

Release Notes: Junos[®] OS Release 17.3R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

30 September 2021

Contents	Introduction 11
	Junos OS Release Notes for ACX Series 11
	New and Changed Features 12
	Release 17.3R2 New and Changed Features 12
	Release 17.3R1 New and Changed Features 12
	Changes in Behavior and Syntax 14
	General Routing 14
	Interfaces and Chassis 14
	Known Behavior 15
	Known Issues 15
	Class of Service 16
	Hierarchical Class of Service 16
	Interfaces and Chassis 17
	Software Installation and Upgrade 17
	Layer 2 Features 17
	Router 17
	Resolved Issues 18
	Resolved Issues: 17.3R2 18
	Resolved Issues: 17.3R1 19

Documentation Updates | 19

Migration, Upgrade, and Downgrade Instructions | 20

 Upgrade and Downgrade Support Policy for Junos OS Releases | 20

Product Compatibility | 21

 Hardware Compatibility | 21

Junos OS Release Notes for EX Series Switches | 22

New and Changed Features | 22

 Release 17.3R2 New and Changed Features | 23

 Release 17.3R1 New and Changed Features | 23

Changes in Behavior and Syntax | 30

 Management | 30

 Network Management and Monitoring | 30

 Services Applications | 31

 VLAN Infrastructure | 31

Known Behavior | 32

 Authentication, Authorization, and Accounting (AAA) (RADIUS) | 33

 Platform and Infrastructure | 33

 Junos Fusion Enterprise | 33

Known Issues | 34

 Infrastructure | 34

 Layer 2 Features | 34

 MPLS | 35

 Network Management and Monitoring | 35

 Platform and Infrastructure | 35

 Virtual Chassis | 36

 VLAN Infrastructure | 36

Resolved Issues | 36

 Resolved Issues: 17.3R2 | 37

 Resolved Issues: 17.3R1 | 39

Documentation Updates | 40

 Traffic Management User Guide for EX4600 Switches | 40

Migration, Upgrade, and Downgrade Instructions | 41

 Upgrade and Downgrade Support Policy for Junos OS Releases | 41

Product Compatibility | 42

Hardware Compatibility | 42

Junos OS Release Notes for Junos Fusion Data Center | 43

New and Changed Features | 43

Changes in Behavior and Syntax | 44

Known Behavior | 44

Junos Fusion Data Center | 44

Known Issues | 45

Resolved Issues | 46

Resolved Issues: Release 17.3R2 | 46

Documentation Updates | 47

Migration, Upgrade, and Downgrade Instructions | 47

Basic Procedure for Upgrading an Aggregation Device | 48

Preparing the Switch for Satellite Device Conversion | 50

Autoconverting a Switch into a Satellite Device | 52

Manually Converting a Switch into a Satellite Device | 55

Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 57

Configuring Satellite Device Upgrade Groups | 58

Converting a Satellite Device to a Standalone Device | 60

Upgrade and Downgrade Support Policy for Junos OS Releases | 62

Downgrading from Release 17.3 | 62

Product Compatibility | 63

Hardware Compatibility | 63

Junos OS Release Notes for Junos Fusion Enterprise | 65

New and Changed Features | 65

Junos Fusion Enterprise | 66

Changes in Behavior and Syntax | 67

Junos Fusion Enterprise | 67

Known Behavior | 67

Junos Fusion Enterprise | 68

Known Issues | 68

Junos Fusion Enterprise | 69

Resolved Issues | 69

Resolved Issues: 17.3R2 | 70

Resolved Issues: 17.3R1 | 70

Documentation Updates | 71

Migration, Upgrade, and Downgrade Instructions | 71

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 72

Upgrading an Aggregation Device with Redundant Routing Engines | 74

Preparing the Switch for Satellite Device Conversion | 74

Converting a Satellite Device to a Standalone Switch | 76

Upgrade and Downgrade Support Policy for Junos OS Releases | 78

Downgrading from Release 17.3 | 79

Product Compatibility | 79

Hardware and Software Compatibility | 80

Hardware Compatibility Tool | 80

Junos OS Release Notes for Junos Fusion Provider Edge | 81

New and Changed Features | 81

Release 17.3R2 New and Changed Features | 82

Release 17.3R1 New and Changed Features | 82

Changes in Behavior and Syntax | 83

Known Behavior | 83

Known Issues | 84

Resolved Issues | 84

Resolved Issues: 17.3R2 | 85

Resolved Issues: 17.3R1 | 85

Documentation Updates | 85

Migration, Upgrade, and Downgrade Instructions | 86

Basic Procedure for Upgrading an Aggregation Device | 86

Upgrading an Aggregation Device with Redundant Routing Engines | 89

Preparing the Switch for Satellite Device Conversion | 89

Converting a Satellite Device to a Standalone Device | 91

Upgrading an Aggregation Device | 93

Upgrade and Downgrade Support Policy for Junos OS Releases | 93

Downgrading from Release 17.3 | 93

Product Compatibility | 94

Hardware Compatibility | 94

Junos OS Release Notes for MX Series 5G Universal Routing Platforms | 95

New and Changed Features | 96

Release 17.3R2 New and Changed Features | 96

Release 17.3R1 New and Changed Features | 96

Changes in Behavior and Syntax | 120

EVPNs | 120

General Routing | 121

Interfaces and Chassis | 121

Management | 121

MPLS | 122

Network Management and Monitoring | 122

Routing Protocols | 123

Security | 124

Services Application | 124

Subscriber Management and Services | 124

VLAN Infrastructure | 125

Known Behavior | 125

Class of Service (CoS) | 126

EVPN | 126

General Routing | 127

High Availability (HA) and Resiliency | 129

Interfaces and Chassis | 129

MPLS | 129

Routing Protocols | 130

Subscriber Management and Services | 130

Known Issues | 131

Application Layer Gateways (ALGs) | 131

Class of Service (CoS) | 131

EVPN | 132

Forwarding and Sampling | 133

General Routing | 133

Infrastructure | 137

- Interfaces and Chassis | **138**
- Layer 2 Ethernet Services | **138**
- Layer 2 Features | **139**
- MPLS | **139**
- Platform and Infrastructure | **140**
- Routing Protocols | **141**
- Services Applications | **143**
- Subscriber Access Management | **143**
- VPNs | **144**

Resolved Issues | **144**

- Resolved Issues: 17.3R2 | **145**
- Resolved Issues: 17.3R1 | **158**

Documentation Updates | **163**

- Subscriber Management Provisioning Guide | **164**

Migration, Upgrade, and Downgrade Instructions | **164**

- Basic Procedure for Upgrading to Release 17.3 | **165**
- Procedure to Upgrade to FreeBSD 10.x based Junos OS | **166**
- Procedure to Upgrade to FreeBSD 6.x based Junos OS | **168**
- Upgrade and Downgrade Support Policy for Junos OS Releases | **169**
- Upgrading a Router with Redundant Routing Engines | **170**
- Downgrading from Release 17.3 | **170**

