] hierarchy level.

[See [Configuring Precision Time Protocol Clocking](#)]

- **Support for defect and event management and SNMP get and walk management for timing (ACX Series)**—Starting in Junos OS Release 17.2R1, the ACX Universal Metro Routers supports defect and event management capabilities for timing features. Defects and events are notified in the form of SNMP traps.

The ACX Universal Metro Routers also supports SNMP get, get-next, and walk management capabilities for the timing features. These capabilities are enabled through the PTP MIB and SyncE MIB objects.

[See [Understanding Timing Defects and Event Management on ACX Series](#) and [Understanding SNMP MIB for Timing on ACX Series](#)]

SEE ALSO

<a href="#">Changes in Behavior and Syntax</a>	<a href="#">15</a>
<a href="#">Known Behavior</a>	<a href="#">17</a>
<a href="#">Known Issues</a>	<a href="#">18</a>
<a href="#">Resolved Issues</a>	<a href="#">19</a>
<a href="#">Documentation Updates</a>	<a href="#">20</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">21</a>
<a href="#">Product Compatibility</a>	<a href="#">22</a>

## Changes in Behavior and Syntax

IN THIS SECTION

- [General Routing](#) | [16](#)
- [Management](#) | [16](#)
- [Subscriber Management and Services](#) | [16](#)

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 17.2R3 for the ACX Series Universal Metro Routers.

## General Routing

- **Support for deletion of static routes when the BFD session goes down (ACX Series)**—Starting with Junos OS 17.2R2, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. Therefore, the static routes are deleted when the BFD receives a session down message.

## Management

- **Junos OS YANG module namespace and prefix changes (ACX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. In earlier releases, Junos OS YANG modules used only a unique identifier to differentiate the namespace for each module, and the prefix for all **juniper-command** modules was **jrpc**.

Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**. The Junos OS YANG extension modules, **junos-extension** and **junos-extension-odl**, use the **junos** device family identifier in the namespace, but the modules are common to all device families.

[See [Understanding Junos OS YANG Modules](#).]

- **Changes to the rfc-compliant configuration statement (ACX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. If you configure the **rfc-compliant** statement at the `[edit system services netconf]` hierarchy level and request configuration data in a NETCONF session on a device running Junos OS Release 17.2 or later, the NETCONF server sets the default namespace for the **<configuration>** element in the RPC reply to the same namespace as in the corresponding YANG model.

[See [Configuring RFC-Compliant NETCONF Sessions](#) and [rfc-compliant](#).]

## Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (ACX Series)**—Starting in Junos OS Release 17.2R3, 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

<a href="#">New and Changed Features</a>	<a href="#">  12</a>
<a href="#">Known Behavior</a>	<a href="#">  17</a>
<a href="#">Known Issues</a>	<a href="#">  18</a>
<a href="#">Resolved Issues</a>	<a href="#">  19</a>
<a href="#">Documentation Updates</a>	<a href="#">  20</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  21</a>
<a href="#">Product Compatibility</a>	<a href="#">  22</a>

## Known Behavior

IN THIS SECTION

- [High Availability \(HA\) and Resiliency](#) | 18

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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.

High Availability (HA) and Resiliency

- **Residual and baseline statistics loss from ISSU**—Using ISSU to upgrade to Junos OS Release 17.2R1 or later will result in a loss of residual and baseline statistics for interfaces, interface set specific statistics, and BBE subscriber service statistics because of an update to the statistics database.

[See [Unified ISSU System Requirements](#).]

SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  12</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  15</a>
<a href="#">Known Issues</a>	<a href="#">  18</a>
<a href="#">Resolved Issues</a>	<a href="#">  19</a>
<a href="#">Documentation Updates</a>	<a href="#">  20</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  21</a>
<a href="#">Product Compatibility</a>	<a href="#">  22</a>

Known Issues

IN THIS SECTION

- [Layer 2 Features](#) | [19](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R3 for the ACX Series Universal Metro Routers.

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

## Layer 2 Features

- Under certain scenarios, if VPLS instances and Layer 3 NNI interfaces are deleted in the same commit, then a traffic duplication is observed for the VPLS traffic. To avoid such instances, it is recommended to delete or deactivate the Layer 3 NNI interfaces and VPLS instances in separate commits. [PR1260156](#)

### SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  12</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  15</a>
<a href="#">Known Behavior</a>	<a href="#">  17</a>
<a href="#">Resolved Issues</a>	<a href="#">  19</a>
<a href="#">Documentation Updates</a>	<a href="#">  20</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  21</a>
<a href="#">Product Compatibility</a>	<a href="#">  22</a>

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 17.2R3](#) | [20](#)
- [Resolved Issues: 17.2R2](#) | [20](#)
- [Resolved Issues: 17.2R1](#) | [20](#)

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: 17.2R3**

There are no fixed issues in the Junos OS Release 17.2R3 for ACX Series.

**Resolved Issues: 17.2R2**

There are no fixed issues in the Junos OS Release 17.2R2 for ACX Series.

**Resolved Issues: 17.2R1**

There are no fixed issues in the Junos OS Release 17.2R1 for ACX Series.

SEE ALSO

<a href="#">New and Changed Features   12</a>
<a href="#">Changes in Behavior and Syntax   15</a>
<a href="#">Known Behavior   17</a>
<a href="#">Known Issues   18</a>
<a href="#">Documentation Updates   20</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   21</a>
<a href="#">Product Compatibility   22</a>

**Documentation Updates**

There are no errata or changes in Junos OS Release 17.2R3 for the ACX Series documentation.

SEE ALSO

<a href="#">New and Changed Features   12</a>
<a href="#">Changes in Behavior and Syntax   15</a>
<a href="#">Known Behavior   17</a>
<a href="#">Known Issues   18</a>
<a href="#">Resolved Issues   19</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   21</a>
<a href="#">Product Compatibility   22</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 21

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Universal Metro 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 16.1, 16.2 and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

For more information about EEOL releases and to review a list of EEOL releases, see <https://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](#) | 12

---

[Changes in Behavior and Syntax](#) | 15

---

[Known Behavior](#) | 17

---

Known Issues   18
Resolved Issues   19
Documentation Updates   20
Product Compatibility   22

## Product Compatibility

### IN THIS SECTION

- Hardware Compatibility | 22

### 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   12
Changes in Behavior and Syntax   15
Known Behavior   17
Known Issues   18
Resolved Issues   19
Documentation Updates   20
Migration, Upgrade, and Downgrade Instructions   21

# Junos OS Release Notes for EX Series Switches

## IN THIS SECTION

- New and Changed Features | 23
- Changes in Behavior and Syntax | 28
- Known Behavior | 32
- Known Issues | 34
- Resolved Issues | 38
- Documentation Updates | 46
- Migration, Upgrade, and Downgrade Instructions | 46
- Product Compatibility | 47

These release notes accompany Junos OS Release 17.2R3 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

## IN THIS SECTION

- Release 17.2R3 New and Changed Features | 25
- Release 17.2R2 New and Changed Features | 25
- Release 17.2R1 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 17.2R3: EX4300, EX4600, and EX9200.

**NOTE:** In Junos OS Release 17.2R3, J-Web is supported on the 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 [Release Notes: J-Web Application Package Release 17.2A1 for EX4300 and EX4600 Switches](#).



## Release 17.2R3 New and Changed Features

### *Restoration Procedures and Failure Handling*

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (EX Series)**—In Junos OS Release 17.2R3, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, an automatic device recovery mode is activated if the system goes into amnesiac mode. In this process, the system automatically retries to boot with the saved rescue configuration. During this process, the system displays a banner "Device is in recovery mode" in the CLI (in both operational and configuration modes). In releases before Junos OS Release 17.2R3, there is no automatic process to recover from amnesiac mode. In those releases, a user with load and commit permission has to log in using the console and fix the issue in the configuration before the system reboots.

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

## Release 17.2R2 New and Changed Features

- There are no new features or enhancements to existing features for EX Series in Junos OS Release 17.2R2.

## Release 17.2R1 New and Changed Features

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

- **Authentication order with priority (EX4300 switches)**—Starting in Junos OS Release 17.2R1, you can configure EX4300 switches not to trigger re-authentication for a client that has been authenticated using MAC RADIUS authentication or captive portal authentication. If the switch receives an EAP-Start message from an authenticated client, the switch typically responds with an EAP-Request message, which triggers re-authentication using 802.1X authentication. You can use the **eapol-block** statement to configure the switch to ignore EAP-Start messages sent by a client that has been authenticated using MAC RADIUS authentication or captive portal authentication, and maintain the existing authentication session for the client.

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

- **Protected Extensible Authentication Protocol (PEAP) for MAC RADIUS authentication (EX4300 switches)**—Starting in Junos OS Release 17.2R1, you can configure the Protected Extensible Authentication Protocol (PEAP) as the authentication method for MAC RADIUS authentication. PEAP is a protocol that encapsulates EAP packets within an encrypted and authenticated Transport Layer Security (TLS) tunnel. The inner authentication protocol, used to authenticate the client's MAC address inside the tunnel, is the Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2). The encrypted exchange of information inside the tunnel ensures that user credentials are safe from eavesdropping.

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

## EVPNs

- **EVPN proxy ARP and ARP suppression (EX9200 switches)**—Starting with Junos OS Release 17.2R1, EX9200 switches that function as provider edge (PE) devices in an Ethernet VPN-MPLS (EVPN-MPLS) or Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) environment support proxy Address Resolution Protocol (ARP) and ARP suppression. The proxy ARP and ARP suppression capabilities are enabled by default. For both features to work properly, the configuration of an integrated and routing (IRB) interface on the PE device is required.

IRB interfaces configured on a PE device deliver ARP requests from both local and remote customer edge (CE) devices. When a PE device receives an ARP request from a CE device, the PE device searches its media access control (MAC)-IP address bindings database for the requested IP address. If the PE device finds the MAC-IP address binding in its database, it responds to the request. If the device does not find the MAC-IP address binding, it swaps the source MAC address in the request with the MAC address of the IRB interface on which the request was received and sends the request to all interfaces.

Even when a PE device responds to an ARP request, ARP packets might still be flooded across the WAN. ARP suppression prevents this flooding from occurring.

[See [EVPN Proxy ARP and ARP Suppression](#).]

## Layer 3 Features

- **Port-based LAN broadcast traffic forwarding (port helpers) for multiple destination servers (EX4300 switches and Virtual Chassis)**—Starting in Junos OS Release 17.2R1, you can configure *port helpers* on EX4300 switches and EX4300 Virtual Chassis on a per-port basis for multiple destination servers. Port helpers are port-based filters that listen on configured UDP ports for incoming LAN broadcast traffic, and forward those packets to configured destination servers as unicast traffic. Configure port helper filters using the **forwarding-options helpers port *port-number*** configuration statement with any of the following scopes:

- Global—Match incoming broadcast traffic on any interface for a configured port, and forward the traffic to the configured server:

```
set forwarding-options helpers port port-number server server-ip-address
```

- VLAN-specific—Match incoming broadcast traffic on an IRB interface for a configured port, and forward the traffic to the configured server:

```
set forwarding-options helpers port port-number interface irb-interface-name
server server-ip-address
```

- Interface-specific—Match incoming broadcast traffic on a Layer 3 interface for a configured port, and forward the traffic to the configured server:

```
set forwarding-options helpers port port-number interface interface-name
server server-ip-address
```

[See [Configuring Port-based LAN Broadcast Packet Forwarding](#).]

### Management

- **Support for device family and release in Junos OS YANG modules (EX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

[See [Understanding Junos OS YANG Modules](#).]

### Multicast

- **Support for static multicast route leaking for VRF and virtual-router instances (QFX5100 and EX4300 switches)**—Starting in Junos OS Release 17.2R1, you can configure your switch to share IPv4 multicast routes among different virtual routing and forwarding (VRF) instances or different virtual-router instances. On EX4300 switches, multicast route leaking is supported only when the switch functions as a line card in a Virtual Chassis, not as a standalone switch. Only multicast static routes with a destination-prefix length of /32 are supported for multicast route leaking. Only Internet Group Management Protocol version 3 is supported. To configure multicast route leaking for VRF or virtual-router instances, include the **next-table routing-instance-name.inet.0** statement at the **[edit routing-instances routing-instance-name routing-options static route destination-prefix/32]** hierarchy level. For **routing-instance-name**, include the name of a VRF or virtual-router instance. This feature was previously introduced in Junos OS Release 14.X53-D40.

[See [Understanding Multicast Route Leaking for VRF and Virtual-Router Instances](#).]

### Network Management and Monitoring

- **SNMP support for monitoring tunnel statistics (EX Series)**—Starting in Junos OS Release 17.2R1, SNMP MIB jnxTunnelStat supports monitoring of tunnel statistics for IPV4 over IPV6 tunnels. This is a new enterprise-specific MIB, Tunnel Stats MIB, that currently displays three counters: tunnel count in rpd, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine. This MIB can be extended to support other tunnel statistics. The MIB is defined in jnx-tunnel-stats.txt. This MIB is attached to jnxMibs.

### System Management

- **Dynamic power management (EX9200 switches)**—Starting with Junos OS Release 17.2R1, EX9200 switches support dynamic power management.

[See [System Services on EX9200 Switches](#).]

SEE ALSO

<a href="#">Known Behavior   32</a>
<a href="#">Known Issues   34</a>
<a href="#">Resolved Issues   38</a>
<a href="#">Documentation Updates   46</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   46</a>
<a href="#">Product Compatibility   47</a>

## Changes in Behavior and Syntax

### IN THIS SECTION

- [General Routing | 29](#)
- [IP Tunneling | 29](#)
- [Management | 29](#)
- [Multicast | 30](#)
- [Network Management and Monitoring | 30](#)
- [Subscriber Management and Services | 31](#)

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 17.2R3 for the EX Series.

## General Routing

- **Support for deletion of static routes when the BFD session goes down (EX Series)**—Starting with Junos OS 17.2R2, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

## IP Tunneling

- **Deprecated no-path-mtu-discovery configuration option for ipip6 tunnels**—Starting in Junos OS Release 17.2R1, the `no-path-mtu-discovery` configuration statement in the `[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]` and `[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]` hierarchies is no longer available for ipip6 tunnels.

## Management

- **Changes to the rfc-compliant configuration statement (EX Series)**—Starting in Junos OS Release 17.2R1, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. If you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level and request configuration data in a NETCONF session on a device running Junos OS Release 17.2R1 or later, the NETCONF server sets the default namespace for the `<configuration>` element in the RPC reply to the same namespace as in the corresponding YANG model.

[See [Configuring RFC-Compliant NETCONF Sessions](#) and [rfc-compliant](#).]

- **Junos OS YANG module namespace and prefix changes (EX Series)**—Starting in Junos OS Release 17.2R1, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each `juniper-command` module uses its own unique module name as the module's prefix. In earlier releases, Junos OS YANG modules used only a unique identifier to differentiate the namespace for each module, and the prefix for all `juniper-command` modules was `jrpc`.

Device families include `junos`, `junos-es`, `junos-ex`, and `junos-qfx`. The Junos OS YANG extension modules, `junos-extension` and `junos-extension-odl`, use the `junos` device family identifier in the namespace, but the modules are common to all device families.

[See [Understanding Junos OS YANG Modules](#).]

## Multicast

- **Support for per-source multicast traffic forwarding with IGMPv3 (EX4300)**—Starting in Junos OS Release 17.2R3, EX4300 switches forward multicast traffic on a per-source basis according to received IGMPv3 INCLUDE and EXCLUDE reports. In releases prior to these releases, EX4300 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](#).]

## Network Management and Monitoring

- **SNMP syslog messages changed (EX Series)**—In Junos OS Release 17.2R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
  - Old Message- **AgentX master agent failed to respond to ping. Attempting to re-register**  
New Message- **AgentX master agent failed to respond to ping, triggering cleanup!**
  - Old Message- **NET-SNMP version %s AgentX subagent connected**  
New Message- **NET-SNMP version %s AgentX subagent Open-Sent!**

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

- **Update to SNMP support of apply-path statement (EX Series)**—In Junos OS Release 17.2R1, SNMP implementation of the **apply-path** configuration statement supports only two lists:
  - **apply-path "policy-options prefix-list <list-name> <\*>"**  
This configuration has been supported from day one.
  - **apply-path "access radius-server <\*>"**  
This configuration is supported as of this release.
- **Change in default log level setting (EX Series)**—In Junos OS Release 17.2R3, the following changes are made to the default logging levels:  
Before this release:
  - SNMP\_TRAP\_LINK\_UP was LOG\_INFO for both the physical (IFD) and logical (IFL) interfaces.
  - SNMP\_TRAP\_LINK\_DOWN was LOG\_WARNING for both the physical and logical interfaces.
 From this release onward:
  - IFD LinkUp -> LOG\_NOTICE (as this is an important message but less frequent)
  - IFL LinkUp -> LOG\_INFO (no change)
  - IFD and IFL LinkDown -> LOG\_WARNING (no change)

See the [MIB Explorer](#).

- **Need to reconfigure SNMPv3 configuration after upgrade (EX4600)**—In Junos OS Release 17.2R1, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. This is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. In releases before Junos OS Release 17.2R1, you need to reconfigure SNMPv3 every time after a reboot. This problem is now fixed. If you upgrade, you must still reconfigure SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3 again. To reconfigure SNMP v3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters.

[See [Configuring the Local Engine ID](#).]

## Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (EX Series)**—Starting in Junos OS Release 17.2R3, 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 IA\_NA and 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	23
Known Behavior	32
Known Issues	34
Resolved Issues	38
Documentation Updates	46
Migration, Upgrade, and Downgrade Instructions	46
Product Compatibility	47

## Known Behavior

### IN THIS SECTION

- General Routing | 32
- High Availability (HA) and Resiliency | 33
- Interfaces and Chassis | 33

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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 EX4600 switches, Zero Touch Provisioning might take more than normal time (or a longer time) to complete because TFTP might take a long time to fetch the required data. [PR980530](#)
- On an EX4300 Virtual Chassis, when you perform a nonstop software upgrade (NSSU), there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On EX4300 switches, when 802.1X single-suplicant authentication is initiated, multiple "EAP Request Id Frame Sent" packets might be sent. [PR1163966](#)
- On EX4300 10-Gigabit Ethernet links, preexisting MACsec sessions might not come up after the following events: 1)The pfex or dot1x process restarts. 2)The system restarts and the link flaps. [PR1294526](#)



High Availability (HA) and Resiliency

- **Residual and baseline statistics loss from ISSU**—Using unified ISSU to upgrade to Junos OS Release 17.2R1 or later will result in a loss of residual and baseline statistics for interfaces, interface set specific statistics, and BBE subscriber service statistics because of an update to the statistics database.  
  
[See [Unified ISSU System Requirements](#).]
- During an NSSU on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading from and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)
- **ISSU restrictions**—Unified ISSU from Junos OS Release 17.2R1 to Junos OS Release 17.2R2 is not supported.

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 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](#)

SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  23</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  28</a>
<a href="#">Known Issues</a>	<a href="#">  34</a>
<a href="#">Resolved Issues</a>	<a href="#">  38</a>
<a href="#">Documentation Updates</a>	<a href="#">  46</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  46</a>
<a href="#">Product Compatibility</a>	<a href="#">  47</a>

## Known Issues

### IN THIS SECTION

- General Routing | 34
- Authentication and Access Control | 36
- EVPN | 36
- Infrastructure | 36
- Layer 2 Features | 36
- Network Management and Monitoring | 37
- Platform and Infrastructure | 37
- Virtual Chassis | 37

This section lists the known issues in hardware and software in Junos OS Release 17.2R3 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-Gbps 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](#)
- On EX Series Virtual Chassis that support PoE, when the master Routing Engine member is rebooted, PoE devices connected to the master might not come back online after the reboot. As a workaround to avoid this issue, when configuring PoE interfaces, use the **set poe interface all configuration** command instead of configuring specific interfaces individually. To recover connections after seeing this issue, disable and reenabling the ports affected by the issue. [PR1203880](#)
- 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** is received for a logical interface that does not exist with the following error message **ERROR-SPECIFIC INFO: nh\_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none**. As such , there is no service impact in MPC2 and MPC cards. [PR1205593](#)

- On an EX9200 switch with MC-LAG, when the **enhanced-convergence** statement is enabled, and when the kernel sends a next-hop message to the Packet Forwarding Engine, the full Layer 2 header is not sent and a packet might be generated with an invalid source MAC address for some VLANs. [PR1223662](#)
- When a configuration that takes a Packet Forwarding Engine offline and another configuration that brings the Packet Forwarding Engine back online are committed in quick succession, there could be a Routing Engine-Packet Forwarding Engine out-of-sync errors logged in the syslog. Most of the time these are benign errors, but sometimes they might result in Packet Forwarding Engine crashes. [PR1232178](#)
- On an EX Series router, if Dynamic Host Configuration Protocol (DHCP) relay or DHCP server is configured along with **bpdu-block**, a memory allocation issue might be seen. This issue can lead to a memory exhaustion for the DHCP process. [PR1259918](#)
- The EX4300 Virtual Chassis might fail to register some jnxOperating SNMP OIDs related to the Routing Engines. This behavior is more likely if Virtual Chassis members 0 and 1 (FPC0 and FPC1) are not selected as Routing Engines. [PR1368845](#)
- On EX4300-48MP, when regression scripts are run, the syslog error **Error in bcm\_port\_sample\_rate\_set(ifl\_cmd) : Reason Invalid port** appears. [PR1376504](#)

## Authentication and Access Control

- This PR is related to Auto-conversion of Network ports for Virtual Chassis ports feature. A network port is automatically converted to a Virtual Chassis port if the following conditions are met: 1) Two ports are connected back to back between two members in a Virtual Chassis 2) LLDP is enabled on the ports 3) Virtual Chassis is configured using Pre-provisioned mode. But, the conversion to VCP does not complete until the Virtual Chassis members are rebooted. This creates a situations where there could be loops caused by these ports. The command will internally not start the state machine for auto-conversion. This will prevent conversion of any further network ports after the command is configured. This configuration is persistent, across reboots, switchover, and restart of VCCPd. Once configured, the VCCPd will not trigger the exchange of 3-way handshake TLVs with the peer. The interface will remain in INIT state. If any interface is in midway of the conversion, this command has no effect on that and the conversion might be successful. Also, this command will not move the converted ports to Network ports. [PR1207566](#)

## EVPN

- When the ESI configuration on an interface is changed from **all-active** to **single-active**, and back to **all-active**, the EVPN split horizon label is not allocated and is shown as 0. [PR1307056](#)

## Infrastructure

- On EX4300 switches, if you configure a firewall filter policer with the **forwarding-class** action on an egress filter, Junos OS might allow the configuration to commit although that action is not supported. [PR1104868](#)
- In a VLAN swap case, the ARP packet processed at SFI contains the original dsa-tag (cvid), which is derived as an invalid hw-token. For this special case, the packet is sent to the kernel. The VLAN classification or regeneration for the invalid hw-token returns zero as the hw-token. [PR1342432](#)

## Layer 2 Features

- The eswd process might crash after a Routing Engine switchover is performed in an EX Series Virtual Chassis scenario. The crash occurs because of disordered processing of VLAN or VLAN members by eswd and L2PT modules. As the order of processing does not remain the same every time, the crash is random across switchovers. [PR1275468](#)
- 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 error is seen in the syslog every few minutes on the ERPS owner. Because the log is caused by ERPS PDUs in an ERPS setup, you can ignore the message. [PR1372422](#)

## Network Management and Monitoring

- The default syslog level is LOG\_NOTICE in the default configuration. SNMP\_TRAP\_LINK\_UP for the physical interface (IFD) was logged as LOG\_INFO from day one. To help debug physical link up issues, SNMP\_TRAP\_LINK\_UP events are now logged by default. [PR1287244](#)
- Trace files are not closed properly; as a result, writing of traceoptions to the log file suddenly stops. [PR1380764](#)

## Platform and Infrastructure

- On EX4300, EX4600, and QFX5100 switches, if a remote analyzer has an output IP address that is reachable through a route learned by BGP, the analyzer might be in a down state. [PR1007963](#)
- On all platforms running Junos OS, the **file copy** CLI command uses `/var/home/<user>` as a temporary staging directory for a nonroot user, and uses `/var/tmp` for the root user. When you issue the **file copy user@x.x.x.x:/dir/ /var/tmp/** CLI command to copy a file to the device, and if the file you are trying to transfer is larger than the temporary staging directory size, the copy operation might fail. [PR1195599](#)
- Every load override and rollback operation increases the refcount by 1. If the count reaches the maximum value of 65,535, the mgd process might crash and the session might be terminated. When mgd crashes, the active lock might remain up preventing any further commits. [PR1313158](#)
- On EX4300 switches, in an RSTP scenario, if you set a wrong bridge ID as the RSTP **bridge-id**, loops might be created in the network. [PR1383356](#)

## Virtual Chassis

- When the FPC in the linecard role is removed and rejoined to the Virtual Chassis immediately, the LAG interface on the master or backup is not reprogrammed in the rejoined FPC. [PR1255302](#)

## SEE ALSO

[New and Changed Features | 23](#)

[Changes in Behavior and Syntax | 28](#)

[Known Behavior | 32](#)

[Resolved Issues | 38](#)

[Documentation Updates | 46](#)

[Migration, Upgrade, and Downgrade Instructions | 46](#)

[Product Compatibility | 47](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 17.2R3 | 38](#)
- [Resolved Issues: 17.2R2 | 43](#)
- [Resolved Issues: 17.2R1 | 45](#)

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: 17.2R3

#### General Routing

- After access is rejected, the dot1x process might crash because of a memory leak. [PR1160059](#)
- An LCD corruption issue occurs while you are booting up EX Series switches. [PR1233580](#)
- The MACsec session cannot be recovered after physically flapping one link of an aggregated Ethernet interface. [PR1283314](#)
- The **show security macsec statistics** command does not show the expected results. [PR1283544](#)
- The EX4300-32F MACsec session stays down on 1-Gigabit and 10-Gigabit Ethernet links after certain events, when events are performed with traffic running. [PR1299484](#)
- The eswd process generates core files if **apply-groups** is configured under **interface-range**. [PR1300709](#)
- An l2ald crash might occur with no apparent trigger. [PR1302344](#)
- The **show snmp mib walk** command used for jnxMIMstMstiPortState does not display any output in Junos OS Release 17.1R2 on EX4600 switches. [PR1305281](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- PEM alarms and I2C failures are observed on MX240, MX480, and MX960 routers, EX9200 switches, and the SRX5000 line of devices. [PR1312336](#)
- The DHCP-security binding table might not get updated. [PR1312670](#)
- A memory leak is seen for dot1xd. [PR1313578](#)
- The vmcore might be seen and the device might reboot after the ICL is changed from an aggregated Ethernet interface to a physical interface. [PR1318929](#)

- The EX Series switches do not send RADIUS requests after the **interface-range** configuration is modified. [PR1326442](#)
- The major alarm about **Fan & PSU Airflow direction mismatch** might be seen when you remove the management cable. [PR1327561](#)
- Traffic going through an aggregated Ethernet interface might be dropped if mastership changes. [PR1327578](#)
- CoS is wrongly applied on the Packet Forwarding Engine leading to egress traffic drop. [PR1329141](#)
- The rpd generates core files on the new backup Routing Engine at **task\_quit,task\_terminate\_timer\_callback,task\_timer\_dispatch,task\_scheduler** after disabling NSR+GRES. [PR1330750](#)
- The STP BPDUs are not sent out on the other active child when the anchor FPC has no active child. [PR1333872](#)
- MQSS errors and alarms might occur with the interface going down. [PR1334928](#)
- The l2cpd process might crash in a VSTP scenario during Routing Engine switchover. [PR1341246](#)
- The statistics process pfed might generate core files on an upgrade between certain releases. [PR1346925](#)
- After an EX9200 FPC comes online, other FPC might increase the CPU usage to 100 percent and result in traffic loss for around 30 seconds. [PR1346949](#)
- The EX4600 switch detects a LATENCY OVER-THRESHOLD event with a wrong value. [PR1348749](#)
- The 40-Gigabit Ethernet might not forward traffic. [PR1349675](#)
- A commit error is observed if the device is downgraded from Junos OS Release 18.2 or Release 18.3 to Junos OS Release 17.3R3 [PR1355542](#)
- On EX4300-48MP, the 802.1x protocol subsystem is taking a longer time to respond to management requests, and the error **the dot1x-protocol subsystem is not responding to management requests**. [PR1361398](#)
- Unexpected **DCD\_PARSE\_ERROR\_SCHEDULER** messages are logged when an MS-MPC or MS-MIC is taken offline or brought online. [PR1362734](#)
- The l2cpd process might crash when MVRP is being configured with private VLAN and with the RSTP **interface all** option enabled. [PR1365937](#)
- MAC refresh packet might not be sent out from the new primary link after an RTG failover. [PR1372999](#)
- BOOTP packets might be dropped if **BOOTP-support** is not enabled at the global level. [PR1373807](#)
- NPC core files are generated when FPCs on the EX9200 line of switches reboot. [PR1374861](#)
- The dot1x does not work with the Microsoft NPS server. [PR1381017](#)

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

- The LLDP-MED cannot forward the correct PoE class. [PR1296547](#)

- The dot1x process might stop authenticating if continuous reauthentication requests clients cannot get processed. [PR1300050](#)
- The dot1xd process might generate core files if you configure the 802.1x interface with EAP-PEAP as an authentication protocol. [PR1322891](#)

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

- When **igmp-snooping** and **bpdu-block-on-edge** are enabled, IP multicast traffic sourced by the kernel, such as OSPF and VRRP traffic, gets dropped at the Packet Forwarding Engine level. [PR1301773](#)

### **Infrastructure**

- Unable to provide management when the em0 interface of an FPC is connected to another FPC Layer 2 interface of the same Virtual Chassis. [PR1299385](#)
- The file system might be corrupted multiple times during an image upgrade or when an operation is committed. [PR1317250](#)
- The upgrade might fail if as a result of file system corruption due to the presence of bad blocks in the flash drive or file system. [PR1317628](#)
- The PFC feature might not work on EX4600. [PR1322439](#)
- The ifinfo process might generate core files on EX4600 Virtual Chassis. [PR1324326](#)
- Support for archiving dmesg file. [PR1327021](#)

### **Interfaces and Chassis**

- On EX4300 Virtual Chassis, an LACP flap is observed after rebooting the master FPC with PDT configurations. [PR1301338](#)
- The interface might not work properly after the FPC restarts. [PR1329896](#)
- The MAC address assigned to an aggregated Ethernet member interface is not the same as that of its parent aggregated Ethernet interface upon master node removal. [PR1333734](#)
- On EX4600, the MC-LAG after reboot of the VRRP master and backup discards the traffic to the downstream switches. [PR1345316](#)
- The MC-LAG peer does not send ARP requests to the host. [PR1360216](#)

### **Platform and Infrastructure**

- The **interface-range** command cannot be used to set speed and autonegotiation properties for a group of interfaces. [PR1258851](#)
- The mismatch of VLAN IDs between a logical interface and a VLAN configuration might result in traffic to be discarded. [PR1259310](#)
- On EX4300 Virtual Chassis, a 10-Gigabit Ethernet VCP might not get a neighbor after a system reboot. [PR1261363](#)



- The IRB interface does not turn down when the master switch is rebooted or halted. [PR1273176](#)
- The CPU utilization for pfex\_junos usage might go high if DHCP relay packets are received continuously. [PR1276995](#)
- Traffic loss might be observed for about 10 seconds if the master member FPC reboots. [PR1283702](#)
- Issuing the **load replace terminal** CLI command and attempting to replace the interface statements might terminate the current CLI session and leave your session hanging. [PR1293587](#)
- Some packets might be dropped after GRE encapsulation on EX4300. [PR1293787](#)
- The **ERROR: /dev/da0s1a is not a JUNOS snapshot** error is seen during the system startup. [PR1297888](#)
- On EX4300 switches, when unknown unicast ICMP packets are received by an interface, packets are routed, so TTL is decremented. [PR1302070](#)
- The FRU PSU removal and insertion traps might not get generated. [PR1302729](#)
- The unknown IPv6 multicast traffic is dropped if **mld-snooping** is enabled. [PR1304345](#)
- Inconsistent IEEE P-bit marking in 802.1Q header for OSPF packets. [PR1306750](#)
- The multicast receiver connected to an EX4300 switch might not be able to get the multicast streaming. [PR1308269](#)
- Traceroute not working in an EX9200 device for routing instances running on Junos OS Release 17.1R3. [PR1310615](#)
- Autonegotiation not working as expected between EX4300 and SRX5800. [PR1311458](#)
- IGMP snooping might not learn the multicast router interface dynamically. [PR1312128](#)
- The interface with 1-Gigabit SFP transceiver might go down if **no-auto-negotiation** is configured. [PR1315668](#)
- IGMPv3 on EX4300 does not have the correct outgoing interfaces in the Packet Forwarding Engine that are listed in the kernel. [PR1317141](#)
- The l2cpd might generate core files if the interface is disabled under VSTP and enabled under RSTP. [PR1317908](#)
- High latency might be observed between the master Routing Engine and other FPCs. [PR1319795](#)
- VLAN might not be processed, which leads to improper STP convergence improperly. [PR1320719](#)
- Multicast traffic might not be forwarded to one of the receivers. [PR1323499](#)
- A MAC learning issue and failure to create VLANs might be experienced by some VLANs on the EX4300. [PR1325816](#)
- The l2cpd might generate a core file. [PR1325917](#)
- Extra EAP request packets might be sent unnecessarily. [PR1328390](#)
- The SNMP trap message are always sent out with the log message **Fan/Blower OK** on an EX4300-VC switch. [PR1329507](#)

- When the TCAM table is being exhausted, the filter might be incorrectly programmed. [PR1330148](#)
- The EX4300 does not generate l2ald storm control action logs if the interface has the RTG configuration. [PR1335256](#)
- IGMP packets are forwarded out of RTG backup interfaces. [PR1335733](#)
- An l2cpd memory leak is seen on EX Series platforms with VoIP configured. [PR1337347](#)
- The **show spanning-tree statistics bridge** command output displays 0 for all VLAN instance IDs. [PR1337891](#)
- The MAC source address filter with **accept-source-mac** does not work if the MAC move limit is configured. [PR1341520](#)
- MSTP might not work normally after permitting commit. [PR1342900](#)
- The filter might not be programmed in the Packet Forwarding Engine even though TCAM entries are available. [PR1345296](#)
- Packet drop might be seen on the logical tunnel interfaces lt-x/2/x or lt-x/3/x. [PR1345727](#)
- On EX4300 or EX4600 switches, the VLAN translation feature does not work for the control plane traffic. [PR1348094](#)
- On EX4300 switches, traffic drop might happen if LLC packets are received with DSAP and SSAP as 0x88 and 0x8e, respectively. [PR1348618](#)
- Running RSI through the console port might cause the system to crash and reboot. [PR1349332](#)
- On EX4300 switches (standalone and Virtual Chassis) running Junos OS Release 16.1R5 or 16.1R6, the firewall filter with the **syslog** option is unable to send syslog messages to the syslog server. [PR1351548](#)
- 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 SFP-T transceiver might still be up after a system halt. [PR1354857](#)
- The FPC might crash because of the memory leak caused by the VTEP traffic. [PR1356279](#)
- Interface flapping is seen on EX4300 switches. [PR1361483](#)
- On EX4300 and EX4600 switches, the l2ald process might crash in an 802.1x scenario. [PR1363964](#)
- The Packet Forwarding Engine might crash if MAC move is encountered frequently. [PR1367141](#)
- The LLDP TLV might be sent with the wrong switch port capabilities. [PR1372966](#)
- Traffic is discarded silently with indirect next hop and load balancing. [PR1376057](#)
- The IRB interface does not go down when the master chassis is rebooted or halted. [PR1381272](#)

### **Routing Protocols**

- The mcsnoopd process generates a core file at `__raise,abort,__task_quit,__task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal` and `(enable_slip_detector=true, no_exit=true)` at `../../../../src/junos/lib/libtask/base/task_scheduler.c:275`. [PR1305239](#)
- The OSPF routes cannot be installed to the routing table until the **lsa-refresh** timer expires. [PR1316348](#)
- The BGP peer is not established after Routing Engine switchover when BFD is enabled and a graceful restart is performed. [PR1324475](#)
- IGMP snooping might be enabled unexpectedly. [PR1327048](#)

### **Virtual Chassis**

- On EX4300 FRU, the removal insertion trap is not generated for nonmaster (backup or line card) FPCs. [PR1293820](#)

## **Resolved Issues: 17.2R2**

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

- On QFX5100, EX4300, or EX4600, traffic might be dropped when there is more than one forwarding class under the **[forwarding-class-sets]** hierarchy. [PR1255077](#)

### **General Routing**

- Clients not getting IP addresses or ports are programmed under an incorrect VLAN. [PR1230073](#)
- The FPC might encounter errors and stop forwarding traffic. [PR1249375](#)
- EX9200: EVPN active/active ARP is not resolving on hosts. [PR1267769](#)
- After MACsec link flaps, traffic stops forwarding across the MACsec link. [PR1269229](#)
- The l2ald memory might leak for every IPv6 ND message it receives from peer the MC-LAG, and it does not free the memory allocated. [PR1277203](#)
- An l2ald crash occurs with no apparent trigger. [PR1302344](#)

### **Infrastructure**

- On an EX4300 egress VLAN-based firewall filter on a Q-in-Q interface, after a switch reboot, firewall counters might not increment as expected. [PR1165450](#)
- The EX4300 aggregated interface goes down when the interface member VLAN is PVLAN and LACP is enabled. [PR1264268](#)

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

- An interface explicitly disabled under RSTP is blocked under some conditions. [PR1266035](#)

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

- EX4300 running Junos OS Release 17.1R1 cannot be converted on satellite mode. [PR1267767](#)
- With **show ethernet-switching table** a few entries are stuck in DLR state after I2-learning restart. [PR1268619](#)
- VRRP split brain in dual access device Junos Fusion. [PR1293030](#)
- An access device without a cascade port cannot reach hosts over ICL link if they are authenticated by dot1x in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)

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

- Layer 3 protocol packets are not being sent out from the switch. [PR1226976](#)
- Preboot Execution Environment (PXE) unicast ACK packet is dropped on EX4300. [PR1230096](#)
- Traffic is not forwarded through GRE tunnel on EX4300 in some cases. [PR1254638](#)
- Unexpected Packet Forwarding Engine manager (pfex) restart is seen on RE switchover. [PR1258863](#)
- The mismatch of vlan-id between an interface IFL and VLAN config might result in traffic blackhole. [PR1259310](#)
- On the EX4300 Virtual Chassis, the FPC might crash and a pfex core file might be generated. [PR1261852](#)
- IPv6 neighbor solicitation messages are dropped when MLD snooping is enabled on EX4300. [PR1263535](#)
- The l2ald process might crash when many dot1x clients are being re-authenticated. [PR1269945](#)
- On EX4300, CPU usage related to pfex\_junos increases because of DHCP relay traffic. [PR1276995](#)

### ***Routing Protocols***

- The BGP session might flap during ISSU, resulting in 40-50 seconds of dropped traffic. [PR1247937](#)

### ***Virtual Chassis***

- When you add an EX4300 switch to the VCF, the following error message is seen: `?ch__map_alarm_id alarm ignored: object 0x7e reason?.` [PR1234780](#)

### ***VLAN Infrastructure***

- VLAN association is not being updated in the Ethernet switching table when the device is configured in single supplicant mode. [PR1283880](#)

## **Resolved Issues: 17.2R1**

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

- MPC might crash during ISSU from Junos OS Release 15.1R1 to a later release when QSFP/CXP/CFP2 optics are present. [PR1216924](#)

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

- After the rebooting of the Virtual Chassis, authentication of SNMPv3 users fails due to the change of the local engine ID. [PR1256166](#)

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

- The egress PE device (EX4300) sends out LLDP frames toward the CE device with the destination MAC address of 01:00:0c:cd:cd:d0, which is a duplicated frame and rewritten by the ingress (PE) device. [PR1251391](#)

### ***Port Security***

- On EX4600 switches and Virtual Chassis, MACsec connections are deleted randomly after a switch reboot, optics removal, deactivation or activation of a MACsec configuration, or fxpc process restart. [PR1234447](#)

### ***Routing Protocols***

- The BGP session might flap during ISSU, resulting in 40-50 seconds of dropped traffic. [PR1247937](#)

## **SEE ALSO**

---

[New and Changed Features | 23](#)

---

[Changes in Behavior and Syntax | 28](#)

---

[Known Behavior | 32](#)

---

<a href="#">Known Issues</a>	<a href="#">34</a>
<a href="#">Documentation Updates</a>	<a href="#">46</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">46</a>
<a href="#">Product Compatibility</a>	<a href="#">47</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 17.2R3 for the EX Series switches documentation.

### SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">23</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">28</a>
<a href="#">Known Behavior</a>	<a href="#">32</a>
<a href="#">Known Issues</a>	<a href="#">34</a>
<a href="#">Resolved Issues</a>	<a href="#">38</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">46</a>
<a href="#">Product Compatibility</a>	<a href="#">47</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [47](#)

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

## Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 16.1, 16.2 and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

### SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">23</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">28</a>
<a href="#">Known Behavior</a>	<a href="#">32</a>
<a href="#">Known Issues</a>	<a href="#">34</a>
<a href="#">Resolved Issues</a>	<a href="#">38</a>
<a href="#">Documentation Updates</a>	<a href="#">46</a>
<a href="#">Product Compatibility</a>	<a href="#">47</a>

## Product Compatibility

### IN THIS SECTION

- [Hardware Compatibility](#) | [48](#)

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

<a href="#">New and Changed Features</a>	<a href="#">  23</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  28</a>
<a href="#">Known Behavior</a>	<a href="#">  32</a>
<a href="#">Known Issues</a>	<a href="#">  34</a>
<a href="#">Resolved Issues</a>	<a href="#">  38</a>
<a href="#">Documentation Updates</a>	<a href="#">  46</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  46</a>

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

- [New and Changed Features](#) | [49](#)
- [Changes in Behavior and Syntax](#) | [64](#)
- [Known Behavior](#) | [64](#)
- [Known Issues](#) | [65](#)
- [Resolved Issues](#) | [66](#)
- [Documentation Updates](#) | [67](#)



- Migration, Upgrade, and Downgrade Instructions | 68
- Product Compatibility | 82

These release notes accompany Junos OS Release 17.2R3 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](https://www.juniper.net/documentation/product/en_US/junos-os)

## New and Changed Features

### IN THIS SECTION

- Release 17.2R3 New and Changed Features | 50
- Release 17.2R2 New and Changed Features | 50
- Release 17.2R1 New and Changes Features | 50

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

## Release 17.2R3 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Data Center in Junos OS Release 17.2R3.

## Release 17.2R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Data Center in Junos OS Release 17.2R2.

## Release 17.2R1 New and Changes Features

### *Junos Fusion Data Center*

- **Junos Fusion Data Center support (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center support is available and brings the Junos Fusion technology to data center networks. Junos Fusion Data Center uses QFX10000 switches in the aggregation device role and allows data center networks to combine numerous switches into a single, port-dense system. The system is managed from a single point (the aggregation devices) and simplifies network topologies because Junos Fusion Data Center is viewed as a single device by the larger network. Junos Fusion Data Center supports the 802.1BR standard.

You can configure the following QFX10000 Series switches as an aggregation device in a Junos Fusion Data Center:

- QFX10002 switches

You can configure the following switches as satellite devices:

- QFX5100 switches—QFX5100-24Q-2P, QFX5100-48S-6Q, QFX5100-48SH-6Q, QFX5100-48T-6Q, QFX5100-48TH-6Q, and QFX5100-96S-8Q
- EX4300 switches—EX4300-24T, EX4300-32F, EX4300-48T, and EX4300-48T-BF

[See [Junos Fusion Data Center Overview](#).]

- **Dual aggregation devices (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can have two aggregation devices in a Junos Fusion Data Center topology to support dual homing from satellite devices.

To configure a dual aggregation device topology, specify a chassis, redundancy group name and ID, peer chassis ID, and interchassis link interface in a redundancy group. All other ICCP parameters are automatically configured as part of the automatic ICCP provisioning of an interchassis link feature, which is enabled by default.

[See [Configuring the Dual Aggregation Device Topology](#).]

## Hardware

- **New satellite device models (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, the new QFX5100-48SH and QFX5100-TH switch models ship from the factory with preinstalled satellite software, allowing you to deploy them in a Junos Fusion Data Center in a plug-and-play manner.

[See [QFX5100 Switch Hardware Guide](#).]

## Class of Service (CoS)

- **Class of service support (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports the standard Junos class of service (CoS) features and operational commands in either a single or dual 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](#).]

## High Availability (HA) and Resiliency

- **Support for Virtual Routing Redundancy Protocol (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports the Virtual Routing Redundancy Protocol (VRRP). You can configure VRRP on dual aggregation devices to provide a common gateway for the hosts connected to the satellite devices and to provide dynamic switchover from one aggregation device to another in the event of failure. Both aggregation devices share the virtual IP address and route upstream packets independently. For protocol control, one of the aggregation devices is elected as the master and the other is placed in the backup role. To configure basic VRRP support, configure VRRP groups on the aggregated interfaces by including the **vrrp-group** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address ip-address]** hierarchy level.

[See [Understanding VRRP](#).]

## Interfaces

- **LACP support (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, LACP is supported on Junos Fusion Data Center. It provides the ability to bundle several physical interfaces to form one logical aggregated Ethernet interface. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode.

You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link

LACP is supported in single and dual aggregation device topologies.

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

- **Increased number of aggregated Ethernet interfaces (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can configure up to 1000 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.
- **Automatic ICCP provisioning of an interchassis link in a Junos Fusion (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, automatic ICCP provisioning of an interchassis link (ICL) simplifies configuration of a Junos Fusion with dual aggregation devices by automatically provisioning the ICCP configuration within the Junos Fusion, instead of requiring the user to manually configure all ICCP parameters.

The configuration of the redundancy group in a Junos Fusion using dual aggregation devices still requires that you specify a chassis, redundancy group name and ID, peer chassis ID, and interchassis link interface as part of the configuration process. All other redundancy group parameters are now automatically set to default values that do not have to be user-configured for a dual aggregation device topology to operate.

Automatic ICCP provisioning is enabled by default. If a user configures a redundancy group parameter that is set by default normally, the user configuration automatically overrides the default parameter. Automatic ICCP provisioning can be disabled by entering the **no-auto-iccp-provisioning** statement at the **[edit chassis satellite-management redundancy-groups redundancy-group-name peer-chassis-id peer-chassis-id-number]** hierarchy level.

[See [Understanding Automatic ICCP Provisioning and Automatic VLAN Provisioning of an Interchassis Link](#).]

**Automatic VLAN provisioning on an interchassis link in a Junos Fusion (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, automatic VLAN provisioning of an interchassis link (ICL) simplifies configuration of a Junos Fusion with dual aggregation devices by allowing the ICL interconnecting the dual aggregation devices to automatically detect all VLAN traffic on the Junos Fusion and seamlessly forward VLAN information between the aggregation devices over the ICL.

When automatic VLAN provisioning is disabled, you have to manually configure the supported VLANs on each ICL to ensure VLAN information is shared between aggregation devices.

Automatic VLAN Provisioning is enabled by default in a Junos Fusion Data Center, and can be disabled using the **set chassis satellite-management redundancy-groups *redundancy-group-name* peer-chassis-id *peer-chassis-id-number* no-auto-vlan-provisioning** statement.

Automatic VLAN Provisioning only works when the ICL is in trunk mode, and when the ICL interfaces are configured into **unit 0 family ethernet-switching**.

[See [Understanding Automatic ICCP Provisioning and Automatic VLAN Provisioning of an Interchassis Link](#).]

- **Configuration synchronization for MC-LAG (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion supports the ability to easily propagate, synchronize, and commit configurations from one MC-LAG peer to another MC-LAG peer. MC-LAG configuration synchronization enables you log into any one of the MC-LAG peers to manage both MC-LAG peers, thus having a single point of management. With MC-LAG configuration synchronization, you can use configuration groups to simplify the configuration process. For example, you can create configuration groups for the local MC-LAG peers, one for the remote MC-LAG peer, and one for the global configuration, which is essentially a configuration that is common to both MC-LAG peers. You can create conditional groups to specify when a configuration is synchronized with another MC-LAG peer. Additionally, you can include the **peers-synchronize** statement at the **[edit system commit]** hierarchy level to synchronize the configurations and commits across the MC-LAG peers by default. NETCONF over SSH provides a secure connection between the MC-LAG peers, and Secure Copy Protocol (SCP) copies the configurations securely between the MC-LAG peers.
- **Uplink port pinning (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, uplink port pinning allows traffic entering an extended port on a Junos Fusion Data Center to select which uplink port or ports are used to carry the traffic from the satellite device to the aggregation device. Uplink port pinning provides more deterministic traffic control by allowing you to select how traffic is forwarded from an extended port to an aggregation device.

When uplink port pinning is not enabled, traffic is forwarded from the satellite device to the aggregation device using all available uplink ports.

Uplink port pinning is configured in the following steps:

1. Create a forwarding policy in a satellite policy that includes an uplink port group by using the **port-group-extended** and **port-group-uplink** statements.
2. Associate the uplink port group with an extended port by configuring a port group alias with the **port-group-alias** statement.
3. Associate the forwarding policy with the Junos Fusion configuration using the **forwarding-policy** statement at the **[edit chassis satellite-management]** hierarchy level.

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

- **Uplink failure detection (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion enables satellite devices to detect link failures on the uplink interfaces used to connect to aggregation

devices. When a host device is multihomed to two satellite devices, and one of the uplink interfaces goes down, the host device can redirect traffic through the other active satellite device. All of the extended ports configured on the satellite device with the uplink interface failure are shut down.

By default, UFD is disabled. To enable UFD for all satellite devices, include the **uplink-failure-detection** statement at the **[edit chassis satellite-management]** hierarchy level. To enable UFD for specific satellite devices, include the **uplink-failure-detection** statement at the **[edit chassis satellite-management fpc]** hierarchy level.

EX4300 and QFX5100 switches configured as satellite devices have a default set of uplink interfaces. [Table 1 on page 54](#) shows the default set of uplink interfaces that UFD selects for failure detection:

**Table 1: UFD Default Uplink Interfaces for Satellite Devices**

Device Type	Default Uplink Interfaces
EX4300-24T (4 ports each on PIC1 and PIC2)	1/0 through 1/3 and 2/0 through 2/3
EX4300-32F	PIC 0 ports 32-35 PIC 1 ports 0-1 PIC 2 ports 0-7
EX4300-48T (4 ports each on PIC1 and PIC2)	1/0 through 1/3 and 2/0 through 2/3
EX4300-48T-BF (4 ports each on PIC1 and PIC2)	1/0 through 1/3 and 2/0 through 2/3
QFX5100-24Q-2P	PIC 0 ports 20-23
QFX5100-48S-6Q or QFX5100-48SH-6Q (6 QSFP+ ports)	0/48 through 0/53
QFX5100-48T-6Q or QFX5100-48TH-6Q (6 QSFP+ ports)	0/48 through 0/53
QFX5100-96S-8Q (8 QSFP+ ports)	0/96 through 0/103

If you choose not to use the default set of uplinks for your satellite devices, you need to specify which uplink interfaces you want to use for UFD. To apply UFD to an uplink interface, include the **ufd-default-policy** statement at the **[edit chassis satellite-management uplink-failure-detection]** hierarchy level. You also need to configure the UFD policy. For example:

```
[edit policy-options]
satellite-policy {
  candidate-uplink-port-profile {
    ufd-default-policy {
```

```

    term qfx5100 {
        product-model QFX5100*;
        uplink-port-group uplink-ports;
    }
}
port-group-alias {
    uplink-ports {
        pic 0 {
            port [1, 2];
        }
        pic 1 {
            port [3,4];
        }
    }
}

```

[See [Overview of Uplink Failure Detection on a Junos Fusion](#).]

- **Supported port types (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports the following port types:
  - Cascade port—Provides a connection to a satellite device. Cascade ports on an aggregation device connect to uplink ports on the satellite device.
  - Uplink port—Provides a connection to an aggregation device. Uplink ports on a satellite device connect to cascade ports on the aggregation device.
  - Extended port—Provides a connection to servers or endpoints. Extended ports are the physical interfaces of the satellite devices. The satellite devices appear as additional FPCs on the aggregation device in a Junos Fusion topology, and extended ports appear as additional interfaces to be managed by the aggregation device.
  - ICL port—Provides a connection between aggregation devices to support a dual-homed topology. ICL interfaces must be configured.

[See [Understanding Junos Fusion Ports](#).]

- **Enhanced interface commands (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center provides information for extended ports and uplink ports on satellite devices through operational mode commands and output. Extended port names include the extended FPC slot number, PIC slot, and port number. For example, a 10-Gigabit Ethernet extended port number might be xe-125/1/8, where 125 is the FPC slot number, 1 is the physical interface card (PIC) slot, and 8 is the extended port number.

The following commands have been enhanced to display the extended ports and uplink ports by using either the slot or the alias. Additionally, you can now use the keyword **satellite** to view information about the satellite device ports:

- **show interfaces satellite-device** (all | *alias*)
- **show interfaces extensive satellite-device** (all | *alias*)
- **show interfaces terse satellite-device** (all | *alias*)

### Layer 2 Protocols

- **Local switching on satellite devices (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can enable local Layer 2 switching at the satellite device level. In local switching mode, all bridging traffic for which the source and destination port are local to a satellite device is forwarded by that satellite device based on the destination MAC address. Each satellite device maintains only the local destination MAC addresses that are directly connected to the device in the bridge forwarding table. Any unknown MAC address on the satellite device is forwarded to the aggregation device for forwarding. To configure a satellite device in a Junos Fusion Data Center into local switching mode, include the **local-switching** statement at the **[edit forwarding-options satellite fpc fpc-slot-number]** hierarchy level on the aggregation device, where *fpc-slot-number* is the FPC slot ID of the satellite device.

[See [Configuring Local Switching on Junos Fusion Data Center](#).]

- **VLAN autosensing (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, VLAN autosensing allows extended ports on satellite devices to provision VLANs dynamically, as needed, to preserve the VLAN memory of the aggregation device with no or minimal impact to the forwarding of VLAN traffic in the Junos Fusion.

You configure VLAN autosensing from the aggregation device on a per-extended port basis by including the **vlan-auto-sense** statement at the **[edit interfaces interface-name unit logical-unit-number family ethernet-switching]** hierarchy level, where *interface-name* is the name of the extended port interface.

For example, to enable VLAN autosensing on extended port xe-101/0/0:

```
[edit]
user@aggregation-device# set interfaces xe-101/0/0 unit 0 family ethernet-switching
vlan-auto-sense
```

Configuration notes for VLAN autosensing:

- VLAN autosensing is supported on extended ports only.
- Only single VLAN tagged packets are autosensed.

[See [Understanding VLAN Autosensing](#).]

- **Loop detection on extended ports (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can configure a Junos Fusion Data Center system to detect and break loops of unicast traffic on downstream extended ports without configuring spanning tree protocols. Typically, the loops are caused



by either miswiring or by misconfiguration. Loop detection transmits special protocol data units (PDUs) periodically, and if a PDU is received on an extended port, the loop is detected and broken. Loop detection blocks the ingress port and issues a loop detection PDU error. When a port is blocked, you need to manually bring up the interface. Loop detection only responds to detect PDUs, not BPDUs.

[See [Understanding Loop Detection and Prevention on a Junos Fusion](#).]

- **Link Layer Discovery Protocol (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Link Layer Discovery Protocol (LLDP) is supported in a Junos Fusion Data Center. Link Layer Discover Protocol (LLDP) allows network devices to advertise their capabilities, identity, and other information onto a LAN. In a Junos Fusion topology, the LLDP protocol running on the satellite port is used for satellite device discovery and also works as a simple hello protocol between the satellite and aggregation devices to establish a two-way adjacency and detect remote-end failures.

[See [Understanding LLDP and LLDP-MED on Junos Fusion](#).]

- **MAC address synchronization (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, aggregation devices synchronize MAC addresses that are learned on the extended ports.

[See [Understanding MAC Address Synchronization in a Junos Fusion](#).]

- **VSTP enhancements (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, VSTP is supported on a QFX10000 switch acting as a single-homed aggregation device. The VSTP configuration can include native ports or extended ports in a Junos Fusion Data Center.
- **Loop detection with BPDU guard on VSTP edge ports (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports bridge protocol data unit (BPDU) protection for VLAN Spanning Tree Protocol (VSTP) on extended ports in a dual aggregation device topology. You can configure an extended port as a VSTP edge interface, and configure BPDU protection on the interface using the **bpdu-block-on-edge** statement. The exchange of BPDUs generated by VSTP prevents loops in network traffic by determining which interfaces block traffic and which interfaces forward traffic. If a BPDU is received on an edge interface with BPDU guard, VSTP will detect a loop and shutdown the interface. Other interfaces in the VLAN remain intact. To clear the interface for forwarding, issue the **clear error bpdu interface** command.

[See [bpdu-block-on-edge](#).]

### **Layer 3 Protocols**

- **Support for Layer 3 protocols (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, the following routing protocols supported on QFX10000 switches have been extended to the satellite devices in a Junos Fusion Data Center topology.

You can configure the following Layer 3 routing protocols on satellite device extended ports using a single aggregation device topology:

- BGP
- BGP for IPv6
- IS-IS

- IS-IS for IPv6
- OSPF
- OSPF version 3

### **Multicast Protocols**

- **Local egress replication for VLAN flooding (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, for a Junos Fusion topology with dual aggregation devices, you can enable egress replication (local replication) using the **local-replication** statement at the **[edit forwarding-options satellite]** hierarchy level. Local replication helps distribute packet replication load and reduce traffic on cascade ports for multicast and flooded VLAN traffic. When local replication is enabled, packet replication behavior for VLAN flooding is as follows:

- The aggregation device sends one copy of the packet to each satellite device that has extended ports in the VLAN.
- The satellite device does replication for each local port in the VLAN.

Use the **show ethernet-switching flood satellite** and **show ethernet-switching flood next-hops satellite** commands to view local replication information for flooded VLAN traffic.

[See [Egress Multicast Replication on the Satellite Devices.](#)]

- **Egress replication for Layer 2 multicast with IGMP Snooping (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, egress multicast replication, also called local replication, is supported for Junos Fusion topologies featuring dual aggregation devices. You can optionally configure local replication for all satellite devices by including the **local-replication** statement at the **[edit forwarding-options satellite]** hierarchy level. For Layer 2 multicast traffic with IGMP snooping configured and local replication enabled, the aggregation device sends only one copy of the packet to each satellite device that has an extended port in the multicast group, and the satellite device does the replication for its local ports that are members of the multicast group. When local replication is not enabled, Junos Fusion defaults to ingress replication, where all replication is done on the aggregation devices and sent to corresponding satellite devices for each extended port receiving the multicast traffic.

Use the following commands to display local replication information:

- **show ethernet-switching satellite device**
- **show multicast ecid-mapping satellite**
- **show multicast next-hops satellite**
- **show multicast snooping next-hops satellite**
- **show multicast snooping route satellite**
- **show multicast statistics satellite**
- **show multicast summary satellite**

Local replication is not compatible with port mirroring, VLAN ID tagging policies, and VPN configurations, and does not take effect (reverts to ingress replication behavior) for IPv6 traffic or Multicast Listener Discovery (MLD) snooping.

[See [Egress Multicast Replication on the Satellite Devices.](#)]

- **Egress replication for Layer 3 multicast IRB interface traffic (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, for a Junos Fusion topology with dual aggregation devices, you can enable egress multicast replication (also called local replication) using the **local-replication** statement at the **[edit forwarding-options satellite]** hierarchy level. Local replication helps distribute multicast packet replication load and reduce traffic on cascade ports, including for Layer 3 multicast traffic being routed between VLANs on IRB interfaces. When local replication is enabled, Layer 3 multicast packet replication behavior is as follows:
  - The aggregation device replicates the data for each IRB interface in the multicast group, and sends copies to each satellite device with member ports—one copy for each VLAN where the satellite device has destination extended ports in the VLAN.
  - Each receiving satellite device replicates the data for its local extended ports in the multicast group for each VLAN.

Local replication is not compatible with interfaces that use VLAN ID tagging policies that add processing overhead to forward egress traffic.

[See [Egress Multicast Replication on the Satellite Devices.](#)]

- **Multicast convergence improvements using enhanced PIM dual designated router mode for dual aggregation devices (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, enhanced PIM dual designated router mode is supported to improve multicast convergence time on a Junos Fusion with dual aggregation devices in the event of designated router (DR) failure and recovery. You can optionally enable this feature by including the **dual-dr enhanced** statement at the **[edit protocols pim interface interface-name]** hierarchy level. With enhanced PIM dual designated router mode enabled, although only one aggregation device is the primary device actively forwarding multicast traffic, both devices join the multicast tree and receive multicast data. As a result, if the primary aggregation device fails, the other aggregation device quickly takes over multicast replication and forwarding. You can enable this feature with egress multicast replication (local replication).

[See [Understanding Multicast Convergence Enhancements for Dual Aggregation Devices in a Junos Fusion.](#)]

- **Support for multicast protocols (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, many of the multicast protocols supported on QFX10000 switches have been extended to the satellite devices in a Junos Fusion topology. You can configure the following multicast protocols on satellite device extended ports:
  - IGMP
  - MLD

- PIM source-specific multicast (SSM)
- PIM sparse mode

### **Network Management and Monitoring**

- **Local port mirroring (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports local port mirroring. 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 local port mirroring to troubleshoot and monitor applications. You can mirror packets per port, and you can configure the source and mirror ports on the same satellite device.

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

- **Analyzers on extended ports (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can use port mirroring (analyzers) on extended ports on satellite devices in a Junos Fusion Data Center. Extended-port port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies 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, correlating events, and so on. When a port is ingress-mirrored, any packet received on that port is mirrored to the user-configured destination. When a port is egress-mirrored, any packet transmitted from that port is mirrored to your configured port-mirroring destination.

In Junos Fusion Data Center, you can use analyzers on extended ports for these purposes:

- Mirror aggregation device ports to extended ports
- Mirror extended ports to extended ports
- Mirror extended ports to aggregation device ports

[See [Understanding Port Mirroring on a Junos Fusion Data Center.](#)]

- **Junos Space Service Now (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center uses Service Now for failure event reporting. Service Now is an application that runs on the Junos Space Network Management Platform to automate fault management and accelerate issue resolution.

[See [Junos Space Service Now User Guide.](#)]

- **Chassis MIB support (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, satellite devices in a Junos Fusion topology are represented in the chassis MIB. Satellite devices are represented as FPC slots (100, 101, 102,...) in the aggregation device. The support is enabled using a range of container indexes, which enable the SNMP process to redirect SNMP requests to the chassis process or SPMD based on the first index entry.

The following tables have been implemented for satellite devices:

- jnxContainersTable
- jnxContentsTable

- `jnxFilledTable`
- `jnxOperatingTable`
- `jnxFRUTable`

alpha supply) is 102 for the power supply of the satellite device. Using these indexes, you can distinguish the satellite device hardware from the aggregation device hardware.

Chassis MIB support is available in single and dual aggregation device topologies.

[See [Chassis MIB Support \(Junos Fusion\)](#).]

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

- **Flow-based uplink selection (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1 on Junos Fusion Data Center, you can configure flow-based uplink selection for satellite devices to achieve better utilization of network resources. To remap specified elephant flows from satellite devices to aggregation devices, you program remapping on all or specific satellite devices to override the default 5-tuple hashing and then distribute those specified flows across uplinks toward aggregation devices. You define specific flows by using flow-based firewall filters statements, and those flows are sent to the uplink port or ports that you define.

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

### ***Storage***

- **Support for DCBX (Junos Fusion Data Center)**—Starting in Junos OS 17.2R1, Junos Fusion Data Center supports Data Center Bridging Capability Exchange Protocol (DCBX), including both DCBX v1.01 and IEEE DCBX. The Junos Fusion Data Center aggregation and satellite devices function as a single logical DCBX capable switch. Configuration for DCBX on Junos Fusion Data Center is performed on the aggregation device and is the same as on a standalone device.

The satellite device acts as a proxy for relaying DCBX messages from the aggregation device to the peer. In a dual-aggregation device setup, the satellite device automatically coordinates DCBX messages from both aggregation devices to relay to the peer, keeping the Junos Fusion Data Center appearing as a single device.

[See [Understanding DCBX](#).]

- **Support for PFC (Junos Fusion Data Center)** — Starting in Junos OS 17.2R1, Junos Fusion Data Center supports priority-based flow control (PFC) for Fibre Channel over Ethernet (FCoE) traffic. The Junos Fusion Data Center aggregation and satellite devices function as a single logical device. Configuration for PFC on Junos Fusion Data Center is performed on the aggregation device and is the same as on a standalone device.

[See [Example: Configuring CoS PFC for FCoE Traffic](#).]

## Software Installation and Upgrade

- **Upgrading and managing the satellite software on satellite devices (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion provides the ability to manage satellite software. To convert a standalone switch to a satellite device, you can use one of the following methods:
  - Autoconversion—Automatically converts a standalone device into a satellite device when it is cabled to a cascade port on the aggregation device.
  - Manual conversion—Installs the satellite software manually from the aggregation device when you issue the **request chassis satellite interface *interface-name* device-mode satellite** command.
  - Preconversion—Installs satellite software onto a device before connecting it to a Junos Fusion topology.

After you convert the switch to a satellite device, you can install satellite software upgrades onto a satellite device through the aggregation device.

**NOTE:** Before you can save satellite software images on a QFX10002 switch acting as an aggregation device, you must issue a one-time command to expand the storage capacity. To expand the storage area on the aggregation device, issue the [request system storage user-disk expand](#) command.

Satellite software upgrade groups are often needed to install satellite software. A satellite software upgrade group is a group of satellite devices that are designated to upgrade to the same satellite software version using the same satellite software package. When you add a satellite to an upgrade group that is not running the same satellite software, the satellite device is automatically updated to the version of satellite software associated with the upgrade group.

You can use the following commands to add and associate a satellite software version with an upgrade group:

- **request system software add upgrade-group**—Add the satellite software and associate it with the specified upgrade group.
- **request system software delete upgrade-group**—Remove the satellite software association from the specified upgrade group.
- **request system software rollback upgrade-group**—Associate an upgrade group with a previous version of satellite software.

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.

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

## Software Licensing

- **Licensing model (Junos Fusion Data Center)**—Starting with Junos OS Release 17.2R1, you need to install a Junos Fusion license in addition to any other feature licenses that you install to track and activate the

following models that are shipped with satellite software. These models can only be used as satellite devices:

- QFX5100-48SH-AFO
- QFX5100-48SH-AFI
- QFX5100-48TH-AFO
- QFX5100-48TH-AFI

**NOTE:** You do not need Junos Fusion licenses for satellite device models that were purchased as Junos OS-based top-of-rack switches.

You install these licenses on the aggregation device. Because the configurations are synchronized between aggregation devices, you only need to purchase one license and install it on one aggregation device regardless of whether you deploy a single or dual aggregation device topology. You can purchase a single-pack license to activate one satellite device, or you can purchase a multipack license to activate multiple satellite devices.

The following Junos Fusion Data Center SKUs are available for purchase:

- QFX10K-C1-JFS-1
- QFX10K-C1-JFS-4
- QFX10K-C1-JFS-8
- QFX10K-C1-JFS-16
- QFX10K-C1-JFS-32
- QFX10K-C1-JFS-64

You can issue the **request system add license**, **request system license delete**, and **request system license save** commands to manage your licenses. You can also issue the **show system license** command to display license information.

[See [Understanding Junos Fusion Licenses](#).]

## SEE ALSO

[Changes in Behavior and Syntax](#) | 64

[Known Behavior](#) | 64

[Known Issues](#) | 65

[Resolved Issues](#) | 66

---

[Documentation Updates | 67](#)

---

[Migration, Upgrade, and Downgrade Instructions | 68](#)

---

[Product Compatibility | 82](#)

## 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 in Junos OS Release 17.2R3 or later for Junos Fusion Data Center.

### Junos Fusion

- **Change in Junos Fusion operational mode syntax (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, the **slot-id** option has been replaced by **fpc-slot** in commands such as **show chassis satellite** and **show chassis environment satellite**. The **slot-id** option, although hidden, remains a valid option to provide backward compatibility for previous versions of Junos Fusion.

### SEE ALSO

---

[New and Changed Features | 49](#)

---

[Known Behavior | 64](#)

---

[Known Issues | 65](#)

---

[Resolved Issues | 66](#)

---

[Documentation Updates | 67](#)

---

[Migration, Upgrade, and Downgrade Instructions | 68](#)

---

[Product Compatibility | 82](#)

## Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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

- When a QFX10002 switch functions as an aggregation device in a Junos Fusion Data Center topology, it only supports cascade port-based slot assignments for satellite devices. In addition, any change in the configuration for a cascade port connected to a satellite device is treated as a catastrophic event and results in the deletion of any related interface state (including the extended ports), which is rebuilt after a period of time. The following additional restrictions also apply:
  - You cannot configure dual-homed satellite device extended ports as pure Layer 3 interfaces. As a result, **family inet** and **family inet6** are not supported on dual-homed extended ports.
  - If the ICL interface goes down, traffic loss will occur. As a workaround, we recommend you configure the ICL interface over an aggregated Ethernet interface with multiple links in the bundle to prevent single-point failures that would cause the ICL interface to shut down.
- On a Junos Fusion Data Center, configuring the following options for CoS forwarding class sets (fc-sets) incorrectly triggers a syslog message but does not result in any commit errors:
  - Priority of strict-high and normal (strict-high mixed with (low and high) queue) mixed in a single fc-set.
  - Total number of strict-high fc-sets configured is more than 1.
  - Transmit rate or guaranteed rate is configured on strict-high fc.

If the incorrect configuration is applied and the aggregation device is restarted, COSD does not start, and the CoS configuration is not sent to the Packet Forwarding Engine. The system will be in an inconsistent state.

SEE ALSO

<a href="#">New and Changed Features   49</a>
<a href="#">Changes in Behavior and Syntax   64</a>
<a href="#">Known Issues   65</a>
<a href="#">Resolved Issues   66</a>
<a href="#">Documentation Updates   67</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   68</a>
<a href="#">Product Compatibility   82</a>

Known Issues

There are no known issues in the Junos OS Release 17.2R3 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.

SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  49</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  64</a>
<a href="#">Known Behavior</a>	<a href="#">  64</a>
<a href="#">Resolved Issues</a>	<a href="#">  66</a>
<a href="#">Documentation Updates</a>	<a href="#">  67</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  68</a>
<a href="#">Product Compatibility</a>	<a href="#">  82</a>

## Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R3](#) | [66](#)
- [Resolved Issues: 17.2R2](#) | [67](#)
- [Resolved Issues: 17.2R1](#) | [67](#)

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: 17.2R3

- The LAG interface might flap if rebooting the aggregation device. [PR1315879](#)
- Duplicated packets might be received on the multicast downstream devices and multicast receivers. [PR1316499](#)

- The aggregate device might show **plus** sign on the ICL link for a satellite device. [PR1335373](#)
- The aggregation device LAG interface might flap during a satellite device upgrade or downgrade. [PR1321575](#)

**Resolved Issues: 17.2R2**

- Native VLAN on an aggregated Ethernet interface terminated on multiple satellite devices. [PR1305698](#)

**Resolved Issues: 17.2R1**

There are no fixed issues in the Junos OS Release 17.2R1 for Junos Fusion Data Center.

SEE ALSO

<a href="#">New and Changed Features   49</a>
<a href="#">Changes in Behavior and Syntax   64</a>
<a href="#">Known Behavior   64</a>
<a href="#">Known Issues   65</a>
<a href="#">Documentation Updates   67</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   68</a>
<a href="#">Product Compatibility   82</a>

**Documentation Updates**

There are no errata and changes in the current Junos Fusion Data Center documentation.

SEE ALSO

<a href="#">New and Changed Features   49</a>
<a href="#">Changes in Behavior and Syntax   64</a>
<a href="#">Known Behavior   64</a>
<a href="#">Known Issues   65</a>
<a href="#">Resolved Issues   66</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   68</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 68
- Preparing the Switch for Satellite Device Conversion | 70
- Autoconverting a Switch into a Satellite Device | 72
- Manually Converting a Switch into a Satellite Device | 75
- Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 78
- Configuring Satellite Device Upgrade Groups | 79
- Converting a Satellite Device to a Standalone Device | 80
- Upgrade and Downgrade Support Policy for Junos OS Releases | 80
- Downgrading from Release 17.2 | 81

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.