Product Compatibility | **171**

- Hardware Compatibility | **171**

Junos OS Release Notes for NFX Series | **172**

New and Changed Features | **172**

- Juniper Device Manager | **173**

Changes in Behavior and Syntax | **173**

Known Behavior | **174**

Known Issues | **174**

Resolved Issues | **175**

Documentation Updates | **175**

Migration, Upgrade, and Downgrade Instructions | **176**

- Upgrade and Downgrade Support Policy for Junos OS Releases | **176**

Product Compatibility | 177

Hardware Compatibility | 177

Junos OS Release Notes for PTX Series Packet Transport Routers | 178

New and Changed Features | 178

Release 17.3R2 New and Changed Features | 179

Release 17.3R1 New and Changed Features | 179

Changes in Behavior and Syntax | 186

Forwarding and Sampling | 187

Interfaces and Chassis | 187

Management | 187

Network Management and Monitoring | 187

Services Application | 188

VLAN-Infrastructure | 189

Known Behavior | 189

General Routing | 189

Multiprotocol Label Switching (MPLS) | 190

Known Issues | 190

General Routing | 191

Interfaces and Chassis | 192

Routing Protocols | 192

Resolved Issues | 192

Resolved Issues: 17.3R2 | 193

Resolved Issues: 17.3R1 | 195

Documentation Updates | 195

Migration, Upgrade, and Downgrade Instructions | 196

Upgrade and Downgrade Support Policy for Junos OS Releases | 196

Upgrading a Router with Redundant Routing Engines | 197

Basic Procedure for Upgrading to Release 17.3 | 197

Product Compatibility | 200

Hardware Compatibility | 200

Junos OS Release Notes for the QFX Series | 201

New and Changed Features | 202

Release 17.3R2 New and Changed Features | 203

Release 17.3R1 New and Changed Features | 205

Changes in Behavior and Syntax | 218

Class of Service (CoS) | 219

EVPNs | 219

Management | 219

Network Management and Monitoring | 220

Virtual Chassis | 221

VLAN Infrastructure | 221

Known Behavior | 221

Class of Service (CoS) | 222

EVPNs | 222

High Availability (HA) and Resiliency | 223

Interfaces and Chassis | 223

Layer 2 Features | 223

Layer 3 Features | 224

Platform and Infrastructure | 224

Routing Protocols | 226

Virtual Chassis | 226

Known Issues | 227

Class of Service (CoS) | 228

EVPN | 228

Infrastructure | 229

Interfaces and Chassis | 229

Layer 2 Features | 229

Multicast | 230

Network Management and Monitoring | 230

Platform and Infrastructure | 230

Routing Protocols | 232

Resolved Issues | 233

Resolved Issues: 17.3R2 | 233

Resolved Issues: 17.3R1 | 237

Documentation Updates | 239

Traffic Management User Guide for the QFX Series | 239

Migration, Upgrade, and Downgrade Instructions | 240

Upgrading Software on QFX Series Switches | 240

Installing the Software on QFX10002 Switches | 243

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 | 243

Installing the Software on QFX10008 and QFX10016 Switches | 245

Performing a Unified ISSU | 249

Preparing the Switch for Software Installation | 250

Upgrading the Software Using Unified ISSU | 250

Product Compatibility | 253

Hardware Compatibility | 253

Junos OS Release Notes for SRX Series | 254

New and Changed Features | 254

Release 17.3R2 New and Changed Features | 254

Release 17.3R1 New and Changed Features | 255

Changes in Behavior and Syntax | 260

Known Behavior | 260

CLI | 261

Known Issues | 261

Authentication and Access Control | 262

Chassis Cluster | 262

Class of Service (CoS) | 263

CLI | 263

Flow-Based and Packet-Based Processing | 263

Interfaces and Routing | 264

J-Web | 264

Network Address Translation (NAT) | 264

Network Management and Monitoring | 264

Platform and Infrastructure | 265

Routing Policy and Firewall Filters | 265

Routing Protocols | 265

System Logs | 265

	VPNs 265
Resolved Issues 267	
	Resolved Issues: 17.3R2 268
	Resolved Issues: 17.3R1 269
Documentation Updates 270	
Migration, Upgrade, and Downgrade Instructions 270	
	Upgrade for Layer 2 Configuration 270
	Upgrade and Downgrade Scripts for Address Book Configuration 271
Product Compatibility 274	
	Hardware Compatibility 274
Upgrading Using Unified ISSU 275	
Compliance Advisor 275	
Finding More Information 275	
Documentation Feedback 275	
Requesting Technical Support 277	
	Self-Help Online Tools and Resources 277
	Opening a Case with JTAC 278
Revision History 278	

Introduction

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

These release notes accompany Junos OS Release 17.3R2 for the ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- New and Changed Features | 12
- Changes in Behavior and Syntax | 14
- Known Behavior | 15
- Known Issues | 15
- Resolved Issues | 18
- Documentation Updates | 19
- Migration, Upgrade, and Downgrade Instructions | 20
- Product Compatibility | 21

These release notes accompany Junos OS Release 17.3R2 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.3R2 New and Changed Features](#) | 12
- [Release 17.3R1 New and Changed Features](#) | 12

This section describes the new features or enhancements to existing features in Junos OS Release 17.3R2 for ACX Series Universal Metro Routers.

Release 17.3R2 New and Changed Features

There are no new features or enhancements to existing features in Junos OS Release 17.3R2 for ACX Series Universal Metro Routers.

Release 17.3R1 New and Changed Features

Hardware

- **Support for 100 MB Optics (ACX5000, ACX1000, ACX1100, ACX2100, ACX2200, ACX4000)**—Starting in Junos OS Release 17.3R1, ACX Series Universal Metro Routers (ACX500, ACX1000, ACX1100, ACX2100, ACX2200, ACX4000) support 100 MB Ethernet optics.

[See [Hardware Compatibility Tool](#)]

Class of Service

- **Support for hierarchical class of service (ACX5000)**—Starting in Junos OS Release 17.3R1, ACX5000 line of routers support hierarchical class of service. You can configure up to 8 queues per logical interface. Scheduling properties can be applied at both the physical and logical interface levels. Service providers will be able to support hierarchical class of service at multiple levels to meet the service level agreements and bandwidth allocations for subscribers.

To enable hierarchical scheduling, include the **hierarchical-scheduler** CLI statement at the physical interface level.

Hierarchical class of service can be enabled for Layer 3 VPN, VPLS, and VPWS services.

[See [Hierarchical Class of Service in ACX5000](#).]

Interfaces and Chassis

- **Support for limiting the number of MAC addresses learned from a logical interface (ACX5000)**—Starting in Junos OS Release 17.3R1, you can limit the number of MAC addresses learned from a logical interface

on the ACX5000 line of routers. The number of MAC entries learned on a logical interface can be limited by configuring a value for **interface-mac-limit**. The logical interface MAC limit allows the MAC address table space to be distributed among the different logical interfaces. The MAC limiting can be done for both VPLS and VLAN networks. The limits for a bridge domain and logical port can also be configured at the same time.

You can configure MAC address limit by enabling the **set protocols l2-learning global-no-hw-mac-learning** CLI command.

You can specify a limit for MAC addresses at a logical interface level by configuring a value for the **interface-mac-limit** CLI command.


[See [Configuring MAC Address Limits on a Logical Interface](#).]

- **Support for receiving multicast traffic in a VRF domain (ACX Series)**—Starting in Junos OS Release 17.3R1, ACX Series routers support multicast traffic to be received in a VRF domain.

[See [Configuring an Interface in the VRF Domain to Receive Multicast Traffic](#).]

Timing and Synchronization

- **Support for PTP grandmaster clock (ACX500)**—Starting in Junos OS Release 17.3R1, ACX500 line of routers supports the PTP grandmaster clock functionality. For an ACX500 router to act as a PTP grandmaster clock, the router needs to receive the timing information from a GPS receiver. ACX500 line of routers supports the integrated GNSS receiver, eliminating the need for an external GPS receiver.

**NOTE:** The grandmaster functionality is supported only on the ACX500 Indoor routers.

[See [Integrated Global Navigation Satellite System \(GNSS\) on ACX500 Series Routers](#) and [IEEE 1588v2 Precision Timing Protocol \(PTP\)](#).]

SEE ALSO

Changes in Behavior and Syntax	14
Known Behavior	15
Documentation Updates	19
Known Issues	15
Resolved Issues	18
Migration, Upgrade, and Downgrade Instructions	20
Product Compatibility	21

Changes in Behavior and Syntax

IN THIS SECTION

- [General Routing | 14](#)
- [Interfaces and Chassis | 14](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 17.3R2 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 Release 17.3R1, 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

- **Support for logical interfaces**—ACX5048 and ACX5096 routers do not support configuring more than 1000 logical interfaces.

SEE ALSO

New and Changed Features 12
Known Behavior 15
Documentation Updates 19
Known Issues 15
Resolved Issues 18
Migration, Upgrade, and Downgrade Instructions 20
Product Compatibility 21

Known Behavior

There are no known limitations in Junos OS Release 17.3R2 for the ACX Series.

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

SEE ALSO

New and Changed Features 12
Changes in Behavior and Syntax 14
Documentation Updates 19
Known Issues 15
Resolved Issues 18
Migration, Upgrade, and Downgrade Instructions 20
Product Compatibility 21

Known Issues

IN THIS SECTION

- [Class of Service | 16](#)
- [Hierarchical Class of Service | 16](#)
- [Interfaces and Chassis | 17](#)
- [Software Installation and Upgrade | 17](#)
- [Layer 2 Features | 17](#)
- [Router | 17](#)

This section lists the known issues in hardware and software in Junos OS Release 17.3R2 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.

Class of Service

- On ACX5000 line of routers, traffic drop is seen after performing ISSU when class of service is configured. [PR1299539](#)

Hierarchical Class of Service

- On ACX5000 line of routers, whenever you make a change to the queue modes and for the changes to take effect you will need to restart the PFE. [PR1256465](#)
- On ACX5000 line of routers, the **show class-of-service scheduler-hierarchy** CLI command is not supported. [PR1261835](#)
- On ACX5000 line of routers, the **show class-of-service interfaces queue *logical-interface-name*** CLI command does not show **Queue Buffer Usage** for a logical interface. As a workaround, you can use the PFE shell **show cos halp mmu buffer ifl** command to see the **Queue Buffer Usage** for a logical interface. [PR1272822](#)
- On ACX5000 line of routers, the **shared-buffer maximum** CLI statement for logical interface hierarchical class of service queues does not work correctly. [PR1275796](#)

Interfaces and Chassis

- On ACX Series routers, when link-speed is configured, the aggregate interface goes down permanently after the router reboots. [PR1022248](#)

Software Installation and Upgrade

- ISSU upgrade fails on the ACX5000 line of routers when an AE interface with VLAN map operation **push**, **push-push**, or **pop** is configured in a bridge with no **VLAN ID**. This occurs when the current running Junos OS image is already having an issue, causing the ISSU upgrade to fail. [PR1318771](#)

Layer 2 Features

- On ACX5000 line of routers, in a normal MAC learning mode, when incremental MAC traffic of higher range than the profile is received and after feb restarts, the MAC entries are not seen in the software CLI, although present in the hardware table. As a workaround, in the hardware MAC learning mode, delete the routing instance and reconfigure the routing instance again. In software MAC learning mode, deactivate the routing instance, clear the pending entries or allow the pending entries to be aged out and then activate the routing instance. [PR1277436](#)

Router

- On ACX500 line of routers, performance issues are seen on the ACX500 Indoor AC router. [PR1290278](#)
- On ACX Series routers, at certain instances, the CLI command and syslog shows FAN failure alarms although the fan is running at high speed.

```
user@host> show chassis alarms no-forwarding
alarms currently active
Alarm time          Class  Description
2010-01-01 00:12:04 UTC  Minor  Single FAN Failure
```

```
user@host> show chassis environment no-forwarding
```

Class	Item	Status
Measurement		
Fans	Fan 1	Check
	Fan 2	OK Spinning at high speed

[PR1127846](#)

SEE ALSO

New and Changed Features	12
Changes in Behavior and Syntax	14
Known Behavior	15
Documentation Updates	19
Resolved Issues	18
Migration, Upgrade, and Downgrade Instructions	20
Product Compatibility	21

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.3R2 | 18
- Resolved Issues: 17.3R1 | 19

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: 17.3R2

General Routing

- On ACX Series routers, the DHCP-RELAY requests with IRB interface are not forwarded after upgrade. [PR1243687](#)
- On ACX Series routers, transit ARP packets coming from logical interfaces that are part of a bridge domain or Layer 2 circuit were being sent ("punted") to the Routing Engine. [PR1263012](#)
- On ACX Series router, syslog error was seen on output/egress firewall filter. [PR1316588](#)
- On ACX5000 line of routers, when the management cable was removed, the **Fan & PSU Airflow direction mismatch** major alarm was seen. [PR1327561](#)

Resolved Issues: 17.3R1

Layer 2 Features

- All the XML duplications and unformatted output are addressed. For Example, histogram was just declared as a element inside pfkey container, with this change a new container is defined for histogram. [PR1271648](#)

SEE ALSO

New and Changed Features	12
Changes in Behavior and Syntax	14
Known Behavior	15
Documentation Updates	19
Known Issues	15
Migration, Upgrade, and Downgrade Instructions	20
Product Compatibility	21

Documentation Updates

There are no errata or changes in Junos OS Release 17.3R2 for the ACX Series documentation.