**NOTE:** For the latest information concerning which hardware and software to select for your Junos Fusion system, see [Junos Fusion Hardware and Software Compatibility](#).

### 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 17.2R3 is different from previous 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 access the Junos Fusion Hardware and Software Compatibility page.
4. Click the **Junos Fusion Data Center (QFX10000)** title to expand the list of supported releases.
5. Click the release number (the software version that you want to download) from the list.
6. Select the aggregation device 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, Canada, and worldwide, use the following command:

```
user@host> request system software add reboot  
source/jinstall-host-qfx-10-f-x86-64-17.2R3.13-secure-signed.tgz
```

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

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

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.

**NOTE:** After you install a Junos OS Release 17.2R3 **jinstall** package, you cannot issue the **request system software rollback** command to return to the previously installed software. Instead you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

## 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](#).

Customers with EX4300 switches, use the following command:

```
user@host> request system software add reboot
source/jinstall-ex-4300-14.1X53-D43.7-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add reboot
source/jinstall-qfx-5-14.1X53-D43.7-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 and set it to a factory-default state:

```
[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.

As an alternative, you can include the **auto-satellite-conversion** statement at the **[edit chassis]** hierarchy level on the target satellite device.

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.

### Autoconverting a Switch into a Satellite Device

Use this procedure to automatically configure a switch into a satellite device when it is cabled into the aggregation device.

You can use the autoconversion procedure to add one or more satellite devices to your Junos Fusion topology. The autoconversion procedure is especially useful when you are adding multiple satellite devices to Junos Fusion, because it allows you to easily configure the entire topology before or after cabling the satellite devices to the aggregation devices.

Before you begin:

- Ensure that your aggregation device is running Junos OS Release 17.2R1 or later, and that the satellite devices are running Junos OS Release 14.1X53-D43 or later.

To autoconvert a switch into a satellite device:

1. Cable a link between the aggregation device and the satellite device, if desired.

**NOTE:** You can cable the aggregation device to the satellite device at any point in this procedure.

When the aggregation device is cabled to the satellite device during this procedure, the process for converting a switch into a satellite device to finalize this process occurs immediately.

If the aggregation device is not cabled to the satellite device, the process for converting a switch into a satellite device to finalize this process starts when the satellite device is cabled to the aggregation device.

2. Log in to the aggregation device.
3. Configure the cascade ports.

For example, to configure interface xe-0/0/1 on the aggregation device into a cascade port:

[edit]



```
user@aggregation-device# set interfaces xe-0/0/1 cascade-port
```

#### 4. Associate an FPC slot ID with each satellite device.

Examples:

- To configure the FPC slot ID of the satellite device that is connected to xe-0/0/1 to 101:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management fpc 101 cascade-ports xe-0/0/1
```

- To map FPC slot ID 101 to the satellite device using the serial number ABCDEFGHIJKL:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management fpc 101 serial-number
ABCDEFGHIJKL
```

- To map FPC slot ID to the satellite device using MAC address 12:34:56:AB:CD:EF:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management fpc 110 system-id
12:34:56:AB:CD:EF
```

#### 5. (Recommended) Configure an alias name for the satellite device:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management fpc slot-id alias alias-name
```

where *slot-id* is the FPC slot ID of the satellite device defined in the previous step, and *alias-name* is the alias.

For example, to configure the satellite device numbered 101 as qfx5100-48s-1:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management fpc 101 alias qfx5100-48s-1
```

#### 6. Configure an FPC slot ID into a software upgrade group.

For example, to add a satellite device with FPC number 101 to an existing software group named group1, or create a software upgrade group named group1 and add a satellite device with FPC slot 101 to the group:

```
[edit]
```

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

If you are creating a new software upgrade group in this step, you also need to associate the group with a satellite software image. You can skip this final step if the software upgrade group has been created and an association already exists.

Before associating a satellite software image with a satellite software group, the configuration with the satellite software upgrade group must be committed:

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

After committing the configuration, associate a satellite software image named **satellite-3.1R1.3-signed.tgz** to the upgrade group named **group1**:

```
user@aggregation-device> request system software add /var/tmp/satellite-3.1R1.3-signed.tgz
upgrade-group group1
```

**NOTE:** Before you can save satellite software images on a QFX10002 switch acting as an aggregation device, you must issue a one-time command to expand the storage capacity. To expand the storage area on the aggregation device, issue the [request system storage user-disk expand](#) command.

#### 7. Enable automatic satellite conversion:

```
[edit]
user@aggregation-device# set chassis satellite-management auto-satellite-conversion satellite
slot-id
```

For example, to automatically convert FPC 101 into a satellite device:

```
[edit]
user@aggregation-device# set chassis satellite-management auto-satellite-conversion satellite
101
```

#### 8. Commit the configuration:

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

The satellite software upgrade on the satellite device begins after this final step is completed, or after you cable the satellite device to a cascade port using automatic satellite conversion if you have not already cabled the satellite device to the aggregation device.

After the satellite software update, the switch operates as a satellite device in the Junos Fusion topology

## Manually Converting a Switch into a Satellite Device

Use this procedure to manually convert a switch into a satellite device after cabling it into the Junos Fusion topology.

This procedure should be used to convert a switch that is not currently acting as a satellite device into a satellite device. A switch might not be recognized as a satellite device for several reasons, including that the device was not previously autoconverted into a satellite device or that the switch had previously been reverted from a satellite device to a standalone switch.

Before you begin:

- Ensure that your aggregation device is running Junos OS Release 17.2R1 or later, and that the switches that will become satellite devices are running Junos OS Release 14.1X53-D43 or later.

To manually convert a switch into a satellite device:

1. Cable a link between the aggregation device and the satellite device.
2. Log in to the aggregation device.
3. Configure the link on the aggregation device into a cascade port, if you have not done so already.

For example, to configure interface xe-0/0/1 on the aggregation device into a cascade port:

```
[edit]
user@aggregation-device# set interfaces xe-0/0/1 cascade-port
```

4. Associate an FPC slot ID with the satellite device.

Examples:

- To configure the FPC slot ID of the satellite device that is connected to xe-0/0/1 to 101:

```
[edit]
user@aggregation-device# set chassis satellite-management fpc 101 cascade-ports xe-0/0/1
```

- To map FPC slot ID 101 to the satellite device using the serial number ABCDEFGHIJKL:

```
[edit]
user@aggregation-device# set chassis satellite-management fpc 101 serial-number  
ABCDEFGHIJKL
```

- To map FPC slot ID to the satellite device using MAC address 12:34:56:AB:CD:EF:

```
[edit]
```

```
user@aggregation-device# set chassis satellite-management fpc 110 system-id
12:34:56:AB:CD:EF
```

5. Configure the interface on the aggregation device into a software upgrade group.

For example, to add a satellite device with FPC number 101 to an existing software group named group1, or create a software upgrade group named group1 and add a satellite device configured with FPC number 101 to the group:

```
[edit]
user@aggregation-device# set chassis satellite-management upgrade-group group1 satellite
101
```

If you are creating a new software upgrade group in this step, you also need to associate the group with a satellite software image. You can skip this final step if the software upgrade group has been created and an association already exists.

Before associating a satellite software image with a satellite software group, the configuration with the satellite software upgrade group must be committed:

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

After committing the configuration, associate a satellite software image named **satellite-3.1R1.3-signed.tgz** to the upgrade group named group1:

```
user@aggregation-device> request system software add /var/tmp/satellite-3.1R1.3-signed.tgz
upgrade-group group1
```

**NOTE:** Before you can save satellite software images on a QFX10002 switch acting as an aggregation device, you must issue a one-time command to expand the storage capacity. To expand the storage area on the aggregation device, issue the [request system storage user-disk expand](#) command.

6. Manually configure the switch into a satellite device:

```
user@aggregation-device> request chassis satellite interface interface-name device-mode
satellite
```

For example, to manually configure the switch that is connecting the satellite device to interface xe-0/0/1 on the aggregation device into a satellite device:

```
user@aggregation-device> request chassis satellite interface xe-0/0/1 device-mode satellite
```

The satellite software upgrade on the satellite device begins after this final step is completed.

After the satellite software update, the switch operates as a satellite device in the Junos Fusion topology.

## Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology

Use this procedure to install the satellite software onto a switch before interconnecting it into a Junos Fusion topology as a satellite device. Installing the satellite software on a switch before interconnecting it to a Junos Fusion topology allows you to more immediately deploy the switch as a satellite device by avoiding the downtime associated with the satellite software installation procedure for Junos Fusion.

Before you begin:

- Ensure that your switch that will become a satellite device is running Junos OS Release 14.1X53-D43 or later.
- Ensure that you have copied the satellite software onto the device that will become a satellite device.

**NOTE:** Ensure there is sufficient space available in the `/var/tmp` directory to be able to copy the software to the switch (especially for EX4300 switches). If there is not enough memory available, issue the **request system storage cleanup** command on the device before attempting to perform the conversion.

In satellite software release 3.1R1, a `satellite-ppc-3.1R1.3-signed.tgz` package is included specifically for converting Junos OS to satellite on EX4300 to address a EX4300 switch space issue. The `satellite-ppc` package is to be used only for configuring a switch into a satellite device before connecting it to a Junos Fusion topology.

- You can manually install the satellite software onto a switch by entering the following command:

```
user@satellite-device> request chassis device-mode satellite URL-to-satellite-software
```

For instance, to install the satellite software package `satellite-3.1R1.3-signed.tgz` stored in the `/var/tmp/` directory on the switch:

```
user@satellite-device> request chassis device-mode satellite
/var/tmp/satellite-3.1R1.3-signed.tgz
```

- To install satellite software onto a QFX5100 switch, use the `satellite-3.1R1.3-signed.tgz` satellite software package.
- To install satellite software onto a EX4300 switch, use the `satellite-ppc-3.1R1.3-signed.tgz` satellite software package.

The device will reboot to complete the satellite software installation.

After the satellite software is installed, follow this procedure to connect the switch into a Junos Fusion topology:

1. Log in to the aggregation device.
2. Configure the link on the aggregation device into a cascade port, if you have not done so already.

For example, to configure interface xe-0/0/1 on the aggregation device into a cascade port:

```
[edit]
user@aggregation-device# set interfaces xe-0/0/1 cascade-port
```

3. Configure the satellite switch into a satellite software upgrade group that is using the same version of satellite software that was manually installed onto the switch.

This step is advisable, but not always required. Completing this step ensures that the satellite software on your device is upgraded to the version of satellite software associated with the satellite software upgrade group when the satellite device connects to the aggregation device.

4. Commit the configuration.

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

5. Cable a link between the aggregation device and the satellite device.

## 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/satellite-3.1R1.3-signed.tgz
```

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 16.1, 16.2 and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

## Downgrading from Release 17.2

To downgrade from Release 17.2 to another supported release, follow the procedure for upgrading, but replace the 17.2 **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

<a href="#">New and Changed Features   49</a>
<a href="#">Changes in Behavior and Syntax   64</a>
<a href="#">Known Behavior   64</a>
<a href="#">Known Issues   65</a>
<a href="#">Resolved Issues   66</a>
<a href="#">Documentation Updates   67</a>
<a href="#">Product Compatibility   82</a>

# Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 82

## Hardware Compatibility

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 on Junos Fusion devices, 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   49
Changes in Behavior and Syntax   64
Known Behavior   64
Known Issues   65
Resolved Issues   66
Documentation Updates   67
Migration, Upgrade, and Downgrade Instructions   68

# Junos OS Release Notes for Junos Fusion Enterprise

## IN THIS SECTION

- New and Changed Features | 84
- Changes in Behavior and Syntax | 86
- Known Behavior | 86
- Known Issues | 88
- Resolved Issues | 89
- Documentation Updates | 90
- Migration, Upgrade, and Downgrade Instructions | 91
- Product Compatibility | 97

These release notes accompany Junos OS Release 17.2R3 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](#) in the [Junos Fusion Enterprise User Guide](#).

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

## New and Changed Features

### IN THIS SECTION

- [Release 17.2R3 New and Changed Features | 84](#)
- [Release 17.2R2 New and Changed Features | 84](#)
- [Release 17.2R1 New and Changed Features | 84](#)

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

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

### Release 17.2R3 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 17.2R3.

### Release 17.2R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 17.2R2.

### Release 17.2R1 New and Changed Features

#### *Interfaces and Chassis*

- **Half-duplex link support on satellite devices (Junos Fusion Enterprise)**—Starting with Junos OS 17.2R1, half-duplex communication is supported on all built-in network copper ports on EX2300, EX3400, and EX4300 satellite devices in a Junos Fusion Enterprise (JFE). *Half-duplex* is bidirectional communication, but signals can flow in only one direction at a time. *Full-duplex* communication means that both ends of the communication can send and receive signals at the same time. The built-in network copper ports are configured by default as full-duplex 1-gigabit links with autonegotiation. If the link partner is set to autonegotiate the link, then the link is autonegotiated to full duplex or half-duplex. If the link is not set

to autonegotiation, then the satellite-device link defaults to half-duplex unless the interface is explicitly configured for full duplex.

To explicitly configure full duplex:

```
[edit]
```

```
user@aggregation-device# set interfaces interface-name link-mode full-duplex
```

To verify a half-duplex setting:

```
user@aggregation-device> show interfaces interface-name extensive
```

[See [Understanding Half-Duplex Links on Satellite Devices in a Junos Fusion Enterprise.](#)]

### Layer 2 Features

- **Private VLANs (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.2R1, Junos Fusion Enterprise (JFE) supports private VLANs (PVLANS). PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the known communication between known hosts. PVLANS can be used for purposes including: to help ensure the security of service providers sharing a server farm; to provide security to subscribers of various service providers sharing a common metropolitan area network; or to achieve isolation within the same subnet in a very large enterprise network. PVLAN is a standard introduced by RFC 5517 to achieve port or device isolation in a Layer 2 VLAN by partitioning a VLAN broadcast domain (also called a *primary VLAN*) into smaller subdomains (also called *secondary VLANs*).

In a JFE PVLAN topology:

- Multiple satellite devices can be clustered into a group and cabled into the JFE as a group instead of as individual satellite devices.
- Aggregation device native ports or satellite device extended ports can act as promiscuous ports, isolated ports, or community VLAN ports.
- The promiscuous port can be attached to a core switch or router through physical interfaces or aggregated links.
- PVLANS are supported in dual aggregation device JFEs.

[See [Understanding Private VLANs on a Junos Fusion Enterprise.](#)]

### SEE ALSO

---

[Changes in Behavior and Syntax](#) | 86

---

[Known Behavior](#) | 86

---

[Known Issues](#) | 88

---

---

[Resolved Issues | 89](#)

---

[Documentation Updates | 90](#)

---

[Migration, Upgrade, and Downgrade Instructions | 91](#)

---

[Product Compatibility | 97](#)

## 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 17.2R3 for Junos Fusion Enterprise.

### SEE ALSO

---

[New and Changed Features | 84](#)

---

[Known Behavior | 86](#)

---

[Known Issues | 88](#)

---

[Resolved Issues | 89](#)

---

[Documentation Updates | 90](#)

---

[Migration, Upgrade, and Downgrade Instructions | 91](#)

---

[Product Compatibility | 97](#)

## Known Behavior

### IN THIS SECTION

- [Junos Fusion Enterprise | 87](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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.

## Junos Fusion Enterprise

- On a Junos Fusion, when using LLDP, the **Power via MDI** and **Extended Power via MDI** TLVs are not transmitted. [PR1105217](#)
- In a Junos Fusion Enterprise topology with dual aggregation devices, firewall statistics are not synchronized across the aggregation devices. [PR1105612](#)
- On a Junos Fusion Enterprise, Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) fast start does not work. [PR1171899](#)
- On a Junos Fusion Enterprise, the **show ethernet-switching table** CLI command takes a few minutes to show entries when an extended port receives with MAC count set to 150K. [PR1117567](#)
- 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](#)
- While applying a loopback filter on aggregation devices in a Junos Fusion Enterprise, Callback Control Protocol (CBCP) packets might be filtered, which might cause CBCP sessions to be dropped and one of the satellite devices in a redundant pair to be in the SplitBrainDn state. To work around this issue, you can add a filter similar to the following to the existing set of loopback filters:

```
set firewall family inet filter accept-icl term accept-icl from source-address
10.0.0.0/30
set firewall family inet filter accept-icl term accept-icl from
destination-address 10.0.0.0/30
```

### [PR1183680](#)

- On a Junos Fusion Enterprise, a loss of connectivity of the link connecting the standalone switch might cause conversion of the switch from Junos OS to SNOS to fail. As a workaround, reboot the standalone switch to restart the conversion process in case of auto-conversion. [PR1232798](#)
- On a Junos Fusion Enterprise, the satellite device might not come online when the systems is converted from cluster to non-cluster mode without accompanying topology changes. As a workaround, ensure the configuration of satellite devices matches the wiring topology: non-cluster devices should not be connected to other clustered devices through default or configured clustering/uplink ports. [PR1251790](#)
- On Junos Fusion Enterprise, when 802.1X is configured in single-secure mode, a firewall counter is created for the default discard term in addition to the configured term. [PR1254503](#)

## SEE ALSO

[New and Changed Features | 84](#)

[Changes in Behavior and Syntax | 86](#)

[Known Issues | 88](#)

<a href="#">Resolved Issues   89</a>
<a href="#">Documentation Updates   90</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   91</a>
<a href="#">Product Compatibility   97</a>

## Known Issues

### IN THIS SECTION

- [Junos Fusion Enterprise | 88](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R3 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, the TCPDUMP command does not capture packets on satellite devices. [PR1125568](#)
- In a Junos Fusion, when a satellite device port is configured with auto-negotiation, and it is operating at a different speed than a link partner device, the port might go back to negotiated speed instead of going down. [PR1247353](#)

### SEE ALSO

<a href="#">New and Changed Features   84</a>
<a href="#">Changes in Behavior and Syntax   86</a>
<a href="#">Known Behavior   86</a>
<a href="#">Resolved Issues   89</a>
<a href="#">Documentation Updates   90</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   91</a>
<a href="#">Product Compatibility   97</a>



## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 17.2R3 | 89](#)
- [Resolved Issues: 17.2R2 | 90](#)
- [Resolved Issues: 17.2R1 | 90](#)

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: 17.2R3

- Mirrored packets are dropped if the analyzer output extended port is reachable through the ICL link. [PR1211123](#)
- On a Junos Fusion Enterprise with dual aggregation devices, while applying Routing Engine lo0 filters and setting the cascade port down on AD2, the SD goes to **ProvSessionDown** state on AD2 while it stays online on AD1. [PR1275290](#)
- In a Junos Fusion environment the satellite device displays U-boot on the LCD screen. [PR1304784](#)
- All the 802.1X sessions are removed when the AUTO ICCP link is disabled. [PR1307588](#)
- LACP aggregated Ethernet interfaces go to a down state when performing **commit synchronize**. [PR1314561](#)
- Packet loss of 2-3 seconds is seen 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 an 802.1X-enabled port is installed on a single-homed satellite device. [PR1328247](#)
- DHCP security binding entries are not synced after the FPC goes offline and comes back online. [PR1332828](#)
- Issue with 802.1X re-authentication in Junos Fusion Enterprise. [PR1345365](#)
- A satellite device does not recover PoE after the device is offline for more than 10 minutes and rejoins the AD. [PR1356478](#)
- The Fusion satellite device reboots after an automatic POE firmware upgrade. [PR1359065](#)
- The ppm-lite process might generate a core file on the Fusion satellite devices. [PR1364265](#)

## Resolved Issues: 17.2R2

### *Junos Fusion Enterprise*

- In dual aggregation device case, when you disable a cascade port, the extended port physical interfaces are marked as being down. [PR1232924](#)
- EX4300 with Junos OS Release 17.1R1 cannot be converted to satellite mode. [PR1267767](#)
- CoS shaping is not happening properly according to the configured shaping rate. [PR1268084](#)
- In a Junos Fusion Enterprise, for **show ethernet-switching table**, a few entries are stuck in DLR state after **I2-learning** restart. [PR1268619](#)
- In a Junos Fusion Enterprise, the DHCP snooping entry is deleted after I2ald restart. [PR1281824](#)
- VRRP split-brain state in dual aggregation device Junos Fusion. [PR1293030](#)
- Aggregation devices without a cascade port cannot reach hosts over an ICL link if they are authenticated by 802.1X authentication in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)
- The 802.1X authentication might fail in a Junos Fusion setup. [PR1299532](#)
- Dot1x might crash in Junos Fusion setup with dual AD. [PR1303909](#)

## Resolved Issues: 17.2R1

There are no resolved issues for Junos Fusion Enterprise in Junos OS Release 17.2R1.

SEE ALSO

<a href="#">New and Changed Features   84</a>
<a href="#">Changes in Behavior and Syntax   86</a>
<a href="#">Known Behavior   86</a>
<a href="#">Known Issues   88</a>
<a href="#">Documentation Updates   90</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   91</a>
<a href="#">Product Compatibility   97</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 17.2R3 for Junos Fusion Enterprise documentation.

## SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  84</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  86</a>
<a href="#">Known Behavior</a>	<a href="#">  86</a>
<a href="#">Known Issues</a>	<a href="#">  88</a>
<a href="#">Resolved Issues</a>	<a href="#">  89</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  91</a>
<a href="#">Product Compatibility</a>	<a href="#">  97</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | [91](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | [93](#)
- [Preparing the Switch for Satellite Device Conversion](#) | [94](#)
- [Converting a Satellite Device to a Standalone Switch](#) | [95](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [95](#)
- [Downgrading from Release 17.2](#) | [96](#)

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 17.2R2:

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. Click the **Junos Fusion EX9200 (Enterprise)** title to expand the list of supported releases.
5. Click the release number (the software version that you want to download) from the list.
6. Select the aggregation device 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 17.2R3 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**.

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

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 the satellite device 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 16.1, 16.2 and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

## Downgrading from Release 17.2

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:** It is not recommended to downgrade the aggregation device from 17.2R1 to 16.1 if there are cluster satellite devices in the setup.

To downgrade a Junos Fusion Enterprise from Junos OS Release 17.2, follow the procedure for upgrading, but replace the 17.2 **junos-install** 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 | 84](#)

[Changes in Behavior and Syntax | 86](#)

[Known Behavior | 86](#)

[Known Issues | 88](#)

[Resolved Issues | 89](#)

[Documentation Updates | 90](#)

[Product Compatibility | 97](#)



# Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 97](#)
- [Hardware Compatibility Tool | 97](#)

## 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](#).

## Hardware Compatibility Tool

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

SEE ALSO

<a href="#">New and Changed Features   84</a>
<a href="#">Changes in Behavior and Syntax   86</a>
<a href="#">Known Behavior   86</a>
<a href="#">Known Issues   88</a>
<a href="#">Resolved Issues   89</a>
<a href="#">Documentation Updates   90</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   91</a>

# Junos OS Release Notes for Junos Fusion Provider Edge

## IN THIS SECTION

- New and Changed Features | 98
- Changes in Behavior and Syntax | 100
- Known Behavior | 100
- Known Issues | 101
- Resolved Issues | 102
- Documentation Updates | 103
- Migration, Upgrade, and Downgrade Instructions | 103
- Product Compatibility | 111

These release notes accompany Junos OS Release 17.2R3 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

## IN THIS SECTION

- Release 17.2R3 New and Changed Features | 99
- Release 17.2R2 New and Changed Features | 99
- Release 17.2R1 New and Changed Features | 99

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 17.2R3 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 17.2R3.

## Release 17.2R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 17.2R2.

## Release 17.2R1 New and Changed Features

### *Class of Service (CoS)*

- **Per-unit scheduler support on extended ports (Junos Fusion Provider Edge)**—Beginning with Junos OS 17.2R1, Junos Fusion Provider Edge supports per-unit schedulers on extended ports. To support per-unit scheduling on an extended port, all cascade ports on the aggregation device for that extended port must have a queueing chip. aggregated Ethernet ports support per-unit schedulers, but all aggregated Ethernet member ports must be on the same satellite device. To enable per-unit scheduling on an extended port, enable the **per-unit-scheduler** option at the **[edit interfaces *interface-name*]** hierarchy level for the extended port.

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

- **Hierarchical CoS support on extended ports (Junos Fusion Provider Edge)**—Beginning with Junos OS 17.2R1, Junos Fusion Provider Edge supports hierarchical CoS (interface set-level scheduling) on extended ports. To support hierarchical CoS on an extended port, all cascade ports on the aggregation device for that extended port must have a queueing chip. aggregated Ethernet ports support hierarchical schedulers, but all aggregated Ethernet member ports must be on the same satellite device. To enable hierarchical CoS on an extended port, enable the **hierarchical-scheduler** option at the **[edit interfaces *interface-name*]** hierarchy level for the extended port.

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

### *Junos Fusion*

- **Support for selective VLAN local switching**—Starting in Junos OS Release 17.2R1, Junos Fusion Provider Edge supports local switching on a service level. When you configure selective VLAN local switching on satellite devices, the other VLANs will continue to follow the default forwarding behavior. Use the **selective-vlan-switching** option for the routing instance at the **[edit forwarding-options *satellite fpc slot*]** hierarchy level to enable selective VLAN local switching for a particular satellite device.
- **Support for an ingress policer**—Starting in Junos OS Release 17.2R1, Junos Fusion Provider Edge supports the use of an ingress policer to filter incoming traffic at the extended port level. This feature supports

a two-color policer that allows you to limit the traffic that is received on an interface. You can configure the Layer 2 ingress policer by using the **input-policer** statement at the **[edit interfaces *interface-name* layer2-policer]** hierarchy level.

#### SEE ALSO

[Changes in Behavior and Syntax | 100](#)

[Known Behavior | 100](#)

[Known Issues | 101](#)

[Resolved Issues | 102](#)

[Documentation Updates | 103](#)

[Migration, Upgrade, and Downgrade Instructions | 103](#)

[Product Compatibility | 111](#)

## Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 17.2R3.

#### SEE ALSO

[New and Changed Features | 98](#)

[Known Behavior | 100](#)

[Known Issues | 101](#)

[Resolved Issues | 102](#)

[Documentation Updates | 103](#)

[Migration, Upgrade, and Downgrade Instructions | 103](#)

[Product Compatibility | 111](#)

## Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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.

SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  98</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  100</a>
<a href="#">Known Issues</a>	<a href="#">  101</a>
<a href="#">Resolved Issues</a>	<a href="#">  102</a>
<a href="#">Documentation Updates</a>	<a href="#">  103</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  103</a>
<a href="#">Product Compatibility</a>	<a href="#">  111</a>

**Known Issues**

There are no known issues in the Junos OS Release 17.2R3 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.

SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  98</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  100</a>
<a href="#">Known Behavior</a>	<a href="#">  100</a>
<a href="#">Resolved Issues</a>	<a href="#">  102</a>
<a href="#">Documentation Updates</a>	<a href="#">  103</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  103</a>
<a href="#">Product Compatibility</a>	<a href="#">  111</a>

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 17.2R3 | 102](#)
- [Resolved Issues: 17.2R2 | 102](#)
- [Resolved Issues: 17.2R1 | 102](#)

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: 17.2R3

#### *Junos Fusion*

- In Junos Fusion, the **show interfaces diagnostics optics satellite** command does not display any output. [PR1327876](#)
- SSH key-based authentication fails after reboot if **chassis satellite-management** is configured. [PR1344392](#)
- Laser receive power of extended ports is higher than the output power of the peer link. [PR1358007](#)
- The shutdown of the cascade port might lead to the invalidation of the MPC. [PR1360876](#)

### Resolved Issues: 17.2R2

#### *Junos Fusion*

- In a Junos Fusion setup, the transit unicast traffic might be discarded on a satellite device when they pass through different IFLs of the same extended port. [PR1264900](#)

### Resolved Issues: 17.2R1

There are no fixed issues in the Junos OS Release 17.2R1 for Junos Fusion Provider Edge.

### SEE ALSO

[New and Changed Features | 98](#)

Changes in Behavior and Syntax	100
Known Behavior	100
Known Issues	101
Documentation Updates	103
Migration, Upgrade, and Downgrade Instructions	103
Product Compatibility	111

## Documentation Updates

There are no errata or changes in Junos OS Release 17.2R3 for Junos Fusion Provider Edge documentation.

### SEE ALSO

New and Changed Features	98
Changes in Behavior and Syntax	100
Known Behavior	100
Known Issues	101
Resolved Issues	102
Migration, Upgrade, and Downgrade Instructions	103
Product Compatibility	111

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 104
- Upgrading an Aggregation Device with Redundant Routing Engines | 106
- Preparing the Switch for Satellite Device Conversion | 106
- Converting a Satellite Device to a Standalone Device | 108
- Upgrading an Aggregation Device | 110
- Upgrade and Downgrade Support Policy for Junos OS Releases | 110
- Downgrading from Release 17.2 | 111

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 **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](#).

The download and installation process for Junos OS Release 16.1R1 and later 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 access the Junos Fusion Hardware and Software Compatibility page.
4. Click the **Junos Fusion MX Series (Provider Edge)** title to expand the list of supported releases.
5. Click the release number (the software version that you want to download) from the list.



6. Select the aggregation device 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.

- For 64-bit software:

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

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

For example:

```
user@host> request system software add validate reboot
source/junos-install-mx-x86-64-17.2R3.9-signed.tgz
```

- For 32-bit software:

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

For example:

```
user@host> request system software add validate reboot
source/junos-install-mx-x86-32-17.2R3.9-signed.tgz
```

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

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

- `http://hostname/pathname`
- `scp://hostname/pathname` (available only for the Canada and U.S. version)

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

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

Rebooting occurs only if the upgrade is successful.

**NOTE:** After you install a Junos OS Release 17.2R3 **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](#)

Customers with EX4300 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-ex-4300-14.1X53-D43.7-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-qfx-5-14.1X53-D43.7-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 the satellite device from the Junos Fusion topology.

**NOTE:** If the satellite device is a QFX5100 switch, you need to install a PXE version of Junos OS. The PXE version of Junos OS is software that includes pxe in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D30 is named `install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz`. If the satellite device is an EX4300 switch, you install a standard `jinstall-ex-4300` version of Junos OS.

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

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:  
<https://www.juniper.net/support/downloads>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** from the drop-down list and select the switch platform series and model for your satellite device.
4. Select the Junos OS Release 14.1X53-D30 software image for your platform.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.

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

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

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

```
[edit]
```

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

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

```
[edit]
```

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

```
[edit]
```

```
user@aggregation-device# commit synchronize
```

Otherwise, commit the configuration to a single Routing Engine:

```
[edit]
```

```
user@aggregation-device# commit
```

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

```
[edit]
```

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

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

```
[edit]
```

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

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

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

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

11. Wait for the reboot that accompanies the software installation to complete.
12. When you are prompted to log back in to your device, uncable the device from the Junos Fusion topology. See *Removing a Transceiver from a QFX Series Device* or *Remove a Transceiver*, as needed. Your device has been removed from Junos Fusion.

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

## Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 17.2R3, you must also upgrade your satellite device to Satellite Device Software version 3.1R3.

## 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 16.1, 16.2, and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

## Downgrading from Release 17.2

To downgrade from Release 17.2 to another supported release, follow the procedure for upgrading, but replace the 17.2 **jinstall** package with one that corresponds to the appropriate release.

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

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

### SEE ALSO

<a href="#">New and Changed Features   98</a>
<a href="#">Changes in Behavior and Syntax   100</a>
<a href="#">Known Behavior   100</a>
<a href="#">Known Issues   101</a>
<a href="#">Resolved Issues   102</a>
<a href="#">Documentation Updates   103</a>
<a href="#">Product Compatibility   111</a>

## Product Compatibility

### IN THIS SECTION

- [Hardware Compatibility | 111](#)

## 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 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. 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

<a href="#">New and Changed Features   98</a>
<a href="#">Changes in Behavior and Syntax   100</a>
<a href="#">Known Behavior   100</a>
<a href="#">Known Issues   101</a>
<a href="#">Resolved Issues   102</a>
<a href="#">Documentation Updates   103</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   103</a>

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

### IN THIS SECTION

- [New and Changed Features | 113](#)
- [Changes in Behavior and Syntax | 144](#)
- [Known Behavior | 162](#)
- [Known Issues | 167](#)
- [Resolved Issues | 185](#)
- [Documentation Updates | 249](#)
- [Migration, Upgrade, and Downgrade Instructions | 250](#)
- [Product Compatibility | 258](#)



These release notes accompany Junos OS Release 17.2R3 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

### IN THIS SECTION

- [Release 17.2R3 New and Changed Features | 114](#)
- [Release 17.2R2 New and Changed Features | 114](#)
- [Release 17.2R1 New and Changed Features | 116](#)

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

## Release 17.2R3 New and Changed Features

### *Interfaces and Chassis*

- **Enhancement to increase the threshold of corrected single-bit errors (MPC7E, MPC8E, MPC9E on MX Series)**—In Junos OS Release 17.2R3, 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 errordetected in HMC 0 - Total count 25

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

[See [Alarm Overview](#).]

### *Restoration Procedures Failure*

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (MX Series)**—In Junos OS Release 17.2R3, 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](#).]

### *Subscriber Management and Services*

- **Controlling search behavior for address allocation from linked pools (MX Series)**—Starting in Junos OS Release 17.2R3, 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](#).]

## Release 17.2R2 New and Changed Features

### *Multicast*

- **Improved multicast performance using distributed IGMP (MX Series)**—Starting in Junos OS Release 17.2R2, you can improve multicast performance by using the distributed Internet Group Management Protocol (IGMP). Distributed IGMP moves IGMP processing from the Routing Engine to the Packet Forwarding Engine. When you configure distributed IGMP, join and leave events are processed across

multiple Modular Port Concentrators (MPCs) on the Packet Forwarding Engine. Instead of being processed through a centralized routing protocol process (rpd) on the Routing Engine, this improves performance and decreases join and leave latency.

For distributed IGMP to function properly, you must configure enhanced IP network services by including the **enhanced-ip** statement at the **[edit chassis network-services]** hierarchy level. To enable distributed IGMP on static interfaces, include the **distributed** statement at the **[edit protocols igmp interface interface-name]** hierarchy level. To enable distributed IGMP on dynamic interfaces, include the **distributed** statement at the **[edit dynamic-profiles profile-name protocols igmp interface \$junos-interface-name]** hierarchy level.

You can optionally configure specific multicast groups to join statically by including the **distributed** option at one of the following hierarchy levels:

- **[edit protocols pim static]**
- **[edit protocols pim static group multicast-group-address]**
- **[edit protocols pim static group multicast-group-address source source-address]**

[See [Understanding IGMP](#).]

### Services Applications

- **Support for disabling the filtering of HTTP traffic with an embedded IP address belonging to a blacklisted domain (MX Series router)**—Starting in Junos OS Release 17.2R2, you can disable the filtering of HTTP traffic that contains an embedded IP address belonging to a blacklisted domain name. To disable the filtering, include the **disable-url-filtering** statement at the **[edit services url-filter profile profile-name template template-name]** hierarchy level when you are configuring URL filtering. However, if the embedded IP address is explicitly identified in the blacklisted URL database, the traffic is still filtered.

[See [Configuring URL Filtering](#).]

- **Maximum number of RPM probes increased (MX Series routers)**—Starting in Junos OS Release 17.2R2, you can configure the maximum allowed number of concurrent real-time performance monitoring (RPM) probes on an MX Series router to be as high as 2000. In Junos OS Release 17.2R1 and earlier, you can configure the maximum number to be as high as 500.

[See [Limiting the Number of Concurrent RPM Probes](#).]

### Subscriber Management and Services

- **Support for excluding tunnel attributes from RADIUS Access-Request messages (MX Series)**—Starting in Junos OS Release 17.2R2, you can use the **exclude** statement at the **[edit access profile profile-name radius attribute]** hierarchy level to exclude the following tunnel attributes from RADIUS Access-Request messages in addition to the previously supported Accounting-Start and Accounting-Stop messages:
  - **acct-tunnel-connection**—RADIUS attribute 68, Acct-Tunnel-Connection
  - **tunnel-assignment-id**—RADIUS attribute 82, Tunnel-Assignment-Id
  - **tunnel-client-auth-id**—RADIUS attribute 90, Tunnel-Client-Auth-Id

- tunnel-client-endpoint—RADIUS attribute 66, Tunnel-Client-Endpoint
- tunnel-medium-type—RADIUS attribute 65, Tunnel-Medium-Type
- tunnel-server-auth-id—RADIUS attribute 91, Tunnel-Server-Auth-Id
- tunnel-server-endpoint—RADIUS attribute 67, Tunnel-Server-Endpoint
- tunnel-type—RADIUS attribute 64, Tunnel-Type

## Release 17.2R1 New and Changed Features

### Hardware

- **RE-S-X6-64G-LT Routing Engine and REMX2K-X8-64G-LT CB-RE Routing Engines(MX Series)**—Starting with Junos OS release 17.2R1, MX Series Routers support the following new Routing Engine and CB-RE:
  - RE-S-X6-64G-LT Routing Engine
  - REMX2K-X8-64G-LT CB-RE

[See [MX240 Routing Engine Description](#), [MX480 Routing Engine Description](#), [MX960 Routing Engine Description](#), and [MX2000 Host Subsystem Description](#).]

**NOTE:** The Routing Engines are equipped with limited encryption support only. 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. See [Junos OS Editions](#).

- **Junos OS support for MX2008 routers**—In Junos OS Release 15.1F7 and 17.2R1, Junos OS supports the MX2008 Universal Routing Platform (model number: CHAS-MX2008). The MX2008 router is a 10-slot half-rack chassis with increased port density, but uses less space and consumes less power. Additionally, with the MX2008, you can scale bandwidth up to 1.6 Tbps per slot by using a chassis that is approximately half a rack in size.

The MX2008 router is an Ethernet-optimized edge router that provides both switching and carrier-class Ethernet routing. The router enables a wide range of business and residential applications and services, including high-speed transport and VPN services, next-generation broadband multiplay services, and high-volume Internet data center networking.

### Class of Service (CoS)

- **Support for user-configurable traffic class map (MX Series routers with MPCs)** — Beginning with Junos OS Release 17.2R1, MX Series routers with MPCs support a user-configurable input priority map, known as a **traffic-class-map**, that enables you to prioritize and classify input traffic entering a Packet Forwarding Engine during ingress oversubscription. You can define traffic class maps for a packet based on DSCP,

IP precedence, MPLS EXP, IEEE 802.1p, and IEEE 802.1ad CoS values and associate these CoS values with **real-time**, **network-control**, and **best-effort** traffic classes.

[See [Managing Ingress Oversubscription at the PFE.](#)]

- **CoS-based forwarding support for up to 16 forwarding classes (MX Series and PTX Series)**—Beginning with Junos OS Release 17.2R1, MX Series routers with MPCs or MS-DPCs, vMX, PTX3000 routers, PTX5000 routers, and VPTX support configuring CoS-based forwarding (CBF) for up to 16 forwarding classes. All other platforms support CBF for up to 8 forwarding classes. To support up to 16 forwarding classes for CBF on MX routers, enable **enhanced-ip** at the **[edit chassis network-services]** hierarchy level.

[See [Forwarding Policy Options Overview.](#)]

- **Propagating CoS shaping rate adjustments that are based on multicast traffic (MX Series)**—Starting in Junos OS Release 17.2R1, you can set up CoS shaping rate adjustments that are based on multicast traffic to be propagated to the parent in the scheduler hierarchy. For service providers that are using interface sets to deliver services such as voice and data and multicast VLANs (M-VLANs) to deliver broadcast television, you can set up CoS so that when a subscriber begins receiving multicast traffic, the shaping rate of the subscriber interface is adjusted to account for the multicast traffic. You can now set up the CoS multicast adjustment to be propagated from the subscriber interface to the interface set, which is the parent in the scheduler hierarchy. This feature prevents oversubscription of the multicast replicator, such as a PON, which can result in dropped traffic and service disruption.

[See [Using Hierarchical CoS to Adjust Shaping Rates Based on Multicast Traffic.](#)]

## EVPNs

- **Support for ARP proxy and suppressing of ARP flooding with EVPN (MX Series routers with MPCs)**—Starting in Junos OS Release 17.2R1, a provider edge (PE) router can function as an Address Resolution Protocol (ARP) proxy with EVPN configured. The ARP proxy/suppression capability is enabled by default. For EVPN instances with IRB interfaces ARP flooding will be suppressed. To disable proxy and suppression of ARP flooding, include the **no-arp-suppression** statement at the **[edit bridge-domains bridge-domain-name]** hierarchy level.

[See [EVPN Proxy ARP and ARP Suppression and Network Discovery Protocol and Network Discovery Protocol Suppression.](#)]

- **NSR and unified ISSU support for EVPN-VPWS and PBB-EVPN**—Starting in Junos OS Release 17.2R1, Junos OS supports NSR and unified ISSU on VPWS with EVPN and provider backbone bridging (PBB) EVPN. NSR and GRES enable the routing system to switch over from a primary Routing Engine to a backup Routing Engine while continuing to forward packets.

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. Unified ISSU upgrade is only supported by dual Routing Engine platforms. Unified ISSU requires both GRES and NSR to be enabled.

To enable GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

To enable NSR, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **commit synchronize** statement at the **[edit system]** hierarchy level.

[See [Overview of VPWS with EVPN Signaling Mechanisms](#) and [Provider Backbone Bridging \(PBB\) and EVPN Integration for Data Center Interconnect Overview](#).]

- **Unified ISSU support for EVPN and VXLAN**—Starting in Junos OS Release 17.2R1, Junos OS supports Unified ISSU on EVPN and VXLAN. 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. Unified ISSU upgrade is only supported by dual Routing Engine platforms. Unified ISSU requires both GRES and NSR to be enabled.

To enable GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

To enable NSR, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **commit synchronize** statement at the **[edit system]** hierarchy level.

[See [NSR and Unified ISSU Support for EVPN Overview](#) and [PIM NSR and Unified ISSU Support for VXLAN Overview](#).]

- **Support for EVPN E-Tree service**—Starting in Release 17.2R1, Junos OS enables you to configure Ethernet VPN E-Tree service. The EVPN E-Tree feature implements E-Tree service as defined by the Metro Ethernet Forum (MEF) in draft-sajassi-l2vpn-evpn-etree-03. The E-Tree service is a rooted-multipoint service that is supported only with EVPN over MPLS in the core. In an EVPN E-Tree service, each customer edge (CE) device attached to the EVPN E-Tree service needs to be designated as either root or leaf. If an interface is not configured for a role, it is assigned the role of “root” by default.

The service adheres to the following forwarding rules:

- A leaf can send or receive traffic only from a root.
- A root can send traffic to another root or any of the leaf devices.
- A leaf or root can be connected to provider edge (PE) devices in single homing mode or multihoming mode.

To configure an Ethernet VPN E-Tree service, use **set evpn-etree** at the **edit routing-instances <routing-instance\_name> protocols evpn** hierarchy level.

To configure an interface as leaf, use **set etree-ac-role leaf** at the **[edit interfaces <interface-name> unit <interface-unit-number>]** hierarchy level.

To configure an interface as root, use **set etree-ac-role root** at the **[edit interfaces <interface-name> unit <interface-unit-number>]** hierarchy level.

[See [EVPN-ETREE Overview](#).]

- **Interconnecting data center networks over WAN (MX Series)**—Starting in Junos OS Release 17.2R1, you can interconnect data center networks running Ethernet VPN (EVPN) with Virtual Extensible LAN (VXLAN) encapsulation through a WAN running MPLS-based EVPN. This feature enables you to:

- Connect data center edge routers over MPLS-based EVPN WAN for data center interconnections.
- Interconnect EVPN-VXLAN and EVPN-MPLS using logical tunnel (lt-) interface on data center edge routers.

[See [EVPN-VXLAN Data Center Interconnect Through EVPN-MPLS WAN Overview](#).]

- **Integrating PBB with EVPN (MX Series with MPCs and MICs)**—Starting in Junos OS Release 17.2R1, the integration of provider backbone bridging (PBB) with Ethernet VPN (EVPN) is supported. With PBB-EVPN, the control plane learning across the core is significantly reduced, allowing a huge number of Layer 2 services, such as data center connectivity, to transit the network in a simplified manner.

In a PBB-EVPN network, the backbone core bridge (BCB) device in the PBB core is replaced with MPLS, while retaining the service scaling properties of the PBB backbone edge bridge (BEB). The B-component (provider routing instance) is signaled using EVPN BGP signaling and encapsulated inside MPLS using provider edge (PE) and provider (P) devices. Thus, PBB-EVPN combines the vast scaling property of PBB with the simplicity of a traditional basic MPLS core network, resulting in significant reduction in the amount of network-wide state information, as opposed to regular PBB.

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

- **NSR and unified ISSU support for EVPN-ETREE**—Starting in Junos OS Release 17.2R1, Junos OS supports NSR and unified ISSU for EVPN-ETREE services. NSR and GRES enables the routing system to switch over from a primary Routing Engine to a backup Routing Engine while continuing to forward packets.

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. Unified ISSU upgrade is only supported by dual Routing Engine platforms. Unified ISSU requires both GRES and NSR to be enabled.

To enable GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

To enable NSR, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **commit synchronize** statement at the **[edit system]** hierarchy level.

[See [EVPN-ETREE Overview](#).]

- **MAC pinning support for PBB-EVPN (MX Series with MPCs)**—Starting in Junos OS Release 17.2R1, the MAC pinning feature is enabled on provider backbone bridging (PBB) and Ethernet VPN (EVPN) integration, including customer edge (CE) interfaces and EVPN over PBB core in both all-active or single-active mode.

To configure MAC pinning for PBB-EVPN, include the **mac-pinning** statement at the **[edit routing-instances pbbn protocols evpn]**, where **pbbn** is the PBB routing instance over backbone port (B-component). With this configuration, the dynamically learned MAC addresses in the PBB I-component (customer routing instance) bridge domain over CE interfaces, as well as PBB-MPLS core interfaces, are pinned. This prevents MAC move on duplicate MAC detection, avoiding loop creation in a network. The duplicate MAC addresses are blocked, and data is dropped if traffic is received on any interface other than the interface on which it is pinned.

[See [PBB-EVPN MAC Pinning Overview](#).]

### **Forwarding and Sampling**

- **Support for multiple server instances under a given interface. (MX Series)**—Starting in Junos OS Release 17.2R1, you can specify multiple Domain Name System (DNS), Trivial File Transfer Protocol (TFTP), or BOOTP servers instances under a given helper port interface. The same packet, with the originator IP address and port requests, is forwarded to the different configured servers; the payload of the UDP packet is not modified.

[See [DNS, Port, and TFTP Service Servers](#).]

- **Improved load balancing for L2TP data transit traffic (MX Series)**—Starting in Junos OS Release 17.2, L2TP load balancing can occur on a per-tunnel basis, or within the same tunnel, on a per-session basis, for better distribution of packets. To enable this feature, enable the **l2tp-tunnel-session-identifier** command at the **[edit forwarding-options hash-key family inet]** hierarchy level.

[See [l2tp-tunnel-session-identifier](#).]

### **General Routing**

- **Support for PTP, Synchronous Ethernet, and hybrid mode over link aggregation group (MX104, MX240, MX480, MX960, MX2010)**—Starting in Junos OS Release 17.2R1, the MPC5E, MPC6E, MPC3E NG, and MPC2E NG 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.

- **OpenConfig: BGP routing table - Support for operational state model (MX Series)**—Starting in Junos OS 17.2R1, the OpenConfig BGP RIB routing table supports local-rib for IPV4 and IPV6. The Openconfig-rib-bgp.yang model supports five logical RIBs per address family. There are five tables for IPv4 routes and five tables for IPv6 routes.
- **Support for PTP over Ethernet, hybrid mode, and G.8275.1 profile (MPC6E, MPC2E NG, MPC3E NG MPCs)**—Starting in Junos OS Release 17.2R1, MPC6E, MPC2E NG, and MPC3E NG MPCs 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.
  - **Hybrid mode**— In hybrid mode, the synchronous Ethernet equipment clock (EEC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP.
  - **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.

**NOTE:** PHY timestamping is supported on MPC2E NG and MPC3E NG only with MIC-3D-20GE-SFP-E.

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

- **Enhancements to Precision Time Protocol feature (MX104)**—Starting in Junos OS Release 17.2R1, the Precision Timing Protocol (PTP) feature in MX104 routers has been enhanced with the following changes:
  - After PTP is phase-aligned, if the system up time is less than 30 minutes and the PTP source is lost before 30 minutes, the PTP state will be moved to **freerun**. On the other hand, if the system up time is more than 30 minutes and the PTP source is lost, the PTP state will move to **holdover**.
  - If PTP is never phase-aligned and PTP source is lost, the PTP state shall move to **freerun**.
  - While operating in PTP Hybrid mode, the state of PTP will be in **holdover** for 8 days after a PTP clock source is lost but a valid high stratum SyncE source is present.
  - PTP state will transition to **holdover** irrespective of the current state of **acquiring** or **phase aligned** as long as PTP was phase-aligned once and system uptime was more than 30 minutes.
- **New command to display upstream and downstream clock information (MX104)**—Starting with Junos OS Release 17.2R1, a new show command, **show ptp all-master-clock**, is introduced to display all the upstream master information and clock advertised to downstream. This command is supported only on MX104 routers.
- **OpenConfig: Supporting for the BGP model in Junos OS (MX Series)**—Starting in Junos OS 17.2R1, the configuration leaf devices defined in the **openconfig-bgp.yang** and **openconfig-bgp-multiprotocol.yang** files are supported.

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

- **Warm standby mode for routing protocols process (MX Series)**—Starting in Junos OS Release 17.2R1, you can set the routing protocol process (rpd) mode to **warm-standby** by using the **set routing-options warm-standby** command. Warm standby mode helps the backup Routing Engine stay synchronized with the master Routing Engine, allowing for faster Routing Engine switchover during GRES.

[See [warm-standby](#).]

- **Support for unified ISSU on MX Series routers and MX Series Virtual Chassis with MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, MPC2E-3D-NG-Q, and MPC5E (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Release 17.2R1, Junos OS supports unified ISSU on MX Series routers and MX Series Virtual Chassis with MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, MPC2E-3D-NG-Q, and MPC5E.

Unified ISSU is supported on MPC5E with the following MICs in non-OTN mode:

- 3X40GE QSFP
- 12X10GE-SFPP OTN
- 1X100GE-CFP2
- 2X10GE SFPP OTN

**NOTE:** Unified ISSU is not supported on MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q with the following MICs:

- MS-MIC-16G
- MIC-3D-8DS3-E3
- MIC-3D-10C192-XFP

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.

- **Kernel synchronization performance and debugging enhancements (MX Series)**—Starting in Junos OS Release 17.2R1, the kernel synchronization process (ksyncd) uses multithreading for increased performance, and you can use new CLI commands for ksyncd debugging and recovery. Use the **set system kernel-replication no-multithreading** command to run ksyncd in single thread mode for debugging purposes. Use the **set system kernel-replication system-reboot recovery-failure** command to configure the automatic reboot of a standby Routing Engine after receiving a ksyncd initialization error.

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

### *Interfaces and Chassis*

- **Software feature support on the MX2008**—In Junos OS Release 15.1F7 and 17.2R1, the MX2008 router supports all software features that are supported by other MX Series routers in Junos OS Release 15.1F6.

The following key Junos OS features are supported:

- Basic Layer 2 features including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
- Class of service (CoS)
- Firewall filters and policers
- Integrated routing and bridging (IRB)
- Interoperability with existing MPCs (excluding the Application Services Modular Carrier Card, or AS-MCC)
- Layer 2 protocols
- Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs
- Layer 3 routing protocols and MPLS

- Layer 3 services supported on MS-MIC and MS-MPC (for example, CGNAT, IP Security, inline active flow monitoring) and inline services
- Multicast forwarding
- Port mirroring
- Spanning-tree protocols, such as STP, MSTP, RSTP, and VSTP
- Synchronous Ethernet and Precision Time Protocol (IEEE 1588)
- Tunneling
- Graceful Routing Engine Switchover (GRES) and Non Stop Routing (NSR)

**MPCs and MICs supported on MX2008 routers**—The MX2008 router (model number: CHAS-MX2008) supports all the MPCs (excluding AS-MCC) and MICs that are supported by the MX2000 line of routers.

MPCs native to the MX2000 line of routers (MPC6E, MPC8E, and MPC9E) are supported without an adapter card, but other MPCs (MS-MPC, MPC1, MPC2, MPC3, MPC4, MPC5, MPC7, MPC2E-NG, MPC3E-NG, and all variants) are supported with an adapter card.

**NOTE:** MX2008 routers do not support the Application Services Modular Carrier Card (AS-MCC).

[See [MPCs Supported by MX240, MX480, MX960, MX2010, and MX2020 Routers](#).]

**Support for centralized clocking on MX2008 routers**—In Junos OS Release 15.1F7 and 17.2R1, the MX2008 router (model number: CHAS-MX2008) uses the centralized Stratum 3 clock module on the Routing and Control Board (RCB) to lock onto Synchronous Ethernet and distribute the frequency to the entire chassis. Supported features include:

- Clock monitoring, filtering, and holdover
- Hitless transition from a distributed to a centralized clocking mode
- Distribution of the selected chassis clock source to downstream network elements by using supported line interfaces

You can view the centralized clock module information by using the **show chassis synchronization clock-module** command.

**NOTE:** The MX2008 supports Precision Time Protocol (PTP) in distributed mode.

**Junos OS support for FRU management of MX2008 routers**—In Junos OS Release 15.1F7 and 17.2R1, Junos OS supports the MX2008 router (model number: CHAS-MX2008). The Junos OS chassis

management software for the MX2008 routers provides enhanced environmental monitoring and field-replaceable unit (FRU) control.

The MX2008 host subsystem consists of two Routing and Control Boards, or RCBs (model number REMX2008-X8-64G). The RCB is an integrated board and a single FRU that provides Routing Engine and Control Board functionality and supports virtualization. The router contains 8 SFBs (fabric cards, model number: MX2008-SFB2) that provides 7+1 redundancy. The router supports a maximum of 10 MPCs including adapter cards, and up to 20 MICS—a maximum of two MICs can be installed in each MPC.

The chassis contains nine power supply modules (PSMs) and two power distribution modules (PDMs) for the power feeds. Each PSM delivers 2500 W of power, and provides 8+1 redundancy. The two PDMs provide feed redundancy, with each PDM connected to primary and backup feeds separately.

The MX2008 cooling system contains two fan trays, with six fans in each. The fan trays can be installed at or removed from the back of the chassis, which allows the space in the front to be used for cable management. The MX2008 supports temperature thresholds for each temperature sensor, which enables the router to precisely control the cooling, raise alarms, and shut down a FRU.

[See [Junos OS for MX Series 5G Universal Routing Platforms](#).]

- **Limited encryption Junos OS image and boot restriction (MX Series)**—Starting with Junos OS Release 17.2R1, the MX240, MX480, MX960, MX2010, and MX2020 routers with the Routing Engines RE-S-X6-64G-LT and RE-MX2K-X8-64G-LT support 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 Routing Engines are restricted to boot only the Junos Limited image.
- **Enhancement to ambient-temperature statement (MX Series)**—In Junos OS Release 15.1F4 and later, the default ambient temperature is set at 40° C on MX480, MX960, MX2010, and MX2020 Universal Routing Platforms. You can override ambient temperature by setting the temperature at 55° C or 25° C.

```
[edit]
user@router# set chassis ambient-temperature ?
Possible completions:
25C                25 degree celsius
40C                40 degree celsius
55C                55 degree celsius
[edit]
```

When a router restarts, the system adjusts the power allocation or the provisioned power for the line cards on the basis of the configured ambient temperature. If enough power is not available, a minor chassis alarm is raised. However, the chassis continues to run with the configured ambient temperature.

You can configure a new higher ambient temperature only after you make more power available by adding new power supply modules or by taking a few line cards offline. By using the provisioned power that is saved by configuring a lower ambient temperature, you can bring more hardware components online.

- **Reordering of MAC addresses after a Routing Engine switchover**—In Junos OS Release 14.2 and later, if you configure multiple aggregated Ethernet interfaces, the MAC address of the aggregated Ethernet interfaces displayed in the **show interfaces ae *number*** command output might get reordered after a Routing Engine switchover or restart.

As a workaround, you can configure static MAC addresses for aggregated Ethernet interfaces. Any external dependency, such as filtering of the MAC addresses that are assigned before the reboot, becomes invalid if the MAC address changes.

### Layer 2 VPN

- **Support for FEC128 and FEC129 in the same routing instance (MX Series)**—Starting in Release 17.2R1, Junos OS supports the configuration of forwarding equivalency class (FEC) 128 mesh groups in a FEC 129 VPN instance. You can configure a FEC 129 VPLS instance to support both BGP autodiscovery as defined in FEC 129 as well as statically configured Label Distribution Protocol (LDP) neighbors as defined by FEC 128. This allows a router to use a common MAC table to forward traffic between a FEC 128 LDP VPLS domain and a FEC 129 domain.

[See [show vpls connections \(with FEC128 and FEC129 in the same routing-instance\)](#).]

### Management

- **Support for fabric statistics sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, you can export fabric statistics through the Junos Telemetry Interface. The types of fabric statistics you can export include those for Packet Forwarding Engine pairs, Flexible PIC Concentrators, and Control Boards and Switch Fabric Boards. To enable a sensor to export fabric statistics include the **resource /junos/system/linecard/fabric/** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level. Only UDP streaming is supported. gRPC streaming is not supported.

[See [Configuring a Junos Telemetry interface Sensor \(CLI Procedure\)](#).]

- **Support for LSP events and properties sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, you can export statistics for LSP events and properties through the Junos Telemetry Interface. Only gRPC streaming for this sensor is supported. You can export statistics for ingress point-to-point LSPs, point-to-multipoint LSPs, bypass LSPs, and dynamically created LSPs. To export data through gRPC, use the **/mpls/lsp/** or **/mpls/signal-protocols/** set of OpenConfig subscription paths. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of the Junos OS with an upgraded FreeBSD kernel, 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 [Guidelines for gRPC Sensors](#).]

- **Support for gRPC streaming for Junos Telemetry Interface firewall filter statistics (MX Series)**—Starting with Junos OS Release 17.2R1, you can use gRPC interfaces to provision sensors to subscribe to and receive firewall filter telemetry data. Hierarchical policer statistics are also collected. Use the `/junos/firewall/firewall-stats/` path to provision a sensor for firewall filter statistics. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent 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. OpenConfig paths are used to define telemetry parameters for data streamed through gRPC. This functionality was previously introduced in Junos OS Release 16.1R4.

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

- **Support for queue statistics for logical interface sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, logical interface sensors also collect egress and ingress queue statistics. Both UDP and gRPC streaming are supported. Queue statistics, including for per-unit queuing and hierarchical queuing, are exported when a queuing structure is configured on a logical interface. To provision a logical interfaces statistics sensor for UDP streaming, include the **resource** `/junos/system/linecard/interface/logical/usage/` statement at the `[edit services analytics sensor sensor-name]` hierarchy level. To provision a sensor for gRPC streaming, include the following resource `/interfaces/interface[name='interface-name']/subinterfaces/` in the subscription path. Use the **telemetrySubscribe** RPC to define telemetry parameters for gRPC streaming. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions.

[See [Overview of the Junos Telemetry Interface](#).]

- **Support for routing protocol processes task memory utilization sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, you can stream telemetry data through gRPC for routing protocol process (rpd) task memory usage. Include the `/junos/task-memory-information/` path to provision a sensor to stream data through gRPC. UDP streaming for this sensor is not supported. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent 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. OpenConfig paths are used to define telemetry parameters for data streamed through gRPC. This functionality was previously introduced in Junos OS Release 16.1R3.

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

- **Support for LSP statistics for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, you can stream telemetry data for LSPs through UDP and gRPC. To provision an LSP statistics sensor for UDP streaming, include the **resource** `/junos/services/label-switched-path/usage/` statement at the `[edit services analytics sensor sensor-name]` hierarchy level. Use the `mpls/lsp/constrained-path/tunnels/tunnel/` path to provision a sensor for streaming LSP statistics through gRPC. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces

to manage gRPC subscriptions. For both UDP and gRPC streaming, you must also configure the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level. Additionally, MX Series routers should operate in enhanced mode. Support for the LSP statistics sensor was previously introduced in Junos OS Release 15.1F6 and Junos OS Release 16.1R4.

[See [Overview of the Junos Telemetry Interface](#).]

- **Support for device family and release in Junos OS YANG modules (MX Series)**—Starting in Junos OS Release 17.2R1, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

[See [Understanding Junos OS YANG Modules](#).]

## MPLS

- **Support for MPLS label types with scale optimization (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the **enhanced-ip** command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can separate the MPLS labels used for different label spaces which provides more flexibility and scalability. The table space in **vrf-table-label** is also increased to at least 16,000, if the platform can support the scale.

For Junos OS Release 17.1 and earlier, MPLS label space was divided into various predefined segments, under **label-space** command, which served different purposes or applications. Due to various restrictions imposed by older platforms with limited capability, the segment allocation was platform dependent and fixed label space.

- **SPRING-TE support in PCEP implementation (MX Series)**—Starting in Junos OS Release 17.2R1, the traffic engineering (TE) capabilities of Source Packet Routing in Networking (SPRING) are supported in Path Computation Element Protocol (PCEP) sessions for the label-switched paths (LSPs) initiated by a Path Computation Element (PCE). Tunnel routes are created in the inet.3 routing table of the Path Computation Client (PCC) corresponding to the SPRING-TE LSPs. Similar to any other tunnel route, the SPRING-TE tunnel routes can be used for resolving indirect next hops for plain IP and service traffic.

To configure SPRING-TE for PCEP:

- Enable external path computing for MPLS and SPRING-TE at the **[edit protocols]** hierarchy level.
- Enable spring capability for the PCE at the **[edit protocols pcep pce pce]** hierarchy level.

[See [Support of SPRING-TE for the Path Computation Element Protocol Overview](#).]

- **Support for empty and loose EROs for PCE-controlled LSPs (MX Series)**—Starting in Junos OS Release 17.2R1, for PCE-initiated and PCC-delegated label-switched paths (LSPs), two Constrained Shortest Path First computation types are introduced for computing constrained paths locally and externally. With this, a Path Computation Client (PCC) can accept an LSP path, or Explicit Route Object (ERO), that includes loose next hops (loose ERO) or does not include a path at all (empty ERO), in addition to strict EROs.

With this enhancement, the existing Junos OS constrained path computation behavior and performance are leveraged, along with the other benefits of external path computing.

[See [PCE-Controlled LSP ERO](#).]

- **IPv6 support for static egress LSPs (MX Series)**—Starting in Junos OS Release 17.2R1, static LSPs on the egress router can be configured with IPv6 as the next-hop address for forwarding IPv6 traffic. Previously, only IPv4 static LSPs were supported. The IPv6 static LSPs share the same transit, bypass, and static LSP features of IPv4 static LSPs.

A commit failure occurs when the next-hop address and destination address of the static LSP do not belong to the same address family (IPv4 or IPv6).

[See [next-hop \(Protocols MPLS\)](#) and [resolution](#).]

- **Scaling optimization of pseudowire service logical interfaces (MX Series)**—Starting in Junos OS Release 17.2R1, the scaling limit for pseudowire service logical interface is increased from 256 to 2000 per Modular Port Concentrator (MPC) and from 2000 to 7000 per device. MX Series routers with Junos Trio based line cards help to imitate and leverage functionality of an Ethernet interface.

#### NOTE:

- Pseudowire service logical interface is supported by MPC with Junos Trio chipset only.
- A *commit check* is performed when you issue the **commit** command at configuration mode. Commit check fails when the scaling limit exceeds the value of 2000 per Flexible PIC Concentrator (FPC) and 7000 per device.

### Network Management and Monitoring

- **MIB enhancement for jnxPPPoESubIfTable and jnxSubscriberTable tables (MX Series)**—Starting in Junos OS Release 17.2R1, you can correlate information between the jnxPPPoESubIfTable and jnxSubscriberTable tables. Prior to Junos OS Release 17.2R1, you could not correlate information between the two tables because they are indexed differently. Now, the jnxPPPoESubIfTable can provide a subscriber session ID, which corresponds to each PPPoE session. This ID can be used to correlate information in the jnxSubscriberTable. Additionally, the physical interface and underlying interface names for a subscriber session are now available in the jnxSubscriberTable.
- **New indicators for the jnxLEDState MIB (MX960, MX2020, and MX2010 routers)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E include the following indicators for the jnxLEDState MIB object in the jnxLEDEntry MIB table:
  - off—Offline, not running
  - blinkingGreen—Entering state of ok, good, normally working
- **Support for kernel features on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E support the following features:



- Addressing the IPv6 NDP DoS issue —You can address the IPv6 Neighbor Discovery Protocol (NDP) denial-of-service (DoS) issue at the Routing Engine by using NDP inspection or protection to prioritize NDP activities on the Routing Engine.
- Maximum period for autogeneration of keepalives by the kernel using precision timer feature—Precision timers in the kernel automatically generate keepalives on behalf of BGP for a specified maximum period of time after a switchover event from standby to master.
- IPv6 support for traceroute with AS number lookup—IPv6 is supported for traceroute with the **as-number-lookup** option. Traceroute is an application used to display a list of routers between the device and a specified destination host.
- Targeted aggregated Ethernet distribution—You can direct traffic through specified links of a logical interface of an aggregate Ethernet bundle that is configured without link protection. By configuring targeted aggregated Ethernet distribution, you can create distribution lists consisting of specific child member links.
- Reduction in the number of IPCs between master agent and subagent- The SNMP GetBulk requests are converted to AgentX GetNext for the repetitions specified in the request. This might result in several inter-process communication (IPCs) between the master agent snmpd and subagent AgentX in proportion to the number of max-repetitions specified in the GetBulk request. The number of IPCs between the master agent and subagent can be reduced by translating GetBulk requests with a high max-repetitions count to a single request between the master agent snmp and the subagent AgentX.
- I3-level liveness detection mechanism for child links of ethernet LAG interface.
- Match-string functionality for efficient syslog message filtering.
- **Support for features on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E support the following features:
  - LDP in an IPv6 network only, and in an IPv6 or IPv4 dual-stack network.
  - The IS-IS protocol can restrict flooding of LSAs to control sharing of routes between multiple level-2 metro ring networks.
  - For routers operating in enhanced IP Network Services mode, you can configure a threshold that triggers fast failover in next-generation MVPNs with hot-root standby on the basis of aggregate flow rate.
  - Control word feature for LDP VPLS and FEC 129 VPLS.
  - You can specify route prefix priority of high or low through the existing import policy in protocols. Through priority, you can control the order in which the routes get updated from LDP/OSPF to RPD, and RPD to kernel.
  - RSVP with traffic engineering (RSVP-TE) protocol extensions for fast reroute (FRR) facility protection to allow greater scalability of LSPs and faster convergence times.
  - The Junos OS implementation of MPLS RSVP-TE is scaled to enhance the usability, visibility, configuration, and troubleshooting of label-switched paths (LSPs).

- Tables and objects defined in RFC 5132, *IP Multicast MIB*, except the `ipMcastZoneTable` table.
- Agent Capabilities MIB provides information about the implementation characteristics of an Agent subsystem in a network management system.
- You can prioritize BGP route updates by using output queues.
- Flow-aware transport (FAT) label for BGP-signaled pseudowires such as Layer 2 VPN and VPLS.
- The NLRI format available for BGP VPN multicast is changing from the existing format of SAFI 128 to SAFI 129 as defined in RFC 6514.
- You can use the **import-labeled-routes** statement at the **[edit routing-instances routing-instance-name protocols vpls]** hierarchy level to specify one or more nondefault routing instances where you want MPLS pseudowire labeled routes to be leaked from the `mpls.0` path routing table in the master routing instance.
- You can configure BGP-ORR with IS-IS as the interior gateway protocol (IGP) on a route reflector to advertise the best path to the BGP-ORR client groups by using the shortest IGP metric from a client's perspective, instead of the route reflector's view.
- **RPM timestamping extension on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E support timestamping of RPM probes in the Packet Forwarding Engine host processor. You can enable this feature by including the **hardware-timestamp** statement at the **[edit services rpm probe probe-name test test-name]** hierarchy level.

[See [hardware-timestamp](#).]

**Support for RPM probes with IPv6 sources and destinations on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address. To specify the destination IPv6 address used for the probes, include the **target (url ipv6-url | address ipv6-address)** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. You can also define the RPM client or the source that sends RPM probes to contain an IPv6 address. To specify the IPv6 protocol-related settings and the source IPv6 address of the client from which the RPM probes are sent, include the **inet6-options source-address ipv6-address** statement at the **[edit services rpm probe owner test test-name]** hierarchy level.

[See [probe-type](#).]

- **SNMP support for monitoring tunnel statistics (MX Series)**—Starting in Junos OS Release 17.2R1, SNMP MIB `jnxTunnelStat` supports monitoring of tunnel statistics for IPV4 over IPV6 tunnels. This is a new enterprise-specific MIB, Tunnel Stats MIB, that currently displays three counters: tunnel count in `rp`, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine. This MIB can be extended to support other tunnel statistics. The MIB is defined in `jnx-tunnel-stats.txt`. This MIB is attached to `jnxMibs`.

### ***Operation, Administration, and Maintenance (OAM)***

- **Support for Ethernet OAM features on MPC7E, MPC8E, and MPC9E (MX Series)**---Starting in Release 17.2R1, Junos OS supports the following Ethernet OAM features on MPC7E, MPC8E, and MPC9E:
  - IEEE 802.3ah standard for OAM
  - IEEE 802.1ag standard for OAM
  - Technical Specification MEF-36-compliant performance monitoring
  - Configuration of multiple maintenance endpoints (MEPs) for a single combination of maintenance association and maintenance domain IDs for interfaces belonging to a particular VPLS service or bridge domain.
- **Enhanced scale support for MIPs and MEPs per chassis (MX Series routers with MPCs)**—Starting in Junos OS Release 17.2R1, Junos OS supports 32000 maintenance intermediate points (MIPs) and maintenance association end points (MEPs) each per chassis for bridge domain and VPLS domain interfaces. Increasing the number of MIPs and MEPs per chassis for specific domains enables effective Ethernet OAM deployment in scaling networks. To increase the number of MIPs and MEPs supported per chassis, enable enhanced connectivity fault management (CFM) by using the **enhanced-cfm-mode** command. To support enhanced CFM, configure the network services mode on the router as **enhanced-ip**. If you do not configure the network services mode, then Junos OS supports only 8000 MIPs and MEPs each per chassis.

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

- **Support for Packet Forwarding Engine features on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 16.1R4 and 17.2R1, MPC7E, MPC8E, and MPC9E support the following features:
  - **Protection against label spoofing or errant label injection across ASBRs**—You can use regular BGP implicit and explicit export policies to restrict VPN ASBR peer route advertisement to a given routing instance.
  - **Policer overhead adjustment at the interface level**—The policer overhead adjustment for ingress and egress policers is defined on a per IFL/direction granularity in order to address MEF CE 2.0 requirements to the bandwidth profile.
  - **Configuration support to improve MC-LAG Layer 2 and Layer 3 convergence**—You can configure multichassis link aggregation (MC-LAG) interfaces to improve Layer 2 and Layer 3 convergence time to subsecond values when a multichassis aggregated Ethernet link goes down or comes up in a bridge domain.
  - **Support for packet-marking schemes on a per-customer basis**—A packet-marking scheme, called policy map, enables you to define rewrite rules on a per-customer basis.

- **MPLS encapsulated payload load-balancing**—Configure the **zero-control-word** option to indicate the start of an Ethernet frame in an MPLS Ethernet pseudowire payload.
- **Latency fairness optimized multicast**—You can reduce latency in the multicast packet delivery by optimizing multicast packets sent to the Packet Forwarding Engines.

### *Routing Protocols*

- **Support for BGP link-state distribution with SPRING extensions (MX Series)**—Starting in Junos OS Release 17.2R1, BGP link-state extensions export source packet routing in networking (SPRING) topology information to software-defined networking controllers. Controllers can get the topology information by either being a part of an interior gateway protocol (IGP) domain or through BGP link-state distribution. BGP link-state distribution is supported on inter-domain networks and provides a scalable mechanism to export the topology information. This feature benefits networks that are moving to SPRING but also have RSVP deployed, and continue to use both SPRING and RSVP in their networks.

[See [Link-State Distribution Using BGP Overview](#).]

- **Support for SRGB in SPRING for IS-IS (MX Series in enhanced IP Mode)**—Starting with Junos OS Release 17.2R1, you can configure the segment routing global block (SRGB) range label used by source packet routing in networking (SPRING). Currently Junos OS allows you to configure only node segment indices. The value of the start label depends on the dynamic label available in the system. The labels from this SRGB range are used for SPRING in the IS-IS domain. The labels advertised are more predictable and deterministic across the segment routing domain.
  - To configure the starting index value of the SRGB label block, use the **start-label start-label-block-value** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.
  - To configure the index range of the SRGB label block, use the **index-range value** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.

[See [source-packet-routing](#)]

- **Support for anycast and prefix segments in SPRING for IS-IS protocols (MX Series)**—Starting in Junos OS Release 17.2R1, there is support for anycast segment identifiers (SIDs) and prefix SIDs in source packet routing in networking (SPRING). Currently there is support for node segments in Junos OS supports node segments for IPv4 and IPv6 when they are explicitly configured under the **[edit protocols isis source-packet-routing node-segments]** hierarchy. Now you can provision prefix SIDs along with node SIDs to prefixes that are advertised in IS-IS protocols through policy configuration. Anycast SID is a prefix segment that identifies a set of routers. You can configure **explicit-NULL** flag on all prefix SID advertisements and configure **shortcut** for SPRING routes using **family inet-mpls** or **family inet6-mpls**.

[See [Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol](#)]

- **FIB scaling and performance enhancements (MX Series)**—Starting in Junos OS Release 17.2R1, the Packet Forwarding Engine is enhanced to scale and support a higher number of routes in the forwarding information base (FIB), also known as forwarding table. However, during graceful Routing Engine switchover (GRES), when there are ten million IPv4 routes in the forwarding table, there is traffic loss.

This traffic loss is not seen when a routing protocol process (rpd) runs in warm standby mode. We currently do not support unified ISSU and NSR at this scale.

- **Support for unique AS path count (MX Series)**—Starting with Junos OS Release 17.2R1, you can configure a routing policy to determine the number of unique autonomous systems (ASs) present in the AS path. The unique AS path count helps determine whether a given AS is present in the AS path multiple times, typically as prepended ASs. In earlier Junos releases it was not possible to implement this counting behavior using the **as-path** regular expression policy. This feature permits the user to configure a policy based on the number of AS hops between the route originator and receiver. This feature ignores ASs in the **as-path** that are confederation ASs, such as **confed\_seq** and **confed\_set**.

To configure AS path count, include the **as-path-unique-count count (equal | orhigher | orlower)** configuration statement at the **[edit policy-options policy-statement policy\_name from]** hierarchy level.

- **TCP IP network stack parallelization for virtual Route Reflector devices**—Starting in Junos OS Release 17.2R1, you can enable TCP IP network stack parallelization on virtual Route Reflector (vRR) devices by using the **set system enable network-stack parallel-mode** command. Network stack parallelization can help increase performance for TCP protocol users, depending on application behavior.

[See [Understanding Virtual Route Reflector](#).]

- **Optimization of rpd resolver module (MX Series)**—Starting in Junos OS Release 17.2R1, the resolver module of the routing protocol process (rpd) is optimized to increase the throughput of inbound processing flow, accelerating the learning rate of the routing information base (RIB) and the forwarding information base (FIB), also known as routing table and forwarding table, respectively.

This enhancement makes the rpd CPU-efficient, and benefits networks with high scale internal BGP (IBGP) routes in the inet.0 and inet6.0 routing tables, internal BGP multipath routes, high RSVP equal-cost multipath routes, and virtual route reflector deployments where a forwarding state is not built.

[See [BGP Route Resolution Overview](#).]

### Services Applications

- **Inline video monitoring for IPv4-over-MPLS flows (MX Series)**—Starting in Junos OS Release 17.2R1, MX Series routers support the inline video monitoring of IPv4-over-MPLS flows to measure media delivery index (MDI) metrics. MDI information enables you to identify devices that are causing excessive jitter or packet loss for streaming video applications.

[See [Configuring Inline Video Monitoring](#).]

- **Configurable interval and threshold values for IKEv2 dead peer detection (MX Series with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.2R1, you can set the dead peer detection (DPD) interval and threshold options in IPsec rules for IKEv2 security associations. The interval is the amount of time that the peer waits for traffic from its destination peer before sending a DPD request packet, and the threshold is the maximum number of unsuccessful DPD requests to be sent before the peer is considered unavailable.

[See [Configuring IPsec Rules](#).]

- **Introducing the Junos OS URL filtering feature (MX Series)**—Starting in Junos OS Release 17.2R1, you can use URL filtering to filter which Web content is accessible to users based on a set of criteria or *template*. Blacklisted URLs are maintained in a URL database file. These URLs are resolved by the URL filtering process (url-filterd) on the Routing Engine to a list of IP addresses that are downloaded to the URL Filter Plugin (jservices-urlf), which is added to the Multiservices PIC management process (msprmand) running on the service PIC.
- **Support for inline 6rd and 6to4 (MX2020)**—Starting in Junos OS Release 17.2R1, you can also configure inline IPv6 rapid deployment (6rd) or IPv6 to IPv4 (6to4) on an MX2020 router on MPC7Es, MPC8Es, and MPC9Es. You can use the inline capability to avoid the cost of using services PICs for required tunneling, encapsulation, and de-encapsulation processes. Anycast is supported for 6to4 using next-hop service interfaces. Hairpinning is also supported for traffic between 6rd domains.

[See [Configuring Inline 6rd](#), [show services inline software statistics](#), and [clear services inline software statistics](#).]

- **Support for Junos Traffic Vision for multiple flow collectors for inline flow monitoring on MX Series routers**—Starting in Junos OS Release 17.2R1, you can export flow records generated by inline flow monitoring to four collectors under a family with the same source IP address. The Packet Forwarding Engine can export the flow record, flow record template, option data, and, option data template packet to all configured collectors. You can configure the multiple collectors at the **[edit forwarding-options sampling instance *instance name*]** hierarchy level.

**NOTE:** You cannot change the source IP address for collectors under the same family. Also, the template mapped across collectors under a family should be same.

[See [Inline Sampling Overview](#)]

- **Support for H.323 gatekeeper mode for NAT64 on MS-MPC and MS-MIC (MX Series routers)**—Starting in Junos OS Release 17.2R1, H.323 gatekeeper mode is supported in NAT-64 rules in addition to NAPT-44 rules and IPv4 and IPv6 stateful firewall rules. H.323 is a legacy VoIP protocol.

[See [ALG Descriptions](#).]

- **IPsec cleanup when local gateway address, MS-MPC, or MS-MIC goes down (MX Series router)**—Starting in Junos OS Release 17.2R1, you can enable an IPsec tunnel's service set to stop sending IKE triggers when the tunnel's local gateway IP address goes down or the MS-MIC or MS-MPC being used in the tunnel's service set goes down. In addition, when the local gateway IP address goes down, the IKE and IPsec security associations (SAs) are cleared for next-hop service sets, and go to the Not Installed state for interface-style service sets. The SAs that have the Not Installed state are deleted when the local gateway IP address comes back up.

[See [Configuring IPsec Service Sets](#).]

- **Support for AMS warm standby on MS-MPC and MS-MIC (MX Series routers)**—Starting in Junos OS Release 17.2R1, you can use the same services interface as the backup in multiple aggregated multiservices

(AMS) interfaces, resulting in an N:1 warm standby option for MS-MPCs and MS-MICs. Each warm standby AMS interface contains two members. One member is the service interface you want to protect, called the primary interface, and the other member is the secondary (backup) interface. You can use the same secondary member interface in multiple warm standby AMS interfaces.

[See [Configuring Warm Standby for Services Interfaces](#).]

- **Vendor-specific logging and reporting function templates**—Starting in Junos OS Release 17.2R1, you see a warning message when committing the configuration of a vendor-specific template for the logging and reporting function (LRF) if you do not identify the vendor with the **vendor-support** statement at the **[edit services lrf profile *profile-name*]** hierarchy level. For Junos OS Release 17.2R1, this restriction only applies to an IBM-specific template.

[See [Configuring an LRF Profile for Subscribers](#).]

- **Exchanging data more efficiently using TCP Fast Open (MX Series)**—Starting in Junos OS Release 17.2, there is an update to TCP, TCP Fast Open (TFO), that significantly improves overall network latency for short Web transfers. The key component of TFO is the TFO cookie, which is a Message Authentication Code (MAC) tag generated by the server. The client requests a TFO cookie in one regular TCP connection, and then uses it for future TCP connections to exchange data *during*, instead of *after*, the three-way handshake, saving up to one full round-trip time (RTT) over standard TCP. TFO support is for MS-MPC and MS-MIC.

- **FlowTapLite support for circuit cross connect traffic (MX Series routers)**—Starting in Junos OS Release 17.2R1, FlowTapLite sampling of circuit cross connect (CCC) traffic is supported. FlowTapLite is a lighter version of Junos Packet Vision, which lets you capture packet flows on the basis of dynamic filtering criteria. While Junos Packet Vision requires a services PIC, FlowTapLite functionality resides in the Packet Forwarding Engine.

[See [Configuring FlowTapLite](#).]

### Software-Defined Networking (SDN)

- **BFD in a VMware NSX Environment with OVSDB and VXLAN (MX Series)**—Within a Virtual Extensible LAN (VXLAN) managed by the Open vSwitch Database (OVSDB) protocol, by default, Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic is replicated and forwarded by one or more software virtual tunnel endpoints (VTEPs) or service nodes in the same VXLAN. (The software VTEPs and service nodes are collectively referred to as *replicators*.)

Starting in Junos OS Release 17.2R1, a Juniper Networks switch or Virtual Chassis that functions as a hardware VTEP in a VMware NSX environment uses the Bidirectional Forwarding Detection (BFD) protocol to prevent the forwarding of BUM packets to a non-functional replicator.

This feature is supported on MX Series routers and enables them to be provisioned in the following ways:

- MX Series router acting as DCI and Layer 2 gateway to translate VLAN traffic coming from an EVPN (a remote data center) to VXLAN traffic
- MX Series router acting as DCI to connect different OVSDB domains through EVPN



- MX Series router acting as a layer 3 gateway to route between an VXLAN domain

By exchanging BFD control messages with replicators at regular intervals, the hardware VTEP can monitor the replicators to ensure that they are functioning and are, therefore, reachable. Upon receipt of a BUM packet on an OVSDB-managed interface, the hardware VTEP can choose one of the functioning replicators to handle the packet.

Feature Explorer family: Software Defined Networking (SDN)

- **Support for Junos node slicing**—Starting in Junos OS Release 17.2R1, Junos node slicing is supported. Junos node slicing allows a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos OS control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF).

The MX Series router functions as the base system (BSYS). The BSYS owns all the physical components of the router, including the line cards and the switching fabric. The BSYS assigns line cards to GNFs.

The Juniper Device Manager (JDM) software orchestrates GNF VMs.

In JDM, a GNF VM is referred to as a virtual network function (VNF).

A GNF thus comprises a VNF and a set of line cards.

JDM and VNFs are hosted on a pair of external industry standard x86 servers.

To set up Junos node slicing, you need an MX960 or MX2020 router and two x86 servers. The server host operating system must be Red Hat Enterprise Linux 7.2 or Ubuntu 16.04 LTS.

### ***Subscriber Management and Services***

- **PIM support for enhanced subscriber management (MX Series)**—Starting in Junos OS Release 17.2R1, MX Series routers support the Protocol Independent Multicast (PIM) protocol for enhanced subscriber management. You can use the **protocols pim** command at the **[edit dynamic-profiles profile-name]** hierarchy level to enable PIM for subscribers within the specified profile. To selectively disable PIM for an individual subscriber, use the **PIM-enable** RADIUS vendor-specific attribute and set the integer value to 0.

The **routing-services** and **protocols pim** commands under the **[edit dynamic-profiles profile-name]** hierarchy level are mutually exclusive and should not be configured together in the same client dynamic profile.

[See [PIM Overview](#).]

- **DHCPv6 support for MAC address in usernames (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the client MAC address to be included in the client username for authentication for both the DHCPv6 local server and the DHCPv6 relay agent. In earlier releases, the MAC address is supported only for DHCPv4 client usernames.

[See [Creating Unique Usernames for DHCP Clients](#).]

- **Support for mapping VLAN session termination cause (MX Series)**—Starting in Junos OS Release 17.2R1, new internal identifiers indicate the reasons that autoconfd initiates termination of individual VLAN



out-of-band subscriber sessions. In earlier releases, the termination cause for a VLAN session is always 6 (administrative reset) and cannot be modified.

The session termination causes map to default code values that are reported in the RADIUS Acct-Terminate-Cause attribute (49) in Acct-Stop messages for the service. You can use the new **vlan** option with the **terminate-code aaa** statement at the **[edit access]** hierarchy level to remap any of the new termination causes to any number in the range 1 through 4,294,967,295.

You can use the new **vlan** option with the **show network-access aaa terminate-code vlan** command to display only the VLAN termination causes and their current code values.

[See [Understanding Session Termination Causes and RADIUS Termination Cause Codes.](#)]

- **Subscriber termination supported in dynamic-bridged GRE tunnels (MX Series)**—Starting in Junos OS Release 17.2R1, dynamic-bridged generic routing encapsulation (GRE) tunnels are created and terminated at the broadband network gateway (BNG) to support the MX Series deployed as a Wi-Fi access gateway model. Dynamic Host Configuration Protocol (DHCP) subscribers are transported through GRE tunnels as either VLAN-tagged or untagged. Subscriber services such as authentication, authorization, and accounting (AAA); address assignment; and class of service (CoS) are supported for individual DHCP subscribers within the GRE tunnels.

[See [Wi-Fi Access Gateway Overview.](#)]

- **Support for per-subscriber application-aware policy control (MX Series with MS-MPCs)**—Starting in Junos OS Release 17.2R1, the MS-MPC supports per-subscriber application-aware policy control based on Layer 7 application identification information for the IP flow (for example, YouTube) or Layer 3 and Layer 4 information for the IP flow (for example, the source and destination IP address). Subscriber application-aware policy actions can include:
  - Redirecting HTTP traffic to another URL or IP address
  - Steering with a routing instance
  - Setting the forwarding class
  - Setting the maximum bit rate
  - Setting the gating status to blocked or allowed
  - Setting the allowed burst size
  - Logging data for subscriber application-aware data sessions and sending that data in an IP Flow Information Export (IPFIX) format to an external log collector, using UDP-based transport.

[See [Understanding Application-Aware Policy Control for Subscriber Management.](#)]

- **New Junos OS predefined variables (MX Series)**—Starting in Junos OS Release 17.2R1, new Juniper Networks predefined variables are available for service sets, service filters, PCEF profiles, and PCC rules in dynamic profiles. These new predefined variables include:
  - \$junos-input-ipv6-service-filter
  - \$junos-input-ipv6-service-set

- `$junos-input-service-filter`
- `$junos-input-service-set`
- `$junos-output-ipv6-service-filter`
- `$junos-output-ipv6-service-set`
- `$junos-output-service-filter`
- `$junos-output-service-set`
- `$junos-pcef-profile`
- `$junos-pcef-rule`

[See [Junos OS Predefined Variables](#).]

- **Reduced time to provision business services with ESSM and increased business services scale (MX Series)**—Starting in Junos Release 17.2R1, Enhanced Subscriber Services Manager (ESSM) can both load and commit configurations into an ephemeral configuration database through an operation (op) script. The ephemeral configuration database is an alternate database that provides a configuration layer separate from both the static configuration database and the configuration layers of other client applications. The ephemeral commit model enables devices running Junos OS to simultaneously commit and merge changes from multiple clients and execute the commits with significantly greater throughput than when committing data to the static configuration database.

Before you commit a configuration, you must validate the op script. Committing to the ephemeral database does not perform a commit check; committing an invalid configuration might result in unexpected behavior.

- **ANCP agent adjustment of downstream data rate and overhead for SDSL, VDSL, and VDSL2 subscriber lines (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the ANCP agent to provide two independent, adjusted values to CoS for downstream subscriber traffic on frame mode DSL types (SDSL, VDSL, and VDSL2), enabling CoS to more accurately adjust the effective shaping rate for the downstream subscriber traffic. You can specify a percentage value that is applied to the actual, unadjusted data rate received in ANCP Port Up messages. You can also specify a number of bytes that is added to or subtracted from the frame overhead for the traffic.

[See [Traffic Rate Reporting and Adjustment by the ANCP Agent](#).]

- **Extended support for service-accounting, service-filter-hit, and force-premium firewall match conditions and actions (MX Series)**—Starting in Junos OS Release 17.2R1, the **service-filter-hit** firewall match condition and the **service-filter-hit**, **force-premium**, **service-accounting**, and **service-accounting-deferred** firewall actions are extended to the family **any** filter on MX Series routers. This means that the filter match conditions and actions can apply to any logical interface independent of protocol. This support is in addition to existing support on the family **inet** and family **inet6** filters. Filter precedence is also supported for family **any**, which with the **service-filter-hit** facilitates filter chaining for service filters.

[See [Firewall Filter Terminating and Nonterminating Actions for Protocol-Independent Traffic in Dynamic Service Profiles](#).]

- **Prevent DHCPv6 and ICMPv6 control packets from affecting idle timeouts (MX Series)**—Starting in Junos OS Release 17.2R1, you can use the terminating filter action **exclude-accounting** to exclude all DHCPv6 and ICMPv6 control traffic from being considered for idle-timeout detection for tunneled subscribers at the LAC.

Include this term at the **[edit firewall family inet6 filter filter-name term term-name then]** hierarchy level. Apply the filter in the dynamic profile as an input and output filter.

In earlier releases, DHCPv6 and ICMPv6 control traffic prevents the idle timeout from ever expiring, leading to incorrect detection of idle periods. When connections are charged based on the time the call is connected, this can result in high call charges.

[See [Firewall Filter Terminating Actions](#).]

- **Support for parameterized filters for protocol-independent packets (MX Series)**—Starting in Junos OS Release 17.2R1, you can use family **any** for parameterized firewall filters in dynamic service profiles. You can also specify a precedence order for family **any** filters when they are attached to a dynamic logical interface. Parameterization enables you to create basic or boilerplate filters under a dynamic profile and have specific values for certain attributes provided only when the dynamic session is activated.

[See [Parameterized Filter Nonterminating and Terminating Actions and Modifiers](#).]

- **Support for inline IP reassembly on an L2TP connection**—Starting in Junos OS Release 17.2R1, you can now configure the service interfaces on MX Series routers with MPC7E-MRATE, MPC7E-10G, MPC8E, and MPC9E to support inline IP packet reassembly on a Layer 2 Tunneling Protocol (L2TP) connection. The IP packet is fragmented over an L2TP connection when the packet size exceeds the maximum transmission unit (MTU) defined for the connection. Depending on the direction of the traffic flow, the fragmentation can occur either at the L2TP access concentrator (LAC) or at the L2TP network server (LNS), and reassembly occurs at the peer interface. (In an L2TP connection, a LAC is a peer interface for the LNS and vice versa.)

You can configure the service interfaces on the LAC or on the LNS to reassemble the fragmented packets inline before they can be further processed on the network. On a router running Junos OS, a service set is used to define the reassembly rules on the service interface. The service set is then assigned to the L2TP service at the **[edit services l2tp]** hierarchy level to configure IP reassembly for L2TP fragments.

You can view the reassembly statistics by using the **show services inline ip-reassembly statistics <fpc fpc-slot | pfe pfe-slot>** command.

[See [IP Packet Fragment Reassembly for L2TP Overview](#).]

- **Support for converged services for Routing Engine-based captive portal (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure converged services at the **[edit dynamic-profiles http-redirect-converged]** hierarchy level. CPCD rules can also be configured under the dynamic profiles stanza to achieve parameterization of the rules. This mechanism provides additional flexibility to customize the different rules on a per-subscriber basis through service attachment.

[See [Subscriber Management HTTP redirect](#).]

- **Support for converged services for MS-MPCs and MS-MICs based captive portal (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure converged services for MS-MPCs and MS-MICs. You can configure captive portal content delivery (CPCD) profiles for MS-MICs and MS-MPCs by including the service interface `ms-fpc/pic/port` statement at the `[edit service-set service set name captive-portal-content-delivery-profile profile name interface-service]` hierarchy level.

[See [Subscriber Management HTTP redirect.](#)]

- **Support for service activation through dynamic profiles at subscriber and underlying interfaces (MX Series)**—Starting in Junos OS Release 17.2R1, service activation can now dynamically apply a full range of CoS parameters to subscriber and underlying (for example, SVLAN) interfaces through dynamic profiles. Dynamic profiles support the attachment of classifiers, traffic control profiles, scheduler maps, and rewrite rules at the `[dynamic-profiles profile-name class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level.

[See [Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS.](#)]

- **Enhanced subscriber management support for external BGP on LNS interfaces (MX Series)**—Starting in Junos OS Release 17.2R1, when enhanced subscriber management is enabled and only for LNS subscribers, you can statically provision a subscriber's client IP address as the BGP neighbor IP address with the existing `neighbor` statement at the `[edit protocols bgp group]` hierarchy level. This is the same method supported in legacy subscriber management; however, as for all routing protocols in enhanced subscriber management, you must also configure the existing `routing-services` statement at the `[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]` hierarchy level.

[See [neighbor \(Protocols BGP\)](#) and [routing-services \(Enhanced Subscriber Management\)](#).]

- **Increased business services scale (MX Series)**—Starting in Junos Release 17.2R1, Enhanced Subscriber Services Manager (ESSM) can support up to 1000 business services per subscriber PPP session and up to 8000 business services per chassis. All combinations of subscribers and services are supported within those limits; for example, 8 subscribers with 1000 services each, 100 subscribers with 80 services each, and so on.

- **Support for bulk CoA (MX Series)**—Starting with Junos OS release 17.2R1, bulk change of authorization CoA is supported for RADIUS-based subscriber services. The two new Radius VSAs introduced are:

- 26-194 (Bulk-CoA-Transaction-Id)
- 26-195 (Bulk-CoA-Identifier)

This functionality enables accumulation of a series of CoA requests (bulk-CoA) and commits all of them together, in bulk, automatically.

[See [AAA Subscriber Access Radius VSA.](#)]

- **Rapid drain mode for DHCP address pools and lease timer enhancements (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the DHCP local server to stop allocating addresses from a local pool and gracefully terminate subscribers that are using addresses from that pool. When a DHCP subscriber attempts to renew the IP address from a pool configured for active drain, the DHCP local server replies with a NAK to the subscriber's T1 renewal messages, forcing a renegotiation, at which

time the server allocates a new IP address from an alternative address pool that is not configured for active drain.

Also, you can now configure the duration for T1 (renewal) and T2 (rebinding) timers for inet and inet6 in seconds. In earlier releases, you can configure the duration of these timers only as percentages. You must use either seconds or percentages for both T1 and T2 for a given pool and address family; you cannot mix the units.

[See [Configuring DHCP Local Address Pool Rapid Drain](#) and [DHCP Lease Timers](#).]

- **Traffic throughput improvements for MPC5 and MPC6 cards (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the **host-prefix-only** statement on the underlying demux interface for static or dynamic VLANs to improve datapath performance for DHCPv4 access models. This statement has the following requirements:
  - All the DHCPv4 subscribers using the underlying interface must be brought up using a 32-bit host prefix.
  - You must configure the **demux-source inet** statement. You must not configure **demux-source inet6** or **demux-source [inet inet6]**.

[See [host-prefix-only](#).]

- **New dynamic variable to create interface sets for a passive optical network (PON) (MX Series)**—Starting in Junos OS Release 17.2R1, you can use the predefined variable `$junos-pon-id-interface-set-name` to extract a portion of the DHCPv4 (Option 82, suboption 2) or DHCPv6 (Option 37) agent remote ID string inserted by the optical line terminal (OLT). The OLT must format the string with a pipe symbol (|) as the delimiter between substrings. The substring consists of the characters following the last delimiter in the agent remote ID string. The contents of the substring are determined by the customer, but can include the name and port of the OLT accessed by the CPE optical network terminal (ONT). After extraction, this substring is used as the name of an interface set and as an identifier to discriminate among individual customer circuits to be aggregated into the interface set.

[See [Extracting an Option 82 or Option 37 Substring to Create an Interface Set](#).]

- **Changes to reporting the effective shaping rate to the LNS (MX Series)**—Starting in Junos OS Release 17.2R1, the methods have changed for deriving the Tx and Rx connect speeds sent by the LAC to the LNS:
  - The **actual** method is deprecated.
  - The **service-profile** method is added to derive the value for the Tx speed from the actual CoS rate that is enforced on the L3 node based on the local policy. The upstream (Rx) speed is the value configured in the dynamic service profile with the **report-ingress-shaping-rate** statement. If this statement is not configured, the Rx speed follows the fallback procedure.
  - The **static** method, previously deprecated in Junos OS Release 15.1 is undeprecated.

[See [Subscriber Access Line Information Forwarding by the LAC Overview](#), [Transmission of Tx Connect-Speed and Rx Connect-Speeds from LAC to LNS](#), and [Configuring the LAC to Report Access Line Information to the LNS](#).]

- **Support for passing Framed-Route attributes from a RADIUS server. (MX Series)**—Starting in Junos OS Release 17.2, for routers running enhanced subscriber management, tagged subscriber host routes from a RADIUS server can be passively imported to the routing table and thus advertised by BGP. The following attributes are included: **tag**, **metric**, and **preference**. To view the attributes, use the **show system subscriber-management route prefix** command.

[See [show system subscriber-management route prefix](#).]

- **MLPPP support for LNS and PPPoE subscribers (MX Series)**—Starting in Junos OS Release 17.2, Multilink PPP (MLPPP) support is provided for static and dynamic LNS (L2TP network server) and PPPoE (Point-to-Point Protocol over Ethernet) terminated and tunneled subscribers running on MX Series with access-facing MPC2 slots. The following features are supported:
  - Mixed mode for customers with both MLPPP and single link PPP subscribers
  - Fragmentation-maps for both static and dynamic inline service **si** interfaces
  - Coexistence support for member-link IFL and the bundle IFL on different lookup engines
  - Link fragmentation and interleaving (LFI) for a single-link bundle
  - Fragment reordering optimization
- **Targeted distribution of subscriber traffic over aggregated Ethernet**—Starting in Junos OS Release 17.2R1, for a demux configuration whose underlying interface is an aggregated Ethernet interface, Junos OS provides targeted distribution of subscriber traffic while also allowing subscriber traffic redundancy. This ensures equal distribution of bandwidth and CoS resources among subscribers.

Service providers can now:

- Provide DPC and port redundancy for subscriber traffic.
- Apply per-subscriber hierarchical QoS and firewall filters on subscriber traffic over LAG.

**NOTE:** The “targeted-distribution” feature needs to be defined on all levels of the profile that require targeted functionality. For example, if you have targeted distribution enabled on **dvlan** profile and you have dynamic client profile. If targeted distribution is required on dynamic client profile, then you have to enable it.

To set targeted distribution in the demux logical interfaces configuration, use the **targeted-distribution** statement at the **[edit interfaces demux0 unit *logical-unit-number*]** hierarchy level.

To schedule an automatic periodic rebalance on an aggregated Ethernet bundle, use the **rebalance-periodic start-time <hh:mm> interval <hours>** option at the **[edit interfaces *aenumber* aggregated-ether-options targeted-options]** hierarchy level.

To provide module redundancy for demux subscribers on aggregated Ethernet bundles configured with targeted distribution, set the **logical-interface-fpc-redundancy** option at the **[edit interfaces aenumber aggregated-ether-options targeted-options]** hierarchy level.

To configure rebalance subscriber granularity, use the **logical-interface-fpc-redundancy rebalance-subscriber-granularity <rebalance-subscriber-granularity>** option or **logical-interface-chassis-redundancy rebalance-subscriber-granularity <rebalance-subscriber-granularity>** option at the **[edit interfaces ae<number> aggregated-ether-options targeted-options]** hierarchy level.

To manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled, use the **request interface rebalance <interface-name>** command.

To display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle, use the **show interfaces targeting aex** command.

To view status information about the specified demux interface, use **show interfaces demux0.<logical-interface-number>** command.

To set targeted distribution in the VLAN logical interface configuration, use the **targeted-distribution** statement at the **[edit interfaces interface-set <interface-set name> demux0 unit logical-unit-number]** hierarchy level.

- **Configurable grace period for unresponsive RADIUS servers (MX Series)**—Starting in Junos OS Release 17.2R1, you can use the **timeout-grace** statement at the **[edit access radius-options]** hierarchy level to configure a grace period that determines when an unresponsive RADIUS authentication server is marked as down or unreachable. When the server fails to respond to any of the attempts made for an authentication request, it times out, the time is noted, and the grace period begins. If the server is unresponsive for subsequent authentication requests, the grace period is checked each time the server times out. When the check determines that the grace period has expired, the server is marked as down or unreachable.

In earlier releases, the grace period is 10 seconds and is not configurable.

[See [Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable.](#)]

- **ANCP agent adjustment of cell overhead for ADSL, ADSL2, and ADSL2+ subscriber lines (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the ANCP agent to adjust the value it reports to CoS for downstream subscriber traffic on cell-mode DSL types (ADSL, ADSL2, and ADSL2+). The adjusted values enable CoS to more accurately adjust the effective shaping rate for the downstream subscriber traffic.

Use the following statements to specify number of bytes that are added to or subtracted from the cell overhead for the traffic: **adsl-bytes**, **adsl2-bytes**, or **adsl2-plus-bytes**. Use the **show ancp cos** command to view the adjustment configuration and the last updated values sent to CoS. The **show class-of-service interface interface-name** command displays the adjusted overhead values CoS has received from the ANCP agent.

[See [Configuring the ANCP Agent to Report Traffic Rates to CoS.](#)]

Virtual Chassis

- **VCP link hashing enhancements(MX Series)**—Starting in Junos OS Release 17.2R1, you can use Virtual Chassis port (VCP) link hashing more effectively. All links are equally utilized no matter how many VCP links are configured. This results in better load balancing and better utilization of VCP links under heavy traffic.  
  
[See [Guidelines for Configuring Virtual Chassis Ports.](#)]
- **Support for MX Series Virtual Chassis environment (MX Series Routers)**—Starting with Junos OS Release 17.2R1, MX240, MX480, and MX960 routers with the Routing Engine RE-S-X6-64G support the MX Series Virtual Chassis environment.

SEE ALSO

<a href="#">Changes in Behavior and Syntax   144</a>
<a href="#">Known Behavior   162</a>
<a href="#">Known Issues   167</a>
<a href="#">Resolved Issues   185</a>
<a href="#">Documentation Updates   249</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   250</a>
<a href="#">Product Compatibility   258</a>

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 145](#)
- [EVPNs | 145](#)
- [Forwarding and Sampling | 146](#)
- [General Routing | 148](#)
- [High Availability \(HA\) and Resiliency | 148](#)
- [Interfaces and Chassis | 148](#)
- [IP Tunneling | 151](#)
- [Management | 151](#)
- [MPLS | 152](#)
- [Network Management and Monitoring | 153](#)



- Routing Protocols | 155
- Services Applications | 157
- Software-Defined Networking | 157
- Software Installation and Upgrade | 157
- Subscriber Management and Services | 157
- User Interface and Configuration | 160
- VPNs | 161

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 17.2R3 for MX Series.

### Class of Service (CoS)

- **Support for 48 classifiers per family (MX Series)**—Starting with Junos OS Release 17.2R2, you can configure up to 48 classifiers per family at the `[edit class-of-service classifiers]` hierarchy level. In earlier releases, you could only configure up to 32 classifiers per family.

[See [CoS Features and Limitations on MX Series Routers](#).]

### EVPNs

- **EVPN E-Tree extended community**—Starting in Junos OS Releases 16.1R5, 17.1R2, 17.2R1 and later releases, the E-Tree leaf indication bit and leaf label in the EVPN E-Tree extended community follows the guidelines defined in the [E-TREE Support in EVPN & PBB-EVPN IET](#) IETF draft. A mixed network environment with routers running versions of Junos OS without this fix and routers with this fix would encounter unexpected forwarding behavior. Previous versions of Junos OS have the incorrect label indication bit and leaf label encoding. Previous versions of Junos OS, including Release 16.1R4, had the incorrect label indication bit and leaf label encoding.
- **EVPN extended community and ISID using standard IANA value**—Starting in Junos OS Release 17.2R1, the router MAC extended community and service identifier (ISID) sub-type values have been corrected to use the Internet Assigned Numbers Authority (IANA) standardized value. In Junos OS Release 17.1R1, when you configure EVPN extended community using a pure type 5 routing mode with VXLAN encapsulation, you might encounter routing issues with the router from another vendor.
- **Changes in the output of show route table command**—Starting in Junos OS Release 17.2R3, the output for `show route table` no longer displays the loopback address as the route distinguisher for MAC address

virtual routing and forwarding (MAC-VRF) routing instances route entries. Instead, the output now displays the route distinguisher for the evpn and virtual switch instance type.

- **Support for LSP on EVPN-MPLS**—Starting in Junos OS Release 17.2R3, 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.

## Forwarding and Sampling

- If a Packet Forwarding Engine (PFE) of an FPC is affected due to fabric path wedge errors, then as part of fabric hardening actions, the affected Packet Forwarding Engine is disabled and the associated fabric also goes offline. Fabric stream wedge occurs when the ASIC of the FPC is in the stuck state, and the ingress Packet Forwarding Engine fails to send traffic to the destination Packet Forwarding Engine. When the Packet Forwarding Engine is wedged, the fabric of the Packet Forwarding Engine goes offline. The output of **show chassis fabric fpcs** and **show chassis fabric plane** commands show a new state for the Packet Forwarding Engine as **Fabric Disabled**.

```
user@router> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    ... PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    ...
  FPC 1
    PFE #0
      Plane 0: Plane enabled
      Plane 1: Plane enabled
      Plane 2: Plane enabled
      ...
    PFE #1
      Plane 0: Plane enabled
      Plane 1: Plane enabled
      Plane 2: Plane enabled
    ...FPC 2
  PFE #0
    : Fabric Disabled
  PFE #1
```

```
Plane 0: Plane enabled  
Plane 1: Plane enabled  
Plane 2: Plane enabled  
...
```

You can use the **request chassis fabric pfe *pfe-number* fpc-*fpc-number* offline** command to offline any Packet Forwarding Engine. There is no *online* option for this statement. To bring the Packet Forwarding Engine back online, you must restart the FPC.

## General Routing

- **Support for deletion of static routes when the BFD session goes down (MX Series)**—Starting with Junos OS 17.2R2, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session-down message.

## High Availability (HA) and Resiliency

- **In Graceful Routing Engine Switchover (GRES) configuration, use only `vmhost reboot` command on MX2008 routers**—In Junos OS Release 17.2R1, you must use the `vmhost reboot` command instead of the `request system reboot` command on MX2008.

## Interfaces and Chassis

- **Support for maximum queues configuration on MPC7E, MPC8E, and MPC9E (MX Series)**—You can configure the maximum number of queues per MPC on MPC7E, MPC8E, and MPC9E. By default, these MPCs operate in per-port queuing mode.

You can use the `set chassis fpc slot-number max-queues queues-per-line-card` command to configure the number of queues per MPC. The possible values for `queues-per-line-card` are 8k, 16k, 32k, 64k, 128k, 256k, 512k, or 1M.

Per-unit scheduling and hierarchical queuing on MPC7E, MPC8E, and MPC9E are licensed features.

You cannot configure the `max-queues` and the `flexible-queuing-mode` statements at the same time.

You use the `flexi-queuing-mode` statement to configure a maximum of 32,000 queues per MPC.

If the `max-queues` statement is *not* configured, which is the default mode, the MPC starts with a message similar to the following:

**FPC 0 supports only port based queuing. A license is required for per-VLAN and hierarchical features.**

If the `max-queues` statement is configured and the value is less than or equal to 32,000, the MPC starts with a message similar to the following:

**FPC 0 supports port based queuing and is configured in 16384 queue mode. A limited per-VLAN queuing license is required for per VLAN and hierarchical queuing features.**

If the `max-queues` statement is configured and the value is greater than 32,000, the MPC starts with a message similar to the following:

**FPC 0 supports port based queuing and is configured in 524288 queue mode. A full scale per-VLAN queuing license is required for per VLAN and hierarchical queuing features.**

[See [Understanding Hierarchical Scheduling for MIC and MPC Interfaces](#) and [Flexible Queuing Mode Overview](#).]

- **Changes to show interfaces *interface-name* extensive Output**—Starting in Junos OS Releases 15.1R7, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the **MAC Control Frames** field of the **show interface *interface-name* extensive** command for a specified 10-Gigabit Ethernet interface displays a value of zero. In previous releases, the value for this field was calculated. Because of continuous traffic and as a result of the calculations, the value displayed for this field changed continuously.
- **Displaying accurate value of estimated BER in show interfaces (10-Gigabit Ethernet) command**—During autorecovery, when the **show interfaces** command for 10-Gigabit Ethernet interface is executed, the **Estimated BER** field displays **Recovery Under Progress** instead of **<= 1E-16**, as the estimated BER is not known during autorecovery.

Before:

```
Physical interface: xe-5/1/0, Enabled, Physical link is Down
  Interface index: 311, SNMP ifIndex: 1503
  Description: XX - ENNI LAG to PE-13 xe-11/3/1
  Link-level type: Flexible-Ethernet, MTU: 9130, MRU: 9138, LAN-PHY mode,
  Speed: 10Gbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: None,
  Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Link-Layer-Down SNMP-Traps Internal: 0x4004000
  CoS queues    : 8 supported, 8 maximum usable queues
  Schedulers    : 0
  Current address: 00:17:cb:d4:67:c7, Hardware address: 00:17:cb:d4:66:c4
  Last flapped  : 2016-01-12 13:37:33 EST (00:06:56 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : LINK
  Active defects: LINK
  PCS statistics
    Bit errors          Seconds
    Errored blocks      8956
  Link Degradate :
    Link Monitoring      : Enable
    Link Degradate Set Threshold : 1E-8
    Link Degradate Clear Threshold : 1E-12
    Link Degradate War Set Threshold : 1E-9
    Link Degradate War Clear Threshold : 1E-11
    Estimated BER        : <= 1E-16
    Link-degrade event    : Seconds          Count
    State
                                5521          2
    Defect Active
  Interface transmit statistics: Disabled
```

After:

```
Physical interface: xe-5/1/0, Enabled, Physical link is Down
  Interface index: 311, SNMP ifIndex: 1503
  Description: XX - ENNI LAG to PE-13 xe-11/3/1
  Link-level type: Flexible-Ethernet, MTU: 9130, MRU: 9138, LAN-PHY mode,
  Speed: 10Gbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: None,
  Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Device flags      : Present Running Down
  Interface flags: Hardware-Down Link-Layer-Down SNMP-Traps Internal: 0x4004000
  CoS queues       : 8 supported, 8 maximum usable queues
  Schedulers       : 0
  Current address: 00:17:cb:d4:67:c7, Hardware address: 00:17:cb:d4:66:c4
  Last flapped    : 2016-01-12 13:37:33 EST (00:06:56 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Active alarms   : LINK
  Active defects  : LINK
  PCS statistics
    Bit errors          Seconds
    Bit errors          7
    Errored blocks      8956
  Link Degradation :
    Link Monitoring      : Enable
    Link Degradation Set Threshold : 1E-8
    Link Degradation Clear Threshold : 1E-12
    Link Degradation War Set Threshold : 1E-9
    Link Degradation War Clear Threshold : 1E-11
    Estimated BER        : Recovery Under Progress
    Link-degradation event : Seconds          Count
  State
    5521                2
  Defect Active
  Interface transmit statistics: Disabled
```

[See [show interfaces \(10-Gigabit Ethernet\)](#).]

- **Aggregate Ethernet IFL (logical interface) targeted distribution feature now provides four level of prioritization**—Starting in Junos OS Release 17.2R1, the aggregate Ethernet logical interface targeted distribution feature supports four levels of prioritization. If you configure all three distribution lists—primary, backup, and standby---then Junos OS will not implicitly add member interfaces to these distribution lists. That is, if any member interface is not defined in either of the configured lists, then it will be assigned a weight higher than the standby list weight and thus will be used only when all the interfaces in all three configured lists are down. This provides four levels of prioritization.

Previously, traffic would fail over to the standby links when both primary and backup links failed.

- **Deprecated maximum transmission unit configuration option for virtual tunnel interfaces**—In Junos OS Release 17.2R2, you cannot configure the maximum transmission unit (MTU) size for virtual tunnel (vt) interfaces because the **mtu bytes** option is deprecated for vt interfaces. Junos OS sets the MTU size for vt interfaces by default to *unlimited*.
- **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)**—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.

## IP Tunneling

- **Deprecated no-path-mtu-discovery configuration option for ipip6 tunnels**—Starting in Junos OS Release 17.2R1, the **no-path-mtu-discovery** configuration statement in the **[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]** and **[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]** hierarchies is no longer available for ipip6 tunnels.

## Management

- **Changes to the rfc-compliant configuration statement (MX Series)**—Starting in Junos OS Release 17.2R1, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. If you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level and request configuration data in a NETCONF session on a device running Junos OS Release 17.2R1 or later, the NETCONF server sets the default namespace for the **<configuration>** element in the RPC reply to the same namespace as in the corresponding YANG model.

[See [Configuring RFC-Compliant NETCONF Sessions](#) and [rfc-compliant](#).]

- **Enhancement to the Junos Telemetry Interface (MX Series)**—Starting in Junos OS Release 17.2R1, the values displayed in the **oper-status** field for data streamed through gRPC for the physical interfaces sensor have changed.

The following values are now displayed to indicate the operational status of an interface:

- operational status up—**UP**
- operational status down—**DOWN**
- operational status unknown—**UNKNOWN**
- **Junos OS YANG module namespace and prefix changes (MX Series)**—Starting in Junos OS Release 17.2R1, Junos OS YANG modules are specific to a device family, and each module's namespace includes

the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. In earlier releases, Junos OS YANG modules used only a unique identifier to differentiate the namespace for each module, and the prefix for all **juniper-command** modules was **jrpc**.

Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**. The Junos OS YANG extension modules, **junos-extension** and **junos-extension-odl**, use the **junos** device family identifier in the namespace, but the modules are common to all device families.

[See [Understanding Junos OS YANG Modules](#).]

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is `/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/`

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

## MPLS

- **Bandwidth underflow sample on LSPs (MX Series)**—Starting in Junos OS Release 16.1R5 and 17.2R2, all zero value bandwidth samples are considered as underflow samples, except for the zero value samples that arrive after an LSP comes up for the first time, and the zero value samples that arrive first after a Routing Engine switchover.
- Prior to Junos OS Release 17.2R1, incoming MPLS labels from the following ranges can be used for static VPLS LSI-based services (Range-R1) and non-LSI-based services (Range-R2), by default:
  - Range-R1: [29696 to 41983]
  - Range-R2: [1000000 to 1048575]

Starting with Junos OS Release 17.2R1 and subsequent releases, any device operating in an enhanced-IP mode cannot use the range R1 for default assignment of incoming static VPLS LSI-based labels. However, range R2 works the same on releases prior to 17.2R1 and subsequent Junos OS Releases.

- Starting in Junos OS Release 16.1R4-S8, 16.1R6-S2, 16.1R7, 16.2R3, 17.1R3, and 17.2R3, the previously hidden configuration statement, **session**, can be configured at the **[edit protocols ldp]** hierarchy level. This statement enables you to configure the LDP session parameters by specifying the session destination address.

[See [session](#).]

- **Support for inet.0 and inet.3 labeled unicast BGP route for protocol LDP (MX Series)**--- Starting in Junos OS Release 17.2R3, LDP egress policy is supported on both inet.0 and inet.3 routing Information bases (RIBs) also known as routing table 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 Release 16.1R1, LDP egress policy was supported in inet.3 RIBs, in addition to inet.0 RIBs, for labeled-unicast BGP routes.

- **Disable M-LDP from using RSVP-TE LSPs for tunneling**—Starting in Junos OS Release 12.3R1, Junos OS provides support for Multipoint LDP (M-LDP) for Targeted LDP (T-LDP) sessions with unicast replication, in addition to link sessions. As a result, the current default behavior of M-LDP over RSVP tunneling is similar to unicast LDP.

However, because T-LDP is chosen over LDP and link sessions to signal point-to-multipoint LSPs, you can enable LDP natively throughout the network, so the point-to-multipoint LSPs take the LDP paths.

[See [p2mp \(Protocols LDP\)](#).]

- **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

- **Hard-coded RFC 3635 MIB OIDs updated (MX Series)**—Starting in Junos OS Release 17.2R1, the following RFC 3635 MIB OIDs have been updated as default values:
  - dot3StatsFCSErrors and dot3HCStatsFCSErrors, framing errors
  - dot3StatsInternalMacReceiveErrors and dot3HCStatsInternalMacReceiveErrors, MAC statistics: Total errors (Receive)
  - dot3StatsSymbolErrors and dot3HCStatsSymbolErrors, code violations
  - dot3ControlFunctionsSupported, flow control
  - dot3PauseAdminMode, flow control
  - dot3PauseOperMode, auto-negotiation
- **MIB buffer overruns can only be counted under ifOutDiscard (MX Series)**---The change done for PR 1140400 introduced a customer-visible behavior change (CVBC) in which qdrops (buffer overruns) were counted under ifOutErrors along with ifOutDiscards. This is against RFC 2863, in which buffer overruns should only be counted under ifOutDiscards and not under ifOutErrors. In Junos OS Release 17.2R1, this is now fixed.
- **Update to SNMP support of apply-path statement (MX Series)**---In Junos OS Release 17.2R1, SNMP implementation for the **apply-path** configuration statement supports only two lists:
  - **apply-path "policy-options prefix-list <list-name> <\*>"**

This configuration has been supported from day 1.

- **apply-path "access radius-server <\*>"**

This configuration is supported as of this release.

- **Enhancement to SNMPv3 traps for contextName field (MX Series)**—Starting in Junos OS Release 17.2R1, the contextName field in SNMPv3 traps generated from a non-default routing instance is populated with the same routing-instance information as is given in SNMPv2 traps. SNMPv2 traps provide the routing-instance information as context in the form of context@community. This information gives the network monitoring system (NMS) the origin of the trap, which is information it might need. But in SNMPv3, until now, the contextName field was empty. For traps originating from a default routing instance, this field is still empty, which now indicates that the origin of the trap is the default routing instance.
- **SNMP syslog messages changed (MX Series)**—In Junos OS Release 17.2R1, 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](#).]

- **Change in default log level setting (MX Series)**—In Junos OS Release, 17.2R3, the following changes were made in default logging levels:

Before this change:

- SNMP\_TRAP\_LINK\_UP was LOG\_INFO for both the physical (IFD) and logical (IFL) interfaces.
- SNMP\_TRAP\_LINK\_DOWN was LOG\_WARNING for both the physical (IFD) and logical (IFL) interfaces.

After this change:

- IFD LinkUp -> LOG\_NOTICE (since this is an important message but less frequent)
- IFL LinkUp -> LOG\_INFO (no change)
- IFD and IFL LinkDown -> LOG\_WARNING (no change)

See the [MIB Explorer](#).

- **Need to reconfigure SNMPv3 configuration after upgrade (MX Series)**—In Junos OS Release 17.2R2, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. This is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. It used to be that customers had to reconfigure SNMPv3 every time after a reboot. This problem was fixed. If you upgrade, you must still reconfigure SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3 again. To

reconfigure SNMP v3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters.

[See [Configuring the Local Engine ID.](#)]

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

## Routing Protocols

- **IPv6 neighbor reachability stale time range modified**—Starting with Junos OS Release 17.2R1, the stale time range of IPv6 neighbor reachability confirmation has changed from [1..1200] to [1..18000]. You can configure **nd6-stale-time** of upto 5 hours at the [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level.
- **Range of flow route rate-limit modified**—Starting with Junos OS Release 17.2R1, the range of flow route rate-limit is modified from [9600..1000000000000] to [0..1000000000000]. The following rate limits trigger the following actions:

Rate limit	Actions
0	discard
1-999	0 kbps
1000-1000000000000	corresponding value in kbps

- **Syslog error message RPD\_ISIS\_PREFIX\_SID\_CNFLCT to resolve conflicting prefix segment advertisement (MX Series)**—Starting in Junos OS Release 17.2R2, the **RPD\_ISIS\_PREFIX\_SID\_CNFLCT** syslog error message is emitted only when the prefix segment advertisement from the remote node is conflicting with an advertisement from the self node. This conflict happens because the same prefix segment index is assigned on different IP addresses or different prefix segment indexes are assigned to the same IP address. To rectify this conflict, identify the remote node in the network originating the conflicting prefix segment advertisement and change the prefix segment index on the local node or on the remote node.

[See [Example: Configuring Anycast and Prefix Segments in SPRING for IS-IS.](#)]

- **New option to force routers running Junos OS to advertise a zero-length next-hop address in BGP routes for flowspec families**—Beginning with Junos OS Release 17.2R1, you can force routers running Junos OS to advertise flow route updates with a zero-length next-hop address even when a valid next-hop address is present in the local routing table. This option provides backward-compatibility with earlier Junos OS releases that flap BGP sessions on receiving a nonzero-length next-hop address. Junos OS assigns a **Fictitious** type next-hop to flowspec routes received with a zero-length next-hop address. To

advertise zero-length next-hop addresses, configure this new option, **strip-nexthop**, at the **[edit protocols bgp family (inet | inet-vpn | inet6 | inet6-vpn) flow]** hierarchy level.

When **strip-nexthop** is not configured, Junos OS advertises a nonzero-length next-hop address (if one exists) for flowspec family routes just as it does for other address families.

[See [strip-nexthop](#).]

- **Format of session up time modified in show bfd session detail output**—Starting in Junos OS 17.2R1, the output of **show bfd session detail** includes the seconds in the session up time field. In earlier Junos OS releases, the session up time was displayed as **1w1d hh:mm**; the seconds were omitted when the up time was more than 24 hours. The modified format of the **session up time** is **1w1d hh:mm:ss**.

[See [show bfd session](#).]

- **Changes to the stitch label operation of transit static LSPs (MX Series)**—Starting in Junos OS Release 17.1R1, 17.1R2, and 17.2, when configuring transit static LSPs with label operation as stitch, the configured next-hop can only be a valid IP address and not an interface name. The stitch next-hop option at the **[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]** hierarchy level has changed from:

```
stitch next-hop (address | interface-name | address/interface-name);
```

to:

```
stitch next-hop (address);
```

## Services Applications

- **Change in behavior of IKE negotiation (MX Series)**—Starting in Junos OS Release 17.2R1, when you commit an IPsec configuration that includes **establish-tunnels immediately** at the **[edit services ipsec-vpn]** hierarchy level, the service set might take up to 30 seconds to initiate IKE negotiations.

## Software-Defined Networking

- 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 17.2R3, the **locally-provisioned** option in the **show mpls lsp ingress** command is behaving as expected.

## Software Installation and Upgrade

- **ZTP is supported on MX Series PPC platforms (MX Series)**—As of Junos OS Release 17.2R3, 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 Series PPC routers.

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

## Subscriber Management and Services

- **Changes to flat-file accounting statistics collection when a service deactivation fails (MX Series)**—Starting in Junos OS Release 17.2, the collection of accounting statistics when an ESSM service is deactivated has changed. When the deactivation is initiated by a Change of Authorization (CoA) message, **essmd** sends a stop request to the accounting daemon (**pfed**), which writes the stop record and marks the statistics values at that time as a new baseline value.

When the commit for the new configuration succeeds, the logical interface on which the service was deactivated is deleted.

When the commit fails, the service is restored rather than deactivated and the logical interface is not deleted. In this case, **essmd** requests the accounting daemon (**pfed**) to resume flat-file accounting for the service. The accounting daemon (**pfed**) writes an accounting start record, then resumes writing interim accounting records, where the interim statistics equal the current value minus the baseline value.

In earlier releases, if the service deactivation fails and the service is restored on the logical interface, no interim accounting statistics are collected for the interval since the stop record was written, resulting in inaccurate values.

- **DNS servers displayed by the show subscribers extensive command (MX Series)**—Starting in Junos OS

17.2, the display of DHCP domain name servers (DNS) by the **show subscribers extensive** command has changed. When DNS addresses are configured at multiple levels, the command displays only the preferred address according to this order of precedence: RADIUS > access profile > global access. The command does not display DNS addresses configured as DHCP local pool attributes.

DNS addresses from RADIUS appear in the following fields: Primary DNS Address, Secondary DNS Address, IPv6 Primary DNS Address, IPv6 Secondary DNS Address.

DNS addresses from the access profile or the global access configuration appear in the following fields: Domain name server inet, Domain name server inet6.

In earlier releases, the command displays only DHCP DNS addresses provided by RADIUS.

- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 17.2R1, the **show subscribers extensive** command displays the **IPv6 Interface Address** field only when the dynamic profile includes the \$junos-ipv6-address predefined variable.

In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

[See [show subscribers](#).]

- **Change to DHCP option 82 suboptions support to differentiate duplicate clients (MX Series)**—Starting in Junos OS Release 17.2R1, only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are considered when this information is used to identify unique clients in a subnet. Other suboptions, such as Vendor-Specific (suboption 9), are ignored.

[See [DHCPv4 Duplicate Client In Subnet Overview](#).]

- **Default L2TP resynchronization method changed and statement deprecated (MX Series)**—Starting in Junos OS Release 17.2R1, the default resynchronization method for L2TP peers in the event of a control connection failure is changed to silent failover. In earlier releases, the default method is failover-protocol-fall-back-to-silent-failover. The silent failover method is preferred because it does not keep tunnels open without traffic flow, waiting for the failed peer to recover and resynchronize. You can use the new **failover-resync** statement at the **edit services l2tp tunnel** hierarchy level to specify either failover protocol or silent failover as the resynchronization method.

Because silent failover is now the default, the **disable-failover-protocol** statement is no longer needed and has been deprecated. If you upgrade to this release with a configuration that includes this statement, it is supported, but the CLI notifies you it is deprecated.

[See [L2TP Failover and Peer Resynchronization](#).]

- **IPv6 link local addresses assigned to underlying static demux interfaces (MX Series)**—Starting in Junos OS Release 17.2R2, when you are using router advertisement for IPv6 subscribers on dynamic demux interfaces that run over underlying static demux interfaces, configure the software to use the same link-local address for both interfaces. In this case, the link-local address for the underlying interface should be based the MAC address of the underlying interface. The following statement causes the system to assign an address using the 64-bit extended unique identifier (EUI-64) as described in RFC 2373:

```

system {
  demux-options {
    use-underlying-interface-mac
  }
}

```

- **Source-specific multicast (SSM) CLI changes for dynamic IGMP and dynamic MLD (MX Series)**—Starting in Junos OS Release 17.2R2, the `ssm-map ssm-map-name` statement at the `[edit dynamic-profiles profile-name protocols (igmp | mld) interface interface-name]` hierarchy level is deprecated and is no longer supported. Instead, you define an SSM map policy with the `policy-statement` statement at the `[edit policy-options]` hierarchy level. Apply the policy for dynamic IGMP or dynamic MLD with the `ssm-map-policy ssm-map-policy-name` statement at the `[edit dynamic-profiles profile-name protocols (igmp | mld) interface interface-name]` hierarchy level.

If you upgrade from a release that does not support enhanced subscriber management (any release earlier than Junos OS Release 15.1R4) with a configuration that includes `ssm-map`, the configuration is allowed. However, the configuration has no effect and subscribers cannot log in.

- **Memory mapping statement removed for Enhanced Subscriber Management (MX Series)**— In Junos OS Release 17.2R3, use the following command when configuring database memory for Enhanced Subscriber Management:

**set system configuration-database max-db-size**

CLI support for the `set configuration-database virtual-memory-mapping process-set subscriber-management` command has been removed to avoid confusion. Using the command for subscriber management now results in the following error message:

**WARNING: system configuration-database virtual-memory-mapping not supported. error: configuration check-out failed.**

[See [Interface Configuring Junos OS Enhanced Subscriber Management](#) for an example of how to use the `max-db-size` command.]

- **Change to ICRQ message inclusion of the ANCP Access Line Type AVP (MX Series)**—Starting in Junos OS Release 17.2R3, 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.
- **Support for IPv6 all-routers address in nondefault routing instance (MX Series)**—Starting in Junos OS Release 17.2R3, 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 17.2R3, 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 \(L2TP\)](#).]

- **Wildcard supported for show subscribers agent-circuit-identifier command (MX Series)**—Starting in Junos OS Release 17.2R3, 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.

- **DHCPv6 lease renewal for separate IA renew requests (MX Series)**—Starting in Junos OS Release 17.2R3, 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.

[See [Using DHCPv6 IA\\_NA with DHCPv6 Prefix Delegation Overview](#).]

## User Interface and Configuration

- **Enhancements to the show chassis fpc errors command to display the PFE enable or disable status (MX Series)**—The **show chassis fpc errors** command output is enhanced to include information about the state of the Packet Forwarding Engine (PFE).

```
user@host> show chassis fpc errors
```

```
FPC  Level Occurred Cleared Threshold Action-Taken Action
1   Minor      0      0      10      0  LOG|
    Major      0      0      1      0  GET STATE|CM ALARM|DISABLE PFE
    Fatal      0      0      1      0  RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
              | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
2   Minor      0      0      10      0  LOG|
    Major      0      0      1      0  GET STATE|CM ALARM|DISABLE PFE
    Fatal      0      0      1      0  RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
              | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
```



```

3   Minor      0      0     10      0   LOG|
    Major      0      0      1      0  GET STATE|CM ALARM|DISABLE PFE
    Fatal      0      0      1      0   RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
| pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
5   Minor      0      0     10      0   LOG|
    Major      0      0      1      0  GET STATE|CM ALARM|DISABLE PFE
    Fatal      0      0      1      0   RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
| pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |

```

- **Junos OS prohibits configuring ephemeral configuration database instances that use the name default (MX Series)**—Starting in Junos OS Release 17.2R3, 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**.

## VPNs

- **Support for ping on a virtual gateway address**—Starting in Junos OS Release 17.2R2, Junos OS supports pinging an IPv4 or IPv6 address on the preferred virtual gateway interface. To set up support for ping, you must include both the **virtual-gateway-accept-data** and the **preferred** statements at the **[edit interfaces irb unit]** hierarchy of the preferred virtual gateway. This enables the interface on the preferred virtual gateway to accept all packets for the virtual IP address, including ping packets.

## SEE ALSO

[New and Changed Features | 113](#)

[Known Behavior | 162](#)

[Known Issues | 167](#)

[Resolved Issues | 185](#)

[Documentation Updates | 249](#)

[Migration, Upgrade, and Downgrade Instructions | 250](#)

[Product Compatibility | 258](#)

## Known Behavior

### IN THIS SECTION

- Flow-Based and Packet-Based Processing | 162
- General Routing | 162
- High Availability (HA) and Resiliency | 163
- Interfaces and Chassis | 163
- Network Management and Monitoring | 163
- Services Applications | 164
- Software-Defined Networking (SDN) | 165
- Software Installation and Upgrade | 165
- Subscriber Management and Services | 166
- User Interface and Configuration | 166

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 for MX Series..

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

### Flow-Based and Packet-Based Processing

- To avoid dropped packets, Juniper Networks recommends that you configure the **maximum-packet-length** equal to or greater than the IP header. For IPv4, set the maximum length to at least 20, and for IPv6, set the maximum length to at least 40.

### General Routing

- **Multiprotocol extensions capability code in notification message**—Starting in Junos OS Release 17.2R1, when a BGP speaker terminates a peering session, because the peer does not support Multiprotocols Extensions for BGP-4, it sends a notification message that contains the multiprotocol extensions capability as per the standard. In earlier releases, the BGP peer sends a notification message that contains internal code for unsupported NLRIs.
- On a Junos OS based platform, sometimes FPC might get stuck in offline state with the reason **Restarted by CLI command** after restarting the FPC immediately after restarting chassisd. This is because of the

fact that it takes some time for the system to stabilize after chassisd restarts. Though chassisd provides the FPC status and accepts the commands, on the back-end device, it is doing many initializations. Therefore, wait until all the PIC status are also available before issuing any command that makes FPC online, offline, and restart. [PR1275530](#)

## High Availability (HA) and Resiliency

- **Residual and baseline statistics loss from ISSU (MX Series)**—Using unified ISSU to upgrade to Junos OS Release 17.2R1 or later will result in a loss of residual and baseline statistics for interfaces, interface set specific statistics, and BBE subscriber service statistics because of an update to the statistics database. [See [Unified ISSU System Requirements](#).]
- **ISSU restrictions**—Unified ISSU is not supported for upgrading Junos OS 17.2R1 to 17.2R2.

## Interfaces and Chassis

- **An additional commit is required when reusing Virtual IP on an interface as an interface address (MX Series)**—When you reuse a virtual IP address on an interface as an interface address, you must first delete the virtual IP address configuration and commit the configuration. You must then add the interface address configuration in a subsequent commit.
- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance). But only one logical interface is assigned with the identical address after commit. During commit, only syslog messages indicating incorrect configuration are seen and no warnings. [PR1221993](#)
- 1. Delay Measurement support for 5-port 100G DWDM PIC and 5-port 100G DWDM MIC is \*ONE TIME Delay Measurement\*. This means that customers intending to measure Delay 2 points should ensure that link is up on both sides and then conduct this test one time. The result value is valid one time once the test is finished. The test result on CLI is not valid after one time measurement as the old result might show up on Routing Engine CLI. 2. Remote-loop-enable should be configured first on remote end. 3. Each time a customer wants to verify this, test has to be \*repeated\*. 4. Processing delays in each mode are different: HGFEC [For 5-port 100G DWDM MIC] being highest, SDFEC in the interim and GFEC being least for the same cable length. 5. In summary, any breakage in transmit/receive path during the delay measurement test will hinder delay measurement. This is true for all FEC modes - GFEC, SDFEC, HGFEC. 6. Currently SNMP walk is not available for delay measurement. [PR1233917](#)

## Network Management and Monitoring

- **SNMP traps for certain interfaces in Admin Down state (MX Series)**—SNMP traps are generated when an interface that supports the Digital Optical Monitoring (DOM) MIB is placed in an administrative down state. This behavior informs the operator of any interface fault, alarm, or threshold condition.

- The **MIB2D\_RTSLIB\_READ\_FAILURE: rtslib\_iflm\_snmp\_pointchange** syslog message occurs during configuration restore. This is because the mib-process sends requests to the kernel to update SNMP ifIndex for the interfaces that it is learning. If this interface is deleted from the kernel, the syslog message is generated. This interface learning by mib-process occur later once the kernel sends the ADD notification for these interfaces. There is no system impact caused by this syslog message during the configuration scenario. [PR1279488](#)

## Services Applications

- 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-Defined Networking (SDN)

- When the BSYS master Routing Engine is rebooted or shut down, the JDM-to-JDM communication, including the commit sync operation, fails. To work around this issue, commit the JDM configurations on server0 and server1 separately.
- If the GNF console remains idle for a long duration (for example, more than 10 minutes), the console might stop responding.
- Pings to the peer JDM might fail even when the connection status is shown to be up. Also, the **show server connections** command might show JDM-to-JDM ping failure issues. These ping failure issues occur when connections from the Control Board to the servers are mapped incorrectly at the JDM. Correct the mapping by verifying the connections.
- The JDM operational command **show virtual-network-functions** might sometimes show the value of the **Liveness** field as **Down** even when the GNF is up and reachable.
- The GNF VM's fxp0 interface might get slower and stop forwarding packets occasionally. When this occurs, disable the fxp0 interfaces and enable it again.

## Software Installation and Upgrade

- **Unified ISSU with active BBE subscribers using advanced services supported only to 17.2R3 and later 17.2 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 17.2 release earlier than 17.2R3. If you perform an ISSU to a 17.2 release earlier than 17.2R3, 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 17.2 release. The warning **ISSU is not supported for RPM configuration** appears.

## Subscriber Management and Services

- The **all** option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the **all** option with the **clear services l2tp destination**, **clear services l2tp session**, or **clear services l2tp tunnel** statements in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

## User Interface and Configuration

- **Modification to configurable link degrade threshold values (MX Series)**—Starting with Junos OS Releases 15.1F7 and 16.1R1, the values of the user-configurable link degrade thresholds have to be configured according to the following guidelines:
  - **set threshold value** must be greater than **warning set threshold value**
  - **set threshold value** must be greater than **clear threshold value**
  - **warning set threshold value** must be greater than **warning clear threshold value**

If the threshold values are not configured according to these guidelines, the configuration fails and a **Commit Error** message is displayed.

## SEE ALSO

[New and Changed Features | 113](#)

[Changes in Behavior and Syntax | 144](#)

[Known Issues | 167](#)

[Resolved Issues | 185](#)

[Documentation Updates | 249](#)

[Migration, Upgrade, and Downgrade Instructions | 250](#)

[Product Compatibility | 258](#)

## Known Issues

### IN THIS SECTION

- Class of Service (CoS) | 167
- EVPN | 168
- Forwarding and Sampling | 169
- General Routing | 170
- High Availability (HA) and Resiliency | 175
- Infrastructure | 175
- Interfaces and Chassis | 176
- J-Web | 176
- Layer 2 Ethernet Services | 176
- Layer 2 Features | 177
- MPLS | 177
- Platform and Infrastructure | 179
- Routing Policy and Firewall Filters | 181
- Routing Protocols | 181
- Services Applications | 184
- Subscriber Access Management | 184
- User Interface and Configuration | 184
- VPNs | 184

This section lists the known issues in hardware and software in Junos OS Release 17.2R3 for MX Series.