SEE ALSO

New and Changed Features	12
Changes in Behavior and Syntax	14
Known Behavior	15
Known Issues	15
Resolved Issues	18
Migration, Upgrade, and Downgrade Instructions	20
Product Compatibility	21

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 20](#)

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. 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://support.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 | 14](#)

[Known Behavior | 15](#)

[Documentation Updates | 19](#)

[Known Issues | 15](#)

[Resolved Issues | 18](#)

[Product Compatibility | 21](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 21](#)

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 | 14](#)

[Known Behavior | 15](#)

[Documentation Updates | 19](#)

[Known Issues | 15](#)

[Resolved Issues | 18](#)

[Migration, Upgrade, and Downgrade Instructions | 20](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- New and Changed Features | 22
- Changes in Behavior and Syntax | 30
- Known Behavior | 32
- Known Issues | 34
- Resolved Issues | 36
- Documentation Updates | 40
- Migration, Upgrade, and Downgrade Instructions | 41
- Product Compatibility | 42

These release notes accompany Junos OS Release 17.3R2 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.3R2 New and Changed Features | 23
- Release 17.3R1 New and Changed Features | 23

This section describes the new features and enhancements to existing features in Junos OS Release 17.3R2 for the EX Series.

NOTE: The following EX Series switches are supported in Junos OS Release 17.3R2: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.3R2, 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 17.3A1 for EX4300 and EX4600 Switches](#).

Release 17.3R2 New and Changed Features

There are no new features or enhancements to existing features for EX Series in Junos OS Release 17.3R2.

Release 17.3R1 New and Changed Features

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

- **Access control and authentication (EX4600 switches)**—Starting with Junos OS Release 17.3R1, EX4600 switches support controlling access to your network using 802.1X authentication and MAC RADIUS authentication.
 - 802.1X authentication provides port-based network access control (PNAC) as defined in the IEEE 802.1X standard. QFX5100 switches support 802.1X features including guest VLAN, private VLAN, server fail fallback, dynamic changes to a user session, RADIUS accounting, and configuration of port-filtering attributes on the RADIUS server using VSAs. You configure 802.1X authentication at the **[edit protocols dot1x]** hierarchy level.
 - MAC RADIUS authentication is used to authentic end devices independently of whether they are enabled for 802.1X authentication. You can permit end devices that are not 802.1X-enabled to access the LAN by configuring MAC RADIUS authentication on the switch interfaces to which the end devices are connected. You configure MAC RADIUS authentication at the **[edit protocols dot1x authenticator interface interface-name mac-radius]** hierarchy level.
- **IPv6 for RADIUS AAA (EX4300 and EX9200)**—Starting in Junos OS Release 17.3R1, EX4300 and EX9200 switches support IPv6 for user authentication, authorization, and accounting (AAA) using RADIUS servers,

in addition to the existing IPv4 support. You can specify which source address Junos OS uses to contact an external RADIUS server. To configure an IPv6 source address for RADIUS authentication, include the `source-address` statement at the `[edit system radius-server server-address]` hierarchy level. To configure an IPv6 source address for RADIUS accounting, include the `source-address` statement at the `[edit system accounting destination radius server server-address]` hierarchy level.

NOTE: If an IPv6 RADIUS server is configured without any source-address, default `::0` is considered to be the source address.

[See [source-address](#).]

- **Port bounce with CoA requests and framed-IPv6-address RADIUS attribute for AAA (EX4300 and EX9200)**—Starting in Junos OS Release 17.3R1, the port bounce feature is supported on EX4300 and EX9200 switches. Change of Authorization (CoA) requests are RADIUS messages sent from the authentication, authorization, and accounting (AAA) server to the switch. They are typically used to dynamically change the VLAN for the host based on device profiling. End devices such as printers do not have a mechanism to detect the VLAN change, so they do not renew the lease for their DHCP address in the new VLAN. The port bounce feature is used to force the end device to initiate DHCP re-negotiation by causing a link flap on the authenticated port. There is no configuration required to enable the port bounce feature. Framed-IPv6-Address is an additional RADIUS attribute to support clients with an IPv6 address. The attribute is included in the Access-Request message sent from the client to the AAA server.

[See [Understanding RADIUS-Initiated Changes to an Authorized User Session](#) and [Understanding 802.1X and RADIUS Accounting on Switches](#).]

EVPNs

- **EVPN type-5 route support (EX9200)**—Starting with Junos OS Release 17.3R1, you can configure type-5 routing in an Ethernet VPN (EVPN) environment. Type-5 routing, which advertises IP prefixes through EVPN, is used when the Layer 2 domain does not exist at the remote data centers or metro network peering points.

On EX9200 switches, two models are supported:

- Pure type-5 route without an overlay next hop and type-2 route (MPLS encapsulation only)
- Type-5 route with a gateway IRB interface as an overlay next hop and type-2 route (MPLS and VXLAN encapsulation)

To enable 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. To enable type-5 routing with a gateway IRB interface, include the `ip-prefix-routes advertise gateway-address` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level. Specify a gateway IRB

interface by including the **gateway-interface *irb-interface-name*** statement at the **[edit routing-instances *routing-instance-name* protocols evpn ip-prefix-routes]** hierarchy level.

[See [ip-prefix-routes](#).]

- **IPv6 support over IRB interfaces for EVPN (EX9200 switches)**—Starting in Junos OS Release 17.3R1, the Ethernet VPN (EVPN) integrated routing and bridging (IRB) solution supports IPv6 and the Neighborhood Discovery Protocol (NDP). NDP is used by IPv6 nodes on the same link to discover each other's presence, determine each other's Link Layer addresses, find routers, and maintain reachability information about the paths to active neighbors. IPv6 addresses over IRB for EVPN is supported for unique VLAN EVPN instances and for virtual switches with protocol EVPN instances.

[See [EVPN with IRB Solution Overview](#).]

- **EVPN multihoming with ESI per logical interface (EX9200)**—In releases before Junos OS Release 17.3R1, for EX9200 switches, you can configure an Ethernet segment identifier (ESI) only on a physical or aggregated Ethernet interface. In an EVPN-MPLS topology where a customer edge (CE) device is multihomed in active-standby or active-active mode to multiple provider edge (PE) devices, if a physical or aggregated Ethernet interface on an EX9200 switch is considered a non-designated forwarder (DF), the logical interfaces configured on the physical or aggregated Ethernet interface cannot be used for other services. Starting with Junos OS Release 17.3R1 for EX9200 switches, you can now configure an ESI on a logical interface. As a result, even if a logical interface is a non-DF, other logical interfaces on the same physical or aggregated Ethernet interface can still be used for other services.

[See [Example: Configuring an ESI on a Logical Interface for EVPN Multihoming](#).]

- **Layer 3 VXLAN gateway in EVPN-VXLAN topology with a two-layer IP fabric (EX9200)**—Starting with Junos OS Release 17.3R1, EX9200 switches can function as a Layer 3 VXLAN gateway, or spine device, in an EVPN-VXLAN topology with a two-layer IP fabric. In this role, the EX9200 switch uses integrated routing and bridging (IRB) interfaces to route traffic between hosts in different virtual networks (VNs) created by the Contrail virtualization software. When physical (bare-metal) servers in one VN need to communicate with other physical servers or virtual machines (VMs) in another VN, you can also configure an IRB interface as a default Layer 3 gateway that handles the inter-VN traffic for physical servers. In an EVPN-VXLAN topology where a provider edge (PE) device such as a Layer 2 VXLAN gateway or a Contrail vRouter is multihomed in active-active mode to two Layer 3 VXLAN gateways, you can configure redundant default gateways on the Layer 3 VXLAN gateways.

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#).]

Layer 2 Features

- **IRB in PVLAN (EX4600)**—Starting with Junos OS Release 17.3R1, you can configure an IRB interface in a private VLAN (PVLAN) so that devices in the community and isolated VLANs can communicate with each other and with devices outside the PVLAN at Layer 3 without requiring you to install a router.

[See [Example: Configuring a Private VLAN Spanning Multiple Switches with an IRB Interface](#).]

- **PVLAN and Q-in-Q configurations co-exist on a physical interface (EX4600)**—Starting with Junos OS Release 17.3R1, a private VLAN (PVLAN) configuration and a Q-in-Q tunneling configuration can co-exist

on the same Ethernet port. Q-in-Q requires a service provider configuration method, and PVLAN requires an enterprise configuration method. To enable both configurations to exist on the same physical interface, you must configure flexible Ethernet services to support dual methods of configuring logical interfaces.

[See [Understanding Flexible Ethernet Services Encapsulation on Switches](#).]

- **L2PT support for tunneling additional protocols (EX9200)**—Starting with Junos OS Release 17.3R1, you can configure Layer 2 protocol tunneling (L2PT) for the following new protocols on EX9200 switches: E-LMI, GVRP, IEEE 802.1X, IEEE802.3AH, LACP, LLDP, MMRP, MVRP, and UDLD.

[See [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#).]

- **L2PT support for tunneling additional protocols (EX4300)**—Starting with Junos OS Release 17.3R1, you can configure Layer 2 protocol tunneling (L2PT) for the following new protocols on EX4300 switches: E-LMI, IEEE 802.1X, MMRP, and UDLD.

[See [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#).]

Layer 3 Features

- **Port-based LAN broadcast traffic forwarding (port helpers) for multiple destination servers (EX9200)**—Starting in Junos OS Release 17.3R1, you can configure port helpers on EX9200 switches with multiple destination servers for a given port. Port helpers listen on configured UDP ports for incoming LAN broadcast traffic, and forward those packets to configured destination servers as unicast traffic. Configure port helpers to listen on a port and forward the traffic to a specified server using the **forwarding-options helpers port port-number** configuration statement with one of the following options:
 - Global—Specify only **server server-ip-address** to listen on *any* interface for the configured port.
 - VLAN-specific—Specify **interface irb-interface-name server server-ip-address** to listen only on a specified IRB interface.
 - Interface-specific—Specify **interface l3-interface-name server server-ip-address** to listen only on a specified Layer 3 interface.

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

Management

- **Support for the Junos Telemetry Interface (EX9200 switches)**—Starting with Junos OS Release 17.3R1, the Junos Telemetry Interface is supported on EX9200 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 EX9200 switches:
 - Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol (gRPC streaming only)
 - Ethernet interfaces enabled with the Link Layer Discovery Protocol (gRPC streaming only)
 - RSVP interface events (gRPC streaming only)

- BGP peers (gRPC streaming only)
- Memory utilization for routing protocol tasks (gRPC streaming only)
- LSP events and properties (gRPC streaming only)
- LSP statistics (UDP and gRPC streaming)
- Network Discovery Protocol table state (gRPC streaming only)
- Address Resolution Protocol table state (gRPC streaming only)
- IPFIX inline flow sampling (UDP streaming only)
- Queue depth statistics for ingress and egress queue traffic (UDP streaming only)
- Logical interfaces (UDP and gRPC streaming)
- Firewall filter statistics (UDP and 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)
- Fabric statistics (UDP streaming only)
- Physical interfaces (UDP and gRPC streaming)
- Chassis components (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 EX9200 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.

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

- **Support for the Junos Telemetry Interface (EX4600 switches)**—Starting with Junos OS Release 17.3R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various network elements without involving polling on EX4600 switches. Only gRPC streaming of statistics is supported on EX4600 switches. UDP streaming is not supported.

The following sensors are supported:

- BGP peers
- RSVP interface events
- Memory utilization for routing protocol tasks

- Label-switched-path events and properties
- Ethernet interfaces enabled with the Link Layer Discovery Protocol

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

- **Support for Two-Way Active Measurement Protocol (TWAMP) (EX4300 Switches)**—Starting in Junos OS Release 17.3R1, you can measure network performance between any two devices that support the TWAMP protocol. You can use the TWAMP-Control protocol to set up performance measurement sessions and the TWAMP-Test protocol to send and receive performance measurement probes.

You can configure TWAMP to start or stop all of the sessions for all of the TWAMP clients, or start or stop a session for a specific TWAMP client. When you start all the test session configured for a particular TWAMP client, the control-client initiates all requested testing with a Start-Sessions message, and the server sends an acknowledgment. If the control connection is not active between the server and the client, the control connection is also established and the test connections are started later. If the control-client name is not specified, all the configured test sessions are commenced.

When you stop the test session, the control connection is closed only after the Stop-sessions message is sent from the TWAMP client to the TWAMP server. If the control-client name is not specified, all the configured test sessions are closed.

Multiprotocol Label Switching (MPLS)

- **Support for resource RSVP (EX9200)**—Starting in Junos OS Release 17.3R1, the EX9200 switch supports RSVP. RSVP is a signaling protocol that reserves resources, such as for IP unicast and multicast flows, and requests QoS parameters for applications. The protocol was extended with MPLS RSVP-TE to enable RSVP to set up label-switched paths (LSPs) that can be used for traffic engineering in MPLS networks. RSVP is automatically enabled on interfaces on which MPLS-TE is configured. You can enable up to 200 RSVP-TE sessions in the EX9200 advanced feature license (AFL).

[See [RSVP Overview](#) .]

Operation, Administration, and Maintenance

- **Junos OS OpenConfig to support operational models for VLANs (EX Series)**—Starting with Junos OS Release 17.3R1, Junos OS supports an OpenConfig YANG model for VLANs via the addition of **openconfig-vlan.yang**, revision 1.0.2. This provides a unified view for the network agent to retrieve an operational state from Junos OS processes (daemons) for VLANs.

Services Applications

- **Support for enhancing the current inline JFlow scale limits for certain line cards (EX9200-6QS, EX9200-12QS, and EX9200-40XS)**—Starting in Junos OS Release 17.3R1, the **ipv4-flow-table-size** and the **ipv6-flow-table-size** allow up to 256 flow-table-size to support 64M flows at the **[edit chassis fpc slot-number inline-services flow-table-size]** hierarchy level. The existing limit on **flow-export-rate** under **inline-jflow** for each family in the sampling instance is increased to 3200 from 400.