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

### Class of Service (CoS)

- In Junos OS Release 17.2R2, when a cascade port is configured, CoS resources are allocated to it and the corresponding CoS parameters applied on extended ports are scaled. This is done irrespective of the cascade port. If a configured cascade port goes down, nothing is done. [PR1262320](#)
- In Junos OS Release 17.2, the egress rate limit at the extended port does not work properly if you have a rate limit configuration applied at the extended port physical interface level by using the

**traffic-control-profile-remaining** and also at some of the extended port logical interfaces by using an explicit **traffic-control-profile** in **hierarchical-scheduler** mode. [PR1271719](#)

- In certain scenarios of congestion, traffic might be dropped due to non-Juniper Networks optics and generate an alarm. [PR1378392](#)

## EVPN

- Routing instances of type **evpn** configured with a VLAN ID advertises 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 must be configured with **vlan-id none** so that the Ethernet tag in advertised MAC routes is set to zero. [PR945247](#)
- A provider edge (PE) device running EVPN IRB with IGP configured in a VRF associated with the EVPN instance is unable to establish an IGP adjacency with a customer edge (CE) device attached to a remote PE device. The IGP instance running in the VRF on the PE device might be able to discover the IGP instance running on the remote CE device through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- On MX Series routers, when an instance type is changed from VPLS to EVPN, and in the same commit an interface is added to the EVPN instance, the newly added EVPN interface might not be able to come up. [PR1016797](#)
- The Layer 2 address learning process (l2ald) might generate a core file in a scaled Layer 2 setup, including bridge-domain, VPLS, EVPN, and so on. The l2ald generates a core file after a kernel page fault. In most cases, the issue recovers after the l2ald core file is generated. In a few cases, the process might need a manual restart 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. Core dumped!. [PR1142719](#)
- In an EVPN scenario with static MAC configured in the EVPN instance, the remote EVPN instance can detect the MAC route information. However, after deactivating and activating the static MAC in the EVPN instance, and then checking the MAC route information in the remote EVPN instance, no such MAC route is found in the EVPN route table. [PR1193754](#)
- In scaled-up EVPN VPWS configurations (approximately 8000 EVPN VPWS), during Routing Engine switchover, rpd scheduler slip messages might be seen. [PR1225153](#)
- When the ESI configuration on an interface is changed from **all-active** to **single-active**, and back to **all-active**, the EVPN split horizon label is not allocated and is shown as 0. [PR1307056](#)
- PBB EVPN cannot flood traffic towards a core layer. Traffic recovers by performing "restart l2-learning". In addition to this, there is a limitation in PBB EVPN active/active (A/A) unicast traffic forwarding. If entropy in the traffic is not sufficient, then uneven load balancing causes a problem on MH peer A/A



routers. This causes a drop for return traffic. These issues are applicable to MAC-in-MAC private network-to-network (PNN)-EVPN and does not affect any other scenario. [PR1323503](#)

- When an EVPN PE device (RR) is configured as single home without ESI, EVPN BGP routes from table `bgp.evpn.0` might leak into the default EVPN table (`__default_evpn__.evpn.0`) causing label leak. Such a leak might lead to all label exhaustion and, as a result, the routing protocol process (rpd) generates a core file. [PR1333944](#)
- When an EVPN route is filtered by using the CLI command **show route evpn-ethernet-tag-id**, it looks for route in all routing tables including `inet.0`. The EVPN route is not present in `inet.0` and the non-evpn route will not have the Ethernet tag, which might result in the rpd process crash. [PR1337506](#)
- In Junos OS platform, the `l2ald` daemon might crash when MAC address is processing. The MAC learning process might impact during the period of `l2ald` crash. The `l2ald` recovers itself. [PR1347606](#)
- When EVPN is configured with class-of-service-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 two seconds. It can take up to 30 seconds for the LSP to come up under either of the following conditions:
  - Creation of the policing filter and its application to the LSP through configuration occurs in the same commit sequence.
  - Load override of a configuration file that has a policing filter and policing filter application to the LSP is followed by a commit. [PR1160669](#)
- When the **push-backup-to-master** statement is configured under the **accounting-options file** section, the corresponding accounting files need to be pushed to the master Routing Engine from the standby Routing Engine. But because of a software defect, the following issues are observed.
  - The files are getting pushed from the standby Routing Engine to the master Routing Engine irrespective of whether the **push-backup-to-master** statement is configured or not.
  - The files are not getting pushed from the standby Routing Engine to the master Routing Engine if the backup option is configured as **master-only**. [PR1236618](#)
- After the **show firewall** command is executed, the **dfwinfo: tvptest:dfwlib\_owner\_create tvp driven policer\_byte\_count support 0** message is seen in message logs. This message is a cosmetic issue and it can be ignored safely. This message can be seen with the following sample configuration: **set interfaces ge-0/0/0 unit 0 family inet filter input test\_filter, set interfaces ge-0/0/0 unit 0 family inet address 100.100.100.1/24, set firewall family inet filter test\_filter term policer then policer policer\_test, set firewall policer policer\_test if-exceeding bandwidth-limit 100m, set firewall policer policer\_test if-exceeding burst-size-limit 125k, and set firewall policer policer\_test then loss-priority low**. [PR1248134](#)

- FreeBSD 10.x based Junos OS is not supported on 32-bit Routing Engines in Junos OS Release 17.1R1. [PR1252662](#)
- In some stress test conditions, the sampled process crashes and generates a core file when connecting to L2BSA and EVPN subscribers aggressively. [PR1293237](#)
- The Junos OS allows the same filter names under different families to be committed. Effective committing without giving a commit error might cause the filter criteria at the **[edit firewall family inet]** hierarchy level not to be applied on an interface. [PR1344506](#)
- On MX960, MX480, MX240, and MX80 routers with EVPN configured, if RSVP and CoS-based forwarding (CBF) are configured, the remote media access control (MAC) address might not be added to the EVPN instance forwarding table, that causes a drop in the traffic. [PR1353555](#)
- Whenever bridge firewall filter is configured and accounting is enabled on it, the filter counter is not written to the accounting file. [PR1392550](#)
- On Junos Fusion, ingress policing on SD is broken. When the statement **set interfaces layer2-policer input-policer <policer-name>** is executed, the ingress policing on AD and SD is not supported. Error is seen where traffic is not getting policed after locally switched for VLAN 100 and 101 while verifying selective local-switching functionality with 4000 VLANs. [PR1395217](#)

## General Routing

- DC-PEMs of MX104 systems might suddenly restart because of high temperature protection and might trigger a system restart. The DC-PEM temperature sensors were not monitored by the fan system algorithm, causing high temperature conditions under certain environmental conditions. MX104 with AC-PEMs are not exposed. [PR1064039](#)
- On MX Series routers with MS-MPC or MS-MIC, memory leak can be seen with `jnx_msp_jbuf_small_oc` object, upon sending millions of Point-to-Point Tunneling Protocol (PPTP) control connections (3 through 5 million) at higher cells per second (cps) (greater than 150,000 cps). This issue is not seen with up to 50,000 control connections at 10,000 through 30,000 cps. [PR1087561](#)
- On MX104 routers, when using `snmpbulkget` or `snmpbulkwalk` on chassisd related component such as `jnxOperatingEntry`, high CPU usage and slow response for the chassisd process might be seen because of a hardware limitation, which might also lead to query timeout on the SNMP client. In addition, the issue might not be seen when you use SNMP query for interface statistics. As a workaround, use `snmpget` or `snmpwalk` instead of `snmpbulkget` or `snmpbulkwalk` and include the `-t30` option when performing a SNMP query—for example, `"snmpget -v2c -c XX -t30"`. Alternatively, use the `"-t30"` option with `snmpbulkget` or `snmpbulkwalk`—for example, `"snmpbulkget -v2c -c XX -t30"`. [PR1103870](#)
- The SIP session fails when the IPv4 SIP client in the public network initiates a SIP call with the IPv6 SIP client in the private network. [PR1139008](#)
- Source-prefix filtering and protocol filtering of the carrier-grade NAT sessions provide incorrect filtering results. For example, **show services sessions extensive protocol udp source-prefix <0:7000::2>** displays incorrect filtering output 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 10-Gigabit Ethernet (xe) interface. The netdev resource developed for interface configuration determines that speed is a configurable parameter that is supported on a Gigabit Ethernet interface but not on an 10-Gigabit Ethernet (xe) interface. Therefore, the netdev interface resource cannot be used to configure an xe interface because of this limitation. [PR1181475](#)
- Junos OS might improperly bind Packet Forwarding Engine ukernel application sockets after a unified ISSU because of a bug in IP >TNP fallback logic. Because of that bug, threads running on the ukernel that relay on UDP sockets can experience connectivity issues with the host, which in turn can lead to various problems. For instance, a Simple Network Time Protocol (SNTP) client might fail to synchronize time, which in turn might lead to other problems such as failure in adjacency formation for HMAC authenticated protocols. [PR1188087](#)
- As described in RFC7130, when LACP is used and considers the member link to be ready to forward traffic, the member link must not be used by the load balancer until all the micro-BFD sessions of the particular member link are in the up state. [PR1192161](#)
- SMID daemon has stopped responding to the management requests after a jl2tpd (L2TP process) 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 logical interface that does not exist ERROR-SPECIFIC INFO: nh\_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none. Any service impact in MPC2 and MPC3 type cards is not seen.** [PR1205593](#)
- This is a rare race condition in which multiple interrupts are not handled properly on MX Series platforms with MPC7E, MPC8E, MPC9E, and PTX Series platforms with FPC3-PTX-U2/FPC3-PTX-U3, which might generate a core file. It is difficult to reproduce. The interrupt code is optimized to avoid the unnecessary call to prevent the issue. [PR1208536](#)
- In certain interface-scaling scenarios, during configuration commit or rollback, you might see an fpcx error message about a problem with **fpcx list\_get\_head list**. You can safely ignore this message because the issue is triggered by the FPGA mechanism on DPC cards for logical interface mapping (ifl\_map). Between the deletion of a physical interface and the monitoring event, the FPGA monitor mechanism checks through the stored logical interfaces. While the mechanism tries to find the family of a recently deleted logical interface that was not cleaned from the logical interface map, harmless messages might populate the log file. [PR1210877](#)
- The PTP master streams on IP and Ethernet are not supported simultaneously. [PR1217427](#)
- A unified ISSU cannot be performed from a Junos OS Release with NPU image size less than 60 MB to a Junos OS Release with NPU image size greater than 60 MB. [PR1222540](#)
- The following MICs in MPC2E-NG and MPC3E-NG are those that do not support timestamping at the physical layer (Layer 1): MIC-3D-4XGE-XFP, MIC3-3D-10XGE-SFP, MIC-3D-2XGE-XFP, and MIC-3D-20GE-SFP. The packet time error can be greater than +/- 100 seconds in these MICs. [PR1226080](#)

- When a configuration that turns the Packet Forwarding Engine off line and another configuration that brings the Packet Forwarding Engine back online are committed in quick succession, there could be RE-PFE out of synchronization errors logged in syslog. Most of the time these are benign errors, but sometimes they might result in Packet Forwarding Engine crashes. [PR1232178](#)
- OSPF is used as routing protocol between the clients and dynamic endpoint (DEP) router with TD configured. The OSPF protocol traffic brings up IPsec on spokes and the DEP router. The IPsec SAs are distributed on the DEP router. The neighbor state between the OSPF peers move to full, but after that it does not stay in that state. The states changes to init, 2-way, ex-start, and to full again. As a result, the data traffic between the routers drops. Thus tunnel distribution with protocol traffic is not supported. [PR1232277](#)
- When a virtual switch type is changed from IRB type to regular bridge, interfaces under the OpenFlow protocol are removed. The OpenFlow process (daemon) fails to program any flows. [PR1234141](#)
- The subinfo core file might be generated or the subscriber database might get stuck on the router with subscriber services during subscriber log in or the log out or any subscriber database access activity in a scaled environment. In a few scenarios, this issue might happen with or without generating a core file, where the subscriber database might get stuck resulting in the following error: **show subscribers client-type pppoe Invalid argument: smid registration failed.** [PR1234746](#)
- Sometimes, when PPPoE subscribers log in and log out from Junos OS Release 16.1, the following messages are generated: **user@devcie> 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 which no longer exists (-7 = SDB\_DATA\_NOT\_FOUND).** These messages, which indicate that the authd process is asking the SDB for records that do not exist, have no impact and can be safely ignored. [PR1236211](#)
- On MX Series with routing protocol process (rpd) in ASYNC mode, if the distributed IGMP is configured, the rpd process might crash, generating a core file. [PR1238333](#)
- The measured noise-transfer gain is around -40 dB, which is higher than the standard metric. [PR1240054](#)
- The following MICs in MPC2E-NG and MPC3E-NG are those that do not support timestamping at the physical layer (Layer 1): MIC-3D-4XGE-XFP, MIC3-3D-10XGE-SFP, MIC-3D-2XGE-XFP, and MIC-3D-20GE-SFP. The packet 2Way/T1/T4 time error can be upto +/-450 nanoseconds in these MICs. [PR1243646](#)
- Following MICs in MPC2E-NG/MPC3E-NG are those that do not support timestamping at the physical layer (Layer 1): MIC-3D-4XGE-XFP, MIC3-3D-10XGE-SFP, MIC-3D-2XGE-XFP, and MIC-3D-20GE-SFP. The packet dynamic time error might be greater than 40 nanoseconds for LF and 70 nanoseconds for HF. [PR1243871](#)
- When a certain route or next hop has been created by an application, it is assumed that it can propagate to the rest of the system. The kernel routing table (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, next hop> pair that looks legal from the application standpoint, but the KRT is not able to download it to the kernel because the kernel rejects the next hop, the <route, next hop> pair get stuck in the routing protocol process (rpd). In the meantime, the previous version of the route (L-IS-IS in this case) that was downloaded still lingers in the kernel and the Packet Forwarding Engine. [PR1253778](#)

- On MX Series routers with the XM chipset (for example, on the MPCs MPC3E, MPC4E, MPC5E, MPC6E, MPC2E-NG, and MPC3E-NG), the MPC might reboot after unified ISSU completion. [PR1256145](#)
- If there are two logical interfaces with the same VLAN ID on the logical tunnel (lt) interface, the bbe-smgd process crashes continuously. The issue is specific to Junos OS Release 15.1F5. [PR1257931](#)
- The 1PPS TE/cTE performance metric can be as high as +/-550 nanoseconds in MPC2E and MPC3E NG QoS/3D 20x 1GE(LAN)-E, SFP with no PHY-Timestamp and non-hybrid mode. [PR1263235](#)
- On an MX Series Virtual Chassis system in a scaled subscriber management scenario, if a unified ISSU is performed while the BGP protocol sessions are active and such BGP sessions are clients of BFD, then these BGP sessions might go down and come back up again, causing traffic loss. [PR1265407](#)
- If the dynamic VLAN profile does not have interface family (IFF) configuration (for example, **family PPPoE** or **family inet**), but has firewall filter configuration, firewall filter indexes are not released after the dynamic VLAN is removed. This eventually leads to the depletion of the available firewall filter indexes. [PR1265973](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- A vMX router does not detect interface link state correctly in SR-IOV mode with i40e driver. [PR1271902](#)
- If **template-referesh-rate** and **option-refresh-rate** are configured with both packets and seconds interval configuration options for inline flow monitoring, the packets interval configuration does not work. [PR1274206](#)
- The **show storm\_cntl halp** database on FPC shell might cause an FPC crash. [PR1127870](#)
- Performance of X710 NIC is lower compared to that of 82599 NIC. A 40G line rate can be achieved at 512-byte packet size for X710 NIC as compared to 256 bytes for 82,599 NIC. [PR1281366](#)
- PPPoE cannot dial in because of all padi dropped as "unknown iif" when an aggregated Ethernet configuration is deactivated or activated. [PR1291515](#)
- With OSPF and BGP route in the same subnet in inet.0 table, if the protocol next hop of the BGP aggregate route falls within the defined destination for the dynamic tunnel, there might be a recursive lookup within the Packet Forwarding Engine. [PR1292425](#)
- IPsec operations are optimized for smaller packet size (up to 1900 bytes approximately) on routers with MS-MPC and MS-MIC, thus yielding higher throughput and lower latency for more common network deployments. Customers might see slightly higher latency if there are jumbo packets in the network. [PR1307867](#)

- FPC crash is observed when a route has unilist next hops, which contain primary or backup paths, while interfaces related to unilist members go down when **set protocol rsvp load-balance bandwidth** is configured. [PR1315228](#)
- Making changes in services **traffic-load-balance** instance for one instance can lead to a refresh of the existing instances. [PR1318184](#)
- When an MX Series router with 100-megabit SFP transceiver is used on MIC-3D-20GE-SFP-E and MIC-3D-20GE-SFP-EH, the transceiver might not work if it is not from Fiberxon or Avago. [PR1344208](#)
- When the MIC is removed from the MPC, the MPC might crash. [PR1350098](#)
- During stress conditions, error log messages regarding route add, change, and delete might be incorrect. [PR1350713](#)
- When an 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 the ephemeral DB maximum size, an ephemeral DB purge might happen. Then it would purge all the commits and roll over. On this purge the mgd gives all the applications a FULL COMMIT view. And on this FULL COMMIT view, IGMP/MLD deletes all configurations and adds them back again. This might cause PIM to prune the groups on those interfaces and send join messages again. Finally, 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 256000 entries. However, in 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 though 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](#)
- On MX Series routers with MPC2E NG and MPC3E NG line cards, if the inline service interfaces are not configured with the explicit bandwidth value (for example, 1 Gbps or 10 Gbps), the default bandwidth value (100 Gbps) will be used. Therefore, only the first two inline service interfaces can be served by available hardware resources. The third and fourth inline interfaces will be not able to send out packets. [PR1355168](#)
- When you use the **show agent sensors verbose** FPC VTY command on the MPC7E, the FPC might crash. [PR1366249](#)
- The interface optic output could be nonzero value even when the port has been administratively disabled. For example, the port xe-1/0/0 has been disabled: `user@router> show configuration interfaces xe-1/0/0 disable`. However, the optic output value is nonzero value: `user@router> show interfaces diagnostics optics xe-1/0/0` Physical interface: xe-1/0/0 Laser bias current : 6.590 mA Laser output power : 0.4940 mW / -3.06 dBm <==== output value is not zero Module temperature : 41 degrees C / 106 degrees F Laser rx power : 0.6477 mW / -1.89 dBm. [PR1376574](#)
- Domain name is not reported as part of the LLDP sysname in the **show lldp neighbor** command. [PR1383295](#)

- During the Zero Touch Provisioning (ZTP) process, the default route is being cleaned up by code. Because of this, if a static default route is configured in the initial configuration (configuration file downloaded from the file server for ZTP), the route will fail to work. This might lead to ZTP failure or device access issue after ZTP. [PR1387724](#)
- On MX2020, MX2010, and MX2008 platforms with SFB2 cards installed, if a newer generation of MPC (for example, MPC type 3, 4, 5, 6, 7, 8 or 9) is installed into a slot that had MPC 3D 16x10GE MPC type 1 or MPC type 2 previously installed, the available fabric bandwidth to the new MPC card would be rate-limited due to residual programming on the fabric planes. Traffic impact is observed during peak utilization. [PR1388780](#)
- In a scaled environment with 32000 subscribers, if the command **show subscriber extensive** is issued from the CLI, and left sitting at the -(more)- prompt, any subsequent CLI session that requests **show subscriber extensive** content will see a delay up to 40 seconds before the prompt is returned. [PR1390762](#)

## High Availability (HA) and Resiliency

- The following error is seen: **error: not enough space in /var on re1**. As a workaround, the space available in **/var** should be twice the size of the target image. This is the basic requirement for unified ISSU to proceed. [PR1354069](#)

## Infrastructure

- The **/var/run** is in storage file system but it should be in memory file system. [PR1198395](#)
- The configuration command **set system ports console log-out-on-disconnect** logs the user out from the console and closes the console connection. If the configuration command **set system syslog console any warning** is used with the earlier configuration and when there is no active Telnet connection to the console, the process tries to open the console and hangs as it waits for a serial connect, which is received only by connecting to the console through Telnet. As a workaround, remove the later configuration by using **set system syslog console any warning**, which solves the issue. [PR1230657](#)
- The syslog messages are observed when one of the following CLI commands is executed: **system syslog file messages kernel any** or **system syslog file messages any any**. These syslog messages do not indicate any functionality breakage or impact. If you need to enable “anyany”, then you need to skip these logs with an appropriate match condition. [PR1239651](#)
- Sometimes OSPF flapping during unified ISSU is observed starting in Junos OS Release 16.2R2 to Junos OS Release 17.2R3. [PR1371879](#)



## Interfaces and Chassis

- During a configuration change and reuse of the VIP address on an interface, you must stop the configuration, perform a commit, and then add the interface address configuration at the next commit. [PR1191371](#)
- In a VPLS multihoming scenario, CFM packets are forwarded over the standby PE device link, resulting in duplicate packets or a loop between the active and standby link. [PR1253542](#)
- Out of sequence packets are seen with the LSQ interface. [PR1258258](#)
- In Junos OS BNG solutions, after commit event, when the configuration contains duplicate vlan-id configured on aggregated and demux interfaces, Junos OS MX Series routers might go into db prompt mode generating a kernel core file. [PR1274038](#)
- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause cfmd to crash after the upgrade. This is because of the old version of `/var/db/cfm.db`. [PR1281073](#)
- In a subscriber management scenario with dynamic demultiplexing (DEMUX) interfaces configured, in the case when subscribers belonging to one aggregated Ethernet interface are migrated to a newly configured aggregated Ethernet interface, subscribers might fail to access the device after deleting the old aggregated Ethernet configuration. [PR1322678](#)
- On MX Series routers, the bbe-smgd reports some error logs because jpppd sent out a Link Control Protocol (LCP) **config-reject** message, but the bbe-smgd misses such messages are sent in the Tx direction. It has no service impact. [PR1378912](#)
- If channelized interface coc1 is configured and FPC restart is performed, then a core file might be generated and DCD restart can be seen. In case of all other interfaces core file is not generated and normal behavior is seen. [PR1387962](#)

## J-Web

- In Junos OS, an integer signedness error occurs in GD Graphics Library (CVE-2016-3074), which results in a heap overflow when compressed data is processed. See <https://kb.juniper.net/JSA10798> for more information. [PR1218092](#)

## Layer 2 Ethernet Services

- When MSTP is configured under a routing instance, both the primary and standby VPLS pseudowires get stuck in ST state because of a bug in the software. [PR1206106](#)
- After changing the underlying physical interface for a static VLAN demux interface, the NAS-Port-ID formed is based on the previous physical interface. [PR1255377](#)
- MX Series routers might display the false positive CB alarm **PMBus Device Fail**. [PR1298612](#)



## Layer 2 Features

- On MX Series routers with MPC or MIC, with scaled VPLS instances configured to use the label-switched interfaces (LSIs) (for example, 102 instances), if the core-facing interface on the PE router flaps (for example, multiple OSPF flaps, sometimes over a period of 2 days), in a rare scenario, VPLS traffic for one instance might be sent out to another instance with an incorrect LSI MPLS label. [PR1013295](#)
- A device running Junos OS with VPLS routing instances configured on one or more interfaces might be susceptible to a buffer memory (mbuf) leak when source and destination MAC addresses of Ethernet frames with the EtherType field of IPv6 (0x86DD) are flooded into the VPLS instance. The Ethernet frames must be injected directly into a connected interface, limiting exposure to directly connected adjacent networks. See <https://kb.juniper.net/JSA10750> for more information. [PR1132568](#)
- If the router is working as a VPLS PE device, because MAC ages every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)
- On routers running Junos OS with GRES enabled, if VPLS is configured with a dynamic profile association, some traffic loss is observed when the Routing Engine switches from master to standby. This traffic loss is due to a change in the underlying database that handles the dynamic profile sessions. As a result, the VPLS connection is destroyed and re-created after a Routing Engine switchover. [PR1220171](#)

## MPLS

- When using the **mpls traffic-engineering bgp-igp-both-ribs** configuration statement 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 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 ABRs or IS-IS Layer 1 or Layer 2 routers and switch RSVP off on the other OSPF area 0/IS-IS Layer 2 routers, or avoid LDP completely and use only RSVP. [PR1048560](#)
- This issue occurs when 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 the routing protocol process (rpd) only in 32-bit mode while the other is capable of running the rpd in 64-bit mode. The situation could 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 on the new master Routing Engine might crash. As a workaround, this issue can be avoided by using the CLI command **set system processes routing force-32-bit**. [PR1141728](#)
- When **minimum-bandwidth** and **bandwidth** statements are both configured, the bandwidth selection of the LSP is inconsistent. [PR1142443](#)
- When Flow-Label (FL) is enabled for PW, the OAM packets are not sent with Flow-Label because rpd is not aware of the Flow-Label values assigned by the Packet Forwarding Engine software. Hence, the packets get dropped by Packet Forwarding Engine at the tail-end PE device. [PR1217566](#)

- In a CE-CE setup, traffic loss might be observed over the secondary LSP when the primary LSP fails over. [PR1240892](#)
- A new configuration, **protocols mpls traffic-engineering bgp-igp-both-ribs**, in the routing instance is required to make cOC work. [PR1252043](#)
- Because of the current way of calculating bandwidth, you see a minimal discrepancy between MPLS statistics and the adjusted bandwidth reported. The algorithm will be enhanced so that both values match. [PR1259500](#)
- The throughput measurement might be inaccurate when doing performance measurement is performed on an MPLS label-switched path. [PR1274822](#)
- With non-stop-routing (NSR), when the routing protocol process (rpd) on the master Routing Engine restarts, the rpd on the backup Routing Engine might also restart. [PR1282369](#)
- In case of CSPF-disabled LSPs, if the primary path ERO is changed to an unreachable strict hop, sometimes the primary path stays up with the old ERO. The LSP does not switch to standby secondary. [PR1284138](#)
- If there are some LSPs for which a router has made link protection available, and when an FPC restart causes primary link failure, a core file might be generated. [PR1317536](#)
- The Packet Forwarding Engines on MX Series platforms follow a certain conversion logic to convert MPLS-VPN labels to certain channel values and then back to MPLS-VPN labels. VPN labels with values 0x7FFFF and above ( 52,4287 and above) are affected by this conversion logic. [PR1323496](#)
- If inet address is not configured for the gr- interface, the gr- interface borrows the address from the loopback interface. Starting in 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 is borrowing an address from the loopback interface. The RSVP path lookup might fail because it gets confused by the node neighbor presence. So, the RSVP LSP might not come up when it goes through the gr- interface, which is borrowing an address from the loopback interface. [PR1340950](#)
- Executing a **restart chassisd** in an MX Series Virtual Chassis router with the following elements configured might result in 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 l3vpn 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 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 aggregated Ethernet IP **equal-cost multipath** MPLS **equal cost multipath**.
- High availability GRES/NSR ISSU fabric redundancy tail-end protection BGP prefix-independent convergence edge
- 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](#)
- Traceroute MPLS from Juniper to Huawei routers does not work as expected due to unsupported TLV. [PR1363641](#)
- If RSVP is disabled and reenabled globally, and in a rare situation, the new RSVP task tries to access a memory allocated by the old RSVP task during a particular RSVP Path State Control Block (PSB) changed path, then the rpd might crash. [PR1366243](#)
- When RSVP link or node protection is deployed and RSVP authentication is used, if the Point of Local Repair (PLR) router and the Merge Point (MP) router run different versions of Junos OS software during local repair, that is, one a  $\geq 16.1$  release and the other a  $< 16.1$  release, the RSVP authentication errors might occur for the bypass MPLS Label Switched Path (LSP) and cause traffic loss. [PR1370182](#)
- 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. As a result, the IP address remains unchanged for the outgoing interface through which the LSP next hop is reachable. As a workaround to avoid this issue, do the following: (1) delete the previous IP address first, then commit; (2) if the system is busy, wait a while; and (3) configure the same IP address with a different network mask, then commit. [PR1372630](#)

## Platform and Infrastructure

- FPC reports the following errors and the FPC is not able to connect any subscriber: **Pkt Xfer:\*\* WEDGE DETECTED IN PFE 0 TOE host packet transfer: %PFE-0: reason code 0x1**. Also, the MQ FI might be wedged and the following log can be seen: **Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) FI Reorder cell timeout Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) FI Enqueuing error, type 1 seq 404 stream 0 Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) MALLOC Pre-Q Reference Count underflow - decrement below zero**. [PR873217](#)
- Starting in Junos OS Release 13.1R1 and later, if **no-fast-sync** is used in configure-private mode, the commit operation might throw errors after the configuration statement is executed under choice (such as **protocol [ ospf pim tcp ]** ) is added or deleted. Also, the configuration statement is executed under choice (such as **protocol [ ospf pim tcp ]** ) is added or deleted, the whole hierarchy is shown as changed when using the **show configuration | compare** command. [PR1042512](#)
- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh\_ucast\_change:291Referenced I2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)

- After changing an outer VLAN tag, the logical interface gets programmed with an incorrect STP state (discarding), so the traffic gets dropped. [PR1121564](#)
- The Junos OS key attribute, which is emitted in the XML format of the configuration, will not be emitted in the JSON format of the configuration. [PR1195928](#)
- The Junos Continuity Software (JAM) might append to the regular format of a Junos OS release, causing PyEZ to fail. [PR1240640](#)
- Because of transient hardware events, the fabric stream might report **CPQ1: Queue underrun indication - Queue <q>** continuously. For such events, all fabric traffic is queued until the Packet Forwarding Engine completes reporting the error, resulting in a high amount of fabric drops. [PR1265385](#)
- When certain hardware transient failures occur on an MQ-chip-based MPC, traffic might be dropped on the MPC, and syslog errors **Link sanity checks** and **Cell underflow** are reported. There is no major alarm or self-healing mechanism for this condition. [PR1265548](#)
- This issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have 10-Gigabit Ethernet (xe) or Gigabit Ethernet links toward the downstream switch and LAG bundles as uplinks towards 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 to avoid loops. Some MAC addresses are not learned on router when LAG bundles that are part of such bridge domains are flapped and other events such as spanning tree root bridge occur. [PR1275544](#)
- 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 the 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 the **[set protocols layer2-control nonstop-bridging]** hierarchy level. This reduces traffic flood to around 10 seconds in a moderate setup. [PR1275621](#)
- The jlaunchd commit-batch is thrashing and does not restart. [PR1284271](#)
- The operational command **show igmp statistics** with no filter does not display the aggregated JOIN/LEAVE/ QUERY statistics from subscribers with the **distributed** statement in the **igmp/ mld** stanza of the dynamic profile. [PR1289415](#)
- Every load override increases the refcount by 1 and after it reaches the maximum value (65,535), the mgd crashes and the session gets terminated. But there is no impact for a new session. [PR1313158](#)
- When chassis control is restarted with an aggregated Ethernet and CoS rewrite configuration, **Platform failed to bind rewrite** messages might be seen in syslog. The issue is specific to aggregated Ethernet interfaces when restart chassis control is done. A timing issue might occur when logical interface deletion is delayed because of the high scale. When logical interfaces come up again after restart, they have different indexes. The issue is only applicable when aggregated Ethernet interfaces are present. [PR1315437](#)
- On MX Series routers with MPC1E, MPC2E, MPC3D with 16 port 10 Gigabit Ethernet ports, MPC3E, MPC4E, or T4000 with Type 5 FPC, if the interface is configured with the **input-vlan-map** option, then

the traffic with more than 2 VLAN tags might be incorrectly rewritten and sent out. As a result, the traffic drops. [PR1321122](#)

- On all platforms, with dual Routing Engines and GRES enabled, if executing switchover, the firewall filter's state might be incorrect and an FPC core file might be seen. [PR1324819](#)
- In a Layer 3 VPN topology, traceroute to a remote PE device for a CE-facing network sees the 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, 16.2R1, and later releases, there is a kernel sysctl value, `icmp.traceroute_l3vpn`. Setting this to one will change the behavior to select an address based on the destination specified in the traceroute command. [PR1358376](#)
- Traffic traversing an IRB is not tagged with a VLAN if the packets goes through an additional routing instance. [PR1377526](#)

## Routing Policy and Firewall Filters

- The **set metric multiplier policy** command had the potential to generate negative values given user-permitted inputs. This might result in values less than 0 being interpreted by rpd as a very large number, leading to unexpected metric values in many protocols. However, replication of this issue might require an unusual configuration and is not normally considered a problematic condition. [PR1349462](#)

## Routing Protocols

- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- Continuous soft core files might be generated due to the bgp-path-selection code. The routing protocol process (rpd) forks a child process and the child process asserts to produce a core file. The problem is with route ordering and it is automatically corrected after the soft assert core file is collected, 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](#)
- With Shared Risk Link Group (SRLG) enabled under corner conditions, after the **clear isis database** command is executed, the rpd might crash because the IS-IS database tree gets corrupted. [PR1152940](#)
- When LDP is deactivated, in a rare case, the result of remote loop free alternate (RLFA) might be computed to go through deactivated LDP node. The situation is self-recovered in the next shortest-path-first (SPF) calculation. [PR1202392](#)

- JTASK\_SCHED\_SLIP for rpd might be seen on doing restart routing or OSPF protocol disable with scaled BGP routes in the MX104 router. [PR1203979](#)
- In the context of a large number of configured VPNs, routes changing in the midst of a BGP path-selection configuration change can sometimes generate an rpd core file. This core file has been seen with the removal of the **always-compare-med** option. [PR1213131](#)
- When an aggregation gateway running Junos OS uses an IPv6 address as the next hop for IPv4 aggregates announced to downstream devices, it might attract traffic prematurely before Packet Forwarding Engines are programmed with more specific IPv4 routes. This happens when the IPv6 address is advertised in the BGP **inet6-labeled-unicast** family. [PR1220235](#)
- When you try to qualify Junos OS Release 16.1X60-D40 on MX960 for BNG/Subscriber Management functionalities, the routing protocol process (rpd) utilization goes up to 100 percent displaying the following output: {master} user@host> show system processes extensive | no-more last pid: 76128; load averages: 1.51, 1.46, 1.68 up 6+04:38:02 14:32:44 198 processes: 2 running, 195 sleeping, 1 waiting Mem: 1415M Active, 5284M Inact, 2441M Wired, 2088M Buf, 6752M Free Swap: 8192M Total, 8192M Free PID USERNAME THR PRI NICE SIZE RES STATE C TIME WCPU COMMAND 10 root 4 155 ki31 0K 64K RUN 3 509.5H 304.10% idle 5207 root 4 20 0 3017M 2140M kqread 0 23.0H 100.00% rpd 4925 root 2 -26 r26 556M 47060K nanslp 1 511:02 5.08% chassisd 5185 root 1 20 0 698M 176M select 2 139:31 0.20% authd 5002 root 1 20 0 455M 7464K select 1 32:43 0.10% license-check 11 root 30 -72 - 0K 480K WAIT 255 888:28 0.00% intr 52981 root 1 35 15 459M 10360K select 1 469:19 0.00% sampled. The following system log show messages are displayed: Dec 7 03:36:56.615 2016 lab31 rpd[5474]: RPD\_KRT\_Q\_RETRIES: route table add: Resource temporarily unavailable Dec 7 03:36:56.615 2016 lab31 rpd[5474]: RPD\_SYSTEM: Get index for rt table failed: Resource temporarily unavailable Dec 7 03:36:56.615 2016 lab31 rpd[5474]: RPD\_KRT\_Q\_RETRIES: route table add: Resource temporarily unavailable Dec 7 03:36:56.615 2016 lab31 rpd[5474]: RPD\_SYSTEM: Get index for rt table failed: Resource temporarily unavailable Dec 7 03:36:56.615 2016 lab31 rpd[5474]: RPD\_KRT\_Q\_RETRIES: route table add: Resource temporarily unavailable. [PR1240273](#)
- BGP NSR replication starts after a delay in certain cases. [PR1256965](#)
- Performance degradation occurs during the computation of LFA and remote LFAs. This has no impact on functionality. [PR1264564](#)
- The BMP session sends both peer down events as well as route withdrawals when a peer monitoring is disabled through a configuration event. After that commit, only the peer down events are sent. [PR1265783](#)
- When **route-distinguisher-id** is configured and a VRF with a route distinguisher is automatically assigned with the **auto-rd** feature configured, the MX Series BNG allows the configuration to be committed, but after the commit the rpd process crashes. [PR1278582](#)
- The backup Routing Engine scheduler slips when a Cisco Rosen7 PE device has MDT-SAFI is enabled; however, the MDT-SAFI update does not include the **route-target** extended community attribute, NSR is enabled, policies are set to import or export the inet-mdt table, but Rosen is not configured. [PR1295712](#)

- In Junos OS, the rpd might crash because of a malformed BGP UPDATE packet (CVE-2018-0020). Refer to <https://kb.juniper.net/JSA10848> for more information. [PR1299199](#)
- When PIM is enabled for multicast traffic, the designated router switchover might lead to multicast traffic getting pruned for random groups. [PR1303050](#)
- An MX104 is connected to SRX1500. IS-IS is running between these devices and BFD has been configured between the IS-IS peers. Unfortunately, BFD does not come up between these devices successfully. [PR1312298](#)
- In a resource public key infrastructure (RPKI) scenario, the validation replication database might have many more entries than the validation database after the RPKI cache server is restarted and the validation session is reestablished. [PR1325037](#)
- When route target filtering (RTF) is configured for VPN routes and multiple BGP sessions flap, there is a slight chance that some of the peers might not receive the VPN routes after the flapped sessions come up. [PR1325481](#)
- With BGP, LDP, and IS-IS configurations, deleted IS-IS routes might still be present in the RIB. The presence of such routes does not impact on-route on route selection or other functionality of routing protocol process (rpd). Just that deleted IS-IS routes do not get removed with specific configurations. [PR1329013](#)
- In a large-scale 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](#)
- In rare cases, rpd might crash during the times of excessive neighbor session instability (flapping). [PR1337304](#)
- When configuring anycast and prefix segments in SPRING for IS-IS, **prefix-segment index 0**, even though the user is allowed to configure 0 as an index. [PR1340091](#)
- From Junos OS Release 16.1, **show bgp neighbor** does not display the correct value for the **Last traffic (seconds)** field anymore. [PR1361899](#)
- On devices running Junos OS platform, 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 crash. [PR1366696](#)
- Ukern memory leak and core crash might be happened when device configured link-node protection with labeled-bgp. [PR1366823](#)
- In BGP scenario with multipath enabled, when import or export policy of IPv6 routes is applied with an IPv4 next hop to a BGP neighbor, the rpd might crash continuously. [PR1390428](#)



## Services Applications

- In an L2TP scenario, when the L2TP network server (LNS) is flooded by high-rate L2TP messages from the LAC, the CPU on the Routing Engine might become too busy to bring up new sessions. [PR990081](#)
- Session counters for cleartext traffic are not updated after decryption. The decrypted packet count can, however, be obtained by running the **show security group-vpn member ipsec statistics** command. [PR1068094](#)
- We recommend that you do not configure an ms- interface when an AMS bundle in one-to-one mode has the same member interface. [PR1209660](#)
- The NAT auto-injected routes might fail to install or when back-to-back commits with changes made to service sets or NAT rules are performed. This issue occurs with a unique configuration where thousands of routes are added by the service PIC process (spd), which manages installation of NAT return routes and destination routes. [PR1223729](#)
- If an L2TP subscriber has static pp0 interface on the LAC side, LCP renegotiation is configured on the LNS side and CPE device has been changed, an issue with successful negotiation of the PPP session between LNS and the CPE device might occur. [PR1235554](#)

## Subscriber Access Management

- In PPPoE subscribers scenario with a large scale of subscribers (for example, 3000), during login and logout, some subscribers might get stuck in the error state of "Terminated". This issue impacts the traffic for these subscribers. [PR1262219](#)
- RAA message has extra AVP with **destination-host** even though it has been not configured under the configuration. [PR1384011](#)

## User Interface and Configuration

- The **max-db-size** configuration does not work on MX5, MX10, MX40, MX80, and MX104. [PR1363048](#)

## VPNs

- For next-generation MVPN, the traffic threshold is ignored if it is configured in a configuration group. As a workaround, apply the group to the MVPN instance. [PR1191002](#)
- In an MVPN scenario with I-PMSI tunnels and a multihomed source, if the link between the source and the PIM-DR PE1 device goes down, the second PE device (PE2) takes the PIM-DR role and starts to advertise Type-5 prefixes. Then, as the link between the source and PE1 comes back up and PE1 reassumes the PIM-DR role, PE1 might not generate Type-5 BGP prefixes for active sources in some multicast groups. Without Type-5 prefixes from the ingress PE device, the receivers' PE devices do not



generate Type-6/7 prefixes and the ingress PE device does not send multicast traffic. As a workaround, clear PIM joins in the affected instance. [PR1242493](#)

- The configuration statement **unicast-umh-election** for NG-MVPN might not work as expected in special cases. This statement is to use the unicast route preference for upstream multicast hop (UMH) selection. However, the nonoptimal route might be selected if the routes have the same IP address value in the route-import community. [PR1315011](#)
- When a C-multicast route (Type 7 or Type 6) for inter-AS non-segmented option C topology is sent with the originator's IP address, Junos OS source PE device does not accept this and thus the PIM join fails. [PR1327439](#)
- A core file is seen on the backup Routing Engine during label allocation and when restarting routing on the master Routing Engine when NSR is enabled. [PR1351425](#)

SEE ALSO

<a href="#">New and Changed Features   113</a>
<a href="#">Changes in Behavior and Syntax   144</a>
<a href="#">Known Behavior   162</a>
<a href="#">Resolved Issues   185</a>
<a href="#">Documentation Updates   249</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   250</a>
<a href="#">Product Compatibility   258</a>

## Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R3 | 186](#)
- [Resolved Issues: 17.2R2 | 232](#)
- [Resolved Issues: 17.2R1 | 242](#)

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: 17.2R3

### *Application Layer Gateways (ALGs)*

- IPsec IKEv2 negotiation fails with IKE ALG enabled. [PR1300448](#)
- IKEv2 negotiation might fail with IKE ESP ALG enabled in IKEv2 redirection scenario. [PR1329611](#)

### *Authentication and Access Control*

- Platform-specific callbacks are not getting initialized. [PR1354855](#)

### *Class of Service (CoS)*

- CoS wildcard configuration is applied incorrectly after the router restarts. [PR1325708](#)
- The Routing Engine might get into amnesiac mode after restarting if **excess-bandwidth-share** is configured. [PR1348698](#)
- When a logical interface is configured with 802.1P rewrite-rules (for both outer and inner VLAN) and fixed classification, after deactivating class of service (CoS) on any other logical interface, the packets sent from this logical interface might still have the original 802.1P bit set in the inner VLAN without being rewritten. [PR1375189](#)
- The Class of Service (CoS) mode per-unit-scheduler is not supported on interface that is an interface-set member, if CoS mode is changed from hierarchical-scheduler to per-unit-scheduler for the interface, the Flexible PIC Concentrator (FPC) card of the interface might crash. [PR1387987](#)

### *EVPN*

- Ethernet A-D per Ethernet Segment Route (Type-1 PER ES) is not generated with new route target after changing vrf-target. [PR1279529](#)
- EVPN traffic mapping to specific LSPs is not working. [PR1281415](#)
- Local preference for EVPN type-5 route might cause unexpected results, if BGP multipaths are configured. [PR1292234](#)
- BGP route refresh request might not be sent after modifying route target. [PR1300332](#)
- The traffic might be dropped after receiving an updated ARP route packet from peer Layer 3 gateway in EVPN-VxLAN scenario. [PR1306024](#)
- The rpd might crash on Junos OS platform with EVPN and NSR enabled after restarting the rpd process in EVPN environment. [PR1320408](#)
- Discard EVPN route is installed on local PE after the connection flap on a remote PE in a multihome EVPN topology. [PR1321125](#)

- When a catastrophic event occurs that leads to the re-creation of the EVPN-VXLAN instance (such as a change in INSTANCE RD, control word, or router-id) or forced cleanup of the EVPN-VXLAN instance (such as simply deleting the EVPN-VXLAN instance configuration), and if there are multiple entries whose vlan-id are zero, the rpd might crash. [PR1321839](#)
- In an EVPN scenario with nonstop-routing (NSR) enabled, the routing protocol process (rpd) crashes and generates a core file on backup Routing Engine while any configuration changes are made on master Routing Engine. [PR1336881](#)
- The rpd might crash if the IRB interface and routing instance are deleted together in the same commit. [PR1345519](#)
- In an EVPN instance with EVPN E-Tree service configured, the rpd might crash if the EVPN instance refers to a vrf-export policy that does not have **then community**. No **then community** in the vrf-export policy is an incorrect configuration. [PR1360437](#)
- Gratuitous ARP request does not update the ARP table when ARP proxy is enabled. [PR1371352](#)

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

- The mib2d process might crash when SNMP walking during commit or rollback. [PR1286448](#)
- The pfed generates a core file in pfed\_process\_session\_state\_notification\_msg is observed, `pfed_timer_manager_c::remove_serv_id, pfed_delete_timer_id_by_serv_sid (serv_sid=0, serv_info=0x0)` at `../../../../src/junos/usr.sbin/pfed/pfed_timer.cc:16`. [PR1296969](#)
- A few accounting files might be missed in case, the remote archive sites are unreachable. [PR1300764](#)
- There is a memory leak on mib2d when polling firewall MIBs. [PR1302553](#)
- Remote CE1 MAC address might take long time to clear post clear MAC. [PR1304866](#)
- Dfwd might crash during execution of **show firewall templates-in-use** command. [PR1305284](#)
- **ACCT\_FORK\_LIMIT\_EXCEEDED** log level is an error even when backup-on-failure feature is enabled for accounting files. [PR1306846](#)
- The second archive site in the accounting-file configuration is not used when the first one uses SFTP and is not reachable. [PR1311749](#)
- Accounting files with no records might be unexpectedly uploaded to the archive-site. [PR1313895](#)
- The commit might fails when the **nexthop-learning** configuration statement is enabled for J-Flow v9. [PR1316349](#)
- The FPC CPU might reach 100 percent constantly, if shared bandwidth policer is configured. [PR1320349](#)
- DHCP service crashes after the router is set to factory default value 0. [PR1329682](#)
- Some firewall filter counters might not be created in SNMP. [PR1335828](#)
- The error logical interface under VPLS might be blocked after MAC moving if the logical interfaces are on the same physical interface. [PR1335880](#)

- On MX Series routers, the l2ald daemon might crash if a duplicate MAC is learned by two different interfaces (CVE-2018-0056). Refer to <https://kb.juniper.net/JSA10890> for more information. [PR1338688](#)
- When the **clear ethernet-switching table** command is issued on all PE devices associated with an EVPN-VXLAN routing instance to clear all learned MACs, it might not work correctly. Some designated router MACs (MACs learned from remote PE devices) are left hanging on some PE devices, as shown by **show ethernet-switching table**. [PR1341328](#)
- Commit failed when attempting to delete any demux0 unit numbers which are greater or equal to 1000,000,000. [PR1348587](#)
- Packet Forwarding Engine process (pfed) creates dummy interface accounting records on the backup Routing Engine. [PR1361403](#)

### General Routing

- The enhanced IP or the enhanced-Ethernet network-services mode and MS-DPC card are not compatible and should not be configured or inserted in the chassis, at the same time. [PR1035484](#)
- On MX Series routers, when packets go through different interfaces with different family types configured, because of the incorrect cookies that are sent from the Packet Forwarding Engine, the packets might get dropped if channelized E1/T1 circuit emulation MIC is equipped as the outgoing interface (for example, receiving IPv4 packets on the incoming interface, and sending out packets with MPLS label on the outgoing or core-facing interface that is e1/t1 interface in a Layer 3 VPN scenario). [PR1064515](#)
- When hybrid timing mode is configured, MX Series MPC2 cards does not interoperate with ACX in VLAN(native-vlan-id). [PR1076666](#)
- The routing protocol process (rpd) memory leak is caused by repeated RSVP reservation state block (RSB) deletes. [PR1115686](#)
- No warning is raised when the bridge family is configured with interface-mode trunk but without **vlan-tagging** or **flexible-vlan-tagging**. [PR1154024](#)
- Ksyncd might crash due to transient replication errors between Routing Engines. [PR1161487](#)
- Unexpected MobileNext gateway activation license alarm is observed when TDF gateway is configured. [PR1162518](#)
- SNMP trap sent for "PEM Input failure" alarm is not generated when single input feed fails on MX960. [PR1189641](#)
- The agentd process crashes and generates a core file when the command **show agent sensors** is executed. [PR1197608](#)
- Checks are missing for confirming the validity of a data structure for platforms that do not use the data structure. Hence, the error message **chassisd[1825]: pvidb\_get\_root\_node: Error(2) retrieving rootnode value** might be seen. [PR1198817](#)
- Stale VBF states occur without sdb sessions. [PR1204369](#)

- 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. That 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 of an upgrade performed with `virsh create`. [PR1220671](#)
- Unable to deregister sub error (131072) for error(0x1b0001) for MIC module. Error messages are seen on MPC5E card. [PR1221337](#)
- The error logs **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](#)
- On MX5, MX10, MX40, and MX80 routers, Power Equipment Module (PEM) failure alarm/SNMP trap sets with status register 0xff, but it is always cleared in the next polling (in 5 seconds). Returning a status of 0xff from PEM firmware is recognized as invalid. You can safely ignore it as a false PEM failure. [PR1231893](#)
- Timestamp conversion within Zarlink stack is causing FPC CPU hog and crash. [PR1232740](#)
- The MS-MPC card might crash when OSPFv3 IPv6 traffic goes through it. [PR1233459](#)
- The **multicast-replication** setting cannot be reflected in the redundancy environment after rebooting both the Routing Engines. [PR1240524](#)
- In a BGP or MPLS scenario, if the next hop type of label route is indirect, disabling and enabling the "family mpls" of the next hop interface that might cause the route to go into a dead state. [PR1242589](#)
- **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 restart. [PR1243364](#)
- The "validation-state:unverified" routing entry might not be shown with proper location in `show route`. [PR1254675](#)
- Prolonged flow-control core is observed for the TFTP ALG traffic (10,000 simulated users). [PR1255973](#)
- The `rpd` might crash during the next hop change, if unicast reverse-path-forwarding (uRPF) is used. [PR1258472](#)
- Temperature reading for TFEB components jumps up and down frequently on MX80. There is no particular trigger needed. By default, the FPC reports temperatures of some components to the Routing Engine or `chassisd` (every 10 seconds). As part of this periodic polling, you can see the issue of the temperature reading for the TBB Packet Forwarding Engine component showing occasional jumps. [PR1259379](#)
- Traffic drops when an MPC have high rate of cell underflow drops after link sanity check. [PR1262868](#)
- vMX FPC generates a core file and panic (format\_string=format\_string@entry=0x9e509c4 "Thread %s attempted to %s with IRQ priority at %d\n"). [PR1263117](#)
- PTP is lost with master when the master-line card is restarted. [PR1264530](#)

- All traffic received from specific fabric stream is dropped with only XMCHIP FI: cell underflow error syslog event. [PR1264656](#)
- PCC controlled LSP metric is not getting updated on the controller, PCE delegated LSPs do not come up. [PR1265864](#)
- On MX Series platforms, 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 card might cause the SIP ALG to generate a core file. [PR1268891](#)
- Messages related to **Logical Addr xxxxxxxx is invalid** seems when FPC restart also passing traffic. [PR1271810](#)
- The IPv6 ping might fail after route leaking policy deployment is done between two L3VPN routing instances. [PR1274339](#)
- When the static link protection mode is configured with backup state as down, the primary port is going to down state instead of the secondary port, while the secondary port remains in up state . [PR1276156](#)
- The routing protocol process (rpd) KRT asynchronous queue might stall, impacting synchronization between RIB and FIB. [PR1277079](#)
- On running certain commands that involve command forwarding, an mgd process is created to retrieve the data. In rare cases, if this command times out or if it is manually terminated (for example, using Ctrl+c), then it might cause the mgd process to utilize high Routing Engine CPU. [PR1297728](#)
- The bbe-smgd might generate a core file in certain cases while using iflsets in universal call admission control policy mode. [PR1278543](#)
- Syslog messages **jnh\_vbf\_flow\_get\_oif\_index: Rollback cmd not found for flow** are generated by MPC during subscriber login. [PR1278580](#)
- On MX104 platform with GRES enabled, the chassis network services might not get set as "Enhanced-IP". [PR1279339](#)
- CoS attachment might be attached to incorrect link if issuing some changes to aggregated Ethernet bundle. [PR1279788](#)
- Syslog messages **CM\_FPC: Error requesting SET BOOLEAN, illegal setting 132,111** are seen after unified ISSU from Junos OS Release 16.2R2 to Junos OS Release 17.1R2.2. [PR1280878](#)
- The kernel might crash when NSR enabled device has BGP peer flapping. [PR1282573](#)
- The rpd process might crash if dynamic interfaces are used by multiple applications. [PR1282854](#)
- The enhancement of reporting total SBE errors when the corrected singlebit errors threshold of 32 is exceeded for MPC7E, MPC8E, and MPC9E. [PR1285315](#)
- LC, PFH, and Packet Forwarding Engine interface is not coming up on RE1. [PR1285606](#)
- With CoS-based forwarding, when the primary path of one of the next hop LSPs flap, traffic carried by the other next hop LSP could get load-balanced across the primary and secondary path. [PR1285979](#)

- Internal latency increases overtime for Packet Forwarding Engine sensors with streaming telemetry. [PR1286286](#)
- The missing statement **Shared bandwidth policer not supported for interface ge-x/x/x** is seen, during a failed commit in Junos OS Release 16.1R3. [PR1286330](#)
- Framed routes might get stuck in KRT queue. [PR1286849](#)
- The oneset or leaf-list configuration might not get deleted with delete operation through JSON. [PR1287342](#)
- During unified ISSU, (FRU upgarde) micro BFD flap is observed. [PR1288433](#)
- Performance issues can be seen when nontranslated traffic is introduced to a service-set using a large number of NAT terms. [PR1288510](#)
- After GRES, the smid is declared thrashing and could not restart after some fatal SDB errors. [PR1288871](#)
- The interfaces might go to a down state after performing GRES. [PR1289493](#)
- The **request system zeroize** command deletes the **/var/db/scripts** directory, which does not get re-created automatically. [PR1289692](#)
- NAT-T and DPD functionality do not work for aggressive mode. [PR1290689](#)
- Incorrect temperature is displayed for MPCP5, MPC7 in **show chassis fpc** output. [PR1290771](#)
- LSP traffic might silently drop and get discarded after a link goes down in the bypass path. [PR1291036](#)
- The routing protocol process (rpd) might generate a core file while restarting the process. [PR1291110](#)
- The switch might incorrectly learn its own IRB MAC address. [PR1291184](#)
- Device going to the DB prompt "db@jsr\_jsm\_send\_ka\_after\_merge,send\_proto\_keepalive" was observed on master Routing Engine. [PR1291247](#)
- The L2TP ICCN fast retransmission occurs after tunnels go down. [PR1291557](#)
- Kernel does not install the route and throws an error. [PR1291917](#)
- Error message might be seen while bringing up the subscriber in a subscriber management environment. [PR1293057](#)
- The flow export rate remains lower than the configured export rate in inline sampling scenario. [PR1294296](#)
- Loss of DHCP or PPPoE subscribers is observed during unified ISSU from Junos OS Release 16.1-20170718\_161\_r4\_s5.0 to 16.1-20170718\_161\_r4\_s5.0. [PR1294709](#)
- During PPPoE subscriber login errors like [ **vbf\_flow\_src\_lookup\_enabled** ] and [ **failed to find iff structure, ifl** ] were seen on FPC. [PR1294710](#)
- The rpd might crash after the interface or BGP flap. [PR1294957](#)
- The KRT queue might get stuck with the error of **RPD\_KRT\_Q\_RETRIES: chain nexthop add: Unknown error: 0**. [PR1295756](#)
- xmlproxyd generates a core file during telemetry streaming. [PR1295831](#)

- The service profile's CoS might be overridden by the client profile's CoS when second family DHCP sessions are added in a dual-stack subscriber scenario. [PR1296002](#)
- The MSPMAND process might crash if you use SCG services on MS-MPC or MS-MIC. [PR1296422](#)
- The jdhcpd might crash when using 'dhcp-security' related command in enhanced subscriber management mode. [PR1296461](#)
- LLDP sensor on telemetry using a lot of bandwidth. [PR1296869](#)
- In ECMP fast reroute scenario, traffic might get silently dropped or discarded because of a next hop in "hold" state. [PR1297251](#)
- Multiple bbe-smgd core files are seen during a subscriber binding configuration with DT CST with as little as 200-300 subscribers and continual core files while scaling. Maximum scale cannot be achieved with multicast-enabled subscribers (related to IPTV profile). [PR1297612](#)
- It is not possible to collect shmlog entries and statistics on MX5, MX10, and MX40 platforms. The code change also includes improvements that should prevent the shmlogctl process from generating a core file because of a timing issue. [PR1297818](#)
- Some random number of ports on MPC7E-10G card might not come up after the remote system or line card restarts or interface flap. [PR1298115](#)
- The log message about shutdown time is incorrect when system exceeds chassis over temperature limit. [PR1298414](#)
- The rpd core files are generated with PPPoE and L2BSA flapping. [PR1298587](#)
- The bbe-smgd process constantly generates core files while ESSM+PPPoE stress test with concurrent GRES is running. [PR1298742](#)
- MX Series BNG does not respond to PADI after GRES on some ports/VLANs. [PR1298890](#)
- The error messages about PEM might be seen in MX Series platform with AC 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](#)
- Flat accounting files are not generated according to the configured timers. [PR1299597](#)
- The bbe-smgd might generate core files after the Routing Engine mastership switch. [PR1299812](#)
- Subscriber database is stuck in "not-ready" state after GRES. [PR1299940](#)
- Chassisd core file generated is seen after insertion of REMX2K-X8-64 in MX2000 platform along with older RE-S-1800x4. [PR1300083](#)
- After IS-IS-TE routes and BGP routes attribute change, traffic loss might be seen because BGP routes point to some stale labels. [PR1300425](#)
- The error **error: the SDN-Telemetry subsystem is not responding to management requests** is seen on issuing the CLI command **show agent sensors**, if traceoptions is enabled for services analytics. [PR1300829](#)



- ICMP, ICMPv6 error messages might be discarded while forwarding through an AMS interface. [PR1301188](#)
- Configured logical interface might not be created correctly after commit. [PR1301823](#)
- In Junos Telemetry Interface setup, the payload maximum transmission unit (MTU) might be much less than 16KB when subscribing to component sensor. [PR1301835](#)
- The rpd might crash by executing the command **show route extensive** when IS-IS configuration is deleted. [PR1301849](#)
- The rpd might crash when NSR is enabled and routing-instance specific configurations are committed. [PR1301986](#)
- Continuous interface flapping might lead to unwanted MIC reset. [PR1302246](#)
- Service cookie data that is sent from Packet Forwarding Engine to service PIC can be corrupted and might lead to unexpected behavior. [PR1302493](#)
- The rpd might crash when toggling the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statement. [PR1302504](#)
- The log message **jam\_cache\_get.636 ERR:entity 0x997 not found, get cache failed** is continuously seen in jam\_chassisd log-file. [PR1302975](#)
- The chassisd crashes during unified ISSU aborted in FRU upgrade phase. [PR1303086](#)
- The multicast resolve-rate value might go back to default after system upgrade or reboot. [PR1303134](#)
- Incorrect MTU might be seen on PPP interfaces when PPP MTU is not defined in the dynamic profile. [PR1303175](#)
- The list of available routing instances is no longer provided for output of **show subscribers routing-instance** command. [PR1303199](#)
- The inline-ka PPP echo requests are not generated for aggregated Ethernet interfaces. [PR1303249](#)
- The command **request auto-configuration reconnect-pending** is no longer available. [PR1303336](#)
- Blocking PPPoE or DHCP to initiate VLAN auto-sensing if VLAN out-of-band connected is in pending state. [PR1303338](#)
- On routers with XM-chip based line cards, log messages might report fan speed changes between full and normal speed continuously, because XM-chip reaches a temperature threshold. [PR1303459](#)
- The kernel log GENCFG messages with Severity 1 (Alert) might be seen. [PR1303637](#)
- 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](#)
- 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** CLI command. [PR1304016](#)

- The fabric planes might go into "check" state after restarting the line cards with SFB2 used on MX2010 or MX2020 platform. [PR1304095](#)
- Effective rate of E3 in framed mode is limited to 30 Mbps on certain channelized MICs. [PR1304344](#)
- After modifying the DSCP value in the classifier, the value is not getting reflected in the LLDP PDU TCP core file. However, the logical interface binding is happening with the modified DSCP value. [PR1304627](#)
- RPF check strict mode causes traffic drop in next generation subscriber management release. [PR1304696](#)
- On MX2000 platform with MPC9E and SFB2 installed, certain high amount traffic volume might cause traffic drops with 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](#)
- In inline J-Flow vMX platform, OIF field of VPLS data records sometimes report SNMP index value of LSI interface instead of egress physical interface. [PR1305411](#)
- MX Series router is sending immediate-interim for the services pushed by SRC. [PR1305425](#)
- Customers running 32-bit Junos OS might generate rpd core file when traceoptions are enabled. [PR1305440](#)
- Improved handling of exit status for JET applications. [PR1305615](#)
- L2BSA subscriber connection attempts failed with VLAN profile-request-error. [PR1305962](#)
- The CLI **start shell pfe network fpc** command is not working on MX960. [PR1306236](#)
- Bbe-smgd might fail to properly add access-internal routes when the router is extremely busy. [PR1306650](#)
- L2BSA subscribers came up, but no new ANCP session got established during the RADIUS disaster backup procedure. [PR1306872](#)
- Smihelperd 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 is not working. [PR1308671](#)
- FPC syslog errors with **pfeman\_inline\_ka\_steering\_gencfg\_handler: nh not found** could mean that steering rules are not installed correctly. [PR1308884](#)
- All the MICs on one FPC, with PWHT subscribers configured, might go offline during the restart of FPC in another slot. [PR1308995](#)
- Error messages might be often seen after MPC restarts. [PR1309013](#)
- Incorrect values are found in the event-timestamp of RADIUS accounting-stop packets for L2BSA subscribers. [PR1309212](#)

- In MX2020 or MX2010, after smooth upgrade from SFB to SFB2, if one plane or SFB is restarted, link training fails between those planes and MPC6 cards. [PR1309309](#)
- The bbe-mibd might generate a core file after Routing Engine mastership switch. [PR1309341](#)
- First access-request is failing for L2BSA subscribers when changing the MTU of LACP aggregate Ethernet A10NSP interface. [PR1309599](#)
- 90 percent subscribers might go down after unified ISSU from Junos OS Release 16.1 to Junos OS Release 17.3. [PR1309983](#)
- In next generation subscriber management release, bbe-smgd process memory leak is seen after deleting or adding the address pool. [PR1310038](#)
- The MS-MIC or MS-MPC memory utilization might stay at high level 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 running SNMP walk on Service PIC NAT OIDs. [PR1310081](#)
- Some harmless syslog messages might be seen. [PR1310678](#)
- Local IPv6 interface address from NDRA prefix is not removed from service interface when subscriber dual-stack session is removed. [PR1310752](#)
- Utilization of **commit check** just after setting master-password can trigger improper decoding of configuration secrets. [PR1310764](#)
- After BSYS reboot rpd is unresponsive on one GNFs sometimes. [PR1310765](#)
- The incorrect error number might be reported for syslog messages with prefix of **%DAEMON-3-RPD\_KRT\_Q\_RETRIES**. [PR1310812](#)
- Fragmented UDP packet might be incorrectly parsed as uBFD packet and dropped. [PR1311134](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- The routing protocol process generates a core file after multiple session flap on scale setup. [PR1312169](#)
- PEM alarms and I2C failures are observed on MX240, MX480, and MX960 Series. [PR1312336](#)
- MIC MRATE might restart after port speed change. [PR1312504](#)
- Counter at PPPoE session logical interface incremented incorrectly cause accounting packet contains incorrect Acct-input-packets value and incorrect Acct-input-octets value. [PR1312998](#)
- False over temperature SNMP trap could be seen when using MPC5, MPC6, MPC7, MPC8, and MPC9 on MX2020. [PR1313391](#)
- On MX Series Virtual Chassis, BNG IPv6 router-solicit (RS) packets are dropped in non-default RI. [PR1313722](#)
- The CLI command **show version detail** gives severity error log message **traffic-dird[20126]: main: swversion pkg: 'traffic-dird' name: 'traffic-dird' ret: 0**. [PR1313866](#)

- The MSPMAND process generates a core file because flow-control is seen while clearing CGNAT+SFW sessions. [PR1314070](#)
- The **show version detail | no-more** CLI hangs for more than 120 seconds in master Routing Engine and more than 60 seconds in 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 crash in MoFRR scenario. [PR1314711](#)
- MPC7E- IR-mode configuration statement commit failure. [PR1314755](#)
- RPC error is seen while committing **system services subscriber-management enable** through NETCONF. [PR1314968](#)
- The L2TP LAC might drop packets that have incorrect payload length while sending packets to the LNS. [PR1315009](#)
- Continuous logs from vhlclient for all the commands executed. [PR1315128](#)
- The RIB and FIB might get out of synchronization because the KRT asynchronous queue might get stuck. [PR1315212](#)
- FPC crashes when a route has unilist next hops in an RSVP scenario. [PR1315228](#)
- **show version detail** gives severity error log **mobiled: main Neither BNG LIC nor JMOBILE package is present,exit mobiled**. [PR1315430](#)
- The command of **show version detail** might generate severity error log **main: name: SRD ret: 0**. [PR1315436](#)
- The FAN speed might frequently keep changing between normal and full for MX Series platform. [PR1316192](#)
- The **show auto-configuration out-of-band** CLI command with different configuration statements show the same output. [PR1316661](#)
- Demux interface sends neighbor solicitation with source link-address of all zeros 00:00:00:00:00:00 MAC. [PR1316767](#)
- The output from **show configuration <> | display json** might not be properly enclosed in double quotes. [PR1317223](#)
- Linux-based micro-kernel might panic because of the concurrent update on mutable objects. [PR1317961](#)
- CoA shaping rate is not applied successfully after unified ISSU, while doing unified ISSU from Junos OS Release 15.1R6.7 to Junos OS Release 16.1R6.2. [PR1318319](#)
- The rpd might crash when the link flap on an adjacent router. [PR1318476](#)
- The daemon bbe-smgd might crash after performing GRES. [PR1318528](#)
- FPC crashes on configuration change for Packet Forwarding Engine sensors. [PR1318677](#)

- MS-MPC and MS-MIC might crash after a new IPsec tunnel is added. [PR1318932](#)
- The MPC with specific failure hardware might impact other MPCs in the same chassis. [PR1319560](#)
- The task replication might not be complete to certain network protocols after multiple GRES. [PR1319784](#)
- The error log message **MIB2D\_COUNTER\_DECREASING: pfes\_stats\_delta: counter** might be seen on vMX. [PR1319996](#)
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX Series Virtual Chassis unified ISSU. [PR1320370](#)
- The **show subscriber summary** command displays incorrect terminated subscriber count. [PR1320717](#)
- PPP inline keepalive does not work fine as expected when CPE aborts the subscriber session. [PR1320880](#)
- MX Series routers send the IPv6 router advertisements and the DHCPv6 advertisements before sending IPCPv6 ACK from CPE. [PR1321064](#)
- Logical interface bind changes is taking more time, many log messages **IFL TCP (38) Bind change notify ran for** are generated by FPC. [PR1321086](#)
- CoS is not applied to Packet Forwarding Engine when VCP link is added. [PR1321184](#)
- The bbe-smgd process generates a core file after massive clients logout and login in PPPoE dual stack subscriber scenario. [PR1321468](#)
- There is CoA-NAK with **Error-Cause = Invalid-Request** sent back to RADIUS server when drop policy under radius-flow-tap is applied in L2TP subscriber scenario. [PR1321492](#)
- The rpd might crash when two next hops are installed with the same next hop index. [PR1322535](#)
- The rpd might crash when OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)
- MS-MIC interface logical interfaces remain down after many iterations of offline or online. [PR1322854](#)
- An incorrect output is observed while verifying the command **show subscribers client-type vlan subscriber-state active logical-system default routing-instance default**. [PR1322907](#)
- The line card might crash upon receipt of a specific MPLS packet. The affected line cards include MPC7E, MPC8E, and MPC9E on MX Series routers, the third-generation FPC on PTX3000 (FPC3-SFF-PTX), the third-generation FPC on PTX5000 (FPC3-PTX-U2 [FPC-P1] and FPC3-PTX-U3 [FPC-P2]) and the built-in FPC on PTX1000. See <https://kb.juniper.net/JSA10864> for details. [PR1323069](#)
- NCP Conf-Ack or Conf-Req packets might be dropped constantly from Cisco MLPPP client on LI interfaces. [PR1323265](#)
- CLI commands in **show system subscriber-management route routing-instance <xxx>** hierarchy show unexpected outputs. [PR1323279](#)
- Memory leaks in MGD-API daemon are observed during get API requests and error handling during set API request. [PR1324321](#)
- Subscribers might fail to log in after the interface is deactivated or activated. [PR1324446](#)

- The memory leakage is seen in mosquito-nossl daemon. [PR1324531](#)
- The SNMP interface filter does not work when "interface-mib" is part of dynamic-profile. [PR1324573](#)
- The VLAN re-write function might put incorrect vlan-id when Ethernet OAM is configured on DPCE cards. [PR1325070](#)
- SNMP values might not be increased monolithically. [PR1325128](#)
- MPC cards might drop traffic under high temperature. [PR1325271](#)
- IS-IS adjacency fails to establish because of packets drop on Packet Forwarding Engine. [PR1325311](#)
- On Junos OS a denial of service vulnerability in MS-PIC, MS-MIC, MS-MPC, MS-DPC, and SRX Series flow daemon (flowd) is related to the SIP ALG (CVE-2018-0051). Refer <https://kb.juniper.net/JSA10885> for more information. [PR1326394](#)
- The VLAN demux interface does not respond to the the ARP request in a subscriber scenario with MX Series routers running Junos OS Release 15.1 or later with subscriber management enabled. [PR1326450](#)
- In MX Series, BNG CoS service object is not deleted properly for TCP and scheduler. [PR1326853](#)
- A few show commands were issued twice when request support information is executed. [PR1327165](#)
- With **auto-installation usb** configured, interface related commits might not take effect because of dcd error. [PR1327384](#)
- Constant logs such as fm\_feacap\_sys\_feature\_get:Attribute DB init is not done, reading from pvid (id: 18) is repeated every 5 seconds in chassisd log. [PR1328868](#)
- If PIC-based sampling is used and the sampling output interface is on the MS-MIC or MS-MPC, a special MPLS packet that is subjected to MPLS sampling might cause unexpectedly prolonged flow control to be triggered on the MS-MIC or MS-MPC and then the MS-MIC or MS-MPC is restarted. [PR1329189](#)
- 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 is not rewriting 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](#)
- **show services nat mappings address-pooling-paired** times out and fails. [PR1330207](#)
- 'Too many supplies missing in Lower/Upper zone' alarm flaps (set/clear) every 20 seconds if a zone does not have minute required PSMs. [PR1330720](#)
- All packets might be dropped if one route is adverted by BGP whose session is established through the subscriber interface. [PR1330737](#)
- The rpd core file is generated on new backup Routing Engine at **task\_quit,task\_terminate\_timer\_callback,task\_timer\_dispatch,task\_scheduler** after disabling NSR+GRES. [PR1330750](#)

- The FPC might be wedged when LSQ interface receives fragmented packets. [PR1330998](#)
- Non-NEBS compliant optics might be disabled when chassis temperature exceeds non-nebs-optics-overheat-trigger. [PR1331186](#)
- When interfaces involved with traffic path are IRB and there is assymetic routing for IPv6 traffic, if the IPv6 packet is egressing an IRB interface that contains an MTU exceeded error or possibly an ICMP6 redirect, the **NH OUT OF SYNC** messages might be seen and traffic might drop. [PR1331911](#)
- On all platforms running Junos OS, the local dhcp6 server might incorrectly respond to confirm messages from clients with existing address bindings with a **NotOnLink** response. This might cause the client to request a new binding. [PR1331995](#)
- The bbe-smgd process might crash after executing the command of **clear ancp access-loop circuit-id <circuit-id>**. [PR1332096](#)
- The rpd core file might be generated in a rare condition in Layer 2 circuit or in a l2vpn environment. [PR1332260](#)
- Inaccurate J-Flow records might be seen for output interface and next hop. [PR1332666](#)
- On all products that support 802.1X, if ports in multisuppllicant mode flap or if the configuration is removed, the memory associated with the dot1xd might leak. As the memory consumption increases, the dot1xd (802.1X process) might crash and restart. [PR1332957](#)
- The subinfo process might crash and it might cause the PPPOE subscribers to get disconnected. [PR1333265](#)
- MX80, MX104, MX240, MX480, and MX960 routers with a DHCPv6 subscriber management environment, might not be able to learn the global IPv6 neighbor address of the DHCPv6 subscriber client if both the neighbor advertisement (NA) source and the destination address are link-local addresses. [PR1333392](#)
- In an AA multihoming EVPN VXLAN, routing protocol daemon shows very high CPU usage. [PR1334235](#)
- Two subscribers cannot reach the online state at the same time if they have an identical Frame-Route attribute value. [PR1334311](#)
- 260G MPC with HQoS went for "restart" after unified ISSU to Junos OS Release 18.2DCB in MX2010 box [PR1334612](#)
- When the MX Series router is used in a subscriber environment, the non-ISSU upgrade might trigger ffp crash. [PR1334745](#)
- The UID limit is reached in large-scale subscriber scenario. [PR1334886](#)
- When using **show subscribers** while FPC has two digits, the interface and IPv6 address get connected together for DHCPv6 PD. [PR1334904](#)
- On MX9200 and MX2000 platforms with MPC7E, MPC8E, MPC9E, when LAG members from different FPCs are unplugged and one member gets plugged back in, MQSS error logs and alarms might be seen. Multiple interfaces might go down and might not come back up until the line card is restarted. [PR1334928](#)

- The IPsec rule might not work if both IPv4 ANY-ANY term and IPv6 ANY-ANY term are configured for it. [PR1334966](#)
- Traffic drops on the MX LNS because of software error/unknown family exception when traffic goes to/coming from MLPPP subscriber if **routing-services** configuration statement is presented in the dynamic-profile used by this subscriber. [PR1335276](#)
- The RIP route updates might be partially dropped when NSR is enabled. [PR1335646](#)
- The MAC\_STUCK might be seen on MS-MPC or MS-MIC. [PR1335956](#)
- Mirrored traffic is not going out through LT interface. [PR1360489](#)
- JET application might not respawn after a normal exit. [PR1336107](#)
- Subscriber might experience SDB DOWN event and drop the clients' connections when issuing **show subscribers** commands. [PR1336388](#)
- 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](#)
- In some corner cases with certificate hierarchy where intermediate CA profiles are not present on the device, the PKI daemon can become busy and stop responding. [PR1336733](#)
- The MACsec AES-GCM-256 hashing algorithm is not compatible with other vendors. The hash value generated for 256-bit key length of AES-GCM-256 algorithm is incorrect. [PR1336834](#)
- Bbe-smgd might crash when performing some interface set-related CoS. [PR1336852](#)
- The command **set protocols lldp neighbour-port-info-display port-id** might not take effect. [PR1336946](#)
- Error log message **sdb\_db\_interface\_remove: del ifl:si-<index> with licnese cnt non zero on** might be seen on LTS during subscriber logout. [PR1337000](#)
- On MX204, MX10003, or MPC7E, MPC8E, MPC9E, or EX9200-40XS or EX9200-12QS, a 100-Gigabit, 40-Gigabit, or 10-Gigabit interface might keep flapping or stay down because of an interoperation issue between the Juniper Networks device and the remote transport device connected. [PR1337327](#)
- On MX2000 platforms with MPC8 and MPC9, if SFB2 goes offline and online, MPC throughput degradation might be seen. [PR1338216](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- Very few of subscribers show incorrect accounting values in a large-scale subscribers scenario. [PR1340512](#)
- There might be traffic loss on some subscriber sessions when more than 32k L2TP subscriber sessions are anchored in ASI interface. [PR1341659](#)
- With discard interfaces (configured with IGMPv3), KRT queue get stuck while deleting multicast next hop (MCNH) with an error **EPERM -- Jtree walk in progress**. [PR1342032](#)
- SNMP walk might fail for LLDP related OIDs. [PR1342741](#)



- In a subscriber management environment, if the commit option **fast-synchronize** is configured, the bbe-smgd process might crash in a rare condition when committing the configuration changes related to dynamic profiles. [PR1342945](#)
- In an MPLS or RSVP environment, LSP might get stuck in Dn state with **Record route: <self> ...incomplete**. [PR1343289](#)
- In Junos OS, memory exhaustion denial-of-service vulnerability is seen in the routing protocol process (rpd) with Juniper Extension Toolkit (JET) support (CVE-2018-0048). Refer to <https://kb.juniper.net/JSA10882> for more information. [PR1344177](#)
- MX Series routers might send IPv6 RA or DHCPv6 advertisement before completing the PPP IPv6CP negotiation. [PR1344472](#)
- The ancpd process generates a core file at **src/junos/usr.sbin/ancpd/ancpd\_smgd.c:2299** in clearing ANCP subscribers in a scaled scenario. [PR1344805](#)
- The framed-route "0.0.0.0/0" might not be installed in MX Series platform with Junos OS enhanced subscriber management releases. [PR1344988](#)
- In an EVPN-VXLAN, ARP packet uses VRRP/virtual-gateway MAC in Ethernet header instead of IRB MAC address. [PR1344990](#)
- The cpcd generates a core file because of the converged services support for Routing Engine-based captive portal used. [PR1345096](#)
- On any product supporting dot1x, as part of authentication of a VoIP phone, its MAC address gets added in both voice and data VLANs. If traffic is received only on the voice VLAN, the MAC address gets aged-out from the data VLAN and because of this the session gets cleared. [PR1345365](#)
- On all platforms, if the **no-propagate-ttl** statement is set in a routing instance that has a route (the route is leaked from one route table to another route table), an rpd crash might be seen. [PR1345477](#)
- New PPPoE users might fail to login. [PR1346226](#)
- **AC system error** counter in **show pppoe statistics** does not work. [PR1346231](#)
- VCCP-ADJDOWN detection is delayed on VC-Bm when deleting one vcp link on VC-Mm. [PR1346328](#)
- On MX Series using MS-MPC, MS-MIC, in an inline NAT scenario, the adaptive services PIC daemon (spd) and eventd might use up the CPU cycles. The spd might crash, resulting in traffic loss of NAT. [PR1346546](#)
- On any platform that does not clear out **/mfs** when installing a new software release (such as EX Series or QFX Series), when upgrading from certain releases to Junos OS Release 18.1R1, the statistics process pfd might generate a core file. This issue does not impact service. [PR1346925](#)
- The twice-napt-44 sessions are not syncing to backup SDG with stateful sync configured. [PR1347086](#)
- IPv6 MAC resolve fails if the DHCPv6 client uses a non-EUI64 link-local address. [PR1347173](#)
- Issue is seen with handling the community\_action ("add") in RPC call. [PR1348082](#)
- The FPC might crash due to MIC error interrupt hogging. [PR1348107](#)