SEE ALSO

Changes in Behavior and Syntax	30
Known Behavior	32
Known Issues	34
Resolved Issues	36
Documentation Updates	40
Migration, Upgrade, and Downgrade Instructions	41
Product Compatibility	42

Changes in Behavior and Syntax

IN THIS SECTION

- Management | 30
- Network Management and Monitoring | 30
- Services Applications | 31
- VLAN Infrastructure | 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 17.3R2 for the EX Series.

Management

- **Changes to custom YANG RPC syntax (EX Series)**—Starting in Junos OS Release 17.3, custom YANG RPCs have the following changes in syntax:
 - The **junos:action-execute** statement is a substatement to **junos:command**. In earlier releases, the **action-execute** and **command** statements are placed at the same level, and the **command** statement is optional.
 - The CLI formatting for a custom RPC is defined within the **junos-odl:format** statement, which takes an identifier as an argument. In earlier releases, the CLI formatting is defined using a container that includes the **junos-odl:cli-format** statement with no identifier.
 - The **junos-odl:style** statement defines the formatting for different styles within the statement. In earlier releases, the CLI formatting for different styles is defined using a container that includes the **junos-odl:cli-format** and **junos-odl:style** statements.

Network Management and Monitoring

- **Enhancement to about-to-expire logic for license expiry syslog messages (EX Series)**—Starting in Junos OS Release 17.3R1, the logic for multiple capacity type licenses and when their expiry raises alarms was changed. Before, the behavior had alarms and syslog messages for expiring licenses raised based on the highest validity, which would mislead users in the case of a license expiring earlier than the highest validity license. The new behavior has the about-to-expire logic based on the first expiring license.
- **Change to default log level setting (EX Series)**—Starting in Junos OS Release 17.3R2, changes were made in default logging levels:

Before the 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 the change:

- IFD LinkUp -> LOG_NOTICE (changed because although this is an important message, it occurs very frequently)
- IFL LinkUp -> LOG_INFO (no change)
- IFD and IFL LinkDown -> LOG_WARNING (no change)

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

- **Changes to SNMP syslog messages changed (EX Series)**—Starting in Junos OS Release 17.3R1, 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](#).]

Services Applications

- **Changes to the show services rpm history-results command (EX Series)**—Starting in Junos OS Release 17.3R2, you must include the **owner** *owner* and **test** *name* options when using the **show services rpm history-results** command.

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

VLAN Infrastructure

- **LAG interface flaps while adding/removing a VLAN**—From Junos OS Release 17.3 or later, the LAG interface flaps while adding or removing a VLAN. The flapping happens when a low speed SFP is plugged into a relatively high speed port. To avoid flapping, configure the port speed to match the speed of the SFP.

SEE ALSO

[New and Changed Features | 22](#)

[Known Behavior | 32](#)

Known Issues	34
Resolved Issues	36
Documentation Updates	40
Migration, Upgrade, and Downgrade Instructions	41
Product Compatibility	42

Known Behavior

IN THIS SECTION

- Authentication, Authorization, and Accounting (AAA) (RADIUS) | 33
- Platform and Infrastructure | 33
- Junos Fusion Enterprise | 33

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 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.

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

- On EX4300 switches, when 802.1X single-suplicant authentication is initiated, multiple "EAP Request Id Frame Sent" packets might be sent. [PR1163966](#)

Platform and Infrastructure

- On EX4600 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 required data. [PR980530](#)

Junos Fusion Enterprise

- In a Junos Fusion Enterprise topology with dual aggregation devices, firewall statistics are not synced across the aggregation devices. [PR1105612](#)
- In a Junos Fusion Enterprise, **show ethernet-switching** table takes a few minutes to show entries when received on an extended port with MAC count set to 150K. [PR1117567](#)
- In a Junos Fusion Enterprise, to use a non-default port as a clustering port in a clustering port policy, the policy must include at least one port that is a default uplink or clustering port for that platform. [PR1241808](#)
- In a Junos Fusion Enterprise, the satellite device link goes down with auto-negotiation enabled when the link partner speed is 100 Mbps. [PR1272107](#)
- In a Junos Fusion Enterprise, automatic medium-dependent interface crossover (auto_MDIX) and energy efficient Ethernet (EEE) are not supported on QFX5100 as a satellite device. The configuration commit succeeds on ge- ports but the features are not enabled. [PR1279928](#)

SEE ALSO

[New and Changed Features | 22](#)

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

[Known Issues | 34](#)

[Resolved Issues | 36](#)

[Documentation Updates | 40](#)

[Migration, Upgrade, and Downgrade Instructions | 41](#)

[Product Compatibility | 42](#)

Known Issues

IN THIS SECTION

- [Infrastructure | 34](#)
- [Layer 2 Features | 34](#)
- [MPLS | 35](#)
- [Network Management and Monitoring | 35](#)
- [Platform and Infrastructure | 35](#)
- [Virtual Chassis | 36](#)
- [VLAN Infrastructure | 36](#)

This section lists the known issues in hardware and software in Junos OS Release 17.3R2 for the EX Series.

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

Infrastructure

- On an EX9200-40XS line card, if you toggle the MACsec encryption option multiple times, encryption and protected MACsec statistics might be updated incorrectly. As a workaround, restart the line card. [PR1185659](#)
- When the configuration statement **set system ports console log-out-on-disconnect** is enabled, the Junos OS eventd process (daemon) blocks the console-open(). However, during this stage with the syslog console configured (always logs on console), any logging continues even if the console session is ended. When the console logging continues to be in the waiting status, the eventd syslog rotation freezes and some processes directly involved in logging in to the system would also go into the wait status, causing undesirable behavior. [PR1253544](#)

Layer 2 Features

- The memory leak might happen because an Ethernet switching process (eswd) is in Marvell based EX Series platform. A message similar to the following one might be seen in syslog: eswd[1330]: JTASK_OS_MEMHIGH: Using 212353 KB of memory, 158 percent of available /kernel: KERNEL_MEMORY_CRITICAL: System low on free memory, notifying init (#2). /kernel: Process (1254,eswd) has exceeded 85% of RLIMIT_DATA: used 114700 KB Max 131072 KB. [PR1262563](#)

- The Ethernet switching process (eswd) might crash after a Routing Engine switchover in an EX Series Virtual Chassis scenario. The crash happens due to disordered processing of VLAN/vmember by eswd and L2PT modules. As the order of processing does not remain the same every time, the crash is random across a switchover. [PR1275468](#)
- Ethernet ring protection switching (ERPS) route update fails during the addition of a new member to the ERPS-configured VLAN. [PR1301595](#)

MPLS

- On chassis-based line cards, the **FI: Protect: Parity error for CP freepool SRAM** SRAM parity error might be seen. It is harmless and can be ignored. [PR1079726](#)

Network Management and Monitoring

- The default syslog level is LOG_NOTICE in default configuration. SNMP_TRAP_LINK_UP for IFD was logged as LOG_INFO from day 1. To help debug physical link UP issues, SNMP_TRAP_LINK_UP events will be logged by default now. [PR1287244](#)

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 an EX9200-12QS line card, interfaces with the default speed of 10-Gigabit Ethernet are not brought down even when the remote end of a connection is misconfigured as 40 Gigabit Ethernet. [PR1175918](#)
- On an EX Series, if Dynamic Host Configuration Protocol (DHCP) relay or DHCP server is configured along with bpdu-block command, a memory allocation issue might be seen. This can lead to a memory exhaustion issue for the DHCP process (daemon). [PR1259918](#)
- A flexible VLAN-tagged interface allows both primary and secondary VLAN configuration on different logical units of the same interface, but might not work as expected. [PR1267160](#)
- On EX4300 10G links, preexisting MACsec sessions might not come up after the following events:
 - 1. Process (pfex, dot1x) restart or system restart
 - 2. Link flaps [PR1294526](#).

Virtual Chassis

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

VLAN Infrastructure

- 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](#)

SEE ALSO

New and Changed Features	 22
Changes in Behavior and Syntax	 30
Known Behavior	 32
Resolved Issues	 36
Documentation Updates	 40
Migration, Upgrade, and Downgrade Instructions	 41
Product Compatibility	 42

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2](#) | [37](#)
- [Resolved Issues: 17.3R1](#) | [39](#)

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

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

Resolved Issues: 17.3R2

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

- MacSec Issue **show security macsec statistics** command does not show expected results. [PR1283544](#)
- The Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) cannot forward correct Packet Ordering Engine class. [PR1296547](#)
- An l2ald crash occurs with no apparent trigger. [PR1302344](#)
- The CLI command **show snmp mib walk** used for jnxMIMstMstiPortState does not display anything in Junos OS Release 17.1R2 on the EX4600 platform. [PR1305281](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- Dhcp-security binding table might not get updated. [PR1312670](#)
- A memory leak is seen for dot1xd. [PR1313578](#)
- The dot1x process might stop authenticating if continuous dot1x clients reauthentication requests can't get processed [PR1300050](#)
- EX series switches do not send radius request after modifying the interface-range configuration. [PR1326442](#)
- QFX5100/EX4600/ACX5k : Major Alarm 'Fan & PSU Airflow direction mismatch' by removing management cable. [PR1327561](#)

Class of Service (CoS)

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

EVPNs

- Split Horizon Label is not allocated after switching configuration of ESI from 'single-active' to 'all-active' [PR1307056](#)

Infrastructure

- On EX Series switches, the file system might get corrupted multiple times during an image upgrade or commit operation. As a result, the image might fail to upgrade because the EX Series switches bypass the file system corruption check when file system is corrupted. [PR1317250](#)
- On EX4600, priority-based flow control (PFC) frames might not work. [PR1322439](#)

Interfaces and Chassis

- In a Virtual Chassis setup with aggregated Ethernet interfaces and multiple protocols configured in the system, intermittently we see LACP flap when the master is rebooted. Workaround is to toggle the interfaces where LACP is flapping. [PR1301338](#)
- The interface might not work properly after FPC restarts. [PR1329896](#)

Layer 2 Features

- Feature swap-swap might not work as expected in a Q-in-Q scenario. [PR1297772](#)

MPLS

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

Platform and Infrastructure

- On EX4300 Virtual Chassis, a 10-Gigabit Ethernet VCP might not get a neighbor after a system reboot. [PR1261363](#)
- CPU utilization for pfex_junos usage might go high if DHCP relay packets are coming continually. [PR1276995](#)
- Traffic loss might be observed for about 10 seconds if master member FPC reboots [PR1283702](#)
- On EX4300 switches, filter-based forwarding (FBF) might not work properly after deactivating or activating. This occurs because stale entries cannot be freed in ternary content addressable memory (TCAM); it leads to insufficient space in TCAM to process filters. [PR1293581](#)
- On an EX4300 switch, for those packets larger than 1452 bytes will be dropped after generic routing encapsulation (GRE) encapsulation, because the "Fragmentation of payload" and "GRE Path MTU discovery" are not supported on EX4300 Series switch. [PR1293787](#)
- On EX4300 some functions of IPv6 Router Advertisement Guard do not work. [PR1294260](#)
- **ERROR: /dev/da0s1a is not a JUNOS snapshot** is seen during system startup. [PR1297888](#)
- On EX4300 switches, when unknown unicast ICMP packets are received by an interface, packets are routed, so TTL is decremented. [PR1302070](#)
- On EX4300 Virtual Chassis, the FRU PSU removal and insertion traps are not generated for master or backup FPCs. [PR1302729](#)
- Inconsistent IEEE P-bit marking in 802.1Q header for OSPF packets. [PR1306750](#)
- Traceroute not working in EX9200 device for routing-instances running on 17.1R3 Junos version. [PR1310615](#)
- IGMP snooping might not learn multicast router interface dynamically. [PR1312128](#)
- On EX4300VC, l2cpd core file might be seen, if the interface is disabled under VSTP and enabled under RSTP [PR1317908](#)
- High latency might be observed between master Routing Engine and other Flexible PIC Concentrator (FPC). [PR1319795](#)
- On EX4300VC, VSTP BPDUs are not getting processed and root-bridge convergence fails for certain vlans [PR1320719](#)
- Multicast traffic might not forward to one of the receivers. [PR1323499](#)
- A Layer 2 Control Protocol process (l2cpd) might generate a core file. [PR1325917](#)

Routing Protocols

- JDI-RCT:M/Mx:Observed mcsnoopd core @
__raise,abort,__task_quit__,task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal
(enable_slip_detector=true,no_exit=true) at ../../../../src/junos/lib/libjtask/base/task_scheduler.c:275
[.PR1305239](#)

Virtual Chassis

- On EX4300 FRU removal/insertion trap not generated for non-master (backup/line card) FPCs.
[PR1293820](#)

Resolved Issues: 17.3R1

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

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

Infrastructure

- EX4300 aggregated interface is down while interface member VLAN is PVLAN and LACP is enabled.
[PR1264268](#)

Interfaces and Chassis

- Junos: EX Series PFE and MX MPC7E/8E/9E PFE crash when fetching interface stats with extended-statistics enabled (CVE-2017-10611); Refer to <https://kb.juniper.net/JSA10814> for more information. [PR1247026](#)

Layer 2 Features

- All the XML duplications and unformatted output are addressed. For Example, histogram was just declared as a element inside pfkey container, with this change a new container is defined for histogram. [PR1271648](#)

Platform and Infrastructure

- Layer 3 protocol packets are not being sent out from the switch. [PR1226976](#)

SEE ALSO

New and Changed Features 22
Changes in Behavior and Syntax 30
Known Behavior 32
Known Issues 34
Documentation Updates 40

Migration, Upgrade, and Downgrade Instructions 41
Product Compatibility 42

Documentation Updates

IN THIS SECTION

- [Traffic Management User Guide for EX4600 Switches | 40](#)

This section lists the errata and changes in Junos OS Release 17.3R2 for the EX Series switches documentation.