- Per-service accounting statistic value is not accurate. [PR1348796](#)
- The chassisd might crash after replacing MPC6E or MPC7E with MPC9E. [PR1348834](#)
- The DHCPv6 solicit packet might be dropped on MX Series Virtual Chassis with L2TP LNS when the packet is received over a VCP port and the anchor si- interfaces exist on the same Packet Forwarding Engine as the VCP port. [PR1348846](#)
- On a single Routing Engine system, after the GRES, the configuration is removed. The Routing Engine mastership keepalive timer is not resumed to the default value. With the unexpected loss of Routing Engine mastership, issues such as chassisd stuck might be seen. [PR1349049](#)
- On all platforms, if any other smid-related daemon crashes, in a rare case, the dcd process might crash. [PR1349154](#)
- A major alarm **Major PEM 0 Input Failure** might be observed for DC PEM. [PR1349179](#)
- The RLT interface setup is broken. By design, the RLT interface is supposed to have a different L1 node and a different stream other than the tunnel stream. This is mentioned in the design specification of RLT and the source code as well. However, on MPC5E or MPC6E line cards and associated MICs, the RLT interface continues to be mapped to same tunnel stream and then on EA. It did not even get set up. [PR1350115](#)
- On platforms running Junos OS, pccd crash is observed in a PCEP scenario. [PR1350240](#)
- The multicast traffic might get dropped due to the "Invalid policy ID" exception. [PR1350380](#)
- The MTU value for subscriber's interface might be programmed incorrectly if the statement **routing-services** or **protocol pim** is configured in dynamic-profile. [PR1350535](#)
- The VCP port might not come back up after removing and adding it again. [PR1350845](#)
- The subinfo process might crash when executing **show subscribers address <> extensive** for a DHCPv6 address. [PR1350883](#)
- PPE Errors and async xtxn errors are seen when FPC restarts. [PR1350909](#)
- If the subscriber or interface statistics are used at large scale (thousands or more), the pfed process might consume high CPU because of the low performance code processing. This applies on all platforms and is primarily observed on PPC-based routers (such as MX104) when large-scale subscribers (such as 8000) log in to a subscriber management environment and accounting is turned on. [PR1351203](#)
- The high CPU usage of bbe-smgd process might be seen when L2BSA subscribers get stuck. [PR1351696](#)
- After GRES, the BGP neighbors at master Routing Engine might reset and the BGP neighbors at backup Routing Engine take long time to establish. [PR1351705](#)
- The bbe-smgd daemon might restart in a subscriber environment. [PR1352546](#)
- In the DHCPv6 relay scenario, there is a conflict with IPv6 relay-reply packet processing when forward-only and the relay-source overrides are configured for the same interface group. This causes the packets to be dropped by the route lookup logic when the packet is sent back toward the client. [PR1352613](#)

- Offline MIC6-100G-CFP2 MIC through the CLI command might trigger FPC card to crash. [PR1352921](#)
- The routing protocol process (rpd) permanently overuses CPU due to logical system configuration commit. [PR1353548](#)
- On platforms running Junos OS, if GRES is not configured, multiple Routing Engine switchover might cause traffic interruption because the old forwarding information base (FIB) state is not getting cleaned up. [PR1354002](#)
- Syslog error: **dfw\_bbe\_filter\_bind:1125 BBE Filter bind type 0x84 index 167806251 returned 1.** [PR1354435](#)
- The rpd process generates a core file when adding an inter-region template in routing instances. [PR1354629](#)
- Starting with the next-generation subscriber management on Junos OS, the static subscribers might not properly update the firewall information on the Packet Forwarding Engine when dynamic configuration changes are made to active subscribers. As a result, complete traffic loss for the client might be seen. [PR1354774](#)
- Memory leak is found in agentd while running valgrind. [PR1354922](#)
- Packets destined to Routing Engine might be dropped in the kernel when LACP is configured. [PR1355299](#)
- The fabric chip failure alarms are observed in GRES scenario. [PR1355463](#)
- The rpd process crashes when issuing the command **show dynamic-tunnels database terse** for RSVP automatic mesh tunnels. [PR1356254](#)
- The I2c messages from PEM/PSM are reported if SNMP is enabled. [PR1356259](#)
- The CLI command **show pppoe underlying-interfaces** in a scaled environment might cause bbe-smgd memory leak. [PR1356428](#)
- The bbe-smgd generates a core file in recursive loop between functions bbe\_autoconf\_if\_I2\_input and bbe\_if\_I3\_input. [PR1356474](#)
- DHCP subscribers fail after reconfiguration of port from tagged to un-tagged mode. [PR1356980](#)
- On all platforms running Junos OS that have dual Routing Engines, if GRES is enabled to provide High Availability (HA) protection, the backup Routing Engine (RE1) might be out of synchronization with the master Routing Engine (RE0), and the kernel state in the backup Routing Engine (RE1) is not cleaned because of a software defect. After staying in such status for a long time, once the keepalive timeout is detected between the master and backup Routing Engine, the backup Routing Engine (RE1) might take over the mastership. All the line cards will be restarted when they are connected to the new master Routing Engine (RE1) after switchover because of the missing master-backup synchronization. Then the new master Routing Engine (RE1) might crash because some data structure field overflows in the kernel because the kernel state has not been cleaned for a long time. After that, the original master Routing Engine might take the mastership back again. This issue causes complete traffic loss. [PR1357427](#)
- On all platforms running Junos OS, when the system runs with a large scale of subscribers (for example, more than 30,000), if the subscriber interfaces are configured with tail/wred drop rules and different

buffer-size values, the PIC concentrator (MPC/FPC) might take too much process time for adding or deleting tail/wred drop rules during binding or releasing subscribers, so that it cannot reply any request messages to the Routing Engine for a long time. Because of this issue, a lot of kernel timeout error logs might also be seen. [PR1358405](#)

- On MX Series routers, if many subscribers are logging in simultaneously, bbe-smgd crash might be seen. [PR1358868](#)
- When an FPC (or an incompatible one) is powered off by configuration or CLI command and the command **show chassis environment fpc** is issued, the status of the FPC might change to **---Bad Voltage---** under **show chassis fpc**. [PR1358874](#)
- The IPv6 subscriber might fail to access network. [PR1359520](#)
- On MX Series routers, if **system services subscriber-management enabled** is configured, bbe-smgd might fail to add members to some of the aggregated Ethernet interfaces randomly when there are many aggregated Ethernet interfaces in the access configuration. [PR1359986](#)
- When the rpd reads next hops from kernel on restart, for the INH -> FWD NH{List NH} -> {Chain NH} scenario, the rpd should not create an old-style list next hop for the forwarding next hop. [PR1360354](#)
- If groups are applied on the top level, when these groups are deleted, modified, and added, all the top level hierarchies that are referred by these groups will be set with a "mark-changed" bit. Everything under these hierarchies is considered as changed. If these groups refer to policy-options and there are policies referring to prefix-list, each prefix in prefix-list is marked as changed even though the prefix-list is actually not changed at all. This causes the duplicate prefix to be added to prefix-list. When the group adding, modifying, and deleting operation is frequently executed, the issue might cause more CPU usage by policy processing, which in turn might cause rpd scheduler slip. [PR1361304](#)
- The DSCP value in customer IP traffic gets rewritten unexpectedly when the **routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp inet6** statement is configured on the core-facing MPLS interface and a certain EXP rewrite rule is applied. [PR1361429](#)
- MX Series routers do not generate a quality level failed alarm (Trap-Id:.1.3.6.1.4.1.2636.3.75.1.1.7) when the transmit SSM-QL is reduced from a valid SSM-QL to a value below the minimum SSM-QAL (SSUB/EEC). [PR1361430](#)
- When a peer is being established when it needs to catch up with other peers that have received many more updates, the merge code might verify the routes that are to be announced. If none of the prefixes are announced before the peer has processed its fair share of entries, the process starts from the beginning again. This situation is most likely to occur when there is zero route churn. [PR1361550](#)
- In a subscriber management environment, because of a timing issue, the bbe-smgd process on the backup Routing Engine might crash either during login of a subscriber with a multicast service or during activation of multicast service for an existing subscriber. [PR1362188](#)
- If the route installation failure case is not handled properly in a BGP multipath scenario, traffic loss might be seen. [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/MS-MIC is brought offline/online. [PR1362734](#)
- On MX Series routers with non-default routing instance subscribers configured, NTP packet might not use the correct non-default routing instance. [PR1363034](#)
- On MX2010 and MX2020 routers equipped with Switch Fabric Board 2 (SFB2), some error messages could be occasionally seen in the logs. There is no operational impact, nor an indication of a real issue caused by these messages. [PR1363587](#)
- The xmlproxycd for internal interfaces is reporting uint32 instead of uint64. [PR1363766](#)
- In Junos OS, during any route change, the kernel and rpd communicate multiple times to update the route and forwarding entries. In large-scale scenarios, where the system contains multicast composite next hops, during any network events that might cause route/next-hop churn at a huge scale, rpd might skip a route operation (DELETE). As a result, krt queue entries for multicast next-hop (MCNH) deletes might get stuck in the krt queue. Therefore the kernel and rpd could go out of synchronization and potentially cause rpd to crash if it encounters a request from the kernel for a route update that is not in line with its own dataset. The MCNH deletes are not sent to the kernel, which indicates improper error handling in the rpd for route DELETES. [PR1363803](#)
- On EX4600 and QFX5100 platforms, if Rapid Spanning Tree Protocol (RSTP) is configured along with aggregated Ethernet, a traffic loop might be seen in a ring topology even though that port is blocked by RSTP. [PR1364406](#)
- The **shmlog** files are not rotated correctly. As a result, there is an increase in file size that consumes most of the disk space. [PR1364775](#)
- The smgd process might restart unexpectedly when stress tests are performed on subscriber management features. [PR1372223](#)
- On MX2010 platforms, if an aggregated Ethernet bundle is configured with Ethernet OAM link fault management (LFM), and at the same time, no Link Aggregation Control Protocol (LACP) is configured for the aggregated Ethernet bundle, the aggregated Ethernet member link flap might cause one member link to be set as "Link-Layer-Down" by LFM even after its physical link is already up. Because of this issue, there are still traffic flows forwarded through the member link in faulty status. Thus, all the traffic affected might be lost, which might lead to service impact. [PR1365263](#)
- Default NIC driver coming as E1000 when vFPC is deployed on VMware uses an OVA image. As a workaround, vmxnet3 is used as the default NIC driver. [PR1365337](#)
- MS-MPC and MS-PIC might crash if two or more service set are configured with the same prefix lists and the SIP ALG is configured in a NAT scenario. [PR1366259](#)
- If an MPLS path uses an IPv6 next hop, the next hop might be stuck in hold state. Initially, the router triggers the IPv6 Neighbor Discovery (ND), but the neighbor advertisement from the peer is not received. Eventually, the neighbor state moves to unreachable state and the next hop of the MPLS path using this

neighbor is rejected. After this, if the router receives a neighbor solicitation message from the peer, the neighbor state might move to reachable state in the IPv6 neighbor table. The IPv6 module should notify the change to MPLS module, but somehow the notification is missed. This causes the next hop of the MPLS path to get stuck in hold state. [PR1366562](#)

- In a Layer 2 bit stream access (L2BSA) subscriber scenario, if there is a misconfiguration on the RADIUS profile for the L2BSA subscriber (for example, the routing instance returned from RADIUS is not configured as VPLS) or an authentication part is missing in the physical interface configuration, the bbe-smgd process might crash during the L2BSA subscribers login. [PR1367472](#)
- After replacing a FPC having more Packet Forwarding Engines with a FPC having less Packet Forwarding Engines (for example, replace DPC having two Packet Forwarding Engines with MPC3E having only one Packet Forwarding Engine), the nonexistent Packet Forwarding Engine might be shown with the command of **show system resource-monitor fpc**. This can be cleared by using restart subscriber-management >restart subscriber-management gracefully. This restart will not affect your services. However, if there are systems actively collecting the interface statistics, services might pause a little bit while the process restarts. [PR1367534](#)
- RTG interface status is shown as incorrect status with **show interface {master:1}[edit] root@host# show switch-options | display set set switch-options redundant-trunk-group group rtg2 interface xe-1/0/5.0 set switch-options redundant-trunk-group group rtg2 interface xe-1/0/6.0 set switch-options redundant-trunk-group group rtg3 interface xe-0/0/2.0 set switch-options redundant-trunk-group group rtg3 interface ae3.0 root@host# run show interfaces terse | match "xe-1/0/6|xe-1/0/5" xe-1/0/5 down down xe-1/0/5.0 up down eth-switch xe-1/0/6 down down xe-1/0/6.0 up down eth-switch {master:1}[edit] root@jtaq-qfx5100-48t-6q-r2284# run show interfaces terse | match rtg rtg2 up up <<<< incorrect status rtg2.16383 up up eth-switch <<<< incorrect status rtg3 up up rtg3.16383 up up eth-switch**. [PR1368006](#)
- On MX Series routers in BBE configurations, receipt of a crafted IPv6 exception packet causes a denial-of-service (CVE-2018-0058). Refer to <https://kb.juniper.net/JSA10893> for more information. [PR1368599](#)
- While performing an SNMP MIB walk for OID jnxIpSecTunnelEntry, the following errors are seen: **May 25 00:30:04 labbox\_re0 kmd[17150]: KMD\_SNMP\_PIC\_NO\_RESPONSE: PIC rsp1 did not respond to SNMP query: No error: 0 May 25 00:30:04 labbox\_re0 kmd[17150]: KMD\_SNMP\_FATAL\_ERROR: Fatal SNMP error occurred: libservicesui: ipc\_pipe\_read() failed - No error: 0**. [PR1369938](#)
- On MX Series routers that support next-generation subscriber management (Apache Tomcat), when the Layer 2 Tunneling Protocol (L2TP) Network Server (LNS) is enabled, if the dynamic profiles are configured with the statement **routing-services** and the firewall filter, the firewall filter might not be removed from the Packet Forwarding Engine after subscriber logout. Due to this issue, the firewall filter index might be used and then no more subscribers can log in. [PR1369968](#)
- In a subscriber management scenario, if an aggregated Ethernet interface is associated as the underlying-interface of a demux0 unit and both the demux0 unit and aggregated Ethernet unit (corresponding to the above aggregated Ethernet interface) are configured with a duplicated vlan-id, the kernel might crash after committing the configuration. [PR1370015](#)

- With an interface-based dynamic GRE tunnel configured, there might be two next hops for a single dynamic GRE tunnel. When a new route is resolved over the dynamic tunnel after Routing Engine switchover is performed or the rpd is restarted, subsequent withdrawal of the routes over that tunnel or master Routing Engine restarting might cause the rpd to crash. [PR1370174](#)
- ALG cannot process IP datagrams exceeding 8000 bytes in size. The packets are dropped by junos-alg plugin. Plugin-related packet drop counter captures these drops. If an IP datagram is not related to ALG sessions, then the junos-alg plugin has nothing to do with them and they are ignored (ALG plugin will not drop). [PR1370582](#)
- An FPC crashes and generates a core file under heavy load, causing the bbe-smgd process to generate a core file. This core file is generated because of the cleanup issues with the VLAN creations in flight. [PR1371926](#)
- MX Series with MPC or MIC-based line card might reach high CPU utilization or might crash because of a defect in handling a memory hot-banking condition. [PR1372193](#)
- 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 be exhausted, which leads to the creation failure of the child next hop. Hence, the Routing Engine might crash. As a workaround, you can avoid this issue by enabling LAG-enhanced. [PR1373079](#)
- The URL filtering feature might not work in Junos OS Release 17.4R2 when the data interfaces participating in the URL filtering functionality move from one routing instance to another routing instance. [PR1373582](#)
- If BOOTP-support is not enabled at the global level, bootstrap protocol (BOOTP) packets might be dropped while they are received on an interface because the device only checks BOOTP-support at the global level. [PR1373807](#)
- Cosmetic log **warning: [---] is protected, 'protocols ---' cannot be deleted** is seen after the commit using **configure private** in a configuration with "protect" flag present. [PR1374244](#)
- On MX Series routers that support the next-generation subscriber management, if the aggregated Ethernet bundle has multiple child interfaces that are located in the same Packet Forwarding Engine complex,(for example, ge-1/0/0 and ge-1/0/1), when the dynamic VLAN subscriber gets online from the aggregated Ethernet bundle, then one physical child interface is removed out of the aggregated Ethernet bundle, (for example, ge-1/0/0). The Flexible PIC Concentrator (FPC) might keep reporting error logs, and the statistics on the dynamic VLAN flow also would not get incremented. Therefore, the Packet Forwarding Engine might be unable to work properly. [PR1374478](#)
- In case of centralized IGMP configuration, the bbe-smgd daemon might restart last subscriber of a multicast group is leaving the group. It is not able to delete this multicast group node from the tree. In this case on daemon restart, in INIT phase, bbe-smgd might again try to delete the multicast group node and its associated multicast group service and restart again. Because of this, the bbe-smgd process might never complete the INIT phase and restart continuously in INIT phase only. [PR1374530](#)
- On MX Series routers with MS-MPC/MS-MIC installed, ICMPv6 packets larger than 1024 might be dropped if **icmp-large-packet-check** is configured on the IDS service. [PR1378852](#)



- The software detects SDB STS lock deadlock and breaks the deadlock itself and system resumes normally processing on its own. [PR1380231](#)
- On MX Series routers with MS-MPC installed, memory leak might be observed when **vty mspdbg-cli** command is executed. [PR1381469](#)
- Rarely, over GRES or Routing Engine reboot, subscribers of all access types were not able to login. As a workaround, restarting the bbe-smgd daemon might solve the issue. [PR1382050](#)
- The export of the J-Flow records is seen at the collector before the configured active timeout value. This export result might not be the expected. [PR1382531](#)
- On M Series, MX Series and T Series platforms, after bringing up IPsec tunnels, when issuing the **show** command, a kmd crash might be seen. [PR1384205](#)
- When dynamic IPsec VPN is rekeyed due to lifetime expiration, IPsec internet key exchange (IKE) phase 1 user datagram protocol (UDP) port 500 and phase 2 UDP port 4500 sessions might be translated into two different public internal protocol (IP) addresses while passing through carrier-grade network address translation (CGNAT), which causes IPsec VPN traffic to fail. This behavior does not cause an issue for Juniper MX Series routers devices with MS-MIC or for SRX Series devices, because for these devices an identity key is used to authenticate the sessions and it is allowed for private IP address to be translated into two different public IP addresses. [PR1386011](#)
- MX Series BNG does not allow two subscribers with the same framed-route prefix and preference values. It allows the first subscriber to log in, while the second subscriber is denied access. When the second subscriber tries to log in, the bbe-smgd daemon crashes and generates a core file. [PR1387690](#)
- In a subscriber management environment, if CoS adjustment is performed for DHCP subscribers-based on DHCP tags, output of the **show class-of-service interface** command for a DHCP subscriber interface might incorrectly show the adjusting application as PPPoE IA tags instead of DHCP tags. [PR1387712](#)
- The bbe-smgd does not respond to NS from the SLAAC client on dynamic VLAN. [PR1388595](#)
- In a subscriber-management environment, it is not possible to control CoS adjustment based on DHCP tags, because the configuration command **class-of-service adjustment-control-profile <profile-name> application dhcp-tags** is ignored. Both CoS adjustments based on PPPoE IA tags and based on DHCP tags were controlled by the command **class-of-service adjustment-control-profile <profile-name> application pppoe-tags**. [PR1390101](#)

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

- With GRES enabled and **set system syslog file messages daemon any** configured, a log message regarding ksyncd might be generated on the backup Routing Engine. [PR1203163](#)
- The ksyncd might crash. [PR1275022](#)
- A node-slicing setup downing the CB ports on both servers might result in one or more GNFs displaying “not ready” under the **show system switchover** command. Performing a NSR in this state might result in protocol flaps and traffic disruption. As a workaround, run the **restart kernel-replication** command on the backup Routing Engine. This will restart ksyncd and make the system GRES ready. [PR1306395](#)



- On MX Series routers with MS-DPC, if sampling or flow-monitoring is configured, the ksyncd on the new backup Routing Engine might crash continuously after performing a GRES. This might cause GRES to not be ready. The ksyncd becomes unrecoverable until you reboot the backup Routing Engine. [PR1329276](#)
- When GRES is configured with large-scale configurations (for example, 20,000 subscribers), if the ksyncd repeatedly runs into replication error, the kernel synchronization process (ksyncd) triggers a "gather-crashinfo" script, which is run by ksyncd internally, to generate debug information into files on both the master and backup Routing Engines. If the files generated run into GB, then it might lead to insufficient available space on the hard disk. And the debug information as well as all the core will be saved in one single .tgz file on the backup Routing Engine. [PR1332791](#)
- The following error is seen during early unified ISSU validation phase: **error: not enough space in /var on re1**. As a workaround, make sure that the space available in /var is twice the size of the target image. This is the basic requirement for unified ISSU to proceed. [PR1354069](#)
- In an MX Series Virtual Chassis scenario, if any events cause the Virtual Chassis to split, then reforms such as VCP port flaps or backup restarts, or the master Routing Engine in the Virtual Chassis backup router (VC-Bm) might not synchronize with the master Routing Engine in the Virtual Chassis master router (VC-Mm). [PR1361617](#)

### Infrastructure

- The rpcbind service opens a nonsecure secure port (111) to the outside world. As workaround, restrict the service only to internal ports. [PR1296262](#)
- The **syscalltrace.sh** script gets installed as part of the Junos OS starting in Junos OS Release 16.1R1 and later releases. The script is triggered whenever there is a replication error on the backup Routing Engine. It logs the system function call to the output file, which provides additional debug information. But it might create large files because of a bug in this script. As a workaround, it is recommended to uninstall this script after Junos OS is upgraded in the production network. The uninstallation of this script will not have any functionality impact on the router. [PR1306986](#)
- Kernel crash (vmcore) occurs during broadcast storm after enabling **monitor traffic interface fxp0**. Refer to <https://kb.juniper.net/JSA10863> for more details. [PR1322294](#)
- The freeBSD kernel creates threads to perform various tasks, but when these threads exit portions of their memories are not released properly. [PR1328273](#)
- On all platforms running Junos OS, on a port configured with both dot1x static MAC by pass and normal authentication, the hosts configured for static MAC by pass might not be able to send traffic. [PR1335125](#)
- A kernel crash is seen and the system will restart after the device issues a race condition in an SNMP query reply scenario. [PR1351568](#)

## Interfaces and Chassis

- On MX240, MX480, and MX960 platforms IPv6 neighborship is not created on the IRB interface. [PR1198482](#)
- If there are optical transport network (OTN) interfaces on the router, the output value is incorrect when you use the CLI and SNMP walk to query the optical power of these interfaces. This is a cosmetic issue with no traffic impact. The displayed value does not represent the actual optical power of the interfaces. [PR1216153](#)
- On MX Series MPC7E, MPC8E, and MPC9E line cards, the Packet Forwarding Engine crashes while fetching interface-statistics with extended-statistics enabled (CVE-2017-10611). Refer to <https://kb.juniper.net/JSA10814> for more information. [PR1247026](#)
- Rate-limit dropped packets are not displayed by the `[show interfaces <ifl> detail]` and `[show interfaces <ifl> extensive]` commands. The drop can be seen with the `show interfaces queue` command. This is cosmetic issue and traffic is passing correctly. [PR1249164](#)
- The jpppd process might report error messages about RLIMIT\_STACK and RLIMIT\_SBSIZE after issuing the command `show version detail`. [PR1262629](#)
- Continuous error messages might be seen when the physical interface quickly flaps on MPC7, MPC8, and MPC9E cards. This might cause egress stream flush failure. [PR1271089](#)
- The BERT test shows the elapsed time "in progress" but gets stuck and never gets completed. [PR1274896](#)
- Starting with Junos OS Release 16.1 and later releases, the monitor interface on the aggregated Ethernet logical interface shows incorrect BPS value compared to `show interface` output. The issue is not visible if taking value of the monitor interfaces on aggregated Ethernet physical interfaces. [PR1283831](#)
- When executing Routing Engine switchover, the dcd process will do a check on the aggregated Ethernet interface. The check will fail if the aggregated Ethernet interface has a member interface with "framing" settings. The failed check will trigger both the aggregated Ethernet interface and its member interface to flap. [PR1287547](#)
- The family inet shows as **Not configured** after adding or deleting the loopback address. [PR1294267](#)
- In a Layer 2 Tunneling Protocol (L2TP) scenario with enhanced subscriber-management mode and an MX Series routers working as L2TP network server (LNS), some L2TP subscribers with fixed-IP returned by RADIUS might not be cleared if the access-internal routes of such subscribers fail to install. [PR1298160](#)
- With this change, you can configure **delay-buffer-rate** on inline LSQ interfaces. [PR1300281](#)
- IRB interface is showing incorrect bandwidth value. [PR1302202](#)
- If one logical interface changes virtual router (VR) state from master to backup, the traffic might get silently dropped and discarded for other logical interfaces that shares the same group ID on a physical interface. [PR1305327](#)
- On MX104 platforms with the `set system process ethernet link-fault-management disable` command configured, AFEB might not come up after restarting the router/AFEB. [PR1306707](#)

- After executing the command **request system reboot both-routing-engines** in a GRES scenario, the jpppd process might become unresponsive and stop handling Point-to-Point Protocol (PPP) control traffic. No subscribers can log in. [PR1310909](#)
- In a PPPoE subscriber management scenario, if subscriber authentication fails, the subscriber logical interface will be in disable state. This will cause the jpppd process to drop the next Link Control Protocol (LCP) termination request packet from the subscriber, instead of answering it with LCP Ack and closing the PPPoE session with a PPPoE Active Discovery Termination (PADT) packet. This might impact the session setup for this subscriber. [PR1311113](#)
- There are two issues regarding this problem report. The first one is that if ufec with OTN is configured, and the physical link goes down, CPU will go to 100 percent. The second issue is that when ufec with OTN is configured on unconnected interfaces, CPU will go to 100 percent. [PR1311154](#)
- The ifinfo process might crash and generate a core file when executing CLI command **show interfaces <name>** with the name greater than 128 characters. [PR1313827](#)
- There is no route to the IP address from the directly connected route on the static VLAN demux interface if the configuration of the static VLAN demux interface is changed from an unnumbered approach to a configuration with an explicit IP address (for example, /30 ). [PR1318282](#)
- On MX Series routers, when PPPoE is tunneled at the MX Series router (LNS - L2TP network server) or the PPPoE session terminates at the MX Series router, the PPP NCP cannot be enabled in the active mode for the multilink PPP session. [PR1319580](#)
- The result of the output is incorrect when executing the command **show interfaces interface-set**. It is a display issue related to logical interface on MX80, MX10, and MX104. [PR1319682](#)
- When running an MX Series router for BNG and subscriber management functionalities, the dual-stacked subscriber "IPv6 framed interface id field" (from **show subscribers extensive** output) is not matching the negotiated one. [PR1321392](#)
- Internet Protocol Control Protocol (IPCP) negotiation might fail for dual-stack PPPoE subscribers. [PR1321513](#)
- In subscriber management scenario with DEMUX configured, in the case where subscribers belonging to one aggregated Ethernet interface are migrated to a new configured aggregated Ethernet interface, subscribers might fail to access the device after deleting the old aggregated Ethernet configuration. [PR1322678](#)
- If a BGP session flaps in a dynamic-tunnels Generic Routing Encapsulation (GRE) scenario, fault log messages might fill log files. The issue does not have an impact on traffic. [PR1326983](#)
- In a PPPoE subscriber environment, continuous fault log messages might be seen on the backup Routing Engine. The issue does not have an impact on service. [PR1328251](#)
- The issue occurs when multiple Virtual Router Redundancy Protocol (VRRP) groups are separately configured on different units of an aggregated Ethernet bundle, the unit 1 of which has both inner and outer VLAN configured. All the other VRRP groups might malfunction with a period of the time configured by "failover-delay" under the VRRP stanza, after deleting aggregated Ethernet bundle unit 1. [PR1329294](#)

- The cfmd process generates a core file and restarts while the cfmd iterator and/or rmep statistics are being retrieved. [PR1329779](#)
- In the case where the interface is configured as a member of interface-set, it might not work properly after an unrelated FPC (not the one where the interface resides at) restarts. The affected FPC is the restarted one. [PR1329896](#)
- In some situations, like multiple commits in a short time with a scaled configuration, dcd memory leak might occur. This could cause the commit to fail. [PR1331185](#)
- In a subscriber management environment, trace logs for jpppd process (configurable in ppp-service stanza) might miss the last digit of the interface name. [PR1332483](#)
- On platforms running Junos OS and with SNMP configured, the SNMP requests optIfOChSinkCurrentExtTable on the valid interface. If the result data is invalid, this might cause transportd to crash. Transportd keeps the interface configuration cache. After it crashes, the function of the interface discontinues and recovers soon. This issue might be seen on all platforms with transportd support for optics. [PR1335438](#)
- On GRES the implicit filters set by DFWD are cleared by DCD. Hence, a momentary dip in traffic is observed. [PR1336455](#)
- On MX Series Junos subscriber management (JSM) environment, when GRES is disabled, restarting chassisd might cause the FPC to restart and some demux interfaces to be deleted. [PR1337069](#)
- When multiple VRRP sessions with the same group-id are configured on the same port (aggregated Ethernet interface or a physical interface), the VRRP virtual IP will be not reachable. [PR1338277](#)
- The MX Series router running in PPPoE subscriber management mode drops the first incoming LCP configure-request message and accepts the subsequent packets. Because of this behavior, the customer might incur a small latency in establishing the subscriber connection. [PR1338516](#)
- The 100G dense wavelength division multiplexing (DWDM) interface might be going down for 15 seconds after a loss of signal event. [PR1343535](#)
- On MX Series routers with PPPoE subscriber scenario, when **on-demand-ip-address** is enabled, a high frequency of on-demand IP address allocation requests might be seen. As a result, authd runs in high CPU usage and subscriber login fails. [PR1348578](#)
- In L2TP scenario when MX Series router functions as L2TP Tunnel Switch (LTS), there is a memory leak in jpppd process running on the backup Routing Engine, which will eventually lead to generating jpppd core file because of an out-of-memory condition. There is no functional impact, because this action occurs on the backup Routing Engine. [PR1350563](#)
- If the multichassis aggregated Ethernet (MC-AE) is configured with enhanced-convergence and the number of logical interfaces under the aggregated Ethernet physical interface are high, the FPC might be stuck at 100 percent during initial configuration load or FPC restart and this might result in other event processing being delayed. This issue only affects MX Series routers with MX Series-based FPC. [PR1353397](#)

- When the link speed of the aggregated Ethernet bundle is configured to oc192, certain sequence operations might lead to the aggregated Ethernet interface flap, which affects the traffic. As a workaround, configure the member links, then, remove a member link from the bundle, and then add a member link back. [PR1355270](#)
- When **on-demand-ip-address** is configured, the PPPoE client remains in an endless loop of continuously sending IPCP configuration requests. [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 subscriber record loss. These subscribers, however, should be able reconnect right away minimizing any service outage. [PR1360870](#)
- The maximum number allowed for subinterfaces of a LAG interface is 2048. However, a software defect introduced in Junos OS Release 17.2R1 does not enforce this. The problem should be fixed to enforce the maximum number of allowable subinterfaces. [PR1361689](#)
- Messages like **dcd[5304]: is\_ih\_chan\_ci\_candidate: 2124 ifname [ds-5/0/2:4:1] is chan ci candidate** are reported by DCD with priority of ERROR during a commit operation. These do not denote any operational impact and can be filtered out safely. [PR1363536](#)
- In a corner case, in which the pfed daemon is still initializing after fresh upgrade, and jpppd is up and processing subscriber login, a subscribers issue might occur. This is because jpppd ends up waiting indefinitely for pfed to respond with a subscriber statistic request. [PR1368650](#)
- On MX Series routers that supports dynamic Multilink Point-to-Point (MLPPP) subscriber, if the dynamic-profile name contains more than 30 characters, MLPPP subscribers might be unable to negotiate sessions with the server, and cannot log in because of this issue. [PR1370610](#)
- Under rare circumstances, MX Series Virtual Chassis unified ISSU might abort with the message **Timed out Waiting for protocol backup chassis master switch to complete**. [PR1371297](#)
- If **vlan-id none** is configured for interface (for example, **set interfaces <xx> unit <xx> vlan-id none**), the dcd process will go down after committing this configuration. A check error should be reported when committing this configuration so as to avoid the dcd crash. [PR1374933](#)
- On MX480 MCLAG, when **parse\_remove\_ifl\_from\_routing\_inst()** command is executed, the following error is seen after l2cpd daemon is restarted: **ERROR : No route inst on et-0/0/16.16386**. [PR1373927](#)
- On MX Series routers, if configuring duplicate IP on the SONET (so-) interface between and another type of interface, the other interface might not get the IP address. [PR1377690](#)
- On MX Series routers, the bbe-subscriber management daemon (bbe-smgd) reports some error logs because of jpppd sent out **LCP config-reject** message. However, the bbe-smgd misses such type message code in Tx direction. It has no service impact. [PR1378912](#)
- Static demux interface stacking over the PS interface is not supported and might cause the dcd process to restart. The commit process should not allow such a configuration. [PR1382857](#)

### **J-Web**

- On Junos OS a denial-of-service vulnerability in J-Web service might allow a remote unauthenticated user to cause denial-of-service, which might prevent other from authenticating or performing J-Web operations. Refer to <https://kb.juniper.net/JSA10897> for more information. [PR1264695](#)

### **Layer 2 Ethernet Services**

- After changing the outer vlan-tags, the logical interface is programmed with incorrect STP state and gets discarded. Hence, the traffic drop is seen. [PR1121564](#)
- When MSTP or VSTP is configured, if GRES is enabled but nonstop bridging (NSB) is not enabled, after Routing Engine switchover, the MAC address might not be learned due to spanning-tree state "discarding" in the kernel table. [PR1205373](#)
- The IA\_PD prefix might be deleted when MX Series routers receives a DHCPv6 IA\_NA request. [PR1286359](#)
- ARP requests are not generated for IRB configured in VPLS over GRE tunnel. [PR1295519](#)
- PPPoE/DHCP clients cannot log in to PPPoE/DHCP dual-stack subscriber scenario. [PR1298976](#)
- A parameter-handling problem might cause the kernel to panic when a neighbor discovery message arrives on an IRB interface. [PR1303415](#)
- Multiple jdncpd core files are generated in jdncpd\_update\_groups at `../..../src/junos/usr/sbin/jdncpd/jdncpd_config.c:2290`. [PR1311569](#)
- DHCPv6 traffic might be dropped in subscriber scenario. [PR1316274](#)
- jdncpd core file might be generated after making DHCP configuration changes. [PR1324800](#)
- The snmpget for OID: dot3adInterfaceName might not work. [PR1329725](#)
- A denial-of-service vulnerability exists in the Juniper Networks Junos OS jdncpd daemon that allows an attacker to generate jdncpd daemon core files by sending a crafted IPv6 packet to the system. Refer to <https://kb.juniper.net/JSA10868> for more information. [PR1334230](#)
- In a DHCPv6 environment, when a DHCPv6 packet with a big client-id option size (more than 255 bytes) is received, the jdncpd process might spike to 100 percent, which results in memory corruption and an unusable DHCP service. [PR1334432](#)
- The memory leak might happen in l2cpd if the l2-learning process is disabled. [PR1336720](#)
- On MX Series platforms with DHCPv6 running over the access model where the underly is a PPP session (LNS or PPPoE), when the customer premises equipment (CPE) sends a separate DHCPv6 solicit message for the IA\_NA and the IA\_PD prefix, the second solicit for the same session might fail. [PR1340614](#)
- ZTP process does not start to load image and configuration for MX Series PPC routers, because there are no ZTP infrastructure scripts. [PR1349249](#)
- 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 discards DHCP request message silently if the requested IP address has been allocated to the other client. [PR1353471](#)
- Restart FPC with homing micro-bfd link causes lacp to generate a core file. [PR1353597](#)
- DHCPv6 relay ignores responses from server when renewed. [PR1354212](#)
- jdhcpd crashes during processing DHCPv6 information request. [PR1368377](#)
- On MX Series routers, if static demux interface over underlying is configured, after subscriber logout, the accounting statistics are not cleared. [PR1383265](#)

### **Layer 2 Features**

- On MX Series routers with MPCs or MICs-based platforms, packets received on the IRB interface in Virtual Private LAN Services (VPLS) will get double tagged. [PR1295991](#)
- The rpd process memory leak is observed upon any changes in VPLS configuration like deleting/re-adding VPLS interfaces. [PR1335914](#)
- VPLS instance stays in NP state after LDP session flap. [PR1354784](#)

### **MPLS**

- Minor difference between mpls.statistics and adjusted BW. [PR1259500](#)
- Potential issues are seen with policy based selection of RSVP LSPs. [PR1261739](#)
- The ingress RSVP LSP fails to come up after **clear rsvp lsp** is performed on egress router. [PR1275563](#)
- The rpd might crash in LDP L2circuit scenario. [PR1275766](#)
- MPLS I2ckt ping packet is incorrectly parsed by the output loopback filter. [PR1288829](#)
- In Junos OS Release 16.1 and later releases, if LDP egress policy is used for inet.3 BGP labeled-unicast route, the route lable might not be installed in the Label Distribution Protocol (LDP) database. [PR1289860](#)
- Received MTU might not get updated in RSVP MTU signaling. [PR1291533](#)
- The process rpd might crash when performing MPLS traceroute. [PR1299026](#)
- The traffic in P2MP tunnel might be lost when next-generation MVPN uses RSVP-TE. [PR1299580](#)
- The rpd process might crash in rare conditions where traffic-engineering is configured. [PR1303239](#)
- On all platforms running Junos OS with the Junos Telemetry Interface, MPLS statistics sensor and GRES are configured. The kysncd process might crash when the backup Routing Engine is removed or inserted. [PR1303491](#)
- The feature "explicit-null" might block host-bound traffic incoming from LSP. [PR1305523](#)
- The RSVP node-hello packet might not work correctly 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 if LDP updates the label for BGP route. [PR1312117](#)
- Delayed **show mpls container-lsp** output. [PR1314960](#)
- RSVP node-neighbor is found even when node-hello has been disabled. [PR1317241](#)
- The rpd might crash after the primary link failure of link protection. [PR1317536](#)
- With dynamic tunnels configured, the rpd might crash when the rpd is restarted or Routing Engine switchover is executed. [PR1319386](#)
- The IPv4 or IPv6 multicast traffic might get dropped in MX Series Virtual Chassis when the traffic comes in through Layer 2 circuit and goes out through aggregated Ethernet member interface across Virtual Chassis members. [PR1320742](#)
- The rpd might crash because of the memory leak in RSVP scenario. [PR1321952](#)
- Receipt of specially crafted UDP/IP packets over MPLS might be able to bypass a stateless firewall filter. The crafted UDP packets must be encapsulated and meet a very specific packet format to be classified in a way that bypasses IP firewall filter rules. The packets themselves do not cause a service interruption (for example, rpd process crash), but receipt of a high rate of UDP packets might be able to contribute to a denial-of-service attack. [PR1326402](#)
- SNMP OID counters for mplsLspInfoAggrOctets show constant value for some LSPs even though traffic is constantly increasing in **show mpls lsp statistics**. [PR1327350](#)
- Receipt of a specific MPLS packet might cause the routing protocol process (rpd) to crash and restart or might lead to remote code execution. By continuously sending specific MPLS packets, an attacker can repeatedly crash the rpd process, causing a sustained denial-of-service. Refer to <https://kb.juniper.net/JSA10877> for more information. [PR1328058](#)
- Packets loss might be observed when auto-bandwidth is enabled for CCC connections. [PR1328129](#)
- The rpd might crash on backup Routing Engine because of the memory exhaustion. [PR1328974](#)
- The statement **install-nexthop lsp/lsp-regex** in policy does not work with dynamic LSPS (Rsvp automesh). For example, an egress PE device is reachable through three automesh lsp **set routing-options dynamic-tunnels LSP-OTHER-AUTOMESH-1 rsvp-te LSP-OTHER label-switched-path-template LSP-OTHER-TEMPLATE**, **set routing-options dynamic-tunnels LSP-OTHER-AUTOMESH-2 rsvp-te LSP-OTHER label-switched-path-template LSP-OTHER-TEMPLATE**, **set routing-options dynamic-tunnels LSP-TEST-AUTOMESH rsvp-te LSP-TEST label-switched-path-template LSP-TEST-TEMPLATE**, [edit] **user@host# run show route 2.2.2.2 inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)** **+ = Active Route, - = Last Active, \* = Both 2.2.2.2/32 \*[IS-IS/18] 1d 02:30:39, metric 20 > to 10.1.1.2 via ge-0/0/1.0 to 10.1.1.6 via ge-0/0/2.0 inet.3: 1 destinations, 2 routes (1 active, 0 holddown, 0 hidden)** **+ = Active Route, - = Last Active, \* = Both 2.2.2.2/32 \*[RSVP/7/3] 00:00:08, metric 20 > to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-1 to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-2 to 10.1.1.6 via ge-0/0/2.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-TEST-AUTOMESH [Tunnel/300] 1d 02:30:44 Tunnel[edit]** **user@host#** . The policy is configured to use LSP-TEST for routes with community 64723:777, but not for other routes: **set policy-options policy-statement USE-TEST-FOR-COMM term 10 from community**



COMM, set policy-options policy-statement USE-TEST-FOR-COMM term 10 then install-nexthop lsp-regex .\*-TEST-.\*, set policy-options policy-statement USE-TEST-FOR-COMM term 10 then accept, set policy-options policy-statement USE-TEST-FOR-COMM term 20 then install-nexthop except lsp-regex . On testing set policy-options policy-statement USE-TEST-FOR-COMM term 20 then accept set policy-options community COMM members 64723:777, set routing-options forwarding-table export EXPORT-PPLB, and set routing-options forwarding-table export USE-TEST-FOR-COMM. This policy does not work as it is clear from the following output: user@host# run show route community 64723:777 Vrf1.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, \* = Both 40.1.1.1/32 \*[BGP/170] 00:27:49, localpref 100, from 2.2.2.2 AS path: I, validation-state: unverified > to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-1 to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-2 to 10.1.1.6 via ge-0/0/2.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-TEST-AUTOMESH [edit] user@host# run show route table Vrf1.inet.0 Vrf1.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, \* = Both 40.1.1.1/32 \*[BGP/170] 00:00:12, localpref 100, from 2.2.2.2 AS path: I, validation-state: unverified > to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-1 to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-2 to 10.1.1.6 via ge-0/0/2.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-TEST-AUTOMESH 50.1.1.1/32 \*[BGP/170] 00:00:12, localpref 100, from 2.2.2.2 AS path: I, validation-state: unverified > to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-1 to 10.1.1.2 via ge-0/0/1.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-OTHER-AUTOMESH-2 to 10.1.1.6 via ge-0/0/2.0, label-switched-path 2.2.2.2:dt-rsvp-LSP-TEST-AUTOMESH. [PR1313185](#)

- With dynamic tunnels configured, the rpd might crash when the rpd is restarted or Routing Engine switchover is executed. [PR1319386](#)
- The rpd might crash while tracing LSP events when MPLS traceoption is configured. [PR1329459](#)
- After RSVP MPLS label switched path (LSP) link flaps (link goes down and comes back up), RSVP tries to create a second MPLS LSP instance. If resv/pathErr message drops for the second MPLS LSP instance, then the second MPLS LSP instance is stuck, and no further optimizations are possible. [PR1338559](#)
- Whenever there is a decrease in the statistics value across an LSP, the mplsLspInfoAggrOctets value takes two intervals to get updated. [PR1342486](#)
- LDP label is generated for serial interface subnet route unexpectedly. [PR1346541](#)
- On all Junos OS platforms that support MPLS, the LSP might not come up after changing the MPLS admin-group mapping in all nodes of the LSP path, because the LSP configuration is not able to update its admin-group when the global admin-group (under the MPLS hierarchy) is changed. [PR1348208](#)
- The rpd might crash in an RSVP setup-protection scenario. [PR1349036](#)
- If interface flapping occurs on downstream device, some LSPs might get stuck on the upstream devices even if the LSP state is UP. [PR1349157](#)
- In a rare scenario, rpd might crash when LDP fails to allocate self-id for the P2MP FEC. [PR1349224](#)
- Non-deterministic load balancing of Routing Engine generated traffic is observed. [PR1354738](#)

- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. [PR1359956](#)
- With **I2-smart-policy** configured for Label Distribution Protocol (LDP), the L2 circuits might flap if the LDP targeted adjacency also has a link hello adjacency and the interface with the link hello adjacency goes down. [PR1360255](#)
- In MPLS scenario with RSVP-signaled point-to-multipoint (P2MP) LSPs configured, rpd might crash during P2MP LSPs churn. [PR1363408](#)
- In Resource Reservation Protocol (RSVP) scenario, the label-switched path (LSP) might remain UP even if no path is acceptable because of constrained shortest path first (CSPF) failure. There are two scenarios that might result in CSPF failure. Scenario 1 with MBB: Optimization timer fires during make-before-break (MBB). Scenario 2 without MBB: A link/IGP flap causes CSPF, but it depends on timing. [PR1365653](#)
- In BGP labeled-unicast (LU) scenario, when labeled BGP route leaked into Label Distribution Protocol (LDP) using the LDP egress policy and either **set protocols mpls traffic-engineering mpls-forwarding** or **bgp-igp-both-ribs** is configured, after the BGP route gets deleted in one routing table (either inet.3 or inet.0), the LDP might spin to allocate and deallocate labels until it run out of labels. This causes rpd to crash. [PR1366920](#)
- When RSVP link or node protection is deployed and RSVP authentication is used, if the point of local repair (PLR) router and the merge point (MP) router run different versions of Junos OS software during local repair, that is, one earlier release of Junos OS 16.1 release and the other Junos OS 16.1 later release. The RSVP authentication errors might occur for the bypass MPLS LSP and cause traffic loss. [PR1370182](#)
- Enhance MPLS LDP traceroute process to accommodate devices that do not support RFC6424 - LSP ping with TLV 20, DDMT. [PR1372924](#)
- When there is more than one RSVP LSP toward the same downstream neighbor and more than one such downstream neighbor exists, if one of the interfaces toward one downstream neighbor is brought down, the weight might become unequal for ECMP and then the traffic might not be load-balanced equally. [PR1373575](#)
- Applying RSVP traceoptions with nsr-synchronization flag or all flag on an NSR-enabled device might cause the rpd process to crash because of memory corruption. The memory corruption occurs when the size of the received RSVP path message being replicated from the master Routing Engine to the standby Routing Engine is greater than 768 characters. [PR1376354](#)
- On Junos OS, receipt of a specifically crafted malicious MPLS packet leads to a Junos kernel crash (CVE-2018-0049). Refer to <https://kb.juniper.net/JSA10883> for more information. [PR1380862](#)

### **Multicast**

- When DHCPv6 relay in a non-default routing instance is configured and there is no server group, packets in UDP might be dropped because the route table index matching check fails. This is a rare configuration. [PR1316210](#)
- Multicast traffic is not forwarded on the newly added P2MP branch/receiver. [PR1317542](#)

- This issue occurs on platforms running Junos OS enabled with Protocol Independent Multicast(PIM) sparse mode and Internet Group Management Protocol(IGMP). In this scenario, the upstream PIM code is connected through two interfaces, A and B. A is the Reverse Path Forwarding (RPF) interface with the multicast traffic flowing through, and B is the non-RPF interface (for example, a Layer 2 integrated routing and bridging interface). If there is some network change that leads to all PIM status cleanup, and causes the multicast traffic starting to flow through B instead of A, some IGMP groups might still have upstream interface A because the discard route is incorrectly installed in the PIM and could not get timed out. [PR1337591](#)

### **Network Management and Monitoring**

- The **show arp no-resolve interface X** for inexistent interface X is showing all unrelated static ARP entries. [PR1299619](#)
- After SNMP configuration activation 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 might crash when SNMP polling on interface mibs and meanwhile FPC restarts or interface flaps. [PR1318302](#)
- The jnxDomLaneAlarmSet trap is sent with an empty interface description. [PR1318913](#)
- SNMP stops or becomes very slow after a very long period of time. [PR1328455](#)
- With **interafce-mib** command, MX Series router is responding with **type : NoSuchInstance** for OIDs when multiple OIDs are polled in one SNMPGET request. [PR1329749](#)
- jnxDcuStatsEntry and jnxScuStatsEntry OIDs are missing post interface configuration change. [PR1354060](#)

### **Platform and Infrastructure**

- The forwarding-class-accounting enhanced feature is not supported in combination with **forwarding-options hyper-mode**. Using both features together results in traffic being silently discarded or dropped. [PR1198021](#)
- Configuration changes under logical system with LSYS user does not take effect on single commit with fast synchronize enabled. [PR1265139](#)
- On MX Series routers, if a large number of routes are processed, then the Packet Forwarding Engine of the MS-MPC might crash. [PR1277264](#)
- Even though multicast appears to be active with the **show multicast route extensive** command, it is not forwarded to the subscriber interface. [PR1277744](#)
- EVPN-VXLAN traffic gets dropped as "Incorrect vxlan fw path executed" due to a sampling configuration on the core interface. [PR1280539](#)
- With MX Series-based MPC, if the IRB index gets an invalid value and the IRB interface is deleted or any configuration change is made for this IRB interface, an MPC crash might be seen. [PR1281107](#)

- In a rare case, error messages might be observed and the IF queue counters will not be incremented on the MPC5E card. [PR1283850](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- Administratively disabling an interface might cause high FPC CPU usage. [PR1285673](#)
- Executing the command of **show services inline ip-reassembly statistics** might cause ukern sheaf memory leak. [PR1285833](#)
- Generate-event time-interval usage now triggers the event only on the actual expiry of the time interval. [PR1286803](#)
- Incorrect load-balancing on the aggregated Ethernet interface might occur if traffic goes from MS-DPC to MPC in enhanced-IP mode. [PR1287086](#)
- The output values of command **show system resource-monitor** are not accurate. [PR1287592](#)
- There might be memory leak on MX Series with MPCs and MICs-based card if the next hop address that is defined in the next-hop-group is reachable through multiple interfaces. [PR1287870](#)
- On Junos OS, the unauthenticated remote root access is possible when RSH service is enabled (CVE-2018-0052). Refer to <https://kb.juniper.net/JSA10886> for more information. [PR1288932](#)
- The mgd process throws an error: **Couldn't open library: /usr/lib/render/libvccpd-render.tlv**. [PR1289158](#)
- When a non-root user accesses the device through a CLI session, issuing the **load replace terminal** CLI command and attempting to replace the interface stanza in the same operation might terminate the current CLI session and and cause the user session to hang. [PR1293587](#)
- The scale-subscriber license might leak on the backup Routing Engine during bulk subscriber logout. [PR1294104](#)
- The mgd process generates a core file after GRES in a subscriber environment. [PR1298205](#)
- **RMOPD\_HW\_TIMESTAMP\_INVALID** is reported 2 to 4 times a day, which raises an alarm when polled through jnxRpmResSumPercentLost MIB. [PR1300049](#)
- Packet corruption with EVPN MPLS double label push with 3 or more ieee 802.1Q VLAN tags. [PR1300211](#)
- Traffic might be dropped in egress Packet Forwarding Engine because of hashing mismatch. [PR1300789](#)
- Packet Forwarding Engine might crash after MPC reset in firewall filter scenario. [PR1300990](#)
- All traffic can be Tail-/RED-dropped on some interfaces when **chassis fpc max-queues** is configured. [PR1301717](#)
- Classifier does not get applied on the aggregated Ethernet member links on DPC (I-chip) based platforms with CoS configured. [PR1301723](#)
- Logs such as **cassis\_alloc\_index\_pool\_create: SVC NH 0x00b00000[0] poolsize 0x000fffc0 is not a multiple of blk\_sz 0x00001000** are seen. The logs are cosmetic and have no service impact. [PR1301924](#)
- MX Series FPC wedges when creating more than 4000 logical tunnel interfaces per Packet Forwarding Engine. [PR1302075](#)

- The interface-mac-limit might fail for aggregated Ethernet interface. [PR1303293](#)
- On an affected FPC type, when traffic is passed through the optimized loopback path (for example, using lt- interfaces) with packet sizes more than 512 bytes, the forwarding through the Packet Forwarding Engine might stall or you might notice performance degradation. The following syslog entry will be reported: **MQSS(0): LI-1: Received a parcel with more than 512B accompanying data**. The MPC that reports this syslog error message needs to be restarted to recover from this condition. The problem is applicable to MX204, MX480, MX960, MX1003, MX2010, and MX2020 using MPC7, MPC8, and MPC9E line cards (MPC1, MPC2, MPC3, MPC4, MPC5, and MPC6 are not affected). Remaining MX Series platforms such as MX5, MX40, MX80, and MX104 are not affected by this issue. [PR1303529](#)
- The TWAMP Request-TW-Session message's Type-P Descriptor format is not RFC-compliant. [PR1305752](#)
- The message **system reaching processes ceiling low watermark** might be seen. [PR1305964](#)
- On MX Series routers with MPCs or MICs, the resource monitor (RSMON) thread might be stuck in a loop, consuming 100 percent of FPC CPU. [PR1305994](#)
- Service cookie opaque data reset incorrectly leading data sent to service pic getting corrupted. [PR1310904](#)
- The MPC1 or MPC2 might crash because of CPU hogging after the chip fails to initialize. [PR1312286](#)
- ICMP error messages are seen in Packet Forwarding Engine and is forwarded to the correct pic in the AMS bundle. [PR1313668](#)
- The statement **rate-limit** configured with small temporal buffer size might cause packet loss. [PR1317385](#)
- Multicast traffic might get duplicated when MoFRR is configured. [PR1318129](#)
- Move XQ\_CMERROR\_XR\_CORRECTABLE\_ECC\_ERR to minor and re-classify remaining XQCHIP CMERROR from FATAL to MAJOR. [PR1320585](#)
- Errors might be observed when "fabric-header-crc-enable" feature is enabled. [PR1320874](#)
- The traffic with more than 2 VLAN tags might be incorrectly rewritten and sent out. [PR1321122](#)
- On MX104 platforms, when sdk-vmmd fails to correctly identify the current platform, the syslog message **is\_platform\_rainier: Platform could not be detected** appears in error severity. However, for the MX104 this behavior is expected, so this should not be in error severity. Hence, reducing log severity from error severity to debug severity. [PR1321622](#)
- The 'no-propagate-ttl' might not take effect if **chained-composite-next-hop ingress l3vpn extended-space** is configured. [PR1323160](#)
- The MAC might not be learned on MX Series MPCs or MICs-based card because of the negative value of the bridge MAC table limit counter. [PR1327723](#)
- The packet might get dropped in LSR if MPLS pseudowire payload does not have control word and its destination MAC starts with '4'. [PR1327724](#)
- Traffic loss might be observed on LT interface. [PR1328371](#)
- Directories and files under **/var/db/scripts** lose 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 tcpdump filter might not work in egress direction on ps and lt logical interfaces. [PR1329665](#)
- A denial-of-service vulnerability in the telnetd service on Junos OS allows remote unauthenticated users to cause high CPU usage, which might affect system performance. [PR1331234](#)
- Router hits database prompt at **netisr\_process\_workstream\_proto**. [PR1332153](#)
- RPM mib pingResultsMinRtt, pingResultsMaxRtt, pingResultsAverageRtt response as "1" while target address is unreachable, should be "0". [PR1333320](#)
- On all Junos OS platforms, execution of Python scripts through enhanced automation does not work on veriexec images. [PR1334425](#)
- Traffic loss might be seen for some flows because of network churn. [PR1335302](#)
- Commit might fail with error reading from commit script handler, **error: commit script failure**. [PR1335349](#)
- On MX104 routers running with dual Routing Engines with GRES enabled, when **family inet6** is configured on the fxp0 interface and you configure **set system management-instance**, a kernel crash occurs. This issue is seen only on the MX104 and not on the modular MX Series routers. [PR1335903](#)
- The MPC might crash after setting max-queues to a very large number. [PR1338845](#)
- Route corruption in Packet Forwarding Engine with CFM enabled on aggregated Ethernet. [PR1338854](#)
- Configuring the same DHCP server in different routing instances is not supported in DHCP relay scenario. [PR1342019](#)
- In a Virtual Router Redundancy Protocol (VRRP) scenario, the backup router resolves the destination to the VRRP virtual media access control address (VMAC), which resides on the master router. When the backup transitions to master it has to own the VMAC now. In this scenario, the kernel is deleting the earlier next-hop entries, which is the VMAC because of the proxy ARP, to program the Packet Forwarding Engine according to the latest VRRP transition. If any user route points to this next hop, it ends up being a route with a dead next hop, which leads to traffic loss to that destination. [PR1342707](#)
- Route corruption in Packet Forwarding Engine with connectivity-fault-management is enabled for Layer 2 circuit. [PR1342881](#)
- Multiple vulnerabilities in NTP have been resolved in Junos OS. Refer <https://kb.juniper.net/JSA10898> for more information. [PR1343195](#)
- When the aggregated Ethernet interface and the child interface are in a configuration group that is applied through "apply-group", if you execute the command **set/delete interface xxx disable** from a non-group, the interface might still stay in down physical state. [PR1343317](#)
- On MX Series routers with MPC5, MPC6, MPC2E-NG, and MPC3E-NG, if the third or fourth logical tunnel (LT) interface is configured (for example, lt-x/2/x or lt-x/3/x) the queuing logic of those LT interfaces will not work properly and therefore packet drop might be seen on them. [PR1345727](#)
- Multiple vulnerabilities in cURL and libcurl have been resolved in Junos OS. Refer to <https://kb.juniper.net/JSA10874> for more information. [PR1347361](#)

- The IPv4 GPRS traffic over aggregated Ethernet interface might be dropped if **gtp-tunnel-endpoint-identifier** is configured. [PR1347435](#)
- FPC CPU utilization with LT interfaces is pegged continuously at 100 percent. [PR1348840](#)
- Running RSI through console port might cause system crash and reboot. [PR1349332](#)
- ICMP error messages are not generated if 'don't fragment' packets exceed the MTU of the multiservice interface. [PR1349503](#)
- The FPC might crash because of the memory leak caused by the VTEP traffic. [PR1356279](#)
- Traffic black hole seen along with JPRDS\_NH:jprds\_nh\_alloc(),651: JNH[0] failed to grab new region for next hop messages. [PR1357707](#)
- A vulnerability in the IP next-hop index database in Junos OS Release 17.3R3 might allow a flood of ARP requests, sent to the management interface, to exhaust the private internal routing interfaces (IRIs) next-hop limit. Once the IRI next-hop database is full, no further next hops can be learned and existing entries cannot be cleared, leading to a sustained denial-of-service (DoS) condition. [PR1360039](#)
- Certain CLI functions are not triggering properly, because of the missing libraries on the router. Affected commands include **set security ssh-known-hosts load-key-file** and **set system master-password**. [PR1363475](#)
- The error **Disconnected after ISSU and before switchover** might be seen and FPC is restarted during a unified ISSU. [PR1364514](#)
- If you try to configure the same VLAN ID on multiple logical interfaces of the same GR interface, commit will fail with error **two IFLs cannot have same vlan-id**. [PR1365640](#)
- On MX Series routers with MPC1, MPC1E, MPC2, and MPC2E, subscribers over aggregated Ethernet interface cannot utilize their bandwidth, because packets larger than 1500 are dropped. [PR1368414](#)
- Forwarding broken after adding EVPN extended-vlan-id protocol. [PR1368802](#)
- If a tunnel interface is anchored on an MX Series-based 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 platforms, after a unified ISSU from Junos OS Release 14.2 to Junos OS Release 16.1, traffic drops on newly added interfaces because of a unified ISSU hardware synchronize phase issue. [PR1371373](#)
- On MX Series routers with multicast-only fast reroute (MoFRR) enabled, any change that results in creation of a new RPF next hop might also result in JNH memory leak. [PR1373631](#)
- When the scaling IFLset members and aggregated Ethernet members are configured on the same FPC, the FPC might crash when it restarts. [PR1380527](#)
- The rpd might crash after issuing the operational command **show route detail** for RIP route. [PR1386873](#)



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

- When the policy condition configurations are used in export policy in BGP add-path scenario, the condition-based policy fails to take action even though condition is matched. [PR1300989](#)
- The rpd might crash if **vrf-target auto** is configured under routing instance. [PR1301721](#)
- The policy configuration might not be evaluated if policy expression is changed. [PR1317132](#)
- Access-internal route might fail to be leaked between routing instances when "from instance" is configured in the policy. [PR1339689](#)

### ***Routing Protocols***

- The **show bgp summary** shows incorrect results while assisting GR. [PR1045151](#)
- The rpd might crash when running rpd for a long time. [PR1092009](#)
- With multipath and the AS-PATH-IGNORE option enabled under BGP, either on global or routing instance, the multipath feature does not work. [PR1163945](#)
- 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](#)
- The rpd might crash after deactivating or activating BGP. [PR1272202](#)
- Few bfd sessions flap while coming up after FPC restarts or reboots. [PR1274941](#)
- After bfdd restart issue with next-generation MVPN and L2VPN route exchange causing MVPN and VPLS traffic drop. [PR1278153](#)
- With NSR enabled, rpd might crash on master Routing Engine during kernel-id change. [PR1278741](#)
- Under some extremely rare condition, OSPF neighbors might not come up if PIM is also configured in the same routing instance. PIM might install a multicast forwarding entry that prevents OSPF from receiving hello protocol messages. This happens in a rare situation when router control plane is under extreme load just after enabling OSPF and PIM. [PR1279682](#)
- Routing loops might be seen after configuring BGP prefix independent convergence (BGP PIC). [PR1282520](#)
- BGP updates might not be advertised to peers completely in certain condition. [PR1282531](#)
- Some BGP-related traceoptions flag settings will not be effective immediately after the configuration commit, until the BGP sessions are flapped. [PR1285890](#)
- With BGP traceoption enabled, executing the rollback and load merge commands for the configuration might cause rpd to crash. [PR1288558](#)
- BGP-RR sends full route updates to its RR-Clients when any **family mpls interface** gets bounced due to any fiber cut or manual events causing high CPU spike. [PR1291079](#)
- the multihop BFD sessions flap continuously. [PR1291340](#)



- BGP monitoring protocol (BMP) might send malformed route monitoring messages. [PR1292848](#)
- Graceful restart helper might lose capabilities during a peer establishment. [PR1293174](#)
- Multicast flow reset might occur on OIF for RPT joined branch when PIM prune comes on another interface. [PR1293900](#)
- The Impd crashes repeatedly when logical-system is configured on the same device. [PR1294166](#)
- Unified ISSU might take more time to complete and the FPC might go offline during unified ISSU reboot. [PR1298259](#)
- The rpd process might crash because of the AS PATH check error that occurs when RIB groups are added first and later the routing instances are added. [PR1298262](#)
- Inline-BFD on IRB will be broken after GRES or NSR switchover, and the anchor FPC subsequent goes offline. [PR1298369](#)
- MSDP sessions might flap when NSR or GRES is enabled. [PR1298609](#)
- Backup rpd crashed because the SNMP index passed from the master is different from the existing SNMP index. [PR1298711](#)
- In Junos OS, the rpd might crash due to malformed BGP UPDATE packet (CVE-2018-0020). [PR1299199](#)
- BGP might send incorrect AS path when alias is enabled and multiple peers are under the BGP group. [PR1300333](#)
- IBGP route damping does not effect on IBGP inet-vpn address family. [PR1301519](#)
- The rpd process might crash with a core file while deleting a multipath route. [PR1302395](#)
- BGP sessions established without SYNC flag. [PR1302426](#)
- Observed mcsnoopd core file 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/libjtask/base/task_scheduler.c:275.` [PR1305239](#)
- The BFD session might flap when querying interface statistics through SNMP or executing show command through a CLI in vMX platform. [PR1305308](#)
- BGP traceoption logs are still written when it is deactivated. [PR1307690](#)
- Junos OS Release 16.2 and later might give the following error: **Request failed: OID not increasing: ospflflpAddress.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](#)
- The BGP session might flap when the connection between the master Routing Engine and the backup Routing Engine keeps flapping with NSR configured. [PR1311224](#)

- The rpd might crash when the neighbor IS-ISv6 router is restarted, causing route churn. [PR1312325](#)
- Unexpected route age refresh might be observed if BGP PIC is configured. [PR1312538](#)
- The IS-IS SPF might be triggered by LSP updates containing changes only in Reservable Bandwidth in TE extension. [PR1313147](#)
- The rpd might crash if RIP neighbor is configured with the local interface IP address. [PR1313712](#)
- The rpd might crash and generate core files with distributed IGMP. [PR1314679](#)
- BGP prefixes with three levels of recursion for resolution will get stuck with a stale next hop at the first level after a link down event. [PR1314882](#)
- The rpd might constantly consume high CPU in BGP setup. [PR1315066](#)
- On a chassis with BMP configured, the rpd might crash when the rpd process is gracefully terminated. [PR1315798](#)
- With an incorrect IP address that is duplicated with the existing address on the common subnet configured, OSPF has the issue to form adjacency, which is expected. After removing the incorrect configuration, OSPF neighbors can form adjacency (full state) and all the database can be received, but the OSPF routes cannot be installed to the routing table, and the corresponding traffic cannot be forwarded, until the lsa-refresh timer expires. [PR1316348](#)
- The primary path of MPLS LSP might switch to other address. [PR1316861](#)
- Isdb entry cleanup might cause rpd crash, if loop free alternative is configured. [PR1317023](#)
- The inactive route cannot be installed in multipath next hop after disabling and enabling the next hop interface in L3VPN scenario. [PR1317623](#)
- With a specific configuration for BGP and Label Distribution Protocol (LDP), if you make some changes to the active LDP route, the MPLS labels next hop for IPv4 labeled unicast route might be incorrect, and this could cause traffic to be silently dropped or discarded. [PR1317800](#)
- BGP-LU update oscillates with BGP-PIC. [PR1318093](#)
- Remove syslog message that got added to code unintentionally. [PR1318458](#)
- IS-IS might choose a sub-optimal path after the metric change in ECMP links. [PR1319338](#)
- Traffic might get silently dropped and discarded temporarily when BGP GR is triggered and the direct interface flap. [PR1319631](#)
- The rpd crash is seen when deactivating static route if the next hop interface is type of P2P. [PR1323601](#)
- When a prefix limit is reached, increasing maximum-prefixes does not take effect. [PR1323765](#)
- BGP peer is not established after Routing Engine switchover when a graceful restart and BFD is enabled. [PR1324475](#)
- The mcsnoopd process has continuous 28-bits memory leak after having igmp-snooping enabled. [PR1326410](#)

- Multiple next hops might not be installed for IBGP multipath route after IGP route update. [PR1327904](#)
- The rpd might crash on backup Routing Engine after BGP peer is deleted. [PR1329932](#)
- Manual GRES with MX Series Virtual Chassis results in some packet loss on core facing interfaces. [PR1329986](#)
- On MX Series and vMX Series platforms with BGP **add-path** configured, the conditional routing cannot withdraw all routes. [PR1331615](#)
- LDP route might disappear from inet.3 but still remains in inet.0 when OSPF rLFA and LFA protections are used. [PR1333198](#)
- When primary interface returns, discard next hop remains until BGP LU neighbor is cleared. This only impacts the cloned route (S=0). [PR1333570](#)
- Junos OS Release 15.1 and later, IGMP joins are not processed with **passive allow-receive** statement is configured on IGMP interface. [PR1334913](#)
- BGP sessions get stuck in active state after remote end (Cisco) restart the device. [PR1335319](#)
- On all platforms with Shared Risk Link Group (SRLG) configured, if the SRLG information is in the IS-IS protocol, the rpd might crash. [PR1337849](#)
- Under certain rare conditions, if the remote BGP peer closes the TCP session, the rpd process might crash. [PR1340379](#)
- The rpd daemon might crash due to receipt of crafted BGP notification messages. Refer to <https://kb.juniper.net/JSA10871> for more details. [PR1340689](#)
- The rpd crash might occur when receiving BGP updates. [PR1341336](#)
- Changes to the displayed value of AIGP in **show route ... extensive** command. [PR1342139](#)
- Traffic black-hole might be seen if local device receiving BFD-down. [PR1342328](#)
- The rpd might crash when eBGP neighbor flap. [PR1342481](#)
- If a rib-group refers to a VRF routing table in a BGP Layer 3 VPN environment, the rpd process might crash when the VRF routing instance is deleted or deactivated, but the rib-group is still kept. [PR1343578](#)
- A bfd process memory leak might be observed when a multihop bidirectional forwarding detection (BFD) session for a static route is enabled with multiple qualified next hops. [PR1345041](#)
- The rpd might crash if a route for RPF uses a qualified-next-hop. [PR1348550](#)
- On all platforms with GRES enabled, if you bring up the scaling number of BGP peers, after executing Routing Engine switchover, rpd crash might be seen. [PR1349167](#)
- In a PIM scenario, during the process of upstream interface shifting from one interface to another, if the device receives the PIM prune packet, it might cause multicast traffic to be dropped for a while. [PR1350806](#)
- The routing protocol process (rpd) might crash when BGP route damping and BGP multipath feature are configured. [PR1350941](#)

- Source AS community is not appended to RP (display issue in **show route** detail output). [PR1353210](#)
- Static route configuration is always parsed after commit even if the configuration is unchanged. For a recursive route, the metric2 value will be overwritten by the resolver. This metric2 comparison causes the route change although there is no configuration change. [PR1366940](#)
- While performing unified ISSU in an MX Series Virtual Chassis deployment, the MX Series Virtual Chassis system might clear TCP connections causing BGP peerings to flap. [PR1368805](#)
- In a the penultimate-hop router in BGP labeled unicast (LU ) scenario using penultimate-hop popping (PHP), a link flap causes the next hop of a label received from the egress router to change. Once the link comes back, the penultimate-hop router might fail to install the clone route (S=0) entry for that label. As a result, the traffic is dropped and discarded. [PR1387746](#)

### **Services Applications**

- PCP mappings cannot be manually cleared when a NAT pool is shared between PCP and standard NAT. [PR1284261](#)
- TLVs in ICRQ for **actual-rate-downstream** and **actual-data-rate-upstream** do not reflect PPPoE-IA value. [PR1286583](#)
- The MSPMAND process generates a core file "@\_arena\_mALLOc" seen in backup SDG's MS70. [PR1291664](#)
- The jl2tpd process might crash in short time span after GRES switchover. [PR1295248](#)
- L2TP subscribers might get stuck in terminating state during login. [PR1298175](#)
- There is a continuous generation of \*jl2tpd\_era\_Ins\* log files occurs even though L2TP is not configured. [PR1302270](#)
- LTS clients experience packets drop for large packets due to fragmentation in LTS. [PR1312691](#)
- AVP 145 is not present in IRQ when ANCP DSL-type = 0. [PR1313093](#)
- L2TP Tunnel Tx and Rx bytes count sometimes decrease when subscriber sessions are reduced within the tunnel. [PR1318133](#)
- SNMP MIBs not yielding data related to sp- interfaces. [PR1318339](#)
- The MRU might be changed to 1492 instead of the default 1500 in L2TP scenario. [PR1319252](#)
- PPP NCP active mode cannot be enabled for MLPPP session on MX Series platforms. [PR1319580](#)
- Long route remains in forwarding table after subscriber session goes down. [PR1322197](#)
- L2TP LTS might drop the first "CHAP Success" packet from LNS due to delayed programming of /136 route on Packet Forwarding Engine. [PR1325528](#)
- The jl2tpd might crash if the RADIUS server returns 32 tunnel-server-endpoints. [PR1328792](#)
- Not all CSURQ replied in case the number of sessions addressed in CSURQ is more than about 107. [PR1330150](#)

- The l2tpd might crash when multiple l2tp related commands are executed together. [PR1337406](#)
- The command **show services stateful-firewall flows count** shows incorrect flow count after services configuration change. [PR1338704](#)
- Output of **show interfaces si-x/y/z.xxxxx extensive** CLI command shows incorrect inet/inet6 MTU value for MLPPP subscriber on MX Series L2TP LNSs. [PR1346049](#)
- The bbe-smgd process might crash if there are 65535 L2TP sessions in a single L2TP tunnel. [PR1346715](#)
- Session limit per tunnel on LAC does not work as expected. [PR1348589](#)
- 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](#)
- JI2tpd process might crash shortly after one of L2TP destinations becomes unavailable. [PR1352716](#)
- If IPsec is configured (even at a low scale of 200 tunnels) with dead peer detection (DPD) enabled, and all the IPsec tunnels are IDLE, when SNMP walk is performed, IPsec tunnels might flap. [PR1353240](#)
- In subscriber management environment where LNS is deployed, if the "local-gateway" of an L2TP tunnel on LNS device is frequently changed using the **replace** command, the gateway might not be operational and the tunnel connection request packets sent by the corresponding LAC devices (having "remote-gateway" matching the LNS's gateway) might get discarded on the LNS device. [PR1362542](#)
- A few tunneled PPPoE subscriber stuck in terminating state in corner case. [PR1363194](#)
- On MX Series and M12/M320 platforms, when the L2TP LNS subscriber is brought down, the accounting stop message might not be sent to the RADIUS server. [PR1368840](#)
- In an IPsec VPN scenario, some special peers (for example, Huawei enodeB) might start a new IPsec-VPN IKE session without clearing the old session upon detecting session failure, which results in the old IKE session getting stuck in "Not Matured" state. There is no impact to service, but these sessions might consume too many memory resources. [PR1369340](#)
- On MX Series routers with MS-DPC used for NAT64, if an ICMPv6 Type 2 packet is received, NAT64 translates the source address and destination address in the packet incorrectly. [PR1374255](#)
- The FTP ALG is not supported with twice-nat even when there is an unsupported translation type with FTP ALG, and a core file is seen. Should display a syslog message instead of generating a core file. [PR1383964](#)

### ***Software Installation and Upgrade***

- Junos Selective Upgrade (JSU) package is not activated after a reboot. [PR1298935](#)
- The new versions of Junos OS does not have the tool for accessing aux port - /usr/libexec/interposer. [PR1329843](#)
- When the device is booted into single-user mode (recovery mode), and any change is made to the configuration(for example, setting the root password), the commit fails. [PR1368986](#)

### **Subscriber Access Management**

- Accounting messages are sent with incorrect timestamp to RADIUS. [PR1262892](#)
- A few IPs might get stuck on a Policy And Charging Rules Function (PCRF) router. [PR1302509](#)
- Service interim for DHCP subscriber is not working in 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 rebooting the MX Series router acting as a BNG. [PR1308966](#)
- Subscriber might get stuck in "Init" state when executing CLI command **test aaa xxx**. [PR1311263](#)
- Memory leak might happen after clearing subscriber with script or manually. [PR1312517](#)
- Service interim missing for random users in JSRC scenario. [PR1315207](#)
- The delegated prefix from RADIUS is incorrectly parsed when the prefix is fewer than 20 bytes long. [PR1315557](#)
- Allowing unified ISSU during accounting suspend. [PR1320038](#)
- IP addresses are assigned discontinuously from the linked IP pools. [PR1323829](#)
- multiple-radius-servers having different dynamic-request-port is not supported. [PR1330802](#)
- Subscriber might get stuck in terminated state when JSRC synchronize state get stuck in "FULL-SYNC in progress". [PR1337729](#)
- On MX Series platforms with L2TP service-rate-limiter service deployed, the transmission of Tx and Rx connection speeds from LAC to LNS might not be updated in the L2TP incoming-call-connected (ICCN) packet when the LAC receives the **Access-Accept** message from the RADIUS server that provides the specific configuration. [PR1338786](#)
- On Junos OS, authd allows assignment of IP address requested by a DHCP subscriber logging in with Option 50 (Requested IP Address) (CVE-2018-0057). Refer to <https://kb.juniper.net/JSA10892> for more information. [PR1351334](#)
- In 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 generates core files. These types of core files are the result of a rare timing deadlock between the SDB secondary table and STS hash lock. [PR1358339](#)
- In a dual Routing Engine system with the enhanced subscriber management feature enabled, if GRES is not configured, the authd process might not be started after executing Routing Engine switchover on the backup Routing Engine. [PR1368067](#)
- Address pool does not correctly cycle to the beginning of the pool when linked-pool-aggregation parameter is defined. Address pool reports **Out of Addresses** even though not all addresses are in use. [PR1374295](#)

- When the RADIUS server sends CoA for the subscriber after the RADIUS server has returned a different dynamic-profile name in access-accept, the subscriber will be updated with the original dynamic-profile. The issue occurs because the new dynamic-profile name sent by the RADIUS is not saved in the subscriber's table. Hence, when the CoA message arrives, the old dynamic-profile name is used. The issue results in CoA updating the subscriber with unexpected values. (The old dynamic-profile instead of the new dynamic-profile is used). [PR1381230](#)
- In a dual-stack PPP/PPPoE-based subscriber scenario, when V4+V6 service is installed with family v4, if some daemon (such as dfwd) fails to add family inet6 IFF during instantiation of the family inet6 portion of some services (such as SRL service), family activation for family inet6 fails. But only the family inet6 portion of service should be removed. The family inet and L2 services such as CoS should remain unchanged, but they are changed. Therefore, some subscribers cannot get some services (such as SRL service) even though the RADIUS messages can be sent and received. It is a timing-specific issue. [PR1381383](#)
- If the default value for the `$junos-routing-instance` predefined variable is configured (that is, `dynamic-profiles <> predefined-variable-defaults routing-instances <>`), the subscriber will come up in the configured default routing-instance even if RADIUS has already supplied the VSA of '26-1 Virtual-Router'. [PR1382074](#)
- When a subscriber is manually logged out using CLI `clear network-access aaa subscriber username <test>` the following log message gets printed (messages file) when the GX-Plus module is clearing/freeing up the subscriber session-id from its table: `Aug 28 12:11:50 jtac-test-node: authd [XXXX]: %DAEMON-3: gx-plus: logout: wrong state for request session-id: <XXX>`. [PR1384599](#)

### *User Interface and Configuration*

- CLI session might die while issuing the command `show configuration | compare rollback 1`. [PR1331716](#)

### *VPNs*

- Next generation MVPN IPv6 RP bootstrap type 3 S-PMSI AD route prefix ff02::d persists after BSR data stop. [PR1269234](#)
- 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 flap repeatedly. [PR1282875](#)
- L2circuits stitched through LT peer interfaces might get stuck in "LD" (local site signaled down) status. [PR1305873](#)
- Un-hide `set protocols pim mvpn family inet6 disable` configuration to allow users to disable inet6 on MVPN. [PR1317767](#)
- While doing Routing Engine switchover in NSR, the deletion of LDP label-related entries on the standby Routing Engine might be not handled correctly. This issue can trigger an rpd crash on standby Routing Engine. [PR1310934](#)
- The rpd might crash 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 rpd crash. [PR1325867](#)
- MVPN sender-site configuration is not allowed with S-PMSI. [PR1328052](#)
- The routing protocol process (rpd) generates a core file on backup Routing Engine with next generation MPVPN and NSR configuration. [PR1328246](#)
- The routing protocol process (rpd) crashes because of the receipt of a specific Draft-Rosen MVPN control packet in a Draft-Rosen MVPN configuration, and restart or might lead to remote code execution. Refer to <https://kb.juniper.net/JSA10879> for more information. [PR1339567](#)
- The 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 L2circuit or VPLS backup neighbor. [PR1340474](#)
- The rpd might crash on backup Routing Engine when changing the L2circuit virtual-circuit-id in an NSR scenario. [PR1345949](#)
- In an L2VPN scenario, rpd might crash if an interface (that is already in downstate) is added to any operating site of an L2VPN instance. [PR1351386](#)
- In a dual-homed next-generation MVPN scenario with spt-only mode enabled, the receipt of type 5 withdrawal removes the downstream join states for some routes when multiple type 5 routes exist and one of them is withdrawn in some cases (such as PE device uplink failure). [PR1368788](#)

## Resolved Issues: 17.2R2

### *Class of Service (CoS)*

- The Routing Engine level **scheduler-hierarchy** command misses a forwarding class when the **per-unit-scheduler** mode is configured. [PR1281523](#)

### *Forwarding and Sampling*

- Aggregated Ethernet interface might move to "down" state after GRES. [PR1233188](#)
- Packet Forwarding Engine mac-learning debug logs are displayed as error logs. [PR1267684](#)
- Unexpected error messages might be seen in logs. [PR1270686](#)
- The sampled process stops collecting data on Routing Engine based sampling supported platforms. [PR1270723](#)
- Firewall filter might not be matched when wildcard (\*.\*) is specified as matching condition. [PR1274507](#)
- Routing-instances information is not being displayed in the flat accounting file. [PR1275225](#)
- Unicast traffic is forwarded out of the logical interface even after the interface is disabled. [PR1277697](#)
- The sampled route reflector (srrd) process might crash in the large routes churn situation. [PR1284918](#)
- The sampled process might crash if traceoptions are enabled. [PR1289530](#)



## General Routing

- ICMP reply traffic might get dropped on MS-MPC line cards. [PR1059940](#)
- With l2tp subscribers, after every subscriber's login attempt, all FPCs except the card that hosts subscribers might report the following log message **jnh\_if\_get\_input\_feature\_list(9723): Could not find ifl state**. [PR1140527](#)
- The FPC might reboot and the error message **Readback error from I2C slave** might be displayed. [PR1174001](#)
- Port block efficiency and unique pool users statistics show incorrect values respectively in the NAT pool, which is being used by the sessions. This issue occurs when adding an address into the NAT pool. Both NAT pools are used in the same service set. [PR1177244](#)
- The CLI command **request vmhost zeroize** or **request vmhost zeroize both** might work only on the local Routing Engine. [PR1197152](#)
- The rpd might crash in the backup Routing Engine after a Routing Engine switchover in an MX Series subscriber environment. [PR1206804](#)
- IPsec phase2 soft lifetime calculation is different between Junos OS Release 11.4R12 and Junos OS Release 14.2R6. [PR1209883](#)
- Continuous error messages **pdb\_open failure** for Routing Engine scope MQTT broker are observed. [PR1224705](#)
- CoS service with reflexive cos-rule should modify CoS values for reverse flow. [PR1227021](#)
- MPC2E-NG and MPC3E-NG generate a core file with specific MIC because of tight loop of PIC Express critical exceptions. [PR1231167](#)
- Major errors related to XQ-chip L4NP parity errors might be reported on MPC. [PR1232952](#)
- With vLNS (vBNG), a commit generates the message **warning: requires 'l2tp-inline-lns' license** even if a valid license is installed. [PR1235697](#)
- Junos Telemetry Interface: Frequent disconnects are seen with the MQTT messaging protocol when the logical interface sensor is provisioned for a longer duration. [PR1238803](#)
- MPC9E might generate an FPC core file when running Junos OS Release 16.1R2.11 if it is configured with "mixed-rate AE bundles" and "adaptive load balancing". [PR1238964](#)
- Half of the Point-to-Point Protocol over Ethernet (PPPoE) subscribers experience keepalive failure on PICs with aggregated Ethernet anchors. [PR1240365](#)
- ANCP neighbors might stream down after commit. [PR1243164](#)
- XM chip-based line card might drop traffic under high temperature. [PR1244375](#)
- A route target per bridge domain for EVPN is not supported. [PR1244956](#)
- Sensors are not reused when the subscriptions have uncommon paths. [PR1245902](#)
- RADIUS accounting statistics of subscribers are doubled after unified ISSU. [PR1250919](#)

- On MX2000 MPC6E, EOAM LFM adjacency flaps when an unrelated MIC accommodated in the same MPC6E slot is online. [PR1253102](#)
- Na-grpcd might crash if openconfig is used for telemetry interface. [PR1254794](#)
- Device control process (dcd) crashes during the ATM-related configuration commit. [PR1258744](#)
- The syslog message **HEAP: Free at interrupt level /Free interrupt violation!** is displayed when interface drops on TRI-RATE SFP-T on MIC-3D-20GE-SFP-E. [PR1259757](#)
- Incorrect egress classification of L3 multicast traffic from ingress VLAN bridge interface after a configuration change. [PR1260413](#)
- Layer 2 control BUS timeout causes SFP thread hogging and an MPC restart. [PR1260517](#)
- On an MX Series platform with an MPC line card, an MPC line card goes offline during a unified ISSU. [PR1260714](#)
- Point-to-Point Protocol over Ethernet (PPPoE) subscribers might not come up while verifying that IPCP renegotiation happens properly for terminated PPPoE subscribers. [PR1260836](#)
- With QSFP optics, Rx loss cleared and set critical messages are logged continuously. [PR1261793](#)
- Extra link transitions might be seen after restarting MPC. [PR1264039](#)
- BGP hold time might be expired after a GRES or NSR switchover. [PR1264436](#)
- Sometimes SDN-Telemetry subsystem does not respond to management requests while issuing **show agent sensors**. [PR1266058](#)
- Unified ISSU related limitation is observed under highly scaled scenarios. [PR1267680](#)
- The openflowd process might get stuck because of 100% CPU memory corruption while deleting and querying the filter. [PR1268527](#)
- The command **show arp interface xe-x/x/x no-resolve | display xml** returns XNM errors. [PR1269170](#)
- MIC error interrupts are more than the threshold (> 2500 per 5 min), so the MIC or FPC is restarted. As a result, MIC error interrupts will hog the CPU when the restart is initiated. [PR1270420](#)
- The multicast blackhole might be seen when the aggregated Ethernet interface flaps with MoFRR enabled. [PR1270939](#)
- When MX Series routers are equipped with a next-generation Routing Engine, the log message **sdk-vmmd: %USER-3: is\_platform\_Next-Gen RE: Platform found as Next-Gen RE** is displayed with error severity. [PR1271134](#)
- The Routing Engine might stop all services after GRES or ISSU. [PR1271306](#)
- Packet Forwarding Engine drops BUM traffic coming from remote PE EVPN instance. [PR1272384](#)
- Virtual forwarding plane failed to load files from virtual control plane if the interconnection has an MTU less than 1500. [PR1273365](#)
- The mspm and log messages about memory zone level are generated incorrectly. [PR1273901](#)

- The l2ald process might crash in an EVPN scenario. [PR1274113](#)
- L2-over-GRE tunnel might use underlying physical interface MTU directly without deducting IP/GRE header length. [PR1274203](#)
- CLI commands fail to execute **show subscribers detail**, **show subscribers extensive**, **show subscribers count client-type <>** and other commands as subscriber management database is unavailable. [PR1274464](#)
- FPC/MPC might crash in EVPN/MPLS or EVPN/VXLAN environment. [PR1274976](#)
- FPC generates a core file when route record with an unknown AS index is received. [PR1275021](#)
- Link stays down after a flap on MPC NG cards with QSFP+-40G direct attach copper (DAC). [PR1275446](#)
- Fixed the default behavior of the configuration statement added for static route's dependency on BFD\_ADMIN\_DOWN, through PR 1070477. [PR1275973](#)
- Routing Engine based captive-portal-content-delivery (CPCD) does not work in vMX or MX86. [PR1276016](#)
- For MPC7E-10G, MPC7E-MRATE, MX2K-MPC8E, and MX2K-MPC9E complete traffic loss is observed when CRC errors are injected on a single plane. [PR1276301](#)
- Junos OS does not use the complete TCP window size and slows the connection when JET application over GRPC is installed on Junos OS. [PR1276443](#)
- On an MX Series platform with MS-MPC or MS-MIC installed, the service PIC daemon (spd) memory leak might be observed after adding or removing a service-set statement. [PR1276809](#)
- Layer 2 control BUS stuck causes SFP+ thread hogging and restarting of MPC. [PR1277467](#)
- MTU configuration option for virtual tunnel interfaces will be removed. [PR1277600](#)
- IS-IS adjacencies over MLPPP links do not connect to the LSQ bundle interface. [PR1278377](#)
- The routing protocol process (rpd) might get stuck 100% when the same BGP prefix routes are learned in different routing instances with multipath and auto-export configured. [PR1279260](#)
- VLAN out-of-band subscriber session fails when it is autoconfigured. This is because the physical interface goes down even if it is physically up. [PR1279612](#)
- When an MS-MPC-PIC is brought offline or online or bounced (because of an AMS configuration change), occasionally, PIC can take approximately 400 seconds to initiate. [PR1280336](#)
- Authenticated subscriber dynamic VLAN interface might get disconnected immediately after successful connection. [PR1280990](#)
- MTU for a Layer 2 over GRE gr- interface should be unlimited. [PR1281173](#)
- The ingress service-accounting-deferred for L2BSA subscribers are not providing the correct IP traffic statistics. [PR1281201](#)
- Establishment of IPsec SAs for link type tunnels might fail under certain conditions. [PR1281223](#)
- DHCP/PPPoE subscribers fail to bind after FPC restart and smgd restart with BBE\_RTsock\_GET\_RTsock\_IFL\_FAIL\_TERMINATED counter going up. [PR1281930](#)

- Optics levels are not sent in Junos Telemetry interface for down interfaces. [PR1281943](#)
- Buffer overflow in sockets library (CVE-2017-2344). [PR1282562](#)
- Inline J-Flow unrelated configuration changes related to a routing-instance results in invalid or incomplete J-Flow data packets. Commit-full resumes proper functionality. [PR1282580](#)
- Variable based flows (VBF) are not programed appropriately on aggregated Ethernet interfaces. [PR1282999](#)
- OAM fails to come up when GRE tunnel source and family inet address are the same. [PR1283646](#)
- When the service-set has both NAT rule and Stateful-Firewall rule configured but a source IP address could not be matched with any NAT rule, but could be matched with Stateful-Firewall rule, the PPTP session from this source IP address might not be able to be established successfully. [PR1285207](#)
- The J-Flow data template sequence number is zero for MPLS flows. [PR1285975](#)
- Unified ISSU is not supported from Junos OS Release 15.1 or later when source release includes one or more BBE features such as logical interface options, CoS fragmentation map, MLPPP, advisory options, advanced services, and multicast distribution. [PR1286507](#)
- The routing protocol process (rpd) crashes during subscriber login or logout with multicast service enabled and while performing GRES switchover. [PR1286653](#)
- A10NSP interface is not getting attached to the Layer 2 routing instance after renaming the routing instance name. [PR1287070](#)
- The routing protocol process (rpd) might generate a core file after changing the **routing-options dynamic-tunnels** configuration. [PR1287109](#)
- LTS functionality does not work on Junos OS Release 16.1R4-S2 if rewrite-rule configuration is applied to the dynamic profile. [PR1287788](#)
- SNMP query for IF-MIB::ifOutQLen reports incorrect type should be Gauge32 or Unsigned32 for a dynamic VLAN demux0 interface. [PR1287852](#)
- The services-oids-ev-policy.slax & services-oids.slax files built in the Junos OS image does not have the latest versions. [PR1287894](#)
- The bbe-smgd process might crash generating a core file on standby Routing Engine during a reboot upgrade with active locally terminated PPPoE subscribers. [PR1288121](#)
- The smg-service process might generate a core file in the backup with a distributed IGMP configuration. [PR1288465](#)
- Kernel "rtdata" memory might leak on an MX Series Virtual Chassis with heartbeat enabled. [PR1289363](#)
- The FPC memory leak might happen in a BBE subscriber environment. [PR1289365](#)
- Memory leak is observed in a bbe-smgd process (daemon) when the subscriber logs out of the multicast group. [PR1290918](#)
- BBE-SMGD generates a core file following a stress test in bbe\_iff\_add\_ifa. [PR1291969](#)

- An error in **vbf\_filter\_add\_orphan\_check** might be seen when the subscribers use filter to log out or log in. [PR1292582](#)
- The syslog **DDR3 TEMP ALARM** messages are logged in chassisd log. [PR1293543](#)
- Login or logout core file is generated using Routing Engine based http-redirect. [PR1293553](#)
- The **show extensible-subscriber-services sessions** reports an incorrect timestamp increase by one hour after a unified ISSU. [PR1293800](#)
- Unable to edit dynamic profiles after scaling up to 400 dynamic profiles. [PR1295446](#)
- The bbe-smgd process generates a core file at bbe\_mcast\_ifl\_vbf\_encoder on service activation or deactivation along with smg-service restarts. [PR1295938](#)
- Routing Engine crashes generating a core file after a loop in rts\_gencfg\_ifstate\_getparent. [PR1296884](#)
- A memory leak is seen when **set protocols mld XXX** stanza is changed and committed. [PR1297454](#)
- The bbe-smgd process crashes when traceoption is enabled due to an invalid username character. [PR1298667](#)

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

- The vmcore files were generated on both VCMm and VCBm at the same time. [PR1274438](#)

#### **Infrastructure**

- The smartd **Offline uncorrectable sectors** critical log keep reporting every 30 minutes. [PR1233992](#)
- The **show system users** CLI command output displays more users who are not using the router. [PR1247546](#)

#### **Interfaces and Chassis**

- IPv6 Neighbor Discovery does not work for DHCPv6 sessions when using static demux VLAN with router advertisement. [PR1250313](#)
- At a high logical interface scale, an ifinfo process (daemon) generates a core file on executing command **<show-interface>**. [PR1254189](#)
- The MRU of aggregated Ethernet interface might reset to default value. [PR1261423](#)
- When adding an additional Data field in a PPP Echo Request packet, keepalive failure might be seen that might disconnect the subscriber. [PR1273083](#)
- The message dot1agCfmMepHighestPrDefect might be reported in the SNMP trap with a value of -1 instead of 0 on recovery after a remote defect indication (RDI). [PR1273278](#)
- The line card hosting an Ethernet OAM LFM session might reboot during a unified ISSU. [PR1283280](#)
- No L2TP sessions come up on some si- interfaces after an MPC restart followed by a Routing Engine switchover. [PR1290562](#)
- A VRRP track interface down did not trigger a mastership election immediately. [PR1294417](#)

### Layer 2 Ethernet Services

- The **show class-of-service fabric statistics** CLI command might fail with a periodic **Error = Operation timed out** message. [PR1228293](#)
- The IPv4 or IPv6 packets originating from a Routing Engine might be corrupted when the bridge domain has 'vlan-id' set to none, but the outgoing L2 interface for the packet is tagged and CoS is enabled. [PR1263590](#)
- DHCP is not using the configured IRB MAC as the source MAC in DHCP offer unicast replies. [PR1272618](#)
- The messages `l2cpd[2486]: task_connect: task MVRP l2ald ipc./var/run/l2ald_control addr /var/run/l2ald_control: No such file or directory` are filling up the syslog. [PR1278189](#)

### Layer 2 Features

- In a scaling VPLS scenario, convergence time is taking more than 10 minutes. [PR1279192](#)
- A misconfiguration that adds an aggregated Ethernet (AE) bundle and its member link to a VPLS instance might cause 100% routing protocol process (rpd) utilization. [PR1280979](#)

### MPLS

- RSVP p2mp sub-LSPs having more than 1 sub-LSP in down state might not get re-optimized after transit path goes down. [PR1174679](#)
- Traffic loss is seen during an auto-BW make-before-break (MBB) on an ingress router as "invalid fabric token". [PR1264089](#)
- When "explicit-null" is configured for LDP, label 0 is assigned as IPv6 explicit null label. [PR1264753](#)
- The routing protocol process (rpd) might crash if egress-policy is configured in LDP. [PR1266358](#)
- Remote targeted LDP session might remain up even though it should not be up. [PR1266802](#)
- Traffic loss will be observed when primary LSP goes down in an LDP-over-RSVP environment. [PR1270877](#)
- JDI-RCT-RPD `rpd core@ bgp_labeled_l2vpn_standby_outmetrics , bgp_rt_ribout_rcv_nlri`: This core file might be generated for subscribers who have configured BGP family L2VPN in Junos OS Release 17.2R1. [PR1271704](#)
- The CLI command **show route extensive** might cause routing protocol process (rpd) to crash. [PR1272993](#)
- RPD core: Assertion failed `rpd[6255]: file src/junos/usr.sbin/rpd/rsvp/rsvp_enh_lp.c`, line 4928: "rsvp\_enh\_lp\_supported\_psb\_type(psb). [PR1276748](#)
- The routing protocol process (rpd) crashes due to LDP defect during NSR-enabled Routing Engine switchover. [PR1290789](#)
- The routing protocol process (rpd) crashes if MPLS LSP path change occurs. [PR1295817](#)

### **Network Management and Monitoring**

- Command ESC-Q does not work when the syslog is disabled. The syslog message is still seen even if it is disabled by ESC-Q. [PR1269274](#)
- MIB2D related syslog message **MIB2D\_RTSLIB\_READ\_FAILURE: rtslib\_iflm\_snmp\_pointchange** is seen during removing and restoring configurations. [PR1279488](#)
- On Junos OS devices with SNMP enabled, a network-based attacker with unfiltered access to the Routing Engine can cause the Junos OS snmpd process (daemon) to crash and restart by sending a crafted SNMP packet. Repeated crashes of the snmpd daemon can result in a partial denial-of-service condition. Additionally, it may be possible to craft a malicious SNMP packet in a way that can result in remote code execution. [PR1282772](#)
- The Management Information Base II process (mib2d) is logging an "RLIMIT curr 1048576000 max 1048576000" message every time a commit is performed, which might confuse the operator into believing that the memory limit of 1GB has been reached. [PR1286025](#)
- If a logical interface of a loopback interface (lo0) is deleted, it will not be deleted in the ifStack tree. It might result in a mib2d crash when polling the object identifier (OID) of ifStackStatus.0. [PR1286351](#)

### **Platform and Infrastructure**

- Traffic drop might occur under a large-scale of firewall filter configuration. [PR1093275](#)
- Kernel might crash on issuing **show arp** or **clear arp** if there is an IPv4 255.255.255.255 address. [PR1120114](#)
- FPC crashes with MAC accounting feature enabled. [PR1173530](#)
- FPC CPU spikes every 6 minutes on MX Series with an MPC or MIC chipset due to a microcode rebalance. [PR1207532](#)
- With a commit script configured, the mgd process might crash when you configure anything in private configuration mode. [PR1244015](#)
- One of the processes (dcd, rpd, dfwd, pfed, cosd, sampled) might generate a core file in a large-scale 8K ESSM login or logout with an ephemeral database. [PR1249979](#)
- GRE tunnel traffic gets dropped after disabling and reenabling the gr-interface. [PR1255706](#)
- **show ephemeral-configuration** has configuration though there are no active client connections. [PR1260124](#)
- Error message **rn timer\_delete\_nh: no pat-node** might be seen when the subscriber logs out. [PR1263983](#)
- FPC might crash with interface-specific firewall filters with policers configured. [PR1267908](#)
- The routing protocol process (rpd) might crash and BGP session flapping might be seen if flapping interfaces or changing configurations. [PR1269116](#)
- Dropping the TCP RST packet incorrectly on Packet Forwarding Engine might cause a traffic drop. [PR1269202](#)

- FPC generates a core file when you are trying to send igmp-membership reports to 16000 subscribers. [PR1270928](#)
- The queued statistics of interface are not correct in CoS scenario on MX Series platform. [PR1271055](#)
- The real-time performance monitoring (RPM) loss percentage values for "overall tests" through SNMP might be incorrect. This is because the RPM probe loss percentage is stored as a 32-bit integer internally but the calculation can exceed a 32-bit boundary, which might lead to a rounding error. [PR1272566](#)
- Ephemeral database configurations are not getting mirrored to the backup Routing Engine. [PR1279653](#)
- **request routing-engine login other-routing-engine** might require a password. [PR1283430](#)
- Incorrect load-balancing occurs for traffic going from MS-DPC to MPC cards. [PR1287086](#)
- Log messages are getting triggered when any non-superuser or non-root user tries to telnet into the router. `// rend_dlinitt: not a proper library: /usr/lib/render/libdcd-render.so: Cannot open "/usr/lib/render/libdcd-render.so" // .` [PR1289974](#)
- The source MAC learned from cross-Packet Forwarding Engine aggregated Ethernet (AE) might bounce between aggregated Ethernet member Packet Forwarding Engines for a long time and might cause MLP-ADD storm. [PR1290516](#)
- RMOPD might get stuck at sbwait upon receiving a specific response from HTTP agent. [PR1292151](#)
- The Broadband Remote Access Server and carrier grade NAT features running on the same MX Series device might trigger transient flow-control asserted by XLP MAC after upgrading the MX Series routers to Junos OS Release 16.1. [PR1293232](#)

### ***Routing Protocols***

- No multicast forwarding in ASM mode after a unified ISSU. [PR1146621](#)
- The routing protocol process (rpd) might crash on platforms with 64-bit X86 Routing Engine if IPv6 is configured. [PR1224376](#)
- Routing protocol process (rpd) on the backup Routing Engine might restart unexpectedly upon the addition of a new L2VPN routing instance. [PR1233514](#)
- Need support for conflict resolution. At times, the same SID might be sent for multiple prefixes, which might cause issues. [PR1239093](#)
- The routing protocol process (rpd) core file might be seen in an MVPN scenario. [PR1240565](#)
- There might be a stale bootstrap rendezvous point (RP) entry in a bootstrap router RP table after deleting static RP configuration from another router. [PR1241835](#)
- When **advertise-from-main-vpn-tables** configuration statement is used under BGP and the router-reflector functionality is added, a refresh message is not sent resulting in some missing routes. [PR1254066](#)
- BGP-LU label might go into "dead" state in forwarding table after the MPLS address family on the next-hop interface is removed and re-added. [PR1262180](#)
- MPLS over UDP tunnel creation failure in the absence of a VRF table. [PR1270955](#)



- "Nexthop AFI=3" is observed in a BGP open message after configuring **family inet unicast extended-nexthop**. [PR1272807](#)
- The BFD down for BGP might cause traffic black holing for customer traffic. [PR1276497](#)
- Error messages are seen when receiving BGP update messages with UNREACH NLRI. [PR1276758](#)
- IS-IS LSPs might be dropped in interop with Cisco in a segment routing (SR) scenario. [PR1280522](#)
- The routing protocol process (rpd) might crash due to a certain chain of events in BGP-LU protection scenario. [PR1282672](#)
- The second multicast packet might be discarded on rendezvous point router. [PR1282848](#)
- The routing protocol process (rpd) might crash while deactivating in a routing instance [protocols pim static]. [PR1284760](#)
- The routing protocol process (rpd) might crash if dynamic Routing Protocol goes down in ECMP topology and also if PIM **join-load-balance automatic** is configured. [PR1288316](#)
- BGP-RR sends full route updates to its RR-Clients when any family MPLS interface gets bounced because of any fiber cut or manual events causing high CPU spike. [PR1291079](#)
- The routing protocol process (rpd) might crash if BGP flap happens. [PR1295062](#)

### **Services Applications**

- L2TP congestion window set to 128 instead of 1 when tunnel is created. [PR1265001](#)
- DTCP non-optimized trigger attributes can delay mirrored traffic forwarding in scaled environments. [PR1269770](#)
- Kernel crash might be seen after performing the CLI command commit. [PR1273357](#)
- Lawful intercept: ingress control packets from the subscriber are mirrored to the mediation device twice. [PR1275592](#)
- Backup Routing Engine goes to the database prompt with a vmcore if the down ASI interface configuration is deleted. [PR1281882](#)
- Layer 2 Tunneling Protocol (L2TP) subscribers are down after a GRES while verifying framed IPv6 route support for L2TP network server (LNS) at a higher scale with a maximum number of Framed-IPv6-Route. [PR1293783](#)
- Each subscriber session gets its own L2TP tunnel without "Tunnel-Client-Endpoint" from RADIUS. [PR1293927](#)

### **Subscriber Access Management**

- Option to exclude tunnel attributes in access-request on L2TP network server (LNS). [PR1264024](#)
- Possible CPS degradation for scaled DHCPv4 or DHCPv6 and PPPoEv4 subscribers. [PR1264052](#)
- Accounting messages are sent with the wrong Event-Timestamp to RADIUS. [PR1270162](#)

- The DHCP subscriber might not get an IP address when the address pool is tight. [PR1274870](#)
- bbe-smgd might spontaneously crash after bbe-smgd daemon restarts from CLI. [PR1277099](#)
- Some RADIUS attributes might not be filtered out of the accounting-on or accounting-off message on an MX Series platform. [PR1279533](#)
- IP assigned by RADIUS is incorrectly counted by local pool after a Virtual Chassis switchover. [PR1286609](#)
- An authd core file is observed while terminating a large number of subscribers. [PR1289215](#)

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

- commitd might generate a core file by removing certain configuration followed by a commit operation. [PR1267433](#)

### ***VPNs***

- The routing protocol process (rpd) crashes after an L2VPN configuration change followed by "ping mpls l2vpn". [PR1272612](#)
- Memory leak in RPD task\_timer, timer 'PIM MVPN Alt KAT Timer'. [PR1276041](#)

## **Resolved Issues: 17.2R1**

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

- The cosd process might crash when you execute the command **show class-of-service queue-consumption**. [PR1066009](#)

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

- Aggregated Ethernet interface might get into "down" state after GRES. [PR1233188](#)
- For certain subscriber types entry in the statistics database is not cleaned up on logout. [PR1251756](#)
- Accounting interim interval is reset after GRES. [PR1261472](#)
- Service statistics are reported in the wrong order. [PR1262876](#)

### ***General Routing***

- The jsscd might crash in a scaled environment. [PR1133780](#)
- When the traffic matches a rule name with junos:rdp, the LRF record has the PCC rule name any-any. [PR1174938](#)
- On MX Series routers, the MS-MIC line card might crash and restore automatically. [PR1183828](#)
- The CPU of processes might get nearly 100% occupied. When SDN-telemetry (the agentd process) is disabled or continuously restarted, certain messages are repeatedly logged in syslog. The agentd process is unable to accept the new subscriptions. As a result, all subscriptions are dropped, triggering agentd to restart several times. [PR1192366](#)

- Error messages are reported during unified ISSU on MX Series routers. [PR1200045](#)
- The command **show subscribers summary port extensive** outputs might have an incorrect tunneled or terminated sessions count. [PR1206208](#)
- Unified ISSU is not supported on MX2008. [PR1213193](#)
- An MS-MPC or MS-MIC service PIC might crash when passing large fragmented traffic through an ALG. [PR1214134](#)
- Syslog message **fpc\_pic\_process\_pic\_power\_off\_config:[xxxx] :No FPC in slot [y]** is incorrectly displayed on an empty FPC slot with no PIC power off configured. [PR1216126](#)
- MPC might crash during unified ISSU from Junos OS Release 15.1R1 to a later release when QSFP, CXP, or CFP2 optics are present. [PR1216924](#)
- Continuous login and logout of PPPoE/DHCP subscribers might cause some subscribers to fail to bind. [PR1221690](#)
- The MX2008 BITS clock module's LED behavior is inconsistent with other platforms. [PR1222041](#)
- The **early/opDel: bad stored heap** messages seen on sending traffic using captive-portal-content-delivery service do not have any affect on functionality. [PR1226782](#)
- MX2008 chassisd process might consume more CPU cycles than the chassisd process running on MX2010 or MX2020. [PR1231333](#)
- Junos Telemetry Interface: Frequent disconnects are seen in MQTT when the logical interface sensor is provisioned for a longer duration. [PR1238803](#)
- BBE CST MX Series Virtual Chassis: Half of PPPoE subscribers KeepAlive failure on MPC5E line card PIC1, if aggregated Ethernet anchors on PIC1. [PR1240365](#)
- ANCP neighbors go down after a commit. [PR1243164](#)
- The **ms90 kernel: kern.maxfiles limit exceeded by uid 0, please see tuning(7)** message is seen after injecting more than 2M routes. [PR1243581](#)
- Route target per bridge domain for EVPN is not supported. [PR1244956](#)
- Sensors are not reused when the subscriptions have non-common paths. [PR1245902](#)
- GNF console hangs after some idle time. [PR1250726](#)
- The rpd might crash when some interfaces and some peers go down. [PR1250978](#)
- KRT queue gets stuck on the Routing Engine, causing RIB and FIB to go out of synchronization. [PR1251556](#)
- Output of **show ancp subscriber detail** might omit certain TLVs. [PR1252747](#)
- Junos OS Release 17.2DCB: High 1PPS phase-transient is seen on physical layer SyncE rearrangements. [PR1253083](#)

- An interoperability is seen between MX Series MPC3E-NG and MS Series MPC2E-NG line cards when connected to third party switch. [PR1254795](#)
- Incorrect data in the output of **show subscribers extensive** . [PR1255029](#)
- Riot (vPFE) process might generate a core file in vMX platform when a lot of subscribers log in or log out when there are a large number of flows (>500K). [PR1255866](#)
- Traffic drop seen on MPC7E cards after rekeying of MACsec. [PR1257041](#)
- The CLI command **show vpls mac-table** does not display all MAC addresses for L2BSA subscribers. [PR1257605](#)
- Unable to run **show subscribers extensive** and some other CLI commands after GRES because subscriber-management database is unavailable. [PR1258238](#)
- DCD process crashes during the ATM-related configuration commit. [PR1258744](#)
- Subscriber management (bbe-smgd) process might crash and generate a core file during Routing Engine mastership switchover. [PR1258817](#)
- When using an AMS interface and running the **show interfaces extensive** command, the subinterfaces will show only 0 for the packet counters. [PR1258946](#)
- Junos Telemetry Interface reporting interval has a skew. [PR1259224](#)
- QSFPP-40GBASE-LR4 might remain down after fiber link flap. [PR1259930](#)
- Incorrect egress classification of L3 multicast traffic from ingress VLAN bridge interface after configuration change. [PR1260413](#)
- I2C BUS timeout causes SFP thread hogging and MPC restart. [PR1260517](#)
- A Packet Forwarding Engine saves only the first multicast IPv4 packet when waiting for a resolve request. [PR1260729](#)
- In MX Series BNG subscriber management environment, there could be a slight deviation in the dynamic profile service accounting statistics when the subscriber session terminates abruptly. [PR1260898](#)
- During multicast activation of dynamic subscribers through a service profile, the bbe-smgd process in the backup Routing Engine could sometimes crash. [PR1261285](#)
- GRPC physical interfaces \*-pkts fields zero suppressed by its own counter. [PR1261589](#)
- The **show auto-configuration** CLI command was mistakenly hidden in Junos OS 15.1 and later releases. [PR1262139](#)
- The dynamic VLAN is removed after 30 seconds if there are no subscribers on it and **remove-when-no-subscribers** is set regardless of its idle-timeout value for the dynamic VLAN. [PR1262157](#)
- Unified ISSU with subscriber-management is enabled. [PR1262877](#)
- ICMP network unreachable message is not sent back when the subscriber is terminated in a routing instance. [PR1263094](#)

- CoS service profile without line rate adjust needs to use "adjust-always" for proper revert behavior. [PR1263337](#)
- After JSD (JET service process) restart, the process is up but it is not listening on any port. [PR1263748](#)
- The smg-service subsystem is not responding to management requests. [PR1264038](#)
- Authd reports pdb\_get\_all\_profiles\_from\_db: Populate full profile tree failed, err:261, and subscribers are unable to connect at the higher number of configured dynamic profiles. [PR1264629](#)
- With the Ethernet frames with more than 2000 bytes of payload, the mspmand process might crash. [PR1264712](#)
- MX Series LAC does not send packets in the l2tp tunnel for some static PPP subscribers. [PR1265414](#)
- PRPD/JET API: BgpRouteMonitorRegister() might not send end-of-rib operation. [PR1265427](#)
- LLDP neighbor ID is captured incorrectly in streaming telemetry output. [PR1265705](#)
- Sometimes the SDN-telemetry subsystem is not responding to management requests while issuing **show agent sensors**. [PR1266058](#)
- BNG accepts IGMPv3/MLDv2 membership reports sent to non-standard multicast addresses. [PR1266309](#)
- Unified ISSU failure might be seen with Junos OS Release 16.1R4-S1. [PR1266317](#)
- ARP requests are hitting AE\_RESERVED\_IFL\_UNIT (AEx.32767) when VSTP is enabled on a double-tagged aggregated Ethernet logical interface. [PR1267238](#)
- The bbe-smgd process generates a core file during subscriber login or logout on the backup Routing Engine under certain boundary conditions. [PR1267646](#)
- The CLI configuration command **set chassis effective-shaping-rate** is enabled for the MX104. [PR1267829](#)
- ANCP Port Up message triggers RADIUS AccessRequest even when a PPP session is established. [PR1267960](#)
- The message **HALP-lbnh\_xlate\_cntr\_db\_get\_stats:250counter id 1573873: Unable to find lbnh xlate counter** is flooding the syslog. [PR1268452](#)
- Router MAC extended community does not use standardized value. [PR1269236](#)
- The Routing Engine might stop all services after GRES or unified ISSU. [PR1271306](#)

### **Infrastructure**

- The smartd **Offline uncorrectable sectors** critical logs keep reporting every 30 minutes. [PR1233992](#)
- A ksyncd crash might be seen on the backup Routing Engine due to stale next hops on the master Routing Engine. [PR1250880](#)
- Legacy Junos OS kernel might generate a core file on userland\_sysctl / sysctl\_root / sysctl\_kern\_proc\_env / panic\_on\_watchdog\_timeout. [PR1254742](#)

- Device reboots due to watchdog timeout. [PR1259616](#)
- Zero suppression does not work for internal interfaces. [PR1260036](#)

### **Interfaces and Chassis**

- T3 interface might not come up due to incorrect subrate. [PR1238395](#)
- The cfmd might crash when CFM filter refers to a firewall policy. [PR1246822](#)
- For CFM over aggregated Ethernet, incorrect Anchor FPC is selected. [PR1258490](#)
- SNMP SET fails when the FPC slot or PIC/port has a value greater than 9. [PR1259155](#)
- Jpppd might crash when traceoptions is enabled under PPPoE. [PR1264000](#)
- On MX Series Virtual Chassis this message is seen: **CHASSISD\_IPC\_WRITE\_ERR\_NULL\_ARGS: FRU has no connection arguments fru\_send\_msg Global FPC 0.** [PR1264647](#)
- Malformed PPP echo reply causes keepalive failure. [PR1273083](#)
- The message **dot1agCfmMepHighestPrDefect** might be reported in the SNMP trap with the value of -1 instead of 0 on recovery after RDI. [PR1273278](#)

### **Layer 2 Ethernet Services**

- The **show class-of-service fabric statistics** CLI command might fail with periodic **Error = Operation timed out** message. [PR1228293](#)
- An MX Series router with MPC/FPC line card might go offline during FRU upgrade phase of unified ISSU. [PR1256940](#)
- The DHCP client key identifier mismatch due to DHCPv4 Option 82 Suboption 9 change during the release time. [PR1257701](#)
- Eliminate the impact of DHCPv6 renegotiation lockout timer for DHCP solicit with rapid commit options. [PR1263156](#)

### **Layer 2 Features**

- In a scaling VPLS scenario, convergence time takes more than 10 minutes. [PR1279192](#)

### **MPLS**

- When the configured metric for one of the LSPs used in ECMP is removed, other LSPs with configured metric might not honor the configured metric value. [PR1261961](#)
- Traffic loss is seen during auto-BW MBB on ingress router as "invalid fabric token". [PR1264089](#)
- TE++ container LSP statistics are showing the same 10 LSPs and looping. [PR1267774](#)
- The core file might be generated for customers who have configured BGP family L2VPN in Junos OS Release 17.2R1. **JDI-RCT-RPD rpd core@ bgp\_labeled\_l2vpn\_standby\_outmetrics , bgp\_rt\_ribout\_rcv\_nlri:** [PR1271704](#)

### **Network Management and Monitoring**

- The eventd process stops sending syslog messages to a configured syslog server. [PR1246712](#)
- SNMPv3 trap does not contain routing instance information in contextName field. [PR1265288](#)

### **Platform and Infrastructure**

- NPC generated a core file. This type of NPC core file might be observed with a dynamic configuration change to the policer. The processing time in attempting to update all associated policers was exceeded. [PR1071040](#)
- Change the default CMERROR actions for the Major Error on MX Series platforms. [PR1186421](#)
- The routing protocol process (rpd) might crash when the ephemeral database is enabled. [PR1214298](#)
- MX Series with MPC or FPC line cards report LUCHIP EDMEM errors during unified ISSU. [PR1249395](#)
- One of the processes (dcd, rpd, dfwd, pfed, cosd, sampled) might generate a core file in large-scale 8000 ESSM login or logout with an ephemeral database. [PR1249979](#)
- The auditd might crash when RADIUS accounting is configured but the RADIUS accounting server is not reachable. [PR1250525](#)
- The bbe-smgd process might crash if you are running a PPPoE login or logout with IGMP distributed enabled. [PR1253036](#)
- After switchover, KRT queue might get stuck on the new master Routing Engine with the error **ENOENT -- Item not found**. [PR1254980](#)
- FPC might crash and generate a core file during unified ISSU because memory is not properly recycled. [PR1258795](#)
- A mismatching in/out pps value is shown with **show pfe statistics traffic detail**. [PR1259427](#)
- The routed traffic going out through IRB/I2 interface with VXLAN-EVPN is getting dropped after I2 interface switch. [PR1259551](#)
- DHCP/BOOTP reply packet for an unnumbered interface might trigger FUD process failure. [PR1260623](#)
- WRED drops on one VLAN when the other VLAN is congested. [PR1260951](#)
- DDRIF checksum error might lead to traffic blackhole. [PR1260983](#)
- FPC might crash with interface-specific firewall filters with policers configured. [PR1267908](#)
- The routing protocol process (rpd) might crash and BGP session flapping might be seen if the interfaces flap or configurations change quickly. [PR1269116](#)

### **Routing Protocols**

- Multicast Source Discovery Protocol (MSDP) source active (SA) messages are sent at irregular intervals. [PR1257668](#)
- Routing protocol process (rpd) might restart unexpectedly with a reference to `ioth_session_delete_internal()` routine. [PR1261970](#)

- The rpd might crash if the IS-IS segment routing is configured but a certain interface is not configured with RSVP. [PR1262612](#)
- MPLS label entry for direct route as BGP-LU route is permanently stuck in KRT queue when vrf-table-label is configured in CoC routing instance. [PR1263291](#)
- When applying an import policy to a BGP neighbor, the rpd process might crash continuously. [PR1265224](#)
- **Nexthop AFI=3** is observed in BGP open message after configuring **family inet unicast extended-nexthop**. [PR1272807](#)

### *Services Applications*

- Traffic is dropped when changing the source address under a NAT rule term for basic NAT translation. [PR1257801](#)
- The kmd process might crash after configuring certain IPsec configuration using the apply-groups method. [PR1265404](#)

### *Subscriber Access Management*

- Possible CPS degradation for scaled DHCP IPv4 or IPv6 and PPPoE IPv4 subscribers. [PR1264052](#)
- An incorrect number of messages in the queue for the RADIUS server is shown in the output for **show network-access aaa statistics radius detail**. [PR1267307](#)
- The CLI command **show network-access requests pending count** keeps increasing the network access requests pending count even if there are no pending authentication requests. [PR1267702](#)

### *VPNs*

- The Routing protocol process (rpd) memory leak is observed in next-generation MVPN environments. [PR1259579](#)

### SEE ALSO

[New and Changed Features | 113](#)

[Changes in Behavior and Syntax | 144](#)

[Known Behavior | 162](#)

[Known Issues | 167](#)

[Documentation Updates | 249](#)

[Migration, Upgrade, and Downgrade Instructions | 250](#)

[Product Compatibility | 258](#)



## Documentation Updates

### IN THIS SECTION

- [Protocol Independent Routing Properties | 249](#)
- [Subscriber Management Access Network Guide | 249](#)
- [Subscriber Management Provisioning Guide | 250](#)

This section lists the errata and changes in Junos OS Release 17.2R3 documentation for MX Series.

### Protocol Independent Routing Properties

- **Support for deletion of static routes when the BFD session goes down (MX Series)**—Starting with Junos OS Release 17.2R2, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session-down message.

### Subscriber Management Access Network Guide

- The “Configuring the L2TP Resynchronization Method” and “disable-failover-protocol (L2TP)” topics have been updated to state that you can configure the LNS to support only silent failover for peer resynchronization. This capability has been supported on both the LAC and the LNS since Junos OS Release 11.2.
- 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

- Support for the packet-triggered subscribers and policy control rule base (PTSP) feature was discontinued starting in Junos OS Release 13.1R1, but this was not reflected in the documentation. Text exclusive to PTSP has been removed from the *Broadband Subscriber Sessions User Guide*. This includes all CLI topics and the following chapters:
  - “Configuring the PTSP Feature to Support Dynamic Subscribers”
  - “Configuring the PTSP Partition to Connect to the External Policy Manager”
  - “Configuring PTSP Services and Rules”
  - “Monitoring and Managing Packet-Triggered Subscribers”

Topics for other features that refer to PTSP are updated to report the end of support.

- The *Broadband Subscriber Sessions User Guide* did not report that you can suspend AAA accounting, establish a baseline of accounting statistics, and resume accounting. This feature was introduced in Junos OS Release 15.1R4.

[See [Suspending AAA Accounting and Baselining Accounting Statistics Overview.](#)]

## SEE ALSO

<a href="#">New and Changed Features   113</a>
<a href="#">Changes in Behavior and Syntax   144</a>
<a href="#">Known Behavior   162</a>
<a href="#">Known Issues   167</a>
<a href="#">Resolved Issues   185</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   250</a>
<a href="#">Product Compatibility   258</a>

# Migration, Upgrade, and Downgrade Instructions

## IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.2 | 252](#)
- [Procedure to Upgrade to FreeBSD 10.x based Junos OS | 252](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 254](#)

- Upgrade and Downgrade Support Policy for Junos OS Releases | 256
- Upgrading a Router with Redundant Routing Engines | 257
- Downgrading from Release 17.2 | 257

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

Starting with Junos OS Release 15.1, in some of the devices, FreeBSD 10.x is the underlying OS for Junos OS instead of FreeBSD 6.x. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.x remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see [Understanding Junos OS with Upgraded 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 10.x-based Junos OS
MX5, MX10, MX40, MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

## Basic Procedure for Upgrading to Release 17.2

**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 10.x based Junos OS

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

To download and install FreeBSD 10.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 comprising 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-x86-32-17.2R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.2R3.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently comprised 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-x86-32-17.2R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.2R3.9-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 10.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 10.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 VM Host Installation topic in the [Installation and Upgrade Guide](#).

**NOTE:** After you install a Junos OS Release 17.2 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

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

Products impacted: MX80, and 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 comprising of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/jinstall-ppc-17.2R3.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently comprising 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-17.2R3.x-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 17.2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

## 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 16.1, 16.2 and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

For more information about EEOL releases and to review a list of EEOL releases, see <https://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 17.2

To downgrade from Release 17.2 to another supported release, follow the procedure for upgrading, but replace the 17.2 package with one that corresponds to the appropriate release.

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

See [Installation and Upgrade Guide](#).

### SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  113</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  144</a>
<a href="#">Known Behavior</a>	<a href="#">  162</a>
<a href="#">Known Issues</a>	<a href="#">  167</a>
<a href="#">Resolved Issues</a>	<a href="#">  185</a>
<a href="#">Documentation Updates</a>	<a href="#">  249</a>
<a href="#">Product Compatibility</a>	<a href="#">  258</a>

# Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 258](#)

## 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

<a href="#">New and Changed Features   113</a>
<a href="#">Changes in Behavior and Syntax   144</a>
<a href="#">Known Behavior   162</a>
<a href="#">Known Issues   167</a>
<a href="#">Resolved Issues   185</a>
<a href="#">Documentation Updates   249</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   250</a>

# Junos OS Release Notes for NFX Series

## IN THIS SECTION

- New and Changed Features | 259
- Changes in Behavior and Syntax | 265
- Known Behavior | 265
- Known Issues | 266
- Resolved Issues | 270
- Documentation Updates | 271
- Migration, Upgrade, and Downgrade Instructions | 272
- Product Compatibility | 276

These release notes accompany Junos OS Release 17.2R3 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

## IN THIS SECTION

- Release 17.2R3 New and Changed Features | 260
- Release 17.2R2 New and Changed Features | 260
- Release 17.2R1 New and Changed Features | 260

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

## Release 17.2R3 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 17.2R3.

## Release 17.2R2 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 17.2R2.

## Release 17.2R1 New and Changed Features

### Hardware

- **NFX250 Platform**—The NFX250 devices constitute Juniper Network’s secure, automated, software-driven customer premises equipment (CPE) devices that deliver virtualized network and security services on demand. Leveraging Network Functions Virtualization (NFV) and built on the Juniper Cloud CPE solution, NFX250 enables service providers to deploy and service chain multiple, secure, high-performance virtualized network functions (VNFs) in a single device.

**Table 2: NFX250 Models**

Product Number	Specifications	Features
NFX250-S1	1.9 GHz 6-core Intel CPU  16 GB of memory and 100 GB of solid-state drive (SSD) storage  Eight 1-GbE network ports, two 1-GbE RJ-45 ports which can be used as either access ports or as uplinks, two SFP ports, two SFP+ ports, one Management port, and two Console ports	Basic Layer 2/Layer 3
NFX250-S2	1.9 GHz 6-core Intel CPU  32 GB of memory and 400 GB of SSD storage  Eight 1-GbE network ports, two 1-GbE RJ-45 ports which can be used as either access ports or as uplinks, two SFP ports, two SFP+ ports, one Management port, and two Console ports	Basic Layer 2/Layer 3

Table 2: NFX250 Models (*continued*)

Product Number	Specifications	Features
NFX250-LS1	<p>1.6 GHz 4-core Intel CPU</p> <p>16 GB of memory and 100 GB of solid-state drive (SSD) storage</p> <p>Eight 1-GbE network ports, two 1-GbE RJ-45 ports which can be used as either access ports or as uplinks, two SFP ports, two SFP+ ports, one Management port, and two Console ports</p>	<p>Supports up to 100 MBPS throughput Secure Router functionality for the following features:</p> <ul style="list-style-type: none"> <li>• IPSec VPN</li> <li>• NAT</li> <li>• Stateful Firewall</li> <li>• Routing services – BGP, OSPF, DHCP, IPv4 and IPv6</li> </ul>

- Transceivers—NFX250 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-ZR
  - 1G-gigabit SFP transceivers: EX-SFP-1GE-SX, EX-SFP-1GE-SX-ET, EX-SFP-1GE-LX, EX-SFP-1GE-LH, EX-SFP-1GE-T, EX-SFP-1GE-LX40K, EX-SFP-GE10KT13R14, EX-SFP-GE10KT14R13, EX-SFP-GE10KT13R15, EX-SFP-GE10KT15R13, EX-SFP-GE40KT13R15, EX-SFP-GE40KT15R13, EX-SFP-GE80KCW1470, EX-SFP-GE80KCW1490, EX-SFP-GE80KCW1510, EX-SFP-GE80KCW1530, EX-SFP-GE80KCW1550, EX-SFP-GE80KCW1570, EX-SFP-GE80KCW1590, EX-SFP-GE80KCW1610
- Direct Attach Copper (DAC) Cables—NFX250 supports the following DAC cables:
  - EX-SFP-10GE-DAC-1M
  - EX-SFP-10GE-DAC-3M
  - EX-SFP-10GE-DAC-5M

### ***Juniper Device Manager***

The Juniper Device Manager (JDM) is a low-footprint Linux container that provides these functions:

**NOTE:** These features were previously supported in the 15.1X53-D40 and 15.1X53-D47 releases of Junos OS.

- Virtual machine (VM) life cycle management
- Device management and isolation of host OS from user installations
- NIC , single-root I/O virtualization (SR-IOV), and virtual input/output (VirtIO) interface provisioning

- Support for the Network Service Orchestrator module to connect to Network Service Activator
- Inventory and resource management
- Internal network and image management
- Service chaining—provides building blocks such as virtual interfaces and bridges for users to implement service chaining policies
- Virtual console access to VNFs including vSRX and vjunos
- Support for outbound SSH connections
- Authentication of users using TACACS+
- Configure SNMP, and handle SNMP queries and traps
- Enhanced CLI to support launching VNFs, service chaining VNFs, and configuring and monitoring various system parameters and statistics
- IPsec—The IPsec implementation for NFX250 platforms has been enhanced to protect the management traffic between JDM, VNFs and the remote SDN controller and other central servers. The IPsec implementation uses AutoKey IKE with preshared keys to authenticate the participants in an IKE session, each side must configure and securely exchange the preshared key in advance. IPsec for NFX250 devices supports only traffic selector based tunnels, multiple IPsec security associations are negotiated based on multiple traffic selectors configured. Configuration of interfaces and static routes is supported.

### ***Junos Control Plane***

Junos Control Plane (JCP) is the Junos VM running on the hypervisor. By default, JCP runs as vjunos0 on NFX250. You can use JCP to configure the network ports of the NFX250 device. You can log in to JCP from JDM by using the SSH service and CLI, which is similar to the Junos OS CLI. The JCP supports the following features:

- Link aggregation—Link aggregation enables you to use multiple network cables and ports in parallel to increase link speed and improve redundancy.
- Support for Layer 3 logical interfaces—A Layer 3 logical interface is a logical division of a physical interface or an aggregated Ethernet interface that operates at the network level and that can receive and forward IEEE 802.1Q VLAN tags. You can use these interfaces to route traffic between multiple VLANs along a single trunk line that connects an NFX250 device to a Layer 2 switch. Only one physical connection is required between the NFX250 device and the switch.
- 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, and to 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. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of

802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.

- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), 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.
- OSPF support—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). NFX devices support OSPFv1 and OSPFv2. You can configure OSPF at the [edit protocols ospf] hierarchy level.
- Bidirectional Forwarding Detection (BFD) support for static routes and the OSPF and RIP protocols—BFD uses control packets and shorter detection time limits to rapidly detect failures in a network. Hello packets are sent at a specified, regular interval by routing devices. A neighbor failure is detected when a routing device stops receiving a reply after a specified interval.
- 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 highly available path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.
- Internet Group Management Protocol (IGMP) support—IGMP manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to multicast routers that are their immediate neighbors. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.
- IGMP Snooping support—IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.
- Protocol Independent Multicast (PIM) sparse mode support—PIM sparse mode enables efficient routing to multicast groups with receivers that are sparsely spread over multiple networks. To configure PIM sparse mode, include the pim statement at the [edit protocols] hierarchy level.
- SNMP support—SNMP includes versions 1, 2, and 3 for monitoring system activity.
- System logging (syslog) support—Syslog enables you to log system messages into a local directory on the switch or to a system log server.
- Port mirroring support—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.
- Firewall filter support—You can provide rules 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.

- Policing support—You can use policing to apply limits to traffic flow and to set consequences for packets that exceed those limits.
- Storm control support—You can enable the switch to monitor traffic levels and take a specified action when a specified traffic level—called the storm control level—is exceeded, preventing packets from proliferating and degrading service. You can configure a switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs.
- Class of service (CoS)—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 the 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.
- Class-of-service (CoS) rewrite rules and classifier support—You can use rewrite rules to set the value of the CoS bits within a packet header, so you can alter the CoS settings of incoming packets. Packet classification maps incoming packets to a particular class-of-service (CoS) servicing level. You can use classifiers to map packets to a forwarding class and a loss priority and to assign packets to output queues based on the forwarding class.
- Secure Boot—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. No action is required to implement Secure Boot.

## vSRX

vSRX offers the same capabilities as Juniper Networks SRX Series Services Gateways in a virtual form factor, providing perimeter security, IPsec connectivity, and filtering for malicious traffic without sacrificing reliability, visibility, and policy control. This virtual security and routing appliance ensures reliability for each application. By default, vSRX version 15.1X49-D75 is preloaded on NFX250 Network Services platform 17.2R1 release. Earlier versions of vSRX is not compatible with the Junos version 17.2R1 release on NFX250.

## SEE ALSO

[Changes in Behavior and Syntax | 265](#)

[Known Behavior | 265](#)

[Known Issues | 266](#)

[Resolved Issues | 270](#)

[Documentation Updates | 271](#)

[Migration, Upgrade, and Downgrade Instructions | 272](#)

[Product Compatibility | 276](#)



## Changes in Behavior and Syntax

There are no changes in default behavior and syntax in Junos OS Release 17.2R3 for the NFX Series.

### SEE ALSO

[New and Changed Features | 259](#)

[Known Behavior | 265](#)

[Known Issues | 266](#)

[Resolved Issues | 270](#)

[Documentation Updates | 271](#)

[Migration, Upgrade, and Downgrade Instructions | 272](#)

[Product Compatibility | 276](#)

## Known Behavior

### IN THIS SECTION

- [Juniper Device Manager | 265](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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.

### Juniper Device Manager

- JDM shell configurations of interfaces override JDM CLI configurations. As a workaround, use the JDM CLI to configure interfaces. [PR1155749](#)
- SR-IOV interfaces do not support more than 64 VLANs on NFX250. [PR1156348](#)

## SEE ALSO

[New and Changed Features | 259](#)[Changes in Behavior and Syntax | 265](#)[Known Issues | 266](#)[Resolved Issues | 270](#)[Documentation Updates | 271](#)[Migration, Upgrade, and Downgrade Instructions | 272](#)[Product Compatibility | 276](#)

## Known Issues

### IN THIS SECTION

- [Infrastructure | 266](#)
- [IPSec | 266](#)
- [Juniper Device Manager | 267](#)
- [Junos Control Plane | 269](#)
- [vSRX | 269](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R3 for the 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.

### Infrastructure

- You might not be able to upgrade from Junos OS releases 15.1X53-D40 and 15.1X53-D45 to Junos OS release 17.2R3. As a workaround, you can use the image file on a USB, configure the NFX device to boot from the USB, and install the upgrade. [PR1252323](#)

### IPSec

- There is no CLI command to clear interface flow-statistics on ipsec-nm. [PR1216474](#)

- Initial allocation of hugepages is not guaranteed when the srpxfe is killed or restarted. [PR1233794](#)

## Juniper Device Manager

- There might be no checks when you configure the IP address on different logical units of interfaces. The commit will go through, and will be displayed in the configuration. [PR1150512](#)
- The following commands are not supported:
  - clear system reboot and clear system commit
  - restart gracefully, restart immediately, restart init, and restart soft
  - show ethernet-switching, show version brief, show version all members, and show system services service-deployment

### [PR1154819](#)

- When you use the netconf command to display system information details such as model and OS, the system OS is displayed as QFX. [PR1160055](#)
- Ubuntu package does not successfully install on the JDM container. As a workaround, install the package passwd by using the sudo apt-get install passwd command, which enables the useradd command again. [PR1168680](#)
- When you configure a static route on JDM in enhanced-orchestration disabled mode, there might not be an explicit check to validate the IP address. [PR1173039](#)
- System Host bridge uses a default MTU of 1500 and does not support Jumbo frames. Currently there is no CLI to configure the MTU on the host bridge. [PR1192169](#)
- The Network Service Orchestrator module commits the configuration on JDM, Junos Control Plane, and IPSec-NM sequentially. If the commit fails on any one of these system VNFs, the Network Service Orchestrator module automatically rolls back to the older configuration on the VNF where the commit error is seen. But, all prior Network Service Orchestrator module configuration commits on the earlier VNFs continue to exist and is not reversed. [PR1196253](#)
- There is no commit check if the PCI address is reused for different interfaces in a VNF. It is recommend to stop the VNF and then add or delete interfaces. [PR1205497](#)
- Certain VNFs support hot plugging of virtio interfaces when the VNF is running. When a VLAN mapped interface is hot plugged to VNFs such as Centos, it is seen that the interface is not reachable from the vjunos0 VM. As a workaround, delete the VNF configuration and re-commit the complete configuration along with the new interface. [PR1213451](#)
- After enabling or disabling the ipsec-nm service on the NFX250 platform, a warning message might not be displayed asking for a consent to reboot the device. The enabling or disabling action will be effective only after the device is rebooted. Similarly, no warning is displayed when Enhanced orchestration is either enabled or disabled. [PR1213489](#)

- Pre-allocation of hugepages might not consider the available memory and proper commit check is required. It is advisable to use the feature based on free system memory availability. By default, the system requires up to 6 to 7 gigabytes of memory for various operations. The system might not function properly if more memory than what is available is allocated. [PR1213944](#)
- While spawning a VNF, there might not be a commit check for the valid image type supported. [PR1221642](#)
- 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](#)
- The following commands are not supported:
  - show host
  - request system software delete
  - request system software rollback
  - request system storage cleanup

[PR1219972](#)

- DHCP service can be configured on custom system bridges for service chaining. There might be no commit check if the lower and higher values of the pool range are swapped. [PR1223247](#)
- If the configured TACACS+ server has an IP that can be accessed from JDM, the tacplus pam might not wait till timeout in case TACACS+ server is unreachable. [PR1224420](#)
- The Swap memory information displays incorrect values in the show system visibility jdm command output for NFX250 platforms with optimized SSD layouts. [PR1227528](#)
- With enhanced-orchestration mode enabled and routing over management configured on vSRX for WAN redundancy for critical traffic, the system CPU utilization will reach 100% if WAN link goes down and traffic routes through out-of-band management. vSRX may not respond to ping or management requests. Egress traffic through management might be throttled. [PR1233478](#)
- Removing the IRB configuration along with the DHCP configuration on JDM and rolling back the configuration might result in the DHCP service not functioning for service chaining of VNFs. [PR1234055](#)
- Hugepages that are pre-configured through CLI are not used if a custom init-descriptor is used. [PR1245330](#)
- When a VLAN tag is configured through a JDM CLI on a VNF that is provisioned to a DPDK enabled VM and the VM is spawned, the VLAN filtering or striping configuration on the VNF stops taking effect. Removing and recommitting the JDM VLAN ID configuration on the VNF can resolve the issue unless the system or the VNF is rebooted. [PR1251596](#)

- **show system visibility cpu** command on JDM has the field values for IOWait and Intr always set to zero. [PR1258361](#)
- Configuring more than the available number of SR-IOV interfaces in Enhanced mode might result in a state where the used MAC addresses for such interfaces are not released back to the system MAC pool on deletion of the VNF. [PR1259975](#)

## Junos Control Plane

- The Alarm LED will be amber for a major alarm instead of red. In the NFX250-S1E model, the Alarm LED does not blink for any alarms. [PR1146307](#)
- Configuring DSCP and DSCPv6 classifiers together on a Layer 2 interface is not supported. [PR1169529](#)
- When the option `accept-source-mac mac-address` is configured on an interface and then deleted, no additional MAC's will be learnt on the interface. Only the MAC's which were earlier configured will be available. [PR1168197](#)
- When LLDP is configured on `vjunos0` on an NFX250 Network Services platform, the system name TLV(5) might not be advertised. [PR1169479](#)
- There might a traffic drop in IPv4 multicast traffic on JCP when flow-control is configured on interfaces and multicast traffic is more than 400pps. [PR1191794](#)
- On an interface with family `inet` configured, you might not be able to configure a classifier or rewrite rules. [PR1262840](#)
- If the traffic in the out-of-band interface is more, the control plane connectivity might get blocked for sometime while the packets are processed. If this interruption persists, the connection between the PFE and control plane is cleared, which results in a PFE restart or shutdown. You must ensure that there is no heavy traffic flow in the management VLAN. [PR1270689](#)

## vSRX

- On an NFX250-S1E platform running vSRX VNF, the performance of SR-IOV with UTM and IDP is lower than VirtIO with UTM and IDP. [PR1214118](#)
- If `per-unit-scheduler` is not configured, the IFD shaping fails and no packet is queued. [PR1264556](#)
- After configuring the IFD shaping, the ingress interface cannot receive packets. [PR1264850](#)
- The current maximum number of concurrent SIP calls is below the specified maximum limit. [PR1273356](#)

SEE ALSO

Changes in Behavior and Syntax	265
Known Behavior	265
Resolved Issues	270
Documentation Updates	271
Migration, Upgrade, and Downgrade Instructions	272
Product Compatibility	276

## Resolved Issues

### IN THIS SECTION

- Resolved Issues: 17.2R3 | [270](#)
- Resolved Issues: 17.2R2 | [270](#)
- Resolved Issues: 17.2R1 | [270](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

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

### Resolved Issues: 17.2R3

There are no resolved issues for NFX Series in Junos OS Release 17.2R3.

### Resolved Issues: 17.2R2

There are no resolved issues for NFX Series in Junos OS Release 17.2R2.

### Resolved Issues: 17.2R1

#### *Juniper Device Manager*

- User-defined login class is not supported on JDM. [PR1155965](#)
- On the device, ping with the record-route option does not work for VirtIO interfaces. [PR1162659](#)

- The default gateway assigned by phone-home client (PHC) for clients connected through the front panel ports is 10.10.10.254. [PR1168284](#)
- The CLI to configure the time zone is not functional. [PR1169675](#)
- SNMP trap is not supported on JDM. [PR1173216](#)

***Junos Control Plane***

- Transmit rate of 0 cannot be configured on schedulers. [PR1158085](#)
- If a cable is not connected to the RJ-45 ports on the front panel, the status LED blinks. [PR1168054](#)
- SFP-T transceivers are not supported. [PR1151575](#), [PR1166808](#), [PR1168203](#)

SEE ALSO

<a href="#">New and Changed Features   259</a>
<a href="#">Changes in Behavior and Syntax   265</a>
<a href="#">Known Behavior   265</a>
<a href="#">Documentation Updates   271</a>
<a href="#">Known Issues   266</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   272</a>
<a href="#">Product Compatibility   276</a>

**Documentation Updates**

There are no errata or changes in Junos OS Release 17.2R3 documentation for NFX Series.

SEE ALSO

<a href="#">New and Changed Features   259</a>
<a href="#">Changes in Behavior and Syntax   265</a>
<a href="#">Known Behavior   265</a>
<a href="#">Known Issues   266</a>
<a href="#">Resolved Issues   270</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   272</a>
<a href="#">Product Compatibility   276</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 272
- Basic Procedure for Upgrading to Release 17.2 | 272

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.

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. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

For more information 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 17.2

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



**NOTE:** Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

**NOTE:** The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [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 17.2R2:

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 17.2R2 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-17.2R2.13-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.2R2.13-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 *Software Installation and Upgrade Guide*.

**NOTE:** After you install a Junos OS Release 17.2 **`jinstall`** package, you cannot return to the previously installed software by issuing the **`request system software rollback`** command. Instead, you must issue the **`request system software add validate`** command and specify the **`jinstall`** package that corresponds to the previously installed software.

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

SEE ALSO

<a href="#">New and Changed Features   259</a>
<a href="#">Changes in Behavior and Syntax   265</a>
<a href="#">Known Behavior   265</a>
<a href="#">Documentation Updates   271</a>
<a href="#">Known Issues   266</a>
<a href="#">Resolved Issues   270</a>
<a href="#">Product Compatibility   276</a>

# Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 276](#)

## 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](#).

SEE ALSO

<a href="#">New and Changed Features   259</a>
<a href="#">Changes in Behavior and Syntax   265</a>
<a href="#">Known Behavior   265</a>
<a href="#">Documentation Updates   271</a>
<a href="#">Known Issues   266</a>
<a href="#">Resolved Issues   270</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   272</a>

# Junos OS Release Notes for PTX Series Packet Transport Routers

## IN THIS SECTION

- New and Changed Features | 277
- Changes in Behavior and Syntax | 294
- Known Behavior | 299
- Known Issues | 301
- Resolved Issues | 305
- Documentation Updates | 312
- Migration, Upgrade, and Downgrade Instructions | 312
- Product Compatibility | 317

These release notes accompany Junos OS Release 17.2R3 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

## IN THIS SECTION

- Release 17.2R3 New and Changed Features | 278
- Release 17.2R2 New and Changed Features | 278
- Release 17.2R1 New and Changed Features | 278

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 17.2R3 New and Changed Features

There are no new features or enhancements to existing features for PTX Series in Junos OS Release 17.2R3.

## Release 17.2R2 New and Changed Features

### *Software Installation and Upgrade*

- **Device serial number added to DHCP option 60 (PTX1000)**—Starting in Junos OS Release 17.2R2, DHCP option 60 (Vendor Class Identifier) includes the serial number of the device when you use zero touch provisioning to automate provisioning of the device configuration and software image. The serial number can uniquely identify the device in a broadcast network. The serial number appears in the format *Juniper-model-number*. For example, a PTX1000 router numbered DA000 appears as *Juniper-ptx1000-DA000*.

## Release 17.2R1 New and Changed Features

### *Hardware*

- **PTX10008 Packet Transport Router**—PTX10008 Packet Transport Router provides 3.0 Tbps per slot forwarding capacity for the service providers and cloud operators. The router provides a smooth transition from 10-Gigabit Ethernet and 40-Gigabit networks to 100-Gigabit Ethernet high-performance networks. This high-performance, 13 rack unit (13RU) modular chassis provides 24 Tbps of throughput and 16 Bpps of forwarding capacity. PTX10008 has eight slots for the line cards that can support a maximum of 1152 10-Gigabit Ethernet ports, 288 40-Gigabit Ethernet ports, or 240 100-Gigabit Ethernet ports. PTX10008 supports two new line cards, LC1101 and LC1102. The LC1101 line card consists of 30 QSFP+ Pluggable Solution (QSFP28) ports and the LC1102 has 36 QSFP+ ports that support 40-gigabit or 100-gigabit Ethernet optical transceivers.

### *Class of Service (CoS)*

- **Support for CoS-based forwarding (PTX10008)**—CoS-based forwarding (CBF) enables the control of next-hop selection based on a packet's class of service field. Starting with Junos OS Release 17.2R1, PTX10008 routers support CBF. CBF can only be configured on a device with eight or fewer forwarding classes plus a default forwarding class. You can implement CBF by specifying **next-hop-map** at the **[edit class-of-service forwarding-policy]** hierarchy level and then applying **next-hop-map** at the **[edit policy-options]** hierarchy level.
- **CoS-based forwarding support for up to 16 forwarding classes (MX Series and PTX routers)**—Beginning with Junos OS Release 17.2R1, MX Series routers with MPCs or MS-DPCs, vMX, PTX3000 routers, PTX5000 routers, and VPTX support configuring CoS-based forwarding (CBF) for up to 16 forwarding classes. All other platforms support CBF for up to 8 forwarding classes. To support up to 16 forwarding classes for CBF on MX routers, enable **enhanced-ip** at the **[edit chassis network-services]** hierarchy level.

[See [Forwarding Policy Options Overview](#).]

- **Support for class of service (CoS) (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, support is extended for class of service (CoS). CoS is the assignment of traffic flows to different service levels. Service providers can use router-based CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows.

On a PTX1000 router, you can divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs.

- **Support for shaping of traffic exiting third-generation FPCs (PTX1000)**—Starting with Junos OS Release 17.2R1, you can shape the output traffic of an FPC3 physical interface on a PTX1000 Packet Transport Router so that the interface transmits less traffic than it is physically capable of carrying. Shaping on all PTX Series router interfaces has a minimum rate of 1 Gbps and an incremental granularity of 0.1 percent of the physical interface speed after that (for example, 10 Mbps increments on a 10 Gbps interface). You can shape the output traffic of a physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces interface-name]** or **[edit class-of-service traffic-control-profiles profile-name]** hierarchy level and applying the traffic control profile to an interface.

[See [shaping-rate \(Applying to an Interface\)](#).]

### **Forwarding and Sampling**

- **Support for Bidirectional Forwarding Detection (BFD) (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, support is extended for Bidirectional Forwarding Detection (BFD). The BFD protocol uses control packets and shorter detection time limits to rapidly detect failures in a network. Hello packets are sent at a specified, regular interval by routing devices. A neighbor failure is detected when a routing device stops receiving a reply after a specified interval.

On a PTX1000 router, you can configure BFD for static routes and for the BGP, IS-IS, OSPF, PIM, and RIP protocols.

### **General Routing**

- **OpenConfig: Supporting the BGP model in Junos OS (PTX Series)**—Starting in Junos OS Release 17.2R1, the configuration leaf devices defined in the **openconfig-bgp.yang** and **openconfig-bgp-multiprotocol.yang** files are supported.
- **OpenConfig: BGP routing table support for operational state model (PTX Series)**—Starting in Junos OS Release 17.2R1, the OpenConfig BGP RIB routing table supports local-rib for IPv4 and IPv6. The **openconfig-rib-bgp.yang** model supports five logical RIBs per address family. There are five tables for IPv4 routes and five tables for IPv6 routes.

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

- **Kernel synchronization performance and debugging enhancements (PTX Series)**—Starting in Junos OS Release 17.2R1, the kernel synchronization process (ksyncd) uses multithreading for increased performance, and you can use new CLI commands for ksyncd debugging and recovery. Use the **set**

**system kernel-replication no-multithreading** command to run ksyncd in single thread mode for debugging purposes. Use the **set system kernel-replication system-reboot recovery-failure** command to configure the automatic reboot of a standby Routing Engine after receiving a ksyncd initialization error.

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

- **Resiliency Support for LC1101 and LC1102 (PTX10008)**—Starting with Junos OS Release 17.2R1, resiliency support is enabled for the following devices:
  - LC1101 and LC1102
  - Switch Interface Boards
- **Chassis management**—In Junos OS Release 16.1X65 and 17.2R1, the following CLI operational mode commands are supported on a PTX1000 router:
  - **show chassis hardware**
  - **show chassis temperature-thresholds**
  - **show chassis environment**
  - **show chassis firmware**

### *Interfaces and Chassis*

- **Support for packet-forwarding features on LC1101 and LC1102 (PTX10008)**—Starting in Junos OS Release 17.2R1, the following key packet-forwarding features are enabled on LC1101 and LC1102 for PTX10008 routers:
  - Basic Layer 2 features and protocols
  - Class of Service (CoS)
  - Firewall filters and policers
  - Hash enhancement
  - Large scaling IPv4 and IPv6 forwarding information base (FIB)
  - Layer 3 VPNs
  - MPLS
  - Sampling and port mirroring
- **Fabric management support (PTX10008)**—Starting in Junos OS Release 17.2R1, you can set up and manage the fabric connections between the Packet Forwarding Engines of LC1101 and LC1102 in the PTX10008 routers. Fabric management includes collecting fabric status and statistics, monitoring health of the hardware, and responding to CLI queries. It also tracks addition and removal of FRUs from the router and monitors faults in the data plane. It is enabled by default and can be monitored by using the following commands:



- **show chassis fabric summary**
- **show chassis fabric fpcs fpc fpc-slot**
- **show chassis fabric sibs**
- **show chassis fabric errors**
- **show chassis fabric reachability**
- **Support for LC1101 and LC1102 with Routing and Control Board (RCB) (PTX10008)**—Starting with Junos OS Release 17.2R1, the PTX10000 Routing and Control Board (RCB) is supported on PTX10008 routers. The PTX10008 chassis can run with one or two RCBs. A fully redundant system requires a second RCB. When two RCBs are installed, one RCB functions as the master and the second as the backup. If the master RCB is removed, the backup starts and becomes the master. The RCB integrates the control plane and Routing Engine functions into a single management unit. The RCB handles system control functions such as environmental monitoring, routing Layer 2 and Layer 3 protocols, alarm and logging functions, and other functions required to manage the operation of a chassis.
- **Support for 10-Gigabit Ethernet on LC1101 - 30C line card (PTX10008)**—Starting in Junos OS Release 17.2R1, PTX10008 routers support 10-Gigabit Ethernet interfaces in addition to 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces on the LC1101 - 30C line card.

When a particular PE chip or Packet Forwarding Engine is working in mode A to support 10-Gigabit Ethernet, ports 6, 7, 16, 17, 26, and 27 at the PE0 to PE5 level are non operational. However, once the PE goes into mode D, these ports become operational and can operate at 40-Gigabit Ethernet, or 100-Gigabit Ethernet speed.

For 10-Gigabit Ethernet, you must configure the port using the channelization command. Because there is no port-groups option for the 100-Gigabit Ethernet line card, you must use individual port channelization commands.

In 30C line card, by default FPC comes up in mode D, when you channelize the first port in any PE, the FPC restarts and the corresponding PE comes up in mode A. Further channelization in that PE does not restart the FPC. However, if you channelize some another ports in another PE, then the whole FPC restarts again. If you undo the channelization of all ports in any PE, then the FPC gets restarted and the corresponding PE comes up in mode D, which is the default mode.

**NOTE:** If any mode changes (A to D or D to A) occur at the PE, the line card automatically performs a cold reboot.

- **Support for channelizing the 40-Gigabit Ethernet ports (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, the PTX1000 Packet Transport Router supports 10-Gbps, 40-Gbps, and 100-Gbps port speeds, enabling service providers to organically distribute peering points throughout the network. You can channelize four 10-Gigabit Ethernet interfaces from the 40-Gigabit Ethernet interfaces. By default, the 40-Gigabit Ethernet interfaces are named **et-fpc/pic/port**. The names of the channelized 10-Gigabit

Ethernet interfaces appear in the format: **et-fpc/pic/port:channel**, where **channel** is a value from 0 through 3.

To configure a port speed of 40-Gbps or 10-Gbps, use the **set chassis fpc slot pic pic-slot port 0..71 channel-speed (10g|40g)** command.

You can also configure 24 out of the 72 ports to operate at 100-Gbps speed.

- **Support for packet-forwarding features (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, the PTX1000 Packet Transport Router supports the following key packet forwarding features:
  - Basic Layer 2 features and protocols
  - Class of service (CoS)
  - Firewall filters and policers
  - Hash enhancement
  - Large scaling IPv4 and IPv6 forwarding information base (FIB)
  - Layer 3 VPNs
  - MPLS
  - Sampling and port mirroring

- **Support for configuring multiple port speeds on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, PTX1000 Packet Transport Router supports 10-Gbps, 40-Gbps, and 100-Gbps port speeds, enabling service providers to organically distribute peering points throughout the network. To configure the port speed, use the **speed [10G | 40G | 100G]** statement at the **[edit chassis fpc slot-number pic pic-number port port-number]** hierarchy level. The default port speed is **10G**.

**Support for configuring local loopback on PTX1000**—In Junos OS Release 16.1X65 and 17.2R1, to enable local loopback, use the **loopback local** configuration statement on PTX1000 interfaces. The PTX1000 supports only local loopback, not remote loopback. Configure the statement at the **[edit interfaces interface-name gigether-options]** hierarchy level.

- **Support for aggregated Ethernet (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, support is extended for aggregated Ethernet. The Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-flow load balancing.

On a PTX1000 router, you can configure the member links of an aggregated Ethernet bundle with any combination of rates—also known as mixed rates. The bandwidth that is provided by an aggregated Ethernet bundle can be utilized completely and efficiently when the links are configured with different rates.

- **Support for DCU accounting and SCU accounting (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, the destination class usage (DCU) accounting and source class usage (SCU) accounting are supported on PTX1000 routers.

You can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets, which are defined as source classes and destination classes. SCU counts packets sent to customers by performing lookups on the source IP address and the destination IP address. SCU accounting enables you to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. DCU counts packets from customers by performing lookups of the IP destination address. DCU accounting enables you to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.

**NOTE:** DCU accounting and SCU accounting are supported only if the enhanced-mode statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Support for unicast RPF (PTX1000)**—Starting in Junos OS Release 17.2R1, you can configure unicast reverse path forwarding (RPF) to reduce the impact of denial-of-service (DoS) attacks on PTX Series routers that have third-generation FPCs installed.

**NOTE:** Unicast RPF is supported only when the enhanced-mode statement is configured at the **[edit chassis network-services]** hierarchy level.

## IPv6

- **IPv6 statistics on PTX1000, PTX3000, PTX5000, and PTX10008 with third-generation FPCs**—Starting in Junos OS Release 17.2R1, you can obtain the transit IPv6 statistics at both the physical interface and logical interface levels on third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008 by using both a CLI command and SNMP MIB counters. Use the **show interfaces statistics** command to display both physical interface and logical interface statistics. You can view only logical interface statistics if you use SNMP MIB counters. However, for aggregated Ethernet interfaces, the accounting is not done at the level of the child links and, thus, IPv6 statistics for child links are not displayed.

To start getting IPv6 statistics on third-generation FPCs, use the **route-accounting** statement at the **[edit forwarding-options family inet6]** hierarchy level. PTX Series routers with first-generation and second-generation FPCs do not display IPv6 statistics for physical interfaces or logical interfaces, and transit statistics on child links in aggregated Ethernet interfaces are also not taken into account.

**NOTE:** Egress IPv6 statistics are not taken into account in case of MPLS POP where IPv6 traffic is encapsulated within MPLS and MPLS is stripped off before the plain IPv6 traffic is forwarded.

[See [route-accounting](#) and [show interfaces extensive](#).]

## Layer 2 Features

- **Support for Layer 2 protocols (PTX10008)**—Starting in Junos OS Release 17.2R1, L2 circuit and L2VPN are supported on PTX10008 routers.

## Layer 3 Features

- **BGP (PTX1000)**—In Junos OS Release 16.1X65 and 17.2R1, BGP is an exterior gateway protocol (EGP) for routing traffic between autonomous systems (AS). You can configure BGP at the **[edit protocols bgp]** hierarchy level.

**OSPF (PTX1000)**—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). PTX1000 routers support OSPFv1, OSPFv2, and OSPFv3. You can configure OSPF at the **[edit protocols ospf]** hierarchy level.

**Synchronization between OSPF and LDP (PTX1000)**—LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by OSPF. If the synchronization between LDP and OSPF is lost, the label-switched path (LSP) goes down. Therefore, LDP and IS-IS synchronization are beneficial.

To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the **ldp-synchronization** statement at the **[edit protocols ospf interface interface-name]** hierarchy.

**IS-IS (PTX1000)**—The IS-IS protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system.

**Synchronization between IS-IS and LDP (PTX1000)**—LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by IS-IS. If the synchronization between LDP and IS-IS is lost, the label-switched path (LSP) goes down. Therefore, LDP and IS-IS synchronization are beneficial.

To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the **ldp-synchronization** statement at the **[edit protocols isis interface *interface-name*]** hierarchy.

- **Support for Layer 3 protocols (PTX10008)**—Starting in Junos OS Release 17.2R1, Layer 3 protocols are supported on PTX10008 routers. Layer 3 protocols include the Multiprotocol Label Switching (MPLS), Layer 3 Virtual Private Network (L3VPN), Bidirectional Forwarding Detection (BFD), Layer 2 Virtual Private Network (L2VPN), Point-to-multipoint (P2MP), Fast ReRoute (FRR), Operations, Administration and Maintenance (OAM), Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP), Adaptive Load Balancing (ALB), and so on.

### Management

- **Support for LSP events and properties sensor for Junos Telemetry Interface (PTX3000 and PTX5000 routers)**—Starting with Junos OS Release 17.2R1, you can export statistics for LSP events and properties through the Junos Telemetry Interface. Only gRPC streaming for this sensor is supported. You can export statistics for ingress point-to-point LSPs, point-to-multipoint LSPs, bypass LSPs, and dynamically created LSPs. To export data through gRPC, use the **/mpls/lsp/** or **/mpls/signal-protocols/** set of OpenConfig subscription paths. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of the Junos OS with an upgraded FreeBSD kernel, 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 [Guidelines for gRPC Sensors](#).]

- **Support for device family and release in Junos OS YANG modules (PTX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

[See [Understanding Junos OS YANG Modules](#).]

- **Support for LSP statistics for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.2R1, you can stream telemetry data for LSPs through UDP and gRPC. To provision an LSP statistics sensor for UDP streaming, include the **resource /junos/services/label-switched-path/usage/** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level. Use the **mpls/lsp/constrained-path/tunnels/tunnel/** path to provision a sensor for streaming LSP statistics through gRPC. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions. For both UDP and gRPC streaming, you must also configure the

**sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level. Support for the LSP statistics sensor was previously introduced in Junos OS Release 15.1F6 and Junos OS Release 16.1R4.

[See [Overview of the Junos Telemetry Interface](#).]

- **Support for routing protocol processes task memory utilization sensor for Junos Telemetry Interface (PTX Series)**—Starting in Junos OS Release 17.2R1, you can stream telemetry data through gRPC for routing protocol process (RPD) task memory usage. Include the **/junos/task-memory-information/** path to provision a sensor to stream data through gRPC. UDP streaming for this sensor is not supported. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent 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. OpenConfig paths are used to define telemetry parameters for data streamed through gRPC. This functionality was previously introduced in Junos OS Release 16.1R3.

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

- **Support for gRPC streaming for Junos Telemetry Interface firewall filter statistics (PTX3000 and PTX5000)**—Starting with Junos OS Release 17.2R1, you can use gRPC interfaces to provision sensors to subscribe to and receive firewall filter telemetry data. Traffic-class counter statistics are also collected. Use the **/junos/firewall/firewall-stats/** path to provision a sensor for firewall filter statistics. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent 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. OpenConfig paths are used to define telemetry parameters for data streamed through gRPC. This functionality was previously introduced in Junos OS Release 16.1R4.

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

- **Support for the Junos Telemetry Interface (PTX1000)**—Starting with Junos OS Release 17.2R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for several network elements without involving polling. You can stream data through UDP or gRPC.

Only the following sensors are supported on PTX1000 routers:

- Physical interfaces statistics
- Label-switched-path (LSP) statistics
- Network processing unit (NPU) memory
- NPU memory utilization
- CPU memory

To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. To provision sensors to stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters for a specified list of OpenConfig command paths. Streaming

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

[See [Overview of the Junos Telemetry Interface](#).]

- **Support for queue statistics for logical interface sensors for Junos Telemetry Interface (PTX3000 and PTX5000 routers)**—Starting with Junos OS Release 17.2R1, logical interface sensors also collect egress and ingress queue statistics. Both UDP and gRPC streaming are supported. Queue statistics, including for per-unit queuing and hierarchical queuing, are exported when a queuing structure is configured on a logical interface. To provision a logical interfaces statistics sensor for UDP streaming, include the **resource /junos/system/linecard/interface/logical/usage/** statement at the **[edit services analytics sensor sensor-name]** hierarchy level. To provision a sensor for gRPC streaming, include **/interfaces/interface[name='interface-name']/subinterfaces/** in the subscription path. Use the **telemetrySubscribe** RPC to define telemetry parameters for gRPC streaming. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions.

[See [Overview of the Junos Telemetry Interface](#).]

## MPLS

- **MPLS inter-AS link protection (PTX Series)**—Starting in Junos OS Release 17.2R1, MPLS inter-AS link protection is supported. Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router will then choose an alternate link through another interface to send traffic to its destination.

For an MPLS inter-AS environment, link protection can be enabled when labeled-unicast is used to send traffic between autonomous systems (ASs). To configure link protection on an interface, the **protection** statement is introduced at the **[edit protocols bgp group group-name family inet labeled-unicast]** hierarchy level.

[See [protection](#).]

- **Support for filter-based GRE for IPv4 and IPv6 tunneling (PTX Series)**—In Junos OS Release 16.1X65 and 17.2R1, the filter-based generic routing encapsulation (GRE) for IPv4 and IPv6 tunneling uses firewall filters to provide de-encapsulation of GRE traffic. The configuration of filter-based GRE de-encapsulation supports the **routing-instance** statement as one of the attributes.

**NOTE:** Configuring filter-based GRE for IPv4 and IPv6 tunneling is supported only when the enhanced-mode statement is configured at the **[edit chassis network-services]** hierarchy level.

- **MPLS inter-AS link protection (PTX1000)**—Starting in Junos OS Release 17.2R1, MPLS inter-AS link protection is supported. Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router will then choose an alternate link through another interface to send traffic to its destination.

For an MPLS inter-AS environment, link protection can be enabled when **labeled-unicast** is used to send traffic between autonomous systems (ASs). To configure link protection on an interface, the **protection** statement is introduced at the **[edit protocols bgp group group-name family inet labeled-unicast]** hierarchy level.

- **LDP support (PTX1000)**—Starting in Release 17.2R1, Junos OS supports LDP on the PTX1000. The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP enables routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths. For more information, see the *MPLS Applications User Guide for Routing Devices*.
- **RSVP support (PTX1000)**—Starting in Release 17.2R1, Junos OS supports RSVP on the PTX1000. RSVP is a resource reservation setup protocol that is used by both network hosts and routers. Hosts use RSVP to request a specific class of service (CoS) from the network for particular application flows. Routers use RSVP to deliver CoS requests to all routers along the datapath. RSVP also can maintain and refresh states for a requested CoS application flow. For more information, see the *MPLS Applications User Guide for Routing Devices*.
- **ECMP (64-way) with configurable Layer 3 hash options (PTX1000)**—Starting in Release 17.2R1, Junos OS supports configuration of 64 equal-cost multipath (ECMP) next hops for RSVP and LDP LSPs on the PTX1000. To configure the maximum limit for ECMP next hops, include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level.

To view the details of the ECMP next hops, issue the **show route** command. The **show route summary** command also shows the current configuration for the maximum ECMP limit.

- **MPLS capabilities (PTX1000)**—Starting in Release 17.2R1, Junos OS supports MPLS capabilities on the PTX1000. MPLS provides both label edge router (LER) and label-switching router LSR and provides the following capabilities:
  - 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 FRR
  - 6PE and 6VPE devices
  - Layer 3 VPNs for both IPv4 and IPv6
- **IPv6 tunneling over an MPLS-based IPv4 network (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, tunneling enables you to connect IPv6 sites over an IPv4 MPLS-enabled backbone. IPv6 packets are carried over an IPv4 MPLS tunnel. To enable this service, you need to deploy provider edge (PE) routers that can run IPv4, MPLS, and BGP toward the core and IPv6 toward the edge.



[ See [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks](#)]

- **Egress peer engineering of service labels (such as BGP and MPLS) and egress peer protection for BGP-LU (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, you can enable traffic engineering of service traffic, such as MPLS LSP traffic between autonomous systems (ASs), by using BGP-labeled unicast for optimum utilization of the advertised egress routes. You can specify one or more backup devices for the primary egress AS boundary router. Junos OS installs the backup path in addition to the primary path in the MPLS forwarding table, which enables MPLS fast reroute (FRR) when the primary link fails. It provides support for the FRR protection backup scheme to perform an IP lookup to determine a new egress interface.

[See [Configuring Egress Peer Traffic Engineering by Using BGP Labeled Unicast and Enabling MPLS Fast Reroute.](#)]

### **Multicast**

- **Support for Internet multicast (PTX Series)**—Starting in Junos OS Release 17.2R1, the `mpls-internet-multicast` routing instance type uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP. Previously this feature was supported only on PTX Series routers with third-generation FPCs installed. Now this feature is supported when first-generation FPCs or second-generation FPCs are installed with third-generation FPCs on a PTX Series router.

[See [Multiprotocol BGP MVPNs Overview.](#)]

**NOTE:** For the third-generation FPCs to interoperate with the previous FPCs, the enhanced-mode statement cannot be configured on the chassis. To support Internet multicast, the MPLS core-facing interfaces must be third-generation FPCs.

### **Network Management and Monitoring**

- **Support for inline active flow monitoring (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, you can use the export capabilities of inline active flow monitoring with IP Flow Information Export (IPFIX) to define a flow record template suitable for IPv4 or IPv6 traffic. The flow record template provides the flexibility for future enhancements and the ability to add new attributes to inline active flow monitoring without changing to a newer version.
- **Support for SNMP (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, SNMP versions 1, 2, and 3 are supported on the PTX1000. SNMP enables you to monitor network devices from a central location. Junos OS includes an SNMP agent that provides remote management applications with access to detailed information about the devices on the network.
- **Junos Space Service Now (PTX1000)**—In Junos OS Release 16.1X65 and 17.2R1, PTX1000 routers support Junos Space Service Now. The Junos Space Service Now is an application that runs on the Junos Space Network Management Platform to automate fault management and accelerate issue resolution.

[ See [Junos Space Service Now.](#)]

- **Support for accounting profiles (PTX1000)**—Starting in Junos OS Release 17.2R1, you can configure accounting profiles to collect data on PTX Series routers that have third-generation FPCs installed.

**NOTE:** Configuring accounting profiles is supported only when the enhanced-mode statement is configured at the **[edit chassis network-services]** hierarchy level.

- **SNMP support for monitoring tunnel statistics (PTX Series)**—Starting in Junos OS Release 17.2R1, SNMP MIB jnxTunnelStat supports monitoring of tunnel statistics for IPv4 over IPv6 tunnels. This is a new enterprise-specific MIB, Tunnel Stats MIB, that currently displays three counters: tunnel count in rpd, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine. This MIB can be extended to support other tunnel statistics. The MIB is defined in jnx-tunnel-stats.txt. This MIB is attached to jnxMibs.

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

- **Hop-limit firewall filter match condition supported (PTX Series)**—Starting in Junos OS Release 17.2R1, you can configure a firewall filter using the hop-limit and hop-limit except match conditions for IP version 6 (IPv6) traffic (family inet6).

**NOTE:** The hop-limit and hop-limit except match conditions are supported on PTX series routers when [enhanced-mode](#) is configured on the router.

[See [Firewall Filter Match Conditions for IPv6 Traffic.](#)]

- **Support for firewall filter with match conditions (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, you can configure a firewall filter with match conditions for IP version 4 (IPv4) traffic.
- **Support for the no-decrement-ttl tunneling attribute (PTX1000)**—Starting in Junos OS Release 17.2R1, you can configure the **no-decrement-ttl** tunneling attribute for filter-based generic routing encapsulation (GRE) for IPv4 and IPv6 tunneling.

**NOTE:** The no-decrement-ttl tunneling attribute is supported only when the enhanced-mode statement is configured at the **[edit chassis network-services]** hierarchy level.

## Routing Protocols

- **Support for BGP link-state distribution with SPRING extensions (PTX Series)**—Starting in Junos OS Release 17.2R1, BGP link-state extensions export segment routing topology information to software-defined networking controllers. Controllers can get the topology information by either being a part of an interior gateway protocol (IGP) domain or through BGP link-state distribution. BGP link-state distribution is supported on inter-domain networks and provides a scalable mechanism to export the topology information. This feature benefits networks that are moving to source packet routing in networking (SPRING) but also have RSVP deployed, and continue to use both SPRING and RSVP in their networks.

[See [Link-State Distribution Using BGP Overview](#).]

- **Support for SRGB in SPRING for IS-IS (PTX Series)**—Starting with Junos OS Release 17.2R1, you can configure the segment routing global block (SRGB) range label used by source packet routing in networking (SPRING). Currently Junos OS allows you to configure only node segment indices. The value of the start label depends on the dynamic label available in the system. The labels from this SRGB range are used for SPRING in the IS-IS domain. The labels advertised are more predictable and deterministic across the segment routing domain.

- To configure the starting index value of the SRGB label block, use the **start-label start-label-block-value** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.
- To configure the index range of the SRGB label block, use the **index-range value** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.

[See [source-packet-routing](#).]

- **Support for anycast and prefix segments in SPRING for IS-IS protocols (PTX Series)**—Starting in Junos OS Release 17.2R1, there is support for anycast segment identifiers (SIDs) and prefix SIDs in source packet routing in networking (SPRING). Currently there is support for node segments in Junos OS supports node segments for IPv4 and IPv6 when they are explicitly configured under the **[edit protocols isis source-packet-routing node-segments]** hierarchy. Now you can provision prefix SIDs along with node SIDs to prefixes that are advertised in IS-IS protocols through policy configuration. Anycast SID is a prefix segment that identifies a set of routers. You can configure **explicit-NULL** flag on all prefix SID advertisements and configure **shortcut** for SPRING routes using **family inet-mpls** or **family inet6-mpls**.

[See [Support for SRGB, Anycast, and Prefix Segments in SPRING for IS-IS Protocol](#).]

- **Support for unique AS path count (PTX Series)**—Starting with Junos OS Release 17.2R1, you can configure a routing policy to determine the number of unique autonomous systems (ASs) present in the AS path. The unique AS path count helps determine whether a given AS is present in the AS path multiple times, typically as prepended ASs. In earlier Junos releases it was not possible to implement this counting behavior using the **as-path** regular expression policy. This feature permits the user to configure a policy based on the number of AS hops between the route originator and receiver. This feature ignores ASs in the **as-path** that are confederation ASs, such as **confed\_seq** and **confed\_set**.

To configure AS path count, include the **as-path-unique-count** *count (equal | orhigher | orlower)* configuration statement at the **[edit policy-options policy-statement policy\_name from]** hierarchy level.

- Support for IS-IS segment routing on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, IS-IS segment routing support is enabled through MPLS. Currently, label advertisements are supported for IS-IS only. IS-IS creates an adjacency segment per adjacency, per level, and per address family (one each for IPv4 and IPv6). Junos OS IS-IS implementation allocates node segment label blocks in accordance with the IS-IS protocol extensions for supporting segment routing node segments and provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:
  - source-packet-routing**—Enable the source packet routing feature.
  - node-segment**—Enable source packet routing at all levels.
  - use-source-packet-routing**—Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and primary IS-IS source packet routing node segments.
  - no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.
- BGP advertises multiple add-paths based on community value (PTX1000)**—Beginning with Junos OS 17.2R1, you can define a policy to identify eligible multiple path prefixes based on community values. BGP advertises these community-tagged routes in addition to the active path to a given destination. If the community value of a route does not match the community value defined in the policy, then BGP does not advertise that route. This feature allows BGP to limit the number of multiple paths that are processed and not advertise more than 20 paths to a given destination. You can limit and configure the number of prefixes that BGP considers for multiple paths without actually knowing the prefixes in advance. Instead, a known BGP community value determines whether or not a prefix is advertised.
- Selective advertising of BGP multiple paths (PTX1000)**—Beginning with Junos OS Release 17.2R1, you can restrict BGP **add-path** to advertise contributor multiple paths only. Advertising all available multiple paths might result in a large overhead of processing on device memory and is a scaling consideration, too. You can limit and configure upto six prefixes that the BGP **multipath** algorithm selects. Selective advertising of multiple paths facilitates Internet service providers and data centers that use route reflector to build in-path diversity in IBGP.
- Support for BGP to carry flow-specification routes (PTX1000)**--Starting in Junos OS Release 17.2R1, BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX1000 routers that have third-generation FPCs installed. Propagating flow routes as BGP NLRI messages in essence enables the propagation of firewall filters which protects the system against denial-of-service (DOS) attacks.

## Security

- **Firewall filter support (PTX10008)**—Starting in Junos OS Release 17.2R1, you can define firewall filters on the PTX10008 routers 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.

## Services Applications

- **Support for inline J-Flow version 9 flow templates (PTX5000 and PTX3000)**—Starting in Junos OS Release 17.2R1, you can use inline J-Flow export capabilities with version 9 flow templates to define a flow record template suitable for IPv4 or IPv6 traffic.

[See [Configuring Flow Aggregation to Use Version 9 Flow Templates on PTX Series Routers](#).]

## Software Installation and Upgrade

- **Zero Touch Provisioning (PTX1000)**—Starting in Junos OS Release 17.2R1, ZTP is supported to automate the provisioning of the device configuration and software image with minimal manual intervention.

When you physically connect a router to the network and boot it with a default configuration, the router attempts to upgrade the Junos OS software image automatically and autoinstall a configuration file from the network. The router uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the router boots with the pre-installed software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version.

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

## User Interface and Configuration

- **Monitoring, detecting, and taking action on degraded physical 100-Gigabit Ethernet links to minimize packet loss (PTX1000)**—Starting with Junos OS Release 17.2R1, you can monitor physical link degradation (indicated by bit error rate (BER) threshold levels) on Ethernet interfaces, and take corrective actions if the BER threshold value drops to a value in the range of  $10^{-13}$  to  $10^{-5}$ .

The following new configurations have been introduced at the `[edit interfaces interface-name]` hierarchy level to support the physical link degrade monitoring and recovery feature on Junos OS:

- To monitor physical link degrade on Ethernet interfaces, configure the **link-degrade-monitor** statement.
- To configure the BER threshold value at which the corrective action must be triggered on or cleared from an interface, use the **link-degrade-monitor thresholds (set value | clear value)** statement.
- To configure the link degrade interval value, use the **link-degrade-monitor thresholds interval value** statement. The configured interval value determines the number of consecutive link degrade events that are considered before any corrective action is taken.
- To configure link degrade warning thresholds, use the **link-degrade-monitor thresholds (warning-set value | warning-clear value)** statement.

- To configure the link degrade action that is taken when the configured BER threshold level is reached, use the **link-degrade action media-based** statement.
- To configure the link degrade recovery options, use the **link-degrade recovery (auto interval value | manual)** statement.

You can view the link recovery status and the BER threshold values by using the **show interfaces interface-name** command.

## VPNs

- **Layer 3 VPN support (PTX1000)**—Starting in Release 17.2R1, Junos OS supports Layer 3 VPN on the PTX1000 router. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing Internet backbone.

In Junos OS, Layer 3 VPNs are based on RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*. This RFC defines a mechanism by which service providers can use their IP backbones to provide Layer 3 VPN services to their customers.

[See [Understanding Layer 3 VPNs.](#)]

## SEE ALSO

[Changes in Behavior and Syntax | 294](#)

[Known Behavior | 299](#)

[Known Issues | 301](#)

[Resolved Issues | 305](#)

[Documentation Updates | 312](#)

[Migration, Upgrade, and Downgrade Instructions | 312](#)

[Product Compatibility | 317](#)

## Changes in Behavior and Syntax

### IN THIS SECTION

- [Forwarding and Sampling | 295](#)
- [General Routing | 295](#)
- [Interfaces and Chassis | 295](#)

- Management | 296
- Network Management and Monitoring | 297
- Routing Protocols | 298
- Subscriber Management and Services | 298

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 17.2R3 for the PTX Series.

## Forwarding and Sampling

- In Junos OS Release 17.2R2, and later, the **SelectorID** field (element id: 302) is sent instead of the **Bytes** field (element id: 1) in the system scope of **version-ipfix** Option template records for all PTX Series Routers. All other elements of the template remain the same.

## General Routing

- **Support for deletion of static routes when the BFD session goes down (PTX Series)**—Starting with Junos OS 17.2R2, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session-down message.

## Interfaces and Chassis

- **Value of *sysObjectID* now displays *jnxProductNamePTX1000* (PTX1000)**—Starting in Junos OS Release 17.2R1, the value of *sysObjectID* is now displayed as *jnxProductNamePTX1000* instead of *jnxProductPTX1000* (which is an incorrect value), as shown in the following example:

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNamePTX1000
```

The *sysObjectID* value is updated to *jnxProductNamePTX1000* to maintain synchronization across devices (or routers) belonging to the PTX Series.

- **Change in command outputs after a health check failure (PTX5000)**—Starting in Junos OS Release 17.2R1, when a health check fail for a PSM is detected on a PTX5000 router, until a system reboot or restart chassisd occurs, the following changes are displayed in the command outputs:

- The output of the **show chassis environment pdu** displays the reason for the health check fail and the following information:

```
Health Check FAILED for PSM PSM_Number
```

```
PSM_Number is Present|Not OK
```

- The status of the PSM which failed the health check is set to offline and the output of **show chassis alarms** command displays the following existing alarm:

```
PDU slot PSM PSM_Number Not OK
```

After a system reboot or restart chassisd, the router checks the PSM register 0x1D bit-0:

- The output of the **show chassis environment pdu** displays the reason for the health check fail and the following information for the PSM:

```
PSM_Number is Present|Not OK
```

- **Restart FPC option supported on PTX1000 router**—In Junos OS Release 17.2R2, you can reboot the FPC gracefully using **request chassis fpc restart slot slot-number** command on a PTX1000 router. Note that **request chassis fpc (online|offline) slot slot-number** command is not supported, which means only restart option is supported, but online and offline options are not supported. See [\[request chassis fpc.\]](#)

## Management

- **Junos OS YANG module namespace and prefix changes (PTX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. In earlier releases, Junos OS YANG modules used only a unique identifier to differentiate the namespace for each module, and the prefix for all **juniper-command** modules was **jrpc**.

Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**. The Junos OS YANG extension modules, **junos-extension** and **junos-extension-odl**, use the **junos** device family identifier in the namespace, but the modules are common to all device families.

[See [Understanding Junos OS YANG Modules.](#)]

- **Changes to the rfc-compliant configuration statement (PTX Series)**—Starting in Junos OS Release 17.2R1, Junos OS YANG modules are specific to a device family, and each module's namespace includes the



module name, device family, and Junos OS release string. If you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level and request configuration data in a NETCONF session on a device running Junos OS Release 17.2R1 or later, the NETCONF server sets the default namespace for the **<configuration>** element in the RPC reply to the same namespace as in the corresponding YANG model.

[See [Configuring RFC-Compliant NETCONF Sessions](#) and [rfc-compliant](#).]

- **Enhancement to the Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.2R1, the values displayed in the **oper-status** key-value field of data streamed through gRPC for the physical interfaces sensor have changed.

The following values are now displayed to indicate the operational status of an interface:

- operational status up—**UP**
- operational status down—**DOWN**
- operational status unknown—**UNKNOWN**
- **Enhancement to NPU memory sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.2R1, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is **/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/**

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

## Network Management and Monitoring

- **SNMP syslog messages changed (PTX Series)**—In Junos OS Release 17.2R1, 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](#).]

- **Update to SNMP support of apply-path statement (PTX Series)**—In Junos OS Release 17.2R1, the SNMP implementation for the **apply-path** configuration statement supports only two lists:
  - **apply-path "policy-options prefix-list <list-name> <\*>"**  
This configuration has been supported from day 1.
  - **apply-path "access radius-server <\*>"**  
This configuration is supported as of this release.

- **Need to reconfigure SNMPv3 configuration after upgrade (PTX Series)**—In Junos OS Release 17.2R2, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. Reconfiguration is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. It used to be that customers had to reconfigure SNMPv3 every time after a reboot. That problem is now fixed. If you upgrade, you must still reconfigure SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3. To reconfigure SNMPv3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters.

[See [Configuring the Local Engine ID.](#)]

## Routing Protocols

- **Syslog error message RPD\_ISIS\_PREFIX\_SID\_CNFLCT to resolve conflicting prefix segment advertisement (PTX Series)**—Starting in Junos OS Release 17.2R2, the **RPD\_ISIS\_PREFIX\_SID\_CNFLCT** syslog error message is emitted only when the prefix segment advertisement from the remote node is conflicting with an advertisement from the self node. This conflict happens because the same prefix segment index is assigned on different IP addresses or different prefix segment indexes are assigned to the same IP address. To rectify this conflict, identify the remote node in the network originating the conflicting prefix segment advertisement and change the prefix segment index on the local node or on the remote node.

[See [Example: Configuring Anycast and Prefix Segments in SPRING for ISIS.](#)]

## Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (PTX Series)**—Starting in Junos OS Release 17.2R3, 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 | 277](#)

[Known Behavior | 299](#)

[Known Issues | 301](#)

[Resolved Issues | 305](#)

[Documentation Updates | 312](#)

[Migration, Upgrade, and Downgrade Instructions | 312](#)

[Product Compatibility | 317](#)

## Known Behavior

### IN THIS SECTION

- [Hardware | 300](#)
- [High Availability \(HA\) and Resiliency | 300](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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.

# Hardware

- **Enhanced resiliency and system snapshot (PTX1000)**—The 2 x 64-GB SSDs installed in the PTX1000 support the **request vmhost snapshot** command, which creates a recovery snapshot of the currently running and active file system partitions, and the **request vmhost snapshot recovery** command, which recovers the primary disk from the snapshot content stored in the backup disk. In addition, the 64-GB SSDs support enhanced hardware resiliency through storage partitioning and redundancy.

Earlier versions of the PTX1000 have 2 x 32-GB M.2 SATA SSDs. PTX1000 routers with 32-GB SSDs do not support the **request vmhost snapshot** and **request vmhost snapshot recovery** commands, and do not support enhanced hardware resiliency. To determine the size of the SSDs installed in your device, issue the **show vmhost hardware** CLI command. The capacity of **Disk1** and **Disk2** is displayed in the output as **32.0 GB** if 32-GB SSDs are installed, and the capacity is displayed as **50.0 GB** if 64-GB SSDs are installed.

[See the [Junos OS Software Installation and Upgrade Guide](#).]

# High Availability (HA) and Resiliency

- **Residual and baseline statistics loss from ISSU**—Using unified ISSU to upgrade to Junos OS Release 17.2R1 or later will result in a loss of residual and baseline statistics for interfaces, interface set specific statistics, and BBE subscriber service statistics because of an update to the statistics database.

[See [Unified ISSU System Requirements](#).]

- **ISSU restrictions**—Unified ISSU is not supported for upgrading Junos OS 17.2R1 to 17.2R2.

# SEE ALSO

<a href="#">New and Changed Features   277</a>
<a href="#">Changes in Behavior and Syntax   294</a>
<a href="#">Known Issues   301</a>
<a href="#">Resolved Issues   305</a>
<a href="#">Documentation Updates   312</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   312</a>
<a href="#">Product Compatibility   317</a>

## Known Issues

### IN THIS SECTION

- [General Routing | 301](#)
- [Interfaces and Chassis | 304](#)
- [Platform and Infrastructure | 304](#)
- [Routing Protocols | 305](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R3 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

- While you upgrade from Junos OS Release 15.1F-based images to Junos OS Release 16.x and later releases or downgrading from Junos OS Release 16.x to Junos OS Release 15.1F images, if the **validate** option is enabled, the chassis process (chassisd) might crash and the upgrade or downgrade might fail. This issue is not seen if both base and target images are from Junos OS Release 15.1F or Junos OS Release 16.x and later. [PR1171652](#)
- Occasionally QSFP28+ interfaces can run into clock stretch and will be disabled with an **i2c-accel sync access failed** error message. [PR1181493](#)
- This is a resiliency feature. If more than 10 FO CRC errors are seen in an interval of 30 seconds, then CMERROR infra raises an alarm and an appropriate action is taken. [PR1197865](#)
- On MX Series routers with MPC7E, MPC8E, MPC9E and on PTX Series routers with FPC3-PTX-U2/FPC3-PTX-U3, which could lead to the generation of core files. It is hard to reproduce. The interrupt code is optimized to avoid the unnecessary call to prevent the issue. [PR1208536](#)
- A PTX Series third-generation FPC might receive noise on the FPC console port and interpret it as valid signals. This might cause a login failure on the console port and generate core files, or even reloads. [PR1224820](#)
- On a PTX Series router with a faulty power supply module (PSM), the PSM might generate excessive interrupt requests. Because hardware interrupt requests are processed by the chassisd, excessive interrupt requests might cause chassisd to restart when the condition persists more than 200 seconds. [PR1226992](#)

- PTX Series with the FPC-PTX-P1-A or FPC2-PTX-P1A FPC might encounter a single event upset (SEU) event, which can cause a linked-list corruption of the TQ CHIP. The following syslog message is reported.  
**Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt\_min\_free\_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero Jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002, Junos OS Chassis Management Error handling does detect such a condition, raises an Alarm, and disables affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, a restart of the FPC is needed. Soft errors are transient or non-recurring. FPCs experiencing such SEU events do not have any permanent damage. Contact your Juniper support representative if the issue is seen after an FPC restart. [PR1254415](#)**
- On rare occasions, upon reboot, the kernel cannot create sysfs entries for the SSDs in the system. This might result in the system entering panic mode and hanging. [PR1261068](#)
- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are seen on the destination FPC. A corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error. [PR1268678](#)**
- On PTX5000 with a third-generation FPC in rare condition, the FPC might crash during lo0.0 inet6 input filter. [PR1268875](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- Interfaces might go down when the Packet Forwarding Engine encounters **TOE::FATAL ERROR** (TOE is a module in Packet Forwarding Engine, and the fatal error can be caused either by a software issue or hardware issues such as memory parity errors or others). Reboot the line card to recover the service when you experience the issue. [PR1300716](#)
- On PTX10000 Series routers with FPC LC1101 - 30C / 30Q / 96X installed, the 10-Gigabit Ethernet interface might flap when the interface is active and it is set to 100-Gbps speed. [PR1315079](#)
- Over a period of time, network events such as route flaps and MBB scenario cause the Packet Forwarding Engine heap memory to get fragmented. This change improves memory management and reduces the chance of memory fragmentation. [PR1318595](#)
- This is an expected behavior for TQ-chip ASICs. It is primarily due to the strict-high priority queue and the shared shaper. Credits that are unused by an output queue (that is, the actual rate for the queues is less than transmit rates) will cause the queue's credit bucket to hit its maximum value. When a queue hits its maximum credit value, the remaining credits will be distributed to other queues. After the other queues get transmit credits, they can start transmitting. Thus, with the TQ chip and the shared shaper

it is virtually impossible to completely shut off a queue through the guaranteed rate mechanism.

[PR1319923](#)

- On PTX Series routers with first-generation and second-generation FPCs and with CoS used, a high-priority queue might not get the entire configured bandwidth. [PR1324853](#)
- PTX3000 reports Chip to Chip Link (CCL) CRC errors while FPC3-SFF-PTX-1X is taken offline through a 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:0x00000037649c2c43e 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: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-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@PTX3000> 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](#)

- When NSR is configured, packets might be dropped because of reverse path forwarding (RPF) during Routing Engine switchover. [PR1354285](#)
- If firewall filter is configured, in a rare condition, the host interface might be wedged on a PTX Series router with third-generation FPCs. [PR1354580](#)
- Traffic loss is greater than 50 ms (around 200 to 300 ms) for IP routes pointing to a unilist of composites with indirect next hops in a link-down scenario . In this case, the Packet Forwarding Engine does not perform the local repair and waits for the rpd to install the new next hops. [PR1383965](#)

## Interfaces and Chassis

- 1. Delay Measurement support for 5-port 100G DWDM PIC and 5-port 100G DWDM MIC is \*ONE TIME Delay Measurement\*. This means that customers intending to measure Delay 2 points must ensure that the link is up on both sides and then conduct this test one time. The result value is valid one time once the test is finished. The test result on CLI is not valid after one time measurement as the old result might show up on the Routing Engine CLI. 2. **remote-loop-enable** must be configured first on the remote end. 3. Each time a customer wants to verify this, the test has to be repeated. 4. Processing delays in each mode are different: HG-FEC [For 5-port 100G DWDM MIC] being highest, SDFEC in the interim and GFEC being least for the same cable length. 5. In summary, any breakage in Transmit or Receive path during the delay measurement test will hinder delay measurement. This is true for all FEC modes - GFEC, SD-FEC, HGFECC. 6. Currently SNMP walk is not available for delay measurement. [PR1233917](#)
- Junos OS upgrade involving Junos OS Release 14.2R5 (and later 14.2 maintenance releases) and Junos OS Release 16.1 (and later mainline releases) with CFM configuration can cause the cfmcmd to crash after the upgrade. This is due to an old version of `/var/db/cfm.db`. [PR1281073](#)

## Platform and Infrastructure

- In scaled FIB setups, IS-IS graceful restart might abort on the restarting node with T3-timer expiry log, because of hold-time expiry on IS-IS GR-helper peers. May 20 01:22:55.992972 T3 Restart timer expired (graceful restart aborted). This is because of the time taken by the rpd to learn routes from the FIB on rpd startup that were installed by a previous instance of the rpd. IS-IS does not get to initialize and send hellos because of this delay, causing holdtimer-expiry at helpers. {master}[edit] user@router# **run show route forwarding-table summary | match user** May 20 03:07:26 user: 801650 routes user: 403638 routes user: 6 routes. [PR1277933](#)
- Every load override and rollback operation increases the refcount by 1, and after it reaches the maximum value of 65,535, an mgd crash might be observed and the session might get terminated. When mgd crashes, the active lock might remain up preventing any further commits. [PR1313158](#)



Routing Protocols

- With Shared Risk Link Group (SRLG) enabled under corner conditions, after executing the command **clear isis database** is executed, the rpd might crash because the IS-IS database tree gets corrupted.  
[PR1152940](#)

SEE ALSO

<a href="#">New and Changed Features</a>	<a href="#">  277</a>
<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  294</a>
<a href="#">Known Behavior</a>	<a href="#">  299</a>
<a href="#">Resolved Issues</a>	<a href="#">  305</a>
<a href="#">Documentation Updates</a>	<a href="#">  312</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  312</a>
<a href="#">Product Compatibility</a>	<a href="#">  317</a>

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R3](#) | [306](#)
- [Resolved Issues: 17.2R2](#) | [310](#)
- [Resolved Issues: 17.2R1](#) | [311](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

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

## Resolved Issues: 17.2R3

### General Routing

- PTX1000 : `ch_get_product_attribute.324`: Cannot find chassisd error message occurs when images are loaded. [PR1217505](#)
- On PTX Series routers, chassisd thread does not get CPU resources for 200 seconds, and, as a result, multiple chassisd core files are generated. [PR1226992](#)
- The **validation-state:unverified** routing entry might not be shown with proper location in the output of the **show route** command. [PR1254675](#)
- Error messages might be seen on PTX5000 FPC3 with 48x10G/12x40G(LWO)QSFP+ PIC inserted. [PR1273575](#)
- 100GBase-ER4 (740-045420) is shown as **UNKNOWN** when the CLI command **show chassis hardware** is executed in Junos OS Release 15.1R5. [PR1280089](#)
- FPC cards might go offline because of fabric healing in a PTX3000 router with a SIB-SFF-PTX-240-S. [PR1282983](#)
- The PTX SPMB might crash after an FPC replacement is followed by a SIB restart. [PR1283553](#)
- Periodic export of IPFIX flow packets with high octet values. [PR1286427](#)
- LSP traffic might silently drop and get discarded after a link goes down in the bypass path. [PR1291036](#)
- The rpd might generate a core file while restarting. [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](#)
- The rpd might crash after interface or BGP flapping. [PR1294957](#)
- The chassisd might run out of memory and restart on a PTX1000 router. [PR1295691](#)
- On a PTX5000 or on an Ethernet Synchronization Message Channel (ESMC), the clock does not get locked when the source interface is a member link of an aggregated Ethernet bundle. [PR1296015](#)
- 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 when sampling is enabled. [PR1296533](#)
- Alarms and syslog errors are seen with the strict-high priority 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 after hybrid memory cube (HMC) fatal errors might have a system-wide impact on PTX Series routers. [PR1300180](#)
- PTX Series FPC3 drops MPLS packets when the maximum transmission unit is less than the MPLS packet size on the outgoing interface with IPv4 traffic. [PR1302256](#)
- Heap memory leak might be observed on PTX Series FPCs during a multicast route installation into the Packet Forwarding Engine. [PR1302303](#)
- On a PTX3000, the powering of a FPC (PTX-IPLC-B-32) card might cause the other FPCs to reboot. [PR1302304](#)
- The third-generation FPC (FPC3-SFF-PTX) might not boot on a PTX3000 router that has the Control Board or Routing Engine. [PR1303295](#)
- On PTX3000 and PTX5000 routers, the 100-Gigabit Ethernet interfaces might not come up. [PR1303324](#)
- 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](#)
- Repeated log messages `%PFE-3 fpcX expr_nh_index_tree_ifl_get` and `expr_nh_index_tree_ipaddr_get` are observed when sampled packets are discarded because of the configured firewall filter **log** (or **syslog**) statement. [PR1304022](#)
- PTX3000 with the RCB-PTX-X6-32G Routing Engine might be unable to recognize the integrated photonic line card (IPLC) or bring it online. [PR1304124](#)
- Some error messages might be observed on an EVPN-VXLAN setup. [PR1307014](#)
- The **interface hold-time down** timer configuration does not work properly on a PTX5000 router with an optical interface. [PR1307302](#)
- PTX10000: Don't bounce FPC without warning or alarm for different port speed settings. [PR1311875](#)
- The rpd generated core files observed after multiple session flaps on a scaled setup. [PR1312169](#)
- Too many logs are generated after many VM-host-client-related commands are related. [PR1315128](#)
- The RIB and FIB might get out of synchronization because the KRT asynchronous queue is stuck. [PR1315212](#)
- The Packet Forwarding Engine on a third-generation FPC on a PTX3000 or PTX5000 or on a PTX10000 line card might be disabled if the interface connecting to the Packet Forwarding Engine goes down. [PR1315823](#)
- The physical interfaces might generate framing errors when ports are connected to an odd interface. [PR1317827](#)
- After an FPC is powered on, (or connected to or disconnected from the network), the output of the `show chassis hardware` command shows No Power. However, the FPC comes up after some time. [PR1319156](#)
- The rpd might crash when the OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)

- JSA10864 2018-07 Security Bulletin: Junos OS: MPC7/8/9, PTX-FPC3 (FPC-P1, FPC-P2), PTX3K-FPC3 and PTX1000: The line card might crash when it received specific MPLS packets (CVE-2018-0030). [PR1323069](#)
- On PTX1000, the local time on an FPC might be different from the local time on Junos OS VM or VM host. [PR1325048](#)
- PTX MKA sessions are not coming up, after CA parameters such as **transmit-interval**, **key-server-priority** are changed. [PR1325392](#)
- MPLS traceroute fails across PTX Series routers. [PR1327609](#)
- On PTX5000 with third-generation FPCs, PTX10000, and PTX1000 routers, output firewall filters that are configured with **syslog** and **discard** actions do not perform the **syslog** action. [PR1328426](#)
- The line card on a PTX10000 might reboot continuously if HMC BIST fails after you upgrade to Junos OS Release 17.2R1 or later. [PR1330618](#)
- Link instability is experienced after a link-down event on a PTX Series device. [PR1330708](#)
- Next-hop programming issue during link flapping on PTX Series routers. [PR1333274](#)
- PTX5000 FPCs might reboot in certain rare scenarios when an interface-specific policer is configured. [PR1335161](#)
- A member link of an IPv4 unicast next hop might be stuck in **Replaced** state after the interface flaps. [PR1336201](#)
- Disabling a breakout 10-Gigabit Ethernet port on et-0/0/5 unexpectedly disables another breakout 10-Gigabit Ethernet port on et-0/0/5. [PR1337975](#)
- The FPC, FPC2, FPCE on PTX Series does not forward traffic. [PR1339524](#)
- Link goes down on PTX3000 or PTX5000 with FPC3 inserted after the router reboots or link flapping occurs. [PR1340612](#)
- The interface might flap continuously after the device reboots. [PR1342681](#)
- MPLS traceroute for P2MP LSPs configured with link protection causes the FPC to crash. [PR1348314](#)
- BFD sessions do not come up on a PTX3000. [PR1352112](#)
- The interfaces on the 15-port 100-Gigabit Ethernet PIC might come up after a delay of around 60 seconds. [PR1357410](#)
- The route might be in a stuck state, and the route installation might fail with traffic loss, after the BGP neighbor and the route experience flapping. [PR1362560](#)
- The rpd might crash in a large-scale environment. [PR1363803](#)
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with the "Link-Layer-Down" flag set. [PR1365263](#)
- Layer 3 VPN traffic is dropped because the selector weight was set to 65,535 after one core-facing interface was down. [PR1380783](#)

### **Infrastructure**

- The PTX Series router might be in an abnormal state because of the malfunction of the protection mechanism for F-Label. [PR1336207](#)

### **Interfaces and Chassis**

- Interface flapping occurs during Routing Engine switchover if the member links of an aggregated Ethernet interface are configured with framing settings. [PR1287547](#)
- 100-Gigabit Ethernet interfaces might not come up when **otn-options laser-enable** is configured on PTX Series routers. [PR1297164](#)

### **MPLS**

- The rpd might crash on backup Routing Engine because of memory exhaustion. [PR1328974](#)
- The rpd might crash when MPLS traceoption is configured. [PR1329459](#)
- MPLS LSP statistics are not shown in the output of the **show mpls lsp ingress statistics** command. [PR1344039](#)
- Some LSPs might be stuck on the upstream devices after interfaces flap on downstream devices. [PR1349157](#)

### **Platform and Infrastructure**

- Continuous log messages are displayed. For example: **tftpd[23724]: Timeout #35593** on DATA block 85. [PR1315682](#)
- Running the request support information (RSI) command through console port might cause system crash and reboot. [PR1349332](#)
- A traffic black-hole condition occurs along with the output of **JPRDS\_NH:jprds\_nh\_alloc(),651: JNH[0]** failed to grab new region for next hop messages. [PR1357707](#)
- Next hop index allocation failed: private index space exhausted through incoming ARP requests to management interface (CVE-2018-0063). [PR1360039](#)

### **Routing Protocols**

- A few BFD sessions flap while coming up after the FPC restarts or reboots. [PR1274941](#)
- The rpd crashes and generates core files multiple times when you receive an OPEN message from an existing BGP peer. [PR1299054](#)
- With BGP LU FRR in an inter-AS scenario, a very high FRR time is visible after the link is up. [PR1307258](#)
- The rpd might constantly consume high CPU in a BGP setup. [PR1315066](#)
- The primary path of an MPLS LSP might switch to another address. [PR1316861](#)
- The rpd might crash after the passive interface is deactivated under IS-IS. [PR1318180](#)
- The rpd might crash if SRLG information is in the IS-IS protocol. [PR1337849](#)

## 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 Layer 2 circuit flaps repeatedly. [PR1282875](#)

## Resolved Issues: 17.2R2

### General Routing

- The **request vmhost zeroize** and **request vmhost zeroize both** commands might work only on the local Routing Engine. [PR1197152](#)
- User-configured TPID is not being applied on a single-tagged VLAN interface. [PR1237687](#)
- An FPC major alarm might be seen with the error messages **DLU: ilp memory cache error** and **DLU: ilp prot1 detected\_imem\_even error**. [PR1251154](#)
- PTX1000 does not match an outer tag if an inner tag exists. [PR1252443](#)
- The kernel log message **mastership: sent other Routing Engine mastership loss signal** might be printed on the backup Routing Engine of the PTX5000 router. [PR1260884](#)
- Sometimes SDN-Telemetry subsystem does not respond to management requests while issuing **show agent sensors**. [PR1266058](#)
- Graceful restart for FPC is provided on PTX1000. [PR1266097](#)
- SPMB ukern panics during ASIC error recovery. [PR1268253](#)
- The log message **sdk-vmmd: %USER-3: is\_platform\_Next-Gen RE: Platform found as Next-Gen RE** is logged with error severity. [PR1271134](#)
- MPLS TTL is reset to 255 on third generation PTX FPCs when the **protocols mpls no-propagate-ttl** command is configured. [PR1287473](#)

### Infrastructure

- The **show system users** CLI output displays more users that are not using the router. [PR1247546](#)

### Layer 2 Ethernet Services

- Messages **l2cpd[2486]: task\_connect: task MVRP l2ald ipc./var/run/l2ald\_control addr /var/run/l2ald\_control: No such file or directory** is filling up syslog. [PR1278189](#)

### MPLS

- The rpd might crash if the MPLS LSP path changes. [PR1295817](#)

### Routing Protocols

- The rpd might crash on platforms with 64-bit X86 Routing Engine, if IPv6 is configured. [PR1224376](#)

## Resolved Issues: 17.2R1

### General Routing

- Junos Telemetry Interface: Frequent disconnections are seen in MQTT when the logical interface sensor is provisioned for a longer duration. [PR1238803](#)
- On PTX Series platform, add 'set' parameter (optional) to CLI command **request system software add**. It provides a way to install multiple software packages and software add-on packages at the same time. [PR1246675](#)
- 10-Gigabit Ethernet interfaces on a QSFP28 PIC might not come up after a system reboot or a PIC restart. [PR1263413](#)
- An incorrect range of voltages is used for proper PE voltages. [PR1263675](#)

### MPLS

- The rpd process terminates and generates a core file if there are a large number of RSVP LSPs. [PR1257367](#)

### SEE ALSO

[New and Changed Features | 277](#)

[Changes in Behavior and Syntax | 294](#)

[Known Behavior | 299](#)

[Known Issues | 301](#)

[Documentation Updates | 312](#)

[Migration, Upgrade, and Downgrade Instructions | 312](#)

[Product Compatibility | 317](#)

## Documentation Updates

### IN THIS SECTION

- [Protocol-Independent Routing Properties | 312](#)

This section lists the errata and changes in Junos OS Release 17.2R3 documentation for the PTX Series.

### Protocol-Independent Routing Properties

- **Support for deletion of static routes when the BFD session goes down (PTX Series)**—Starting with Junos OS Release 17.2R2, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

### SEE ALSO

---

[New and Changed Features | 277](#)

---

[Changes in Behavior and Syntax | 294](#)

---

[Known Behavior | 299](#)

---

[Known Issues | 301](#)

---

[Resolved Issues | 305](#)

---

[Migration, Upgrade, and Downgrade Instructions | 312](#)

---

[Product Compatibility | 317](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.2 | 313](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 315](#)



- Upgrading Using Unified ISSU | 316
- Upgrading a Router with Redundant Routing Engines | 316

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

### Basic Procedure for Upgrading to Release 17.2

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

**NOTE:** Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

**NOTE:** The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [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 17.2R3:

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:** We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-17.2R3.9.tgz
```

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-17.2R3.9-limited.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://hostname/pathname**
  - **http://hostname/pathname**
  - **scp://hostname/pathname** (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 17.2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

## Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases

provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 16.1, 16.2, and 17.1 are EEOL releases. You can upgrade from Junos OS Release 16.1 to Release 16.2 or even from Junos OS Release 16.1 to Release 17.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

## Upgrading Using Unified ISSU

Unified 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. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [Understanding High Availability Features on Juniper Networks Routers](#).

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

## Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

SEE ALSO

New and Changed Features   277
Changes in Behavior and Syntax   294
Known Behavior   299
Known Issues   301
Resolved Issues   305
Documentation Updates   312
Product Compatibility   317

# Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 317

## 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   277
Changes in Behavior and Syntax   294
Known Behavior   299

---

[Known Issues | 301](#)

---

[Resolved Issues | 305](#)

---

[Documentation Updates | 312](#)

---

[Migration, Upgrade, and Downgrade Instructions | 312](#)

## Junos OS Release Notes for the QFX Series

### IN THIS SECTION

- [New and Changed Features | 318](#)
- [Changes in Behavior and Syntax | 338](#)
- [Known Behavior | 344](#)
- [Known Issues | 347](#)
- [Resolved Issues | 351](#)
- [Documentation Updates | 363](#)
- [Migration, Upgrade, and Downgrade Instructions | 364](#)
- [Product Compatibility | 376](#)

These release notes accompany Junos OS Release 17.2R3 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## New and Changed Features

### IN THIS SECTION

- [Release 17.2R3 New and Changed Features | 319](#)
- [Release 17.2R2 New and Changed Features | 319](#)
- [Release 17.2R1 New and Changed Features | 319](#)

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 17.2R3: QFX5100, QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016.

## Release 17.2R3 New and Changed Features

### *Restoration Procedure Failure*

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (QFX Series)**—In Junos OS Release 17.2R3, 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 17.2R2 New and Changed Features

- There are no new features or enhancements to existing features for QFX Series in Junos OS Release 17.2R2.

## Release 17.2R1 New and Changed Features

### *Hardware*

- **QFX5110-32Q**—The QFX5110 line of switches is Juniper Network's versatile fixed-configuration solution for hybrid cloud deployments. The model QFX5110-32Q is a flexible configuration switch allowing either 32 ports of 40-Gigabit Ethernet quad small form-factor pluggable plus (QSFP+) or 20 ports of QSFP+ and 4 ports of high-density 100-Gigabit Ethernet quad small form-factor pluggable solution (QSFP28). Each QSFP+ port can operate as a native 40-Gigabit Ethernet port, or as four independent 10-Gigabit ports when using breakout cables. The four QSFP28 ports are available either as access ports or as uplinks. The QFX5110-32Q provides full duplex throughput of 960 Gbps. The QFX5110-32Q has a 1 U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow. The model is available with either AC or DC power supplies.

- **QFX10000-60S-6Q Line Card (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.2R1, QFX10000-60S-6Q line cards support 1 Gbps speeds on the 10 Gigabit Ethernet SFP+ ports.

[See [QFX10000-60S-6Q Line Card](#).]

- **QFX10K-12C-DWDM Coherent Line Card (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.2R1, QFX10008 and QFX10016 modular switch chassis support the QFX10K-12C-DWDM Coherent Line Card. The QFX10K-12C-DWDM Coherent Line Card provides up to 1.2 Tbps packet forwarding for cloud providers, service providers, and enterprises that need coherent dense wavelength-division multiplexing (DWDM) with MACsec security features. The six-port line card, with built-in optics, supports flexible rate modulation at 100 Gbps, 150 Gbps, and 200 Gbps speeds. A maximum of four QFX10K-12C-DWDM Coherent Line Cards are supported in either the QFX10008 switch chassis or the QFX10016 switch chassis.

[See [QFX10K-12C-DWDM Coherent Line Card](#).]

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

- **Access control and authentication (QFX5100 switches)**—Starting in Junos OS Release 17.2R1, QFX5100 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 vendor-specific attributes (VSAs). 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 access control and authentication features at the the `[edit protocols dot1x]` hierarchy level. This feature was previously supported in an “X” release of Junos OS.

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

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

- **Support for class of service on QFX5200 switches**—Starting in Junos OS Release 17.2R1, the QFX5200 supports class-of-service (CoS). When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default 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.

[See [Traffic Management User Guide for the QFX Series](#).]

- **Support for FIP snooping and DCBX on QFX5200 switches**—Starting in Junos OS Release 17.2R1, the QFX5200 supports both FIP snooping and DCBX. FIP snooping filters prevent an FCoE device from gaining unauthorized access to a Fibre Channel (FC) storage device or to another FCoE device. Data Center Bridging Capability Exchange Protocol (DCBX) discovers the data center bridging (DCB) capabilities



of connected peers. DCBX advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and values (TLVs).

[See [Traffic Management User Guide for the QFX Series.](#)]

### *Dynamic Host Configuration Protocol (DHCP)*

- **User-defined interface description for DHCP relay (QFX5100, QFX5110, and QFX5200 switches)**--Starting in Junos OS Release 17.2R1, you can define an interface description to be included in DHCP relay option 82 that is independent of the textual interface description configured at the **[edit interfaces interface-name]** hierarchy level.

[See [user-defined.](#)]

### *EVPNs*

- **Support for IGMP snooping for EVPN-VXLAN in a multihomed environment (QFX10000 switches)**--Starting in Junos OS Release 17.2R1, QFX10000 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. 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. The following new operational commands are also supported: **show evpn igmp snooping database extensive**, **show igmp snooping evpn database**, **show igmp snooping evpn membership**, and **show evpn multicast-snooping next-hops**.

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

- **Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network (QFX5100 switches)**--Starting in Junos OS Release 17.2R1, QFX5100 switches that function as Layer 2 VXLAN tunnel endpoints (VTEPs) can tunnel single- and double-tagged Q-in-Q packets through an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network. In addition to tunneling Q-in-Q packets, the ingress and egress VTEPs can perform the following Q-in-Q actions:
  - Delete, or pop, an outer service provider VLAN (S-VLAN) tag from an incoming packet.
  - Add, or push, an outer S-VLAN tag onto an outgoing packet.
  - Map a configured range of customer VLAN (C-VLAN) IDs to an S-VLAN.

**NOTE:** The QFX5100 switch does not support the pop and push actions with a configured range of VLANs.

The ingress and egress VTEPs support the tunneling of Q-in-Q packets and the Q-in-Q actions in the context of specific traffic patterns.

To enable the tunneling of the Q-in-Q packets on the VTEPs, you must configure a flexible VLAN tagging interface, which can transmit 802.1Q VLAN single- and double-tagged packets, on ingress and egress

VTEPs. It is also important to configure the interface to retain the inner C-VLAN tag while a packet is tunneled.

[See [Examples: Configuring QFX5100 Switches to Tunnel Q-in-Q Traffic Through an EVPN-VXLAN Overlay Network](#).]

- **EVPN-VXLAN support of Virtual Chassis and Virtual Chassis Fabric (QFX5100, QFX5100 Virtual Chassis, and Virtual Chassis Fabric)**—Ethernet VPN (EVPN) supports multihoming active-active mode, which enables a host to be connected to two leaf devices through a Layer 2 link aggregation group (LAG) interface. In previous Junos OS releases, the two leaf devices had to be QFX5100 standalone switches. Starting in Junos OS Release 17.2R1, the two leaf devices can be QFX5100 standalone switches, QFX5100 switches configured as a Virtual Chassis (VC), QFX5100 switches configured as a Virtual Chassis Fabric (VCF), or a mix of these options.

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

[See [EVPN-VXLAN Support of Virtual Chassis and Virtual Chassis Fabric](#).]

- **EVPN pure type-5 route support (QFX10000 switches)**—Starting in Junos OS Release 17.2R1, you can configure pure type-5 routing in an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which is used to carry the MAC address of the sending switch and to provide next hop reachability for the prefix. This router MAC extended community provides next-hop reachability without requiring an overlay next hop or supporting type-2 route. To configure pure type-5 routing, include the **ip-prefix-routes advertise direct-nexthop** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. Pure type-5 routing was previously introduced in Junos OS Release 15.1x53-D60.

[See [ip-prefix-routes](#).]

## Infrastructure

- **Secure Boot (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, 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. No action is required to implement Secure Boot.

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

## Interfaces and Chassis

- **Resilient hashing support for link aggregation groups and equal cost multipath routes (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, resilient hashing is now supported by link aggregation groups (LAGs) and equal cost multipath (ECMP) sets.

Resilient hashing enhances LAGs by minimizing destination remapping when a new member is added to or deleted from the LAG.

Resilient hashing works in conjunction with the default static hashing algorithm. It distributes traffic across all members of a LAG by tracking the flow's LAG member utilization. When a flow is affected by a LAG member change, the packet forwarding engine (PFE) rebalances the flow by reprogramming the flow set table. Destination paths are remapped when a new member is added to or existing members are deleted from a LAG.

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

[See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups.](#)]

- **Multichassis link aggregation groups (MC-LAG) (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, 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 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.

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

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

- **Channelizing 100-Gigabit Ethernet QSFP28 interfaces (QFX5200 switches)**—This feature enables you to 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.

To channelize the ports, manually configure the port speed using the **set chassis fpc slot-number port port-number channel-speed speed** command, where the speed can be set to 10G, 25G, or 50G. The ports do not support autochannelization.

**NOTE:** If a 100G transceiver is connected to the switch, channelize the port only to 25G or 50G. If a 40G transceiver is connected, channelize the port only to 10G. Note that there is no commit check for these options.

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

[See [Channelizing Interfaces on QFX5200 Switches](#).]

- **IRB interface in a PVLAN (QFX5110 switches)**—Starting with Junos OS Release 17.2R1, you can configure an integrated routing and bridging (IRB) interface in a private VLAN (PVLAN) on QFX5110 switches so that devices within community VLANs and isolated VLANs can communicate with each other and with devices outside the PVLAN at Layer 3 without requiring you to install a router. This feature was previously supported in an "X" release of Junos OS.

[See [Example: Configuring a Private VLAN Spanning Multiple Switches with an IRB Interface](#).]

## IPv4

- **Generic routing encapsulation (GRE) support (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, you can use GRE tunneling services on QFX5110 switches 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.

## IPv6

- **IPv6 feature support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can configure the Neighbor Discovery Protocol, the Virtual Router Redundancy Protocol (VRRP) for IPv6, and Protocol Independent Multicast (PIM) for IPv6. You can also configure BGP and IS-IS for IPv6 as well as OSPFv3. Additionally, unicast IPv6 is supported for virtual router instances. DHCPv6 is also supported. IPv6 feature support for QFX5110 and QFX5200 switches was previously introduced in "X" releases of Junos OS.

[See [Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery](#) and [Verifying and Managing DHCPv6 Local Server Configuration](#).]

## Layer 2 Features

- **Layer 2 features (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, the following features are supported:
  - VLAN support—Enables you to divide one physical broadcast domain into multiple virtual domains.
  - 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—Allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.
  - Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) support – Provides Layer 2 loop prevention.

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

- **Q-in-Q tunneling support (QFX5200 switches)**—Starting in Junos OS Release 17.2R1, QFX5200 switches support Q-in-Q tunneling, which enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the

customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag. This feature was previously supported in an "X" release of Junos OS.

[See [Understanding Q-in-Q Tunneling](#).]

### Layer 3 Features

- **Support for hierarchical ECMP groups (QFX5200 switches)**—Starting in Junos OS Release 17.2R1, hierarchical equal-cost multipath (ECMP) groups are enabled by default at system start. Hierarchical ECMP provides for two-level route resolution automatically through the Packet Forwarding Engine. Two-level route resolution through ECMP groups enhances load balancing of traffic. This feature was previously introduced in Junos OS Release 15.1X53-D30.

[See [Overview of Hierarchical ECMP Groups](#).]

- **Support for 64 next-hop gateways for ECMP (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, you can configure as many as 64 equal-cost-multipath (ECMP) next hops for RSVP and LDP LSPs or external BGP peers. The following Layer 3 protocols are supported as ECMP gateways for both IPv4 and IPv6 traffic: OSPF, ISIS, EBGp, and IBGP (resolving over IGP routes). Include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level. This feature was previously introduced on QFX5110 switches in Junos OS Release 15.1X53-D210.

[See [maximum-ecmp](#).]

- **Support to disable hierarchical ECMP (QFX5200 switches)**—Starting with Junos OS Release 17.2R1, you can disable hierarchical equal-cost multipath (ECMP) groups at system start time. Hierarchical ECMP is enabled by default. Disabling this feature effectively increases the number of ECMP groups. Include the **no-hierarchical-ecmp** statement at the **[edit forwarding-options]** hierarchical level. Disabling hierarchical ECMP causes the Packet Forwarding Engine to restart. To reenab le hierarchical ECMP, issue the following command: **delete forwarding-options no-hierarchical-ecmp**. This feature was previously introduced in Junos OS Release 15.1X53-D210.

[See [no-hierarchical-ecmp](#).]

### Management

- **Support for device family and release in Junos OS YANG modules (QFX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

[See [Understanding Junos OS YANG Modules](#).]

- **Support for the Junos Telemetry Interface (QFX10000 switches)**—Starting with Junos OS Release 17.2R1, the Junos Telemetry Interface is supported on QFX10000 switches. Both UDP and gRPC streaming of statistics are supported. Junos Telemetry Interface enables you to provision sensors to export telemetry data for various network elements without involving polling.

The following sensors are supported on QFX10000 switches:

- Logical interfaces (UDP and gRPC streaming)
- Physical interfaces (UDP and gRPC streaming)
- Firewall filters, including traffic-class counters (UDP and gRPC streaming)
- LSP statistics (UDP and gRPC streaming)
- LSP events and properties (gRPC streaming)
- Optical interfaces (UDP and gRPC streaming)
- Network processing unit (NPU) memory (UDP and gRPC streaming)
- NPU memory utilization (UDP and gRPC streaming)
- CPU memory (UDP and gRPC streaming)
- Chassis components (gRPC streaming only)
- RSVP interface events (gRPC streaming only)
- BGP peers (gRPC streaming only)
- Memory utilization for routing protocol tasks (gRPC streaming only)
- Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol (gRPC streaming only)
- Ethernet interfaces enabled configured with the Link Layer Discovery Protocol (gRPC streaming only)
- Network Discovery Protocol table state (gRPC streaming only)
- Address Resolution Protocol table state (gRPC streaming only)

To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. To provision sensors to stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters for a specified list of OpenConfig command paths. Because QFX10000 switches run a version Junos OS with an upgraded FreeBSD kernel, 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.

The LSP events and properties sensor is supported in Junos OS Release 17.2R1 for the first time. You can export statistics for ingress point-to-point LSPs, point-to-multipoint LSPs, bypass LSPs, and dynamically created LSPs. To export data through gRPC, use the **/mpls/lsp/** or **/mpls/signal-protocols/** set of OpenConfig subscription paths.

[See [Overview of the Junos Telemetry Interface.](#)]

- **Support for the Junos Telemetry Interface (QFX5200 switches)**—Starting with Junos OS Release 17.2R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various

network elements without involving polling. On QFX5200 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
- Ethernet interfaces enabled configured with the Link Layer Discovery Protocol
- BGP peers
- RSVP interface events
- Memory utilization for routing protocol tasks
- Address Resolution Protocol table state
- Network Discovery Protocol table state

To provision sensors to stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters for a specified list of OpenConfig commands paths. 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](#).]

## MPLS

- **TE++ dynamic bandwidth management using container LSPs (QFX5100)**—Starting with Junos OS Release 17.2R1, a new type of label-switched path (LSP), called a container LSP, is introduced to enable load balancing across multiple point-to-point member LSPs between the same ingress and egress routers. Each member LSP takes a different path to the same destination and can be routed along a different interior gateway protocol (IGP) cost path. Based on the configuration and aggregate traffic, a container LSP provides support for dynamic bandwidth management by enabling the ingress router to dynamically add and remove member LSPs through a process called LSP splitting and LSP merging, respectively. Member LSPs can also be re-optimized with different bandwidth values in a make-before-break way. The feature was previously supported in a "X" release of Junos OS.

[See [Dynamic Bandwidth Management Using Container LSP Overview](#).]

- **Entropy labels for LSPs (QFX10000 switches)**—Starting with Junos OS Release 17.2R1, you can configure entropy labels for label-switched paths (LSPs). An entropy label is a special load-balancing label that 0 enhances the ability of the switch to load-balance traffic across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs). The entropy label allows the switch to efficiently load-balance traffic using just the label stack rather than deep packet inspection (DPI). To configure entropy labels, include the entropy-label statement at the **[edit protocols mpls labeled-switched-path labeled-switched-path-name]** hierarchy level.



[See [Understanding Entropy Label for BGP Labeled Unicast LSPs](#) and [Automatic Bandwidth Allocation for LSPs](#).]

- **Support for a label stack for BGP label unicast for MPLS advertisements (QFX10000 switches)**—Starting with Junos OS 17.2R1, QFX10000 switches implement RFC 3701, which supports a stack of labels in BGP label unicast for both IPv4 and IPv6 traffic. Previously, only one label per prefix was supported in the BGP unicast label. You can now specify to include up to five labels per prefix in the BGP labeled unicast updates. This feature enables the use of the BGP label unicast stack to program a stack of labels to control packet forwarding in a network configured with hierarchical MPLS label-switched paths. To configure as many as five labels to advertise through MPLS, include the **maximum-labels *number*** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family mpls]** hierarchy level. The **show route receive-protocol bgp *neighbor-address* detail** and **show route advertising-protocol *neighbor-address* detail** operational commands are enhanced to display multiple labels for one prefix in the Labels field.

[See [Configuring the Maximum Number of MPLS Labels](#).]

- **MPLS support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, MPLS is supported on the QFX5110 and QFX5200 switches. MPLS supports 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) support, a component of MPLS local protection for both one-to-one and many-to-one local protection.
  - Loop-free alternate (LFA)
  - 6PE 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](#).]

- **Support for equal cost multipath (ECMP) routing on label-switching routers (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can configure ECMP 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 [Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing](#).]

- **Ethernet over MPLS (Layer 2 circuit) support (QFX5100 Virtual Chassis and Virtual Chassis Fabric)**—Starting in Junos OS Release 17.2R1, a QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF) supports Ethernet over MPLS (Layer 2 circuit). The Virtual Chassis or VCF can act as a provider edge switch on which you configure MPLS and LDP for the interfaces that will carry the Layer 2 circuit traffic. The Layer 2 circuit can be port-based (pseudo-wire) or VLAN-based. These features were previously supported for a QFX5100 Virtual Chassis or VCF in an “X” release of Junos OS.

[See [Understanding Ethernet-over-MPLS \(L2 Circuit\)](#) and [Configuring Ethernet over MPLS \(L2 Circuit\)](#).]

### **Multicast**

- **Layer 3 multicast support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, IGMP, including versions 1,2, and 3, IGMP snooping, PIM sparse mode, and PIM source-specific multicast are supported. You can also configure IGMP, IGMP snooping, and PIM in virtual router instances. Multicast Source Discovery Protocol (MSDP) is also supported. Configure IGMP at the **[edit protocols igmp]** hierarchy level. Configure IGMP snooping at the **[edit protocols igmp-snooping]** hierarchy level. Configure PIM at the **[edit protocols pim]** hierarchy level. Configure MSDP at the **[edit protocols msdp]** hierarchy level. Layer 3 multicast support was previously introduced in “X” releases of Junos OS.

[See [Multicast Overview](#).]

- **Support for static multicast route leaking for VRF and virtual-router instances (QFX5100 and EX4300 switches)**—Starting in Junos OS Release 17.2R1, you can configure your switch to share IPv4 multicast routes among different virtual routing and forwarding (VRF) instances or different virtual-router instances. On EX4300 switches, multicast route leaking is supported only when the switch functions as a line card in a Virtual Chassis, not as a standalone switch. Only multicast static routes with a destination-prefix length of /32 are supported for multicast route leaking. Only Internet Group Management Protocol version 3 is supported. To configure multicast route leaking for VRF or virtual-router instances, include the **next-table routing-instance-name.inet.0** statement at the **[edit routing-instances routing-instance-name routing-options static route destination-prefix/32]** hierarchy level. For **routing-instance-name**, include the name of a VRF or virtual-router instance. This feature was previously introduced in Junos OS Release 14.X53-D40.

[See [Understanding Multicast Route Leaking for VRF and Virtual-Router Instances](#).]

## Network Management and Monitoring

- **sFlow enhancements (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 17.2R1, sFlow IPv4 and IPv6 packets support extended router information, including the IP address of the next-hop router, the outgoing VLAN ID, the source IP address prefix length, and the destination IP address prefix length. This information is collected only if BGP is configured on the switch.

In addition, a configuration statement was introduced that allows the sFlow sampling rate to stay within the maximum sampling rate of 1 out of 64,000 packets. Packet-based sampling is implemented in the hardware, so all of the interfaces can be monitored with very little overhead. However, if traffic levels are unusually high, the hardware generates more samples than it can handle. The extra samples are dropped by the software rate-limiting algorithm and can cause inaccurate results. You can include the **disable-sw-rate-limiter** statement at the **[edit protocols sFlow]** hierarchy to disable the software, allowing the hardware sampling rate to stay within the maximum sampling rate for sFlow.

[See [Understanding How to Use sFlow Technology for Network Monitoring on a Switch.](#)]

- **sFlow technology support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, the QFX5110 and QFX5200 switches support 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.](#)]

- **Port mirroring (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can use port mirroring on QFX5110 and QFX5200 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, correlating events, and so on. This feature was previously supported in an "X" release of Junos OS.

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

- **SNMP support for monitoring tunnel statistics (QFX Series)**—Starting in Junos OS Release 17.2R1, SNMP MIB jnxTunnelStat supports monitoring of tunnel statistics for IPv4 over IPv6 tunnels. This is a new enterprise-specific MIB, Tunnel Stats MIB, that currently displays three counters: tunnel count in rpd, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine. This MIB can be extended to support other tunnel statistics. The MIB is defined in jnx-tunnel-stats.txt. This MIB is attached to jnxMibs.

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

## Port Security

- **Media Access Control Security (MACsec) support (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 17.2R1, MACsec is supported on all six interfaces of the QFX10K-12C-DWDM line card when it is installed in a QFX10008 or QFX10016 switch. MACsec is an 802.1AE IEEE industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is capable of identifying and preventing most security threats, and can be used in combination with other security protocols to provide end-to-end network security. MACsec can be enabled only on domestic versions of Junos OS software.

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

- **Access security support (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, the following access security features are supported on QFX5110 switches:
  - DHCP snooping—DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information, which it uses to build and maintain a database of valid IP-address-to-MAC-address (IP-MAC) bindings called the DHCP snooping database. Clients on untrusted ports are only allowed to access the network if they can be validated against the database.
  - DHCPv6 snooping—DHCP snooping for DHCPv6.
  - DHCP option 82—You can use DHCP option 82, also known as the *DHCP relay agent information* option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
  - DHCPv6 option 37—Option 37 is the DHCPv6 equivalent of the remote ID suboption of DHCP option 82. It is used to insert information about the network location of the remote host into DHCPv6 packets.
  - Dynamic ARP inspection (DAI)—DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as *ARP poisoning* or *ARP cache poisoning*). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons.
  - IPv6 neighbor discovery (ND) inspection—IPv6 ND inspection mitigates attacks based on the Neighbor Discovery Protocol by inspecting neighbor discovery messages and verifying them against the DHCPv6 snooping table.
  - MAC limiting—You can configure a MAC limit per interface and per VLAN, and set an action to take on the next packet the interface or VLAN receives after the limit is reached.

- **MAC move limiting**—You can configure MAC move limiting to track MAC address movements on the switch, so that if a MAC address changes more than the configured number of times within one second, the changes to MAC addresses are dropped, logged, or ignored, or the interface is shut down.
- **Persistent MAC learning**—Persistent (also called *sticky*) MAC addresses help restrict access to an access port by identifying the MAC addresses of workstations that are allowed access to a given port. Secure access to these workstations is retained even if the switch is restarted.

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

### Routing Protocols

- **Support for BGP Monitoring Protocol (BMP) Version 3 (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, you can configure BMP, which sends BGP route information from the switch to a monitoring application, or station, on a separate device. To deploy BMP in your network, you need to configure BMP on each switch and at least one BMP monitoring station. Only version 3 is supported. To configure BMP, include the **bmp** set of statements at the **[edit routing-options]** hierarchy level. To configure a BMP monitoring station, include the **station-address ip-address** and the **station-port number** statements at the **[edit routing-options bmp]** hierarchy level.

[See [Configuring BGP Monitoring Protocol Version 3](#).]

- **Support for segment routing for IS-IS (QFX5100 switches and QFX10000 switches)**—Starting with Junos OS Release 17.2R1, you can advertise MPLS labels through IS-IS to support segment routing. IS-IS advertises a set of segments, which enables an ingress device to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the path to take. Two types of segments are supported: node and adjacency. A node segment represents a shortest-path link to a node. An adjacency segment represents a specific adjacency to a node. To enable segment routing, include the **source-packet-routing** statement at the **[edit protocols isis]** hierarchy level. By default, segment routing is enabled on all IS-IS levels. To disable advertising of the adjacency segment for a specified interface, include the **no-advertise-adjacency-segment** statement. You can also specify an interval for maintaining adjacency segments by including the **adjacency-segment hold-time milliseconds** statement.

To enable node segments, include the **node-segment** statement at the **[edit protocols isis source-packet-routing]** hierarchy level. You have two options for advertising a range of indices for IPv4 or IPv6 addresses. Use the **index-range** statement to specify a dynamic label range managed by MPLS. To specify a specific block of indices, also known as a segment routing global block, include the **start-label <number> index-range <number>** statements at the **[edit protocols isis source-packet-routing srgb]** hierarchy level. This configuration enables MPLS to reserve the specified label range.

Segment routing in IS-IS also supports provisioning prefix segment indices (SIDs) and anycast SIDs for both IPv4 and IPv6 prefixes. These SIDs are provisioned through a routing policy for each prefix. Include the **then prefix-segment index number** statement at the **[edit policy options policy-statement policy-name]** hierarchy level. You can also enable IPG shortcuts for prefix segment routes. Include the **shortcuts** statement at the **[edit protocols isis traffic-engineering family (inet-mpls | inet6-mpls)]** hierarchy level.

[See [source-packet-routing](#).]

- **Support for segment routing for OSPF (QFX5100 switches and QFX10000 switches)**—Starting with Junos OS Release 17.2R1, you can advertise MPLS labels through OSPF to support segment routing. Only IPv4 is supported. OSPFv3 is not supported. OSPF advertises a set of segments, which enables an ingress device to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the path to take. Two types of segments are supported: node and adjacency. A node segment represents a shortest-path link to a node. An adjacency segment represents a specific adjacency to a node. To enable segment routing, include the **source-packet-routing** statement at the **[edit protocols ospf]** hierarchy level. By default, segment routing is enabled for all OSPF areas. To disable for a specific area, include the **no-source-packet-routing** statement at the **[edit protocols ospf area area-id]** hierarchy level. To enable node segments, include the **node-segment** statement. You can specify a range for IPv4 addresses to advertise, which MPLS manages dynamically. To disable advertising of the adjacency segment for a specified interface, include the **no-advertise-adjacency-segment** statement.

[See [source-packet-routing](#).]

- **Support for unique AS path count (QFX Series)**—Starting with Junos OS Release 17.2R1, you can configure a routing policy to determine the number of unique autonomous systems (ASs) present in the AS path. The unique AS path count helps determine whether a given AS is present in the AS path multiple times, typically as prepended ASs. In earlier Junos releases it was not possible to implement this counting behavior using the **as-path** regular expression policy. This feature permits the user to configure a policy based on the number of AS hops between the route originator and receiver. This feature ignores ASs in the **as-path** that are confederation ASs, such as **confed\_seq** and **confed\_set**.

To configure AS path count, include the **as-path-unique-count count (equal | orhigher | orlower)** configuration statement at the **[edit policy-options policy-statement policy\_name from]** hierarchy level.

## Security

- **Support for filter-based decapsulation over an IP-IP interface (QFX10000 switches)**—Starting in Junos OS Release 17.2R1, you can use a firewall filter over an IP-IP interface to de-encapsulate traffic on the switch, without the need to create any tunnel interfaces. IP-in-IP packets are special IP tunneling packets with no GRE header. With this feature, you can define a filter with filtering terms to classify packets based on packet fields such as destination IP address and protocol type. This provides significant benefits in terms of scalability, performance, and flexibility.

[See [Configuring a Firewall Filter to De-Encapsulate IP-in-IP Traffic](#).]

- **Policing support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can use policing (or rate-limiting) to apply limits to traffic flow and to set consequences for packets that exceed those limits. The device polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing 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](#).]

- **Storm control support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, 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](#).]

- **Firewall filters support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can provide rules 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 [Overview of Firewall Filters](#).]

- **Generic routing encapsulation (GRE) support (QFX5100 and QFX5200 switches)**—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. This feature was previously supported in an “X” release of Junos OS.

[See [Configuring a Firewall Filter to De-Encapsulate GRE Traffic](#).]

### **Services Applications**

- **Support for IPFIX templates for flow aggregation (QFX10002 switches)**—Starting with Junos OS Release 17.2R1, you can define a flow record template for unicast IPv4 and IPv6 traffic in IP Flow Information Export (IPFIX) format. Templates are transmitted to the collector periodically. To define an IPFIX template, include the **version-ipfix template *template-name*** set of statements at the **[edit services flow-monitoring]** hierarchy level.

You must also perform the following configuration:

- Sampling instance at the **[edit forwarding-options]** hierarchy level.
- Associate the sampling instance with the FPC at the **[edit chassis]** hierarchy level and with a template configured at the **[edit services flow-monitoring]** hierarchy level.
- Firewall filter for the family of traffic to be sampled at the **[edit firewall]** hierarchy level.

[See [Configuring Flow Aggregation to Use IPFIX Flow Templates](#).]

### **Software Defined Networking (SDN)**

- **OVSDB-VXLAN support with Contrail (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, the Open vSwitch Database (OVSDB) management protocol provides a means through



which a Contrail controller can communicate with QFX5110 and QFX5200 switches to provision them as Layer 2 VXLAN gateways. In an environment in which Contrail Release 2.22 or later is deployed, a Contrail 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.

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

[See [Understanding the OVSDb Protocol Running on Juniper Networks Devices.](#)]

- **Layer 2 VXLAN gateway (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can implement a QFX5110 or QFX5200 switch as a Layer 2 Virtual Extensible LAN (VXLAN) gateway. 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 (VMs) 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. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

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

[See [Understanding VXLANs.](#)]

- **BFD in a VMware NSX environment with OVSDb and VXLAN (QFX5100 switches, QFX5100 Virtual Chassis)**

Within a Virtual Extensible LAN (VXLAN) managed by the Open vSwitch Database (OVSDb) protocol, by default, Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic is replicated and forwarded by one or more software virtual tunnel endpoints (VTEPs) or service nodes in the same VXLAN. (The software VTEPs and service nodes are collectively referred to as *replicators*.)

Starting in Junos OS Release 17.2R1, a Juniper Networks switch or Virtual Chassis that functions as a hardware VTEP in a VMware NSX environment uses the Bidirectional Forwarding Detection (BFD) protocol to prevent the forwarding of BUM packets to a non-functional replicator.

By exchanging BFD control messages with replicators at regular intervals, the hardware VTEP can monitor the replicators to ensure that they are functioning and reachable.



### **Software Installation and Upgrade**

- **Support for FreeBSD 10 kernel for Junos OS (QFX5200 and QFX5110 switches)**—Starting with Junos OS Release 17.2R1, FreeBSD 10 is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. Because installation restructures the file system, logs and configurations are lost unless precautions are taken. Now there are Junos OS and OAM volumes, that provide the ability to boot from the OAM volume upon failures. Some system commands display different output and a few others are deprecated.

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

[See [Understanding Junos OS with Upgraded FreeBSD.](#)]

### **Software Licensing**

- **Integrated software feature licenses (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, 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.

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

[See [Software Features That Require Licenses on the QFX Series.](#)]

### **System Management**

- **Support for Precision Time Protocol (PTP) transparent clock (QFX5100 and QFX5110 switches)**—Starting in Junos OS Release 17.2R1, PTP synchronizes clocks throughout a packet-switched network. With a transparent clock, the PTP packets are updated with residence time as the packets pass through the switch. There is no master/slave designation. End-to-end transparent clocks are supported. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that the timestamps are sent in one packet. In a two-step process, estimated timestamps are sent in one packet, and additional packets contain updated timestamps. In addition, User UDP over IPv4 and IPv6 and unicast and multicast transparent clock are supported.

You can configure the transparent clock at the **[edit protocols ptp]** hierarchy.

[See [Understanding Transparent Clocks in Precision Time Protocol.](#)]

- **Zero Touch Provisioning (QFX5100, QFX5110, and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, Zero Touch Provisioning allows you to provision new Juniper Networks switches in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information,

the switch boots with the preinstalled software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version.

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

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

**VLAN Infrastructure**

- **Double VLAN tags on Layer 3 subinterfaces (QFX10000 switches)**—Starting in Junos OS Release 17.2R1, you can configure double VLAN tags on Layer 3 subinterfaces (also called “Layer 3 logical interfaces) on QFX10000 switches. Layer 3 double-tagged logical interfaces support **inet**, **inet6**, and **mpls** families.

Support for double-tagging VLANs on Layer 3 logical interfaces includes:

- Configuration of an IPv4, an IPv6, or an **mpls** family on the logical interface
- Configuration over an aggregated Ethernet interface
- Configuration of multiple logical interfaces on a single physical interface

[See [Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces](#).]

SEE ALSO

<a href="#">Changes in Behavior and Syntax</a>	<a href="#">  338</a>
<a href="#">Known Behavior</a>	<a href="#">  344</a>
<a href="#">Known Issues</a>	<a href="#">  347</a>
<a href="#">Resolved Issues</a>	<a href="#">  351</a>
<a href="#">Documentation Updates</a>	<a href="#">  363</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  364</a>
<a href="#">Product Compatibility</a>	<a href="#">  376</a>

**Changes in Behavior and Syntax**

**IN THIS SECTION**

- [Class of Service \(CoS\)](#) | [339](#)
- [General Routing](#) | [340](#)
- [EVPNs](#) | [340](#)
- [Interfaces and Chassis](#) | [340](#)

- Network Management and Monitoring | 341
- Management | 342
- Routing Protocols | 343
- Virtual Chassis | 343

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 17.2R3 for the QFX Series.

## Class of Service (CoS)

- The following CoS options are hidden under the Traffic control profiles for QFX10002 and QFX10008 products:

- > delay-buffer-rate
- > excess-rate
- > excess-rate-high
- > excess-rate-low
- > excess-rate-medium-high
- > excess-rate-medium-low

The shaping-rate option is hidden on QFX10008 but not on QFX10002, as shaping rate configurations are used in the QFX10002 satellite solution setup. [PR1261988](#)

- When you configure the **transmit-rate**, you must also configure the **guaranteed-rate** under **traffic-control-profiles**. If you commit a **transmit-rate** configuration without having configured **guaranteed-rate**, a warning message is displayed and the default scheduler map is applied.

## General Routing

- **Support for deletion of static routes when the BFD session goes down (QFX Series)**—Starting with Junos OS 17.2R2, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

## EVPNs

- On QFX10000 switches running Junos OS Release 17.2R3 or later, the local preference setting for an Ethernet VPN (EVPN) pure type-5 route is inherited by IP routes that are derived from the EVPN type-5 route. Further, when selecting an IP route for incoming traffic, the QFX10000 switches consider the local preference of the route. A benefit of the QFX10000 switches including the local preference in their route selection criteria is that you can set up a policy to manipulate the local preference, thereby controlling which route the switch selects.

## Interfaces and Chassis

- **Changes to the show interface interface-name command (QFX10002)**—Two additional CLI fields, **FEC Corrected Errors Rate** and **FEC Uncorrected Errors Rate** are added to the `show interface interface-name` command. For example:

```
user@router> show interfaces et-0/0/35
```

```
Physical interface: et-0/0/35, Enabled, Physical link is Up
  Interface index: 658, SNMP ifIndex: 541
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 100Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled,
  Media type: Fiber
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None
  CoS queues       : 8 supported, 8 maximum usable queues
    Last flapped    : 2017-02-07 21:35:08 PST (00:08:07 ago)
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)
  Active alarms    : None
  Active defects   : None
  PCS statistics
    Bit errors      Seconds
    Errorred blocks 2
```

```

FEC MODE                                FEC91
FEC Corrected Errors                    193929
FEC Uncorrected Errors                  2075
FEC Corrected Errors Rate                0
FEC Uncorrected Errors Rate              0
Interface transmit statistics: Disabled

```

## Network Management and Monitoring

- **SNMP syslog messages changed (QFX Series)**—In Junos OS Release 17.2R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:

- Old message: **AgentX master agent failed to respond to ping. Attempting to re-register**  
New message: **AgentX master agent failed to respond to ping, triggering cleanup!**
- Old message: **NET-SNMP version %s AgentX subagent connected**  
New message: **NET-SNMP version %s AgentX subagent Open-Sent!**

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

- **Update to SNMP support of apply-path statement (QFX Series)**—In Junos OS Release 17.2R1, SNMP implementation of the **apply-path** configuration statement supports only two lists:

- **apply-path "policy-options prefix-list <list-name> <\*>"**

This configuration has been supported from the first release.

- **apply-path "access radius-server <\*>"**

This configuration is supported as of this release.

- **Change in default log-level setting (QFX Series)**—In Junos OS Release 17.2R3, the following changes are made to the default logging levels:

Before this release:

- SNMP\_TRAP\_LINK\_UP was LOG\_INFO for both the physical (IFD) and logical (IFL) interfaces.
- SNMP\_TRAP\_LINK\_DOWN was LOG\_WARNING for both the physical (IFD) and logical (IFL) interfaces.

From this release onward:

- IFD LinkUp -> LOG\_NOTICE (as this is an important message but less frequent)
- IFL LinkUp -> LOG\_INFO (no change)
- IFD and IFL LinkDown -> LOG\_WARNING (no change)

See the [MIB Explorer](#).

- **Need to reconfigure SNMPv3 configuration after upgrade (QFX Series)**—In Junos OS Release 17.2R2, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. This is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. In releases before Junos OS Release 17.2R1, you need to reconfigure SNMPv3 every time after a reboot. This problem is now fixed. If you upgrade, you must still reconfigure SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3 again. To reconfigure SNMP v3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters. Platforms affected are QFX5100, QFX10002, QFX10008, and QFX10016.

[See [Configuring the Local Engine ID](#).]

## Management

- **Junos OS YANG module namespace and prefix changes (QFX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command** module uses its own unique module name as the module's prefix. In earlier releases, Junos OS YANG modules used only a unique identifier to differentiate the namespace for each module, and the prefix for all **juniper-command** modules was **jrpc**.

Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**. The Junos OS YANG extension modules, **junos-extension** and **junos-extension-odl**, use the **junos** device family identifier in the namespace, but the modules are common to all device families.

[See [Understanding Junos OS YANG Modules](#).]

- **Changes to the rfc-compliant configuration statement (QFX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. If you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level and request configuration data in a NETCONF session on a device running Junos OS Release 17.2R1 or later, the NETCONF server sets the default namespace for the **<configuration>** element in the RPC reply to the same namespace as in the corresponding YANG model.

[See [Configuring RFC-Compliant NETCONF Sessions](#) and [rfc-compliant](#).]

- **Enhancement to the Junos Telemetry Interface (QFX10000 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, the values displayed in the **oper-status** key-value field of data streamed through gRPC for the physical interfaces sensor have changed.

The following values are now displayed to indicate the operational status of an interface:

- operational status up—**UP**
- operational status down—**DOWN**

- operational status unknown—**UNKNOWN**
- **Enhancement to NPU memory sensors for Junos Telemetry Interface (QFX10000 switches)**—Starting with Junos OS Release 17.2R1, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is `/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/`  
[See [Guidelines for gRPC Sensors](#).]

## Routing Protocols

- **Syslog error message RPD\_ISIS\_PREFIX\_SID\_CNFLCT to resolve conflicting prefix segment advertisement (QFX Series)**—Starting in Junos OS Release 17.2R2, the **RPD\_ISIS\_PREFIX\_SID\_CNFLCT** syslog error message is emitted only when the prefix segment advertisement from the remote node is conflicting with an advertisement from the self node. This conflict happens because the same prefix segment index is assigned on different IP addresses or different prefix segment indexes are assigned to the same IP address. To rectify this conflict identify the remote node in the network originating the conflicting prefix segment advertisement and change the prefix segment index on the local node or on the remote node.  
[See [Example: Configuring Anycast and Prefix Segments in SPRING for ISIS](#).]

## Virtual Chassis

- **Adaptive load balancing (ALB) feature (Virtual Chassis Fabric)**—Starting in Junos OS Release 17.2R2, the adaptive load balancing (ALB) feature for Virtual Chassis Fabric (VCF) is being deprecated to avoid potential VCF instability. The **fabric-load-balance** configuration statement in the **[edit forwarding-options enhanced-hash-key]** hierarchy is no longer available to enable and configure ALB in a VCF. When upgrading a VCF to a Junos OS release where ALB is deprecated, if the configuration has ALB enabled, you should delete the **fabric-load-balance** configuration item before initiating the upgrade.  
[See [Understanding Traffic Flow Through a Virtual Chassis Fabric](#) and [fabric-load-balance](#).]

## SEE ALSO

[New and Changed Features | 318](#)

[Known Behavior | 344](#)

[Known Issues | 347](#)

[Resolved Issues | 351](#)

[Documentation Updates | 363](#)

[Migration, Upgrade, and Downgrade Instructions | 364](#)

## Known Behavior

### IN THIS SECTION

- [EVPNs | 344](#)
- [General Routing | 344](#)
- [High Availability \(HA\) and Resiliency | 345](#)
- [MPLS | 346](#)
- [Layer 2 Features | 346](#)
- [Routing Protocols | 346](#)
- [Virtual Chassis | 346](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R3 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.

### EVPNs

- A PE device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE. The IGP instance running in the VRF on the PE may be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE.  
**Workaround:** Run IGP sessions between a PE and locally attached CEs. Use L3VPN to distribute the IGP-learned routes between PEs across the core. [PR977945](#)

### General Routing

- On EX4600 and QFX5100 switches, the amount of time that it takes for Zero Touch Provisioning to complete might be lengthy because TFTP might take a long time to fetch the required data. [PR980530](#)
- On EX Series or QFX Series switches, nonstop software upgrade (NSSU) cannot be used to upgrade from Junos OS Release 14.1X53 to Junos OS Release 15.1 or later. [PR1087893](#)



- On an QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- With multihop BFD, traffic loss of around 5 to 10 seconds is observed when an intermediate interface is brought down. After 5 to 10 seconds, the traffic recovers and no action is needed. [PR1150695](#)
- For a LAG interface, the Packet Forwarding Engine populates only the bundle statistics and not the child's logical interface statistics. A value of zero (0) is always returned for logical interface statistics. A limitation in the hardware does not allow the correct statistics to be returned. [PR1250870](#)
- On QFX5200 switches, if the port speed is changed from 25 Gbps to 100 Gbps, or if there are repeated changes in the port speed settings, then the link might remain down. This is SDK limitation and has been addressed in SDK versions 6.5.8 and later. [PR1250891](#)
- On QFX10000 switches, a part of the fabric management cell-drop error detection and correction feature is not supported as part of Junos OS Release 17.1. [PR1252448](#)
- On the QFX10K-12C-DWDM coherent line card, when an interface is configured in 8QAM mode, pull out of fiber on the second "OT" interface in the same AC400 module brings both the "OT" interfaces down. This does not affect any functionality. [PR1258539](#)
- When fiber is pulled out of the OT interface and plugged back in, the Et interfaces go up and after 5 seconds the OT interface show as up. But actually both OT and ET are up at the same time. Only reporting of OT up is delayed by a maximum of 5 seconds. This does not affect the functionality. [PR1258551](#)
- Currently, a maximum of 64 LAG members under a single aggregated Ethernet bundle is supported for the QFX Series. [PR1259515](#)
- On QFX10008 and QFX10016 with the QFX10K-12C-DWDM line card installed, protocols (such as BFD and BGP) running over aggregated Ethernet interfaces might flap when the aggregated Ethernet member link is disabled. [PR1289703](#)
- ERPS convergence takes time after a GRES and thus, traffic loss is observed for a brief period. [PR1290161](#)

## High Availability (HA) and Resiliency

- **On QFX5100 switches, residual and baseline statistics loss from unified ISSU**—Using unified ISSU to upgrade to Junos OS Release 17.2R1 or later will result in a loss of residual and baseline statistics for interfaces, interface set specific statistics, and BBE subscriber service statistics because of an update to the statistics database.

[See [Unified ISSU System Requirements](#).]

- During an NSSU on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)

## MPLS

- On QFX5100, QFX5110, QFX5200 switches with Layer 2 circuit configured on the PE switches, enabling VLAN bridge encapsulation on a CE interface drops packets if flexible Ethernet services and VLAN CCC encapsulation are configured on the same logical interface. You can configure only one encapsulation type on a particular logical interface, for example, use either **set interfaces xe-0/0/18 encapsulation flexible-ethernet-services** or **set interfaces xe-0/0/18 encapsulation vlan-ccc**. [PR1329451](#)

## 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](#)
- After you delete and re-add 1000 LAG interfaces, traffic drops might be seen for some time even though all LAG interfaces comes up. [PR1289546](#)

## Routing Protocols

- During a graceful Routing Engine switchover (GRES) on QFX10000 switches, some IPv6 groups might experience momentary traffic loss. This issue occurs when IPv6 traffic is running with multiple paths to the source, and the **join-load-balance** statement for PIM is also configured. [PR1208583](#)
- On QFX10000 switches, the outbound firewall filter on aggregated Ethernet to block micro-BFD packets (destination port 6784) does not work. In this case, the micro-BFD sessions continue to stay up although, the inbound firewall filter to block micro-BFD packets works fine. [PR1248504](#)

## Virtual Chassis

- If a QFX5100 switch running Junos OS Release 17.2R1 or later is in the same Virtual Chassis or Virtual Chassis Fabric (VCF) as a Juniper Networks device that does not support Virtual Extensible LAN (VXLAN) for example, an EX4300 switch, then the Junos OS CLI of the EX4300 switch supersedes the Junos OS CLI of the QFX5100. As a result, the **vxlan** configuration statement at the **[edit vlans vlan-name]** hierarchy level does not appear. [PR1176054](#)

## SEE ALSO

[New and Changed Features | 318](#)

[Changes in Behavior and Syntax | 338](#)

[Known Issues | 347](#)

[Resolved Issues | 351](#)

---

[Documentation Updates | 363](#)

---

[Migration, Upgrade, and Downgrade Instructions | 364](#)

---

[Product Compatibility | 376](#)

---

## Known Issues

### IN THIS SECTION

- [General Routing | 347](#)
- [EVPN | 350](#)
- [Interfaces and Chassis | 350](#)
- [Layer 2 Features | 350](#)
- [Network Management and Monitoring | 350](#)
- [Platform and Infrastructure | 350](#)
- [Routing Protocols | 350](#)

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.2R3.

### General Routing

- While SSH is used to log in to a VNF, the following error message appears **Unrecognized command is seen**. This error has no impact on the functionality. [PR1108785](#)
- After sending leave and rejoin, in a few seconds Layer 3 multicast traffic might not converge up to 100 percent and a certain amount of traffic drops might be seen continuously. This behavior is observed when you scale beyond 2000 VLANs or 2000 IRBs with VLAN replication in the system. [PR1135045](#)
- While scaling beyond 2000 VLANs or 2000 IRBs, Layer 3 multicast traffic does not converge to 100 percentage and continuous drops are observed after the downstream interface is brought down or brought up or while an FPC comes online after a FPC restart. [PR1161485](#)
- When per-packet load balancing is removed or deleted, the next-hop index might change. [PR1198092](#)
- An ICCP session is maintained by multihop BFD (non-distributed mode). The time interval for BFD keepalive messages is similar to that in the GRES configuration (for example, keepAlive = 8 seconds). [PR1230576](#)

- On QFX10008 switches, the IPv6 packets or bytes counter shows higher values than the total packets or bytes of the interface if LAG child members belong to the same PE device. As a workaround, if you monitor IPv6 statistics over the LAG, choose LAG child members across all PE devices. [PR1232388](#)
- On a QFX5110-48S switch, a Gigabit Ethernet interface goes down and comes back up once on a peer as part of a reboot. [PR1237572](#)
- When displaying the unified forwarding table on QFX5200-VC, only the local member's table details are displayed. [PR1243758](#)
- On the QFX10000 line of switches, sFlow monitoring technology output might display a negative number of samples after a long run. As a workaround, issue the **clear sflow collector** command to show or reset the count. [PR1244080](#)
- On the QFX10000-12C-DWDM coherent line card, link flapping is sometimes experienced when MACsec is enabled on the Ethernet interfaces. [PR1253703](#)
- The 50-Gigabit Ethernet(channelized from 100-Gigabit Ethernet) interfaces might not come up between QFX5200 switches running Junos OS Releases 17.2 and 15.1X53-D210.7 until the FEC is disabled on QFX5200 switches running Junos OS Release 17.2. [PR1258524](#)
- The management process (daemon) might crash if the OpenConfig package is installed immediately or within minutes of Network Agent package installation. This is a transient issue and does not impact any functionality. There is no action needed from your side in response to the crash. As a workaround, install OpenConfig before installing Network Agent. [PR1265815](#)
- For QFX10000 switches that do not support a discontinuous mask within the source address or the destination address of a firewall filter, when you commit a firewall filter with such a discontinuous mask prefix (for example, x.x.x.x/255.255.0.240), the commit is successful but the filter does not take effect (the firewall compilation returns an error because a discontinuous IP address mask is not supported and the filter is not programmed in hardware). [PR1267498](#)
- On disabling interfaces with MPLS or LSP configuration, the PathErr message is not being received on link failure under **show mpls container-lsp name**. [PR1275392](#)
- After the analytic configurations are committed and traffic is started, the analytic statistics might not work properly and might result in a core file being generated. [PR1277030](#)
- The mgd process might panic after you modify the aggregated Ethernet interface members under the **ethernet-switching vlan** configuration statement. As a result of the panic, the remote session might be terminated. [PR1325736](#)
- The QFX5200 and QFX10002 devices that have been shipped with Junos OS Releases 15.1X53-D21, 15.1X53-D30, 15.1X53-D31, 15.1X53-D32, 15.1X53-D33 and 15.1X53-D60 or have been upgraded to these releases using the .bin or .iso images might contain an unintended additional Open Network Install Environment (ONIE) partition. [PR1335713](#)
- On QFX10000 switches, the syslog error messages might be seen after you configure multiple interfaces that include LAG Interfaces under the protocol sFlow. Example of error messages: **Mar 13 12:04:24 host1 fpc0 expr\_dfw\_asic\_action\_update\_sflow\_sample\_id:2578 dfw inst lookup failed**

IFD\_EGRESS\_IMPL\_FILTER Mar 13 12:04:24 host1 fpc0 Sflow prds\_sflow\_add\_sample\_in\_hw(442):  
 Sample class (60): Implicit-filter binding set error Mar 13 12:04:24 host1 fpc0 Sflow  
 prds\_sflow\_handle\_int\_event(927): Error(1000) while enabling sflow in hw for intf 560 [PR1346493](#)

- On QFX10000 switches in a DDoS scenario, incorrect DDoS counter values and syslog messages might be seen after a specific protocol statistics is manually cleared. [PR1351212](#)
- On QFX5110, the FEC value for 100-gigabit optics is not being displayed when the expected behavior is for the FEC to be shown as **NONE**. On QFX10002-72Q, the FEC for 40-gigabit optics is being displayed as **NONE** when expected behavior is for the FEC not to be displayed. On QFX10008, the FEC for 40-gigabit optics is being displayed as **NONE** when the expected behavior is for the FEC not to be displayed. [PR1360948](#)
- When MC-LAG is configured with **Force-Up** enabled on MC-LAG nodes, the value of the LACP admin key must not match with that of the access or customer-edge device. [PR1362346](#)
- On QFX5000 switches, the tagged traffic is not passed through the untagged interface in the EVPN-VXLAN scenario if the **ethernet-bridge** configuration statement is configured. [PR1366336](#)

## EVPN

- In a EVPN-VXLAN scenario, a previously learned MAC address from a remote Ethernet segment Identifier (ESI) cannot be changed to local even if it is connected directly. The MAC address of the host might remain as learned from the ESI instead from the local interface until the MAC address is aged out.

[PR1303202](#)

## Interfaces and Chassis

- On QFX5100 switches, with MAC and ARP configurations inside an interface address configuration, an error message that says an IRB interface and an aggregated Ethernet logical interface do not belong to the same routing instance might be displayed, even though they do belong to the same routing instance.

[PR1239191](#)

## Layer 2 Features

- After deleting and re-adding 1000 LAG interfaces, the traffic drop is experienced for some time even though all LAG interfaces come up. [PR1289546](#)

## Network Management and Monitoring

- The default syslog level is LOG\_NOTICE in the default configuration. The SNMP\_TRAP\_LINK\_UP for the physical interface (IFD) was logged as LOG\_INFO from day one. To help debug physical link up issues, SNMP\_TRAP\_LINK\_UP events will be logged by default. [PR1287244](#)

## Platform and Infrastructure

- On all devices running Junos OS, the **file copy** command uses **/var/home/<user>** as a temporary staging directory for a non-root user, and uses **/var/tmp** for a root user. When you issue the **file copy user@x.x.x.x:/dir/ /var/tmp/** to copy a file to the device, and if the file you are trying to transfer is larger than the size of the temporary staging directory, then the copy operation fails. [PR1195599](#)
- Every load override and rollback operation increases the refcount by 1, and after it reaches the maximum value of 65,535, an mgd crash might be observed and the session might be terminated. When mgd crashes, the active lock might remain up, preventing any further commits. [PR1313158](#)

## Routing Protocols

- On the QFX10000 line of switches, traffic drop is seen with IS-IS version 6 traffic during convergence in either of the following two scenarios: 1) While bringing up the ports after bringing them down. 2.

- While the FPC comes online after an FPC restart. This behavior is seen when one of the IS-IS version 6 session flaps. [PR1190180](#)
- On EX4300, EX4600, QFX5100, or QFX10000 switches, traffic drops might occur in MC-LAG configurations. This occurs when an interchassis data link (ICL) interface and then the MC-LAG interface are brought up. The traffic drop occurs because the ARP next-hop update is not recognized on the Packet Forwarding Engine. To recover the traffic path over the MC-LAG interfaces, issue the **clear arp** command. To avoid the issue, enable ICL interfaces and MC-LAG interfaces at the same time. [PR1236201](#)

SEE ALSO

<a href="#">New and Changed Features   318</a>
<a href="#">Changes in Behavior and Syntax   338</a>
<a href="#">Known Behavior   344</a>
<a href="#">Resolved Issues   351</a>
<a href="#">Documentation Updates   363</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   364</a>
<a href="#">Product Compatibility   376</a>

## Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R3 | 352](#)
- [Resolved Issues: 17.2R2 | 359](#)
- [Resolved Issues: 17.2R1 | 362](#)

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: 17.2R3

### General Routing

- On QFX5110 switches, dcpfe might generate core files after the applied lo0 FF term is changed in scaled conditions. [PR1241733](#)
- The LAG interface input bytes counter continuously decrements when no packets come in. [PR1266062](#)
- When the backup state configured for the static link protection mode is down, the primary port goes to the down state instead of the secondary port, and the secondary port remains in the up state. [PR1276156](#)
- The **show security macsec statistics** command does not show the expected results. [PR1283544](#)
- After upgrading the QFX5100 or the EX4600 switch to Junos OS Release 16.1 or later from Junos OS Release 15.1, the commit warning **/boot/ffp.cookie+** might be seen. [PR1283917](#)
- On QFX5100 switches, an aggregated Ethernet interface might flap upon commit if an explicit speed is configured on an aggregated Ethernet member interface. [PR1284495](#)
- The BFD sessions might flap if BFD is configured over IRB interfaces. [PR1284743](#)
- OVSDB and OpenFlow have some limitations on QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016 switches running Junos OS Releases 17.1R1, 17.1R2, and 17.2R1. [PR1288227](#)
- Protocols might flap when you disable the aggregated Ethernet member link. [PR1289703](#)
- The storm-control flags are not set after a Routing Engine switchover. [PR1290246](#)
- On QFX-5100, the fxpc process generates a core file. [PR1294033](#)
- The 1-Gigabit Ethernet port on QFX10008 becomes unusable after inserting a third-party SFP-T optic. [PR1294394](#)
- The received ARP reply packet, whose destination MAC address is the same as the MAC address of the IRB interface, might be flooded on the VLAN. [PR1294530](#)
- The 40-Gigabit Ethernet interface might not come up if a specific vendor's DAC cable is used. [PR1296011](#)
- The DHCP client is not working on the replacement build release. [PR1296774](#)
- On QFX Series platforms, the connectivity of IPv4 might be lost if the logical interface (IFL) gl2d-property (eth) bit is set to 0. [PR1297594](#)
- On QFX Series platforms with ZTP environment, the DHCP clients are not getting an IP address if the DHCP pool with the /31 subnet is configured. [PR1298234](#)
- The dcpfe process might crash and restart on MC-LAG active and standby nodes when there is an ARP or NDP next-hop change. [PR1299112](#)
- The disabled 10-Gigabit Ethernet interfaces might stay up on the QFX10000 line of switches. [PR1300775](#)
- On QFX10008 or QFX10016 switches, a commit error is seen when mixed speed is configured. [PR1301923](#)



- The rpd might crash when the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statements are toggled. [PR1302504](#)
- The sFlow records are missing the **extendedType ROUTER** fields as well as an outbound interface for traffic that is using the BGP multipath. [PR1303236](#)
- If MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic, and display an error message: **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)
- Switches running 32-bit Junos OS might generate an rpd core file when traceoptions are enabled. [PR1305440](#)
- Digital optical monitoring statistics cannot be received through the CLI in Junos OS Release 15.1X53 through Release 17.x. [PR1305506](#)
- On QFX5200 switches, the new apply group is not being applied to the Virtual Chassis after a reboot. [PR1305520](#)
- The QFX5100 switch crashes and the fxcp process generates an core file. [PR1306768](#)
- Some error messages can be observed in a EVPN-VXLAN setup. [PR1307014](#)
- The QSFP+4x10G-IR channelized interface is down between QFX5200 and PTX5000. [PR1307400](#)
- The QFX5200 switch does not send out any frames. [PR1308443](#)
- The runtime PPS statistics value might show zero for a subinterface of an aggregated Ethernet interface. [PR1309485](#)
- Traffic loss might be seen if you send traffic through the 40-Gigabit Ethernet interface. [PR1309613](#)
- Some log messages are seen on the QFX5110 platform when plugging in an SFP-SX. [PR1311279](#)
- One of the aggregated Ethernet members 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](#)
- A memory leak is seen for the dot1xd process. [PR1313578](#)
- The AOC link between QFX5200 and its peer might stay down after the QFX5200 switch reboots. [PR1314323](#)
- Certain IGMP join packets cannot be processed correctly at a high rate. [PR1314382](#)
- Transit traffic over a GRE tunnel might hit the CPU and trigger a DDoS violation on the Layer 3 next hop. [PR1315773](#)
- The Packet Forwarding Engine might crash after changing analyzer configuration, if the output includes a LAG interface. [PR1316245](#)
- On an Layer 2 next-generation switch platform (EX4300/EX4600/EX9200/QFX5100/QFX10000), l2cpd might generate core files repeatedly if an interface is connected to a VoIP product with LLDP and LLDP-MED is enabled. [PR1317114](#)

- After zeroizing, the QFX5100 switch treats the 40-gigabit AOC uplink as 4x10-gigabit breakout with **auto-channelization** enabled. [PR1317872](#)
- Packets such as TDLS without the IP header are looped between virtual gateways. [PR1318382](#)
- The optic interface transmits power even after it has been administratively shut down. [PR1318997](#)
- Packets might be dropped for 4-60 seconds when the master Routing Engine is rebooted in a Virtual Chassis. [PR1319146](#)
- The chassis MIB SNMP OIDs for VC-B member chassis are not available after MX Series Virtual Chassis undergoes unified ISSU. [PR1320370](#)
- The FPCs go offline in some situations. [PR1321198](#)
- The OpenFlow session cannot be established correctly with controller and interfaces options configured on QFX5100 switches. [PR1323273](#)
- Update the new firmware versions for jfirmware package issues of 100G-PSM4 and 100G-AOC line cards. [PR1323321](#)
- The VLAN or VLAN bridge might not be added or deleted if there is an IFBD HW token limit exhaustion. [PR1325217](#)
- ARP request packets might not be flooded on QFX5110. [PR1326022](#)
- Poor performance when clearing scaled DHCP relay bindings [PR1326922](#)
- The major alarm **Fan & PSU Airflow direction mismatch** might be seen after the management cable is removed. [PR1327561](#)
- In an multihoming EVPN-VXLAN scenario with Service Provider style interface, deleting one VXLAN might cause traffic to loop on another VXLAN . [PR1327978](#)
- On QFX5100 series platforms, in some cases, CoS (class of Service) configuration is not properly applied in the Packet Forwarding Engine, leading to unexpected egress traffic drop on some interfaces.[PR1329141](#)
- The etherStatsCRCAlignErrors counters might disappear in the SNMP tree. [PR1329713](#)
- After a commit, members of the Virtual Chassis or VCF are split and some members might get disconnected.. [PR1330132](#)
- After an IP move, the ARP table information is not in synchronization between the two spines. [PR1330663](#)
- The rpd generates core files on the new backup Routing Engine at **task\_quit,task\_terminate\_timer\_callback,task\_timer\_dispatch,task\_scheduler** after NSR and GRES are disabled. [PR1330750](#)
- On QFX10002 switches, the **out of HMC range** and **HMC READ failed** error messages seen. [PR1332251](#)
- The DHCPv6 SOLICIT message is dropped. [PR1334680](#)
- The SNMP jnxBoxDescr OID returns different values when upgrading to Junos OS Release 17.2. [PR1337798](#)

- The analyzer status might show as down when port mirroring is configured to mirror packets from an aggregated Ethernet member. [PR1338564](#)
- The DDoS counters for OSPF might not increment. [PR1339364](#)
- The l2ald generates core files at `../../../../src/junos/usr.sbin/l2ald/l2ald_vxlan_evpn.c:1603`, when the host is moved between two multihomed interfaces. [PR1339543](#)
- On QFX10000 switches, broadcast frames might be modified with the ethertype 0x8850. [PR1343575](#)
- The fxpc process might generate a core file when you are removing the VXLAN configuration. [PR1345231](#)
- The statistics process pfed might generate core files during an upgrade between certain releases. [PR1346925](#)
- The QFX5100-48T 10-Gigabit Ethernet interface might be autonegotiated at 100 Mbps speed instead of 10 Gbps. [PR1347144](#)
- On QFX5110-48S-4C switches , part numbers and serial numbers are not displayed for any of the 10-Gigabit optics/DAC connected. [PR1347634](#)
- The 40-Gigabit Ethernet port on a QFX5100 has an interoperability issue with some other vendors. [PR1349664](#)
- The pfed process might consume high CPU if subscriber or interface statistics are used at a large scale. [PR1351203](#)
- The GTP traffic might not be hashed correctly for the aggregated Ethernet interface. [PR1351518](#)
- On QFX5000 switches, ARP learning might fail after the interface MAC address is changed. [PR1353241](#)
- The SFP-LX10 on QFX5110 might fail to connect with another device because it is not being able to negotiate the port speed. [PR1353677](#)
- The major alarms might be seen when a QFX10000 switch is booting up. [PR1354582](#)
- A commit error is observed if the device is downgraded from Junos OS Release 18.2 or Release 18.3 to Junos OS Release 17.3R3. [PR1355542](#)
- On QFX5110 platforms, the transceiver needs to be reinserted after autonegotiation is enabled or disabled. [PR1355746](#)
- Unable to create QFX5200 Virtual Chassis with 100-gigabit DACs. [PR1360721](#)
- The VME interface might be unreachable after link flap of em0 on the master FPC. [PR1362437](#)
- Traffic might not be forwarded when the member link of the aggregated Ethernet is added or deleted. [PR1362653](#)
- The 1-Gigabit Ethernet interface might stop working when **auto-negotiation** is disabled by default. [PR1362977](#)
- Traffic loss is observed when ISSU is performed with aggregated Ethernet interfaces configured with LACP. [PR1365316](#)

- On QFX5110, QFX5200, and QFX10000 switches, the root password recovery process does not work. [PR1365740](#)
- The chassisd might crash after the **show chassis hardware** CLI command is issued. [PR1366746](#)
- On QFX Series switches, IS-IS adjacency with a Cisco device might go down. [PR1368913](#)
- In certain routing topologies with sFlow configured, sampled packets might be duplicated and sFlow records are not sent to the collector. [PR1370464](#)
- MAC refresh packets might not be sent out from the new primary link after an RTG failover. [PR1372999](#)
- BOOTP packets might be dropped if **BOOTP-support** is not enabled at the global level. [PR1373807](#)
- On QFX5100-48F-6Q, the Packet Forwarding Engine might display **DISCARD next-hop for overlay-bgp-lo0-ip** in a leave-spine topology. [PR1380795](#)
- The master Virtual Chassis is copying **/var/db/ovsdatabase** to the backup every 10 seconds that causes a high write IO and shortens the SSD lifetime in an Open vSwitch Database (OVSDb) environment. [PR1381888](#)
- The Packet Forwarding Engine might crash if the GRE destination IP address is resolved over another GRE tunnel. [PR1382727](#)

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

- For some of the frame sizes throughput is not 100 percent. [PR1256671](#)
- Unable to filter packets with destination IP address as 224/4 and **DST MAC = QFX\_intf\_mac** on the loopback interface using a single match condition for source address 224.0.0.0/4. [PR1354377](#)

#### **EVPN**

- Next-hop installation error messages are seen on the QFX10000 line of switches. [PR1258930](#)
- The dynamic routing protocols might not work correctly over the IRB interface in an EVPN-VXLAN scenario with ECMP. [PR1301521](#)
- The EVPN proxy ARP cannot work properly if **system arp passive-learning** is configured. [PR1312672](#)
- A VXLAN traffic loss is observed after deleting and adding the VLANs. [PR1318045](#)
- The remote ARP entry might cause an error in an EVPN-VXLAN Layer 3 gateway scenario with multihoming mode. [PR1326691](#)
- The MAC movement between a remote VTEP and a local VTEP might cause traffic to be transmitted incorrectly in an EVPN-VXLAN scenario. [PR1335431](#)
- The ARP entry might be deleted in a redundant Layer 3 gateway EVPN-VXLAN scenario after a IP address move happens. [PR1336185](#)
- Configuring **encapsulate-inner-vlan** on the partial VXLANs might cause a traffic impact. [PR1337953](#)

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

- When **igmp-snooping** and **bpdu-block-on-edge** are enabled, the IP protocol multicast traffic sourced by the kernel, such as OSPF and VRRP traffic, gets dropped at the Packet Forwarding Engine level. [PR1301773](#)

### **Interfaces and Chassis**

- Multicast data packets are looping in an MC-LAG scenario. [PR1281646](#)
- There is an ARP reply drop in an MC-LAG scenario. [PR1282349](#)
- Upgrading might encounter commit failure if **redundancy-group-id-list** is not configured under ICCP. [PR1311009](#)
- Packets might be dropped on an ICL of MC-LAG peer where MC-LAG is up. [PR1345316](#)
- If the CVLAN's range is 16, then traffic might not pass through the 16 VLANs in a Q-in-Q scenario. [PR1345994](#)
- The MC-LAG peer does not send the ARP request to the host. [PR1360216](#)

### **Layer 2 Features**

- The QFX10000 line of switches transmit packets that exceed the interface MTU. [PR1306724](#)
- The **bpdu-block-on-edge** configuration does not work correctly when **fast-tune** is enabled [PR1307440](#)
- On the QFX10000 line of switches, the NLB heartbeat packets might be dropped. [PR1322183](#)
- The ARP entry might be learned on STP blocking ports. [PR1324245](#)
- MAC learning might fail for the device on extended port of a satellite device after MAC move occurs in a Junos Fusion scenario. [PR1324579](#)
- The DHCP Discover packets might be looped in an MC-LAG and DHCP-Relay scenario. [PR1325425](#)
- On QFX5100, with multiple logical units configured on an interface, **input-vlan-map POP** does not remove the outer VLAN tag when Q-in-Q and VXLAN are involved. [PR1331722](#)
- The operation of pushing a VLAN tag does not work for VXLAN local switching tunneled Q-in-Q traffic. [PR1332346](#)
- The **flexible-vlan-tagging** and **family ethernet-switching** interface configurations do not work on QFX10000. [PR1337311](#)
- Broadcast frames might be modified with the ethertype 0x8850. [PR1343575](#)
- On random initialization of a QFX5100, the programming of storm control profile is missed within hardware on random interfaces. [PR1354889](#)
- When **native-vlan-id** is configured, LACP packets are dropped after a reboot. [PR1361054](#)
- Hashing does not work for the IPv6 packet encapsulated in a VXLAN scenario. [PR1368258](#)

- When **native-vlan-id** is configured for aggregated Ethernet interfaces, the LACP session to multihomed server goes down. [PR1369424](#)
- The DHCP discover packets might be dropped if VXLAN is configured. [PR1377521](#)

### **Layer 2 Ethernet Services**

- The `jdhcpd` process generates core files after DHCP configuration changes are made. [PR1324800](#)

### **Multiprotocol Label Switching (MPLS)**

- On QFX5100, unified ISSU is not supported with MPLS configuration. [PR1264786](#)
- The DHCP clients cannot get IP addresses over a BGP Layer 3 VPN. [PR1303442](#)
- The LSP stops transferring or passing traffic after an MPLS route is changed. [PR1309058](#)
- The MPLS forwarding might not happen properly for some LSPs. [PR1319379](#)
- The `rpd` might crash on the backup Routing Engine because of memory exhaustion. [PR1328974](#)
- The hot standby for Layer 2 circuit does not work on QFX5100. [PR1329720](#)
- The LSP might remain up even if no path is acceptable because of the CSPF failure. [PR1365653](#)
- The NO-PROPAGATE-TTL flag acts on an MPLS swap operation. [PR1366804](#)

### **Platform and Infrastructure**

- Traffic loss might be observed for about 10 seconds if the master member FPC reboots. [PR1283702](#)
- The `dexp` process might crash after the **set system commit delta-export** command is committed. [PR1284788](#)
- An `l2ald` crash is seen while changing the configuration. [PR1294075](#)
- The OSPFv3 authentication using IPsec SA does not work if you are using IPsec to authenticate OSPFv3 neighbors on some QFX Series platforms. [PR1301428](#)
- The directories and files under `/var/db/scripts` lose execution permission or the `jet` directory is missing under `/var/db/scripts`, causing the **error: Invalid directory: No such file or directory** error during the commit. [PR1328570](#)
- Traffic is silently discarded with indirect next hop and load balancing. [PR1376057](#)

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

- The `rpd` might crash if **vrf-target auto** is configured under **routing-instance**. [PR1301721](#)

### **Routing Protocols**

- On the QFX10000 line of switches, filter-based forwarding (FBF) with the **next-ip**, **next-ip6**, and **next-interface** configurations is not working. [PR1289642](#)
- In a data center environment with EVPN-VXLAN and proxy MAC plus IP advertisement enabled on a Layer 3 gateway, the state for some MACs might be lost during MAC moves. [PR1291118](#)

- The IPv6 multicast traffic drop occurs in a PIM SSM scenario. [PR1292519](#)
- The dcpfe process might crash after a period of idle time on QFX10000 switches. [PR1294055](#)
- The mcsnoopd generates a core file 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/libjtask/base/task_scheduler.c:275.` [PR1305239](#)
- On QFX5100 switches, the Packet Forwarding Engine is unable to delete the next-hop HW token for reject route, leading the **brcm\_nh\_l3\_hw\_install (-6 Table full)** error. [PR1307009](#)
- Packet drop is seen when programming for GRE traffic. [PR1308438](#)
- Some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update. [PR1320723](#)
- IS-IS Layer 2 hello packets are dropped when they come from a Brocade device. [PR1325436](#)
- The looped back IRB interface is not accessible to the remote network. [PR1333019](#)
- Loopback filter does not work on QFX5100 and DFW errors are seen in the logs. [PR1336137](#)
- On QFX5100 switches, parity errors in the Layer 3 IPv4 table in the Packet Forwarding Engine memory might cause traffic to be discarded silently. [PR1364657](#)

### ***Software Installation and Upgrade***

- Commit might fail in single-user mode. [PR1368986](#)

## **Resolved Issues: 17.2R2**

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

- On QFX5100, EX4300, or EX4600, traffic might be dropped when there is more than one forwarding class under the **[forwarding-class-sets]** hierarchy. [PR1255077](#)
- Storm control might not be programmed correctly in the Packet Forwarding Engine if it is applied with a port-speed configuration in a single commit. [PR1255562](#)

### **Dynamic Host Configuration Protocol (DHCP)**

- DHCP reply packets are not relayed by the DHCP relay when there is a GRE tunnel. [PR1198982](#)

### **EVPNs**

- Route target per bridge domain for EVPN is not supported. [PR1244956](#)

### **General Routing**

- QFX100002 and QFX5110 generated an L2ALD core file for an unknown reason at: l2ald\_mac\_process\_update\_fwd\_entry\_mask , l2ald\_mclag\_update\_change\_for\_learn\_mask , logging , vlogging , vlogging\_event. [PR1264432](#)
- The jdncpd process might crash and DHCP does not work if scaling prefixes are configured under the [policy-options prefix-list \*] hierarchy. [PR1272646](#)
- The l2ald memory might leak for every IPv6 ND message it receives from peer MC-LAG and it is not freeing the memory allocated. [PR1277203](#)
- Multicast Listener Discovery (MLD) messages are seen continuously on QFX switches if the management ports are connected through a network. [PR1277618](#)
- Analytics json data format reporting incorrect value for 'rxbps' counter. [PR1285434](#)
- OVSDB and Openflow are caveated for QFX 5110, 5200, 10002, 10008, 10016 platforms in Junos OS Release 17.1R1, 17.1R2, and 17.2R1. [PR1288227](#)
- DCPFE might crash and restart on MC-LAG active and standby node when ARP/NDP next-hop change. [PR1299112](#)

### **Hardware**

- ULC-60S-6Q LC on QFX10008: the port becomes unusable after inserting non Juniper SFP-T optic. [PR1294394](#)

### **Infrastructure**

- On QFX10000 switches, match "pps"} O/P is not returning any values and sometimes it is completely stuck. [PR1250328](#)
- Disabled 10G interfaces might stay up on QFX10000 switches. [PR1300775](#)

### **Interfaces and Chassis**

- The traffic might be dropped in some rare conditions. [PR1241297](#)
- FPC Major Alarm might be seen with error messages "DLU: ilp memory cache error" & "DLU: ilp prot1 detected\_imem\_even error". [PR1251154](#)
- QFX5110: MC-LAG VRRP: Multicast traffic is not forwarded to MC-ae interface after deactivating and reactivating that interface. [PR1257586](#)
- Interfaces do not come up randomly after a line card rebooted. [PR1262839](#)



- Description for 40G-AOC cable in **show chassis hardware** shows UNKNOWN. [PR1269018](#)
- The 40G interface might flap between QFX5100 and other product. [PR1273861](#)
- QFX10000: Observed ot- link flap whenever an optics tca alarm is raised, but there is no loss of service and no traffic loss observed. [PR1279351](#)
- MAC pause frames might increase when SXE interfaces are erroneously configured. [PR1281123](#)
- Traffic might not be received on a 1G interface if autonegotiation is disabled and speed/duplex is configured on QFX and peer end. [PR1292275](#)
- High heap memory utilization might be seen if multiple SFP-T optics are inserted or **set interface <> link-mode full-duplex** is enabled. [PR1294208](#)

#### ***Junos Fusion Provider Edge***

- In a dual access device scenario, when you disable a cascade port, the extended port physical interfaces are marked as being down. [PR1232924](#)

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

- Native VLAN on an aggregated Ethernet interface terminated on multiple satellite devices. [PR1305698](#)

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

- Action-shutdown in storm-control does not bring physical interface down. [PR1240845](#)
- Packets are getting dropped if outer TPID is set with 0x9100. [PR1267178](#)

#### ***Multiprotocol Label Switching (MPLS)***

- Resolving static LSPs next hops. [PR1259238](#)
- QFX5110 MPLS: dcpfe core noticed during the MPLS ingress and egress scale tests. [PR1263201](#)

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

- Dropping the TCP RST packet incorrectly on PFE might cause traffic drop. [PR1269202](#)

#### ***Routing Protocols***

- After running restart routing in the master Routing Engine, the PIM join states of VXLAN multicast groups in the backup Routing Engine are not in sync with the master Routing Engine. [PR1255480](#)
- BGP session failed to establish over IPv6 link-local address. [PR1267565](#)
- IPv4 traffic drops when changing the member interface of the LAG. [PR1270011](#)
- The fxpc process might crash and restart when the fxpc process tries to access already freed up memory. [PR1271825](#)

- GRE tunnel traffic doesn't switch over to the alternate path if the primary path to tunnel destination changes. [PR1287249](#)
- UDP traffic with destination port 520 and 521 is discarded on QFX5110 switches after a Junos OS upgrade. [PR1287271](#)

### **Software Installation and Upgrade**

- When upgrading from 15.1X53-D62 to 17.1R1 or 17.2R1, protocols evpn vni-options vni vrf-target configuration is missing and customer needs to add the missing configuration. [PR1243105](#)

### **Virtual Chassis**

- When you add a QFX5100 switch to the VCF, the following error message is seen: `?ch__map_alarm_id alarm ignored: object 0x7e reason?.` [PR1234780](#)
- VCF - NSSU : the next member/group begin to reboot before the previous one ready caused service down [PR1272240](#)

### **VLAN Infrastructure**

- VLAN association is not being updated in the Ethernet switching table when the device is configured in single supplicant mode. [PR1283880](#)

## **Resolved Issues: 17.2R1**

### **General Routing**

- DHCP Reply packets are not relayed by the DHCP Relay when there is a GRE tunnel. [PR1198982](#)
- On QFX10008 and QFX10016-60x10G ULC 1G mode is not supported in Junos OS Release 17.1R1. [PR1239091](#)
- sFlow might show a negative count for a number of samples after a long run. [PR1244080](#)
- On QFX5100, **show interface** incorrectly displays an interface as 'Link-mode: Auto Speed: Auto' even though the interface is configured for, and up at, 100M/Full. [PR1260986](#)
- On QFX5200, the error log **ifd ifd-number; does not exist** might appear during an SNMP query and the SNMP query might be delayed. [PR1263794](#)
- QFX5100 VCF: Removing force-up causes return-traffic to be dropped by leaf (to spine). [PR1264650](#)
- Description for 40G-AOC cable in **show chassis hardware** shows UNKNOWN. [PR1269018](#)

### **Layer 2 Features**

- If RTG and VSTP are configured on the same VLAN, communication doesn't work over RTG interfaces. [PR1230750](#)
- On QFX10000 Series switches, in a multichassis link aggregation (MC-LAG) scenario, single-homed link (S-Link) MAC might not be learned before the MAC timeout on remote MC-LAG peer. [PR1260316](#)

- Flexible tagged LAG interface might go down when configuring native VLAN. [PR1262529](#)
- The QFX5100, QFX5110, and QFX5200 switches do not transfer BPDU packets though xSTP is disabled. [PR1262847](#)

**Routing Protocols**

- VCF does not forward BUM after fabric-tree-root is configured. [PR1257984](#)
- IPv4 traffic drops when changing the member interface of the LAG. [PR1270011](#)

SEE ALSO

<a href="#">New and Changed Features   318</a>
<a href="#">Changes in Behavior and Syntax   338</a>
<a href="#">Known Behavior   344</a>
<a href="#">Known Issues   347</a>
<a href="#">Documentation Updates   363</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   364</a>
<a href="#">Product Compatibility   376</a>

**Documentation Updates**

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.2R3.

SEE ALSO

<a href="#">New and Changed Features   318</a>
<a href="#">Changes in Behavior and Syntax   338</a>
<a href="#">Known Behavior   344</a>
<a href="#">Known Issues   347</a>
<a href="#">Resolved Issues   351</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   364</a>
<a href="#">Product Compatibility   376</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Upgrading Software on QFX Series Switches | 364
- Installing the Software on QFX10002 Switches | 366
- 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 | 366
- Installing the Software on QFX10008 and QFX10016 Switches | 368
- Performing a Unified ISSU | 372
- Preparing the Switch for Software Installation | 373
- Upgrading the Software Using Unified ISSU | 373

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 [Junos OS 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 **17.2** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 17.2 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new `jinstall` package on the device.

**NOTE:** We recommend that you upgrade all software packages out of band using the console, because in-band connections are lost during the upgrade process.

Customers in the United States and Canada use the following command:

```
user@host> request system software add source/jinstall-host-qfx-5-17.2R3.n-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 17.2 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

## 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-D43. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D43 or Junos OS Release 17.2R1.

**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 not compatible 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-10-17.2R3.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-10-17.2R3.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-10-17.2R3.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-10-17.2R3.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 373](#)
- [Upgrading the Software Using Unified ISSU on page 373](#)

## 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.

## 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-132\_x51\_vjunos.domestic.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-qfx-5-13.2X51-D15.4-domestic ...
Install jinstall-qfx-5-13.2X51-D15.4-domestic 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.

SEE ALSO

[New and Changed Features | 318](#)

Changes in Behavior and Syntax	338
Known Behavior	344
Known Issues	347
Resolved Issues	351
Documentation Updates	363
Product Compatibility	376

# Product Compatibility

## IN THIS SECTION

- Hardware Compatibility | 376

## 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	318
Changes in Behavior and Syntax	338
Known Behavior	344
Known Issues	347
Resolved Issues	351



## 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**—Determine the features supported on MX Series, PTX Series, QFX Series devices. The Juniper Networks Feature Explorer is a Web-based app that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. <https://pathfinder.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. [prsearch.juniper.net](https://prsearch.juniper.net).
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](https://apps.juniper.net/hct/home)

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

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. [apps.juniper.net/compliance/](https://apps.juniper.net/compliance/).

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post sales technical support, you can access our tools and resources online or open a case with JTAC.

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- **JTAC policies**—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- **Product warranties**—For product warranty information, visit <https://support.juniper.net/support/warranty/>.
- **JTAC hours of operation**—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://support.juniper.net/support/>
- Search for known bugs: <https://kb.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>

- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://support.juniper.net/support/downloads/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://forums.juniper.net>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at <https://support.juniper.net/support/requesting-support/>.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to support@juniper.net. For documentation issues, fill out the bug report form located at <https://www.juniper.net/documentation/feedback/>.

## Revision History

3 September 2020—Revision 9, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

24 October 2019—Revision 8, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

6 June 2019—Revision 7, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

25 April 2019—Revision 6, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

14 February 2019—Revision 5, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

7 February 2019—Revision 4, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

10 January 2019—Revision 3, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

3 January 2019—Revision 2, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

26 November 2018—Revision 1, Junos OS Release 17.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

17 May 2018—Revision 11, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

12 April 2018—Revision 10, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

30 March 2018—Revision 9, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

15 February 2018—Revision 8, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion

11 January 2018—Revision 7, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

14 December 2017—Revision 6, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

16 November 2017—Revision 5, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

12 October 2017—Revision 4, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

6 October 2017—Revision 3, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

5 October 2017—Revision 2, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

28 September 2017—Revision 1, Junos OS Release 17.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

20 July 2017—Revision 6, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

6 July 2017—Revision 5, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

29 June 2017—Revision 4, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

20 June 2017—Revision 3, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

13 June 2017—Revision 2, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

6 June 2017—Revision 1, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, and Junos Fusion.

Copyright © 2017, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.