Traffic Management User Guide for EX4600 Switches

- **Consolidation of the Traffic Management User Guide for QFX Series and EX4600 Switches (EX4600)**—Starting in Junos OS Release 17.3R1, the following three traffic management guides are consolidated into one user guide:
 - Traffic Management User Guide for QFX Series
 - Traffic Management User Guide for QFX 10000 Series
 - Traffic Management User Guide for EX4600 Switches

[See [Traffic Management User Guide for QFX Series and EX4600 Switches](#).]

SEE ALSO

New and Changed Features 22
Changes in Behavior and Syntax 30
Known Behavior 32
Known Issues 34
Resolved Issues 36
Migration, Upgrade, and Downgrade Instructions 41
Product Compatibility 42

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 41](#)

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. 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://support.juniper.net/support/eol/junos.html>

SEE ALSO

[New and Changed Features | 22](#)

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

[Known Behavior | 32](#)

[Known Issues | 34](#)

[Resolved Issues | 36](#)

Documentation Updates 40
Product Compatibility 42

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 42

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 22
Changes in Behavior and Syntax 30
Known Behavior 32
Known Issues 34
Resolved Issues 36
Documentation Updates 40
Migration, Upgrade, and Downgrade Instructions 41

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

- New and Changed Features | 43
- Changes in Behavior and Syntax | 44
- Known Behavior | 44
- Known Issues | 45
- Resolved Issues | 46
- Documentation Updates | 47
- Migration, Upgrade, and Downgrade Instructions | 47
- Product Compatibility | 63

These release notes accompany Junos OS Release 17.3R2 for the Junos Fusion Data Center. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

There are no new features in Junos OS Release 17.3R2 for Junos Fusion Data Center.

SEE ALSO

Changes in Behavior and Syntax 44
Known Behavior 44
Known Issues 45
Resolved Issues 46
Documentation Updates 47
Migration, Upgrade, and Downgrade Instructions 47
Product Compatibility 63

Changes in Behavior and Syntax

There are no changes in behavior and syntax for Junos Fusion Data Center in Junos OS Release 17.3R2.

SEE ALSO

New and Changed Features 43
Known Behavior 44
Known Issues 45
Resolved Issues 46
Documentation Updates 47
Migration, Upgrade, and Downgrade Instructions 47
Product Compatibility 63

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 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.

SEE ALSO

New and Changed Features 43
Changes in Behavior and Syntax 44
Known Issues 45
Resolved Issues 46
Documentation Updates 47
Migration, Upgrade, and Downgrade Instructions 47
Product Compatibility 63

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.3R2 for the 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

New and Changed Features 43
Changes in Behavior and Syntax 44
Known Behavior 44
Resolved Issues 46
Documentation Updates 47
Migration, Upgrade, and Downgrade Instructions 47
Product Compatibility 63

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 17.3R2](#) | 46

This section lists the issues fixed in Junos OS Release 17.3R2.

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

Resolved Issues: Release 17.3R2

Junos Fusion Data Center

- Native VLAN on an aggregated Ethernet interface terminated on multiple satellite devices. [PR1305698](#)
- In a Junos Fusion topology with LAG on extended ports from satellite devices which are dual-homed to aggregation devices, the LAG interface might flap if rebooting one of the aggregation devices. [PR1315879](#)
- On a Junos Fusion topology with the QFX10002 platform as Aggregate Devices and the Aggregate Devices have dual cascade links to each Satellite Devices as redundancy, duplicated multicast traffic might be seen on downstream devices and multicast receivers if the multicast traffic pass through the Aggregate Devices. As a workaround, please deactivate and re-activate the VLAN in which duplicated multicast traffic is seen. [PR1316499](#)

SEE ALSO

[New and Changed Features](#) | 43

[Changes in Behavior and Syntax](#) | 44

[Known Behavior](#) | 44

[Known Issues](#) | 45

[Documentation Updates](#) | 47

[Migration, Upgrade, and Downgrade Instructions](#) | 47

[Product Compatibility](#) | 63

Documentation Updates

This section lists the errata or changes in Junos OS Release 17.3R2 for Junos Fusion Data Center documentation.

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

SEE ALSO

[New and Changed Features | 43](#)

[Changes in Behavior and Syntax | 44](#)

[Known Behavior | 44](#)

[Known Issues | 45](#)

[Resolved Issues | 46](#)

[Migration, Upgrade, and Downgrade Instructions | 47](#)

[Product Compatibility | 63](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 48](#)
- [Preparing the Switch for Satellite Device Conversion | 50](#)
- [Autoconverting a Switch into a Satellite Device | 52](#)
- [Manually Converting a Switch into a Satellite Device | 55](#)
- [Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 57](#)
- [Configuring Satellite Device Upgrade Groups | 58](#)
- [Converting a Satellite Device to a Standalone Device | 60](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 62](#)
- [Downgrading from Release 17.3 | 62](#)

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

The download and installation process for Junos OS Release 17.3R1 is different from previous Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the page.
5. Select the **Software** tab.

6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the aggregation device.

NOTE: We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

Customers in the United States and Canada, use the following command, replacing *n* with the spin number.

```
user@host> request system software add reboot
source/jinstall-host-qfx-10-f-17.3R2.n-secure-signed.tgz
```

All other customers, use the following command, replacing *n* with the spin number.

```
user@host> request system software add reboot source/jinstall-host-qfx-10-f-17.3R2
.n-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)

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

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

Rebooting occurs only if the upgrade is successful.

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos Fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device conversion software requirements, please refer to the [Junos Fusion Hardware and Software Compatibility Matrices](#).

NOTE: The following conditions must be met before a Junos switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The Junos switch can only be converted to SNOS 3.1 and higher.
- The Junos switch must be either set to factory-default configuration using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

Customers with EX4300 switches, use the following command, replacing *n* with the spin number:

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

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

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

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

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

```
user@satellite-device> request system zeroize
```

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

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

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

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

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

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

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

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

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration.

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 17.3R1 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-3.1R1.6-signed.tgz** to the upgrade group named **group1**:

```
user@aggregation-device> request system software add /var/tmp/satellite-3.1R1.6-signed.tgz
upgrade-group group1
```

NOTE: Before issuing **request system software add /var/tmp/satellite-3.1R1.6-signed.tgz** **upgrade-group group1**, you must issue a one-time command to expand the storage capacity. Use the **request system storage user-disk expand** command to increase the size of /user partition.

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 17.3 R1 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-3.1R1.6-signed.tgz** to the upgrade group named group1:

```
user@aggregation-device> request system software add /var/tmp/satellite-3.1R1.6-signed.tgz
upgrade-group group1
```

NOTE: Before issuing **request system software add /var/tmp/satellite-3.1R1.6-signed.tgz upgrade-group group1**, you must issue a one-time command to expand the storage capacity. Use the **request system storage user-disk expand** command to increase the size of /user partition.

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-3.1R1.6-signed.tgz** package is included specifically for converting Junos OS to a satellite device 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.

1. 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-3.1R1.6-signed.tgz** stored in the **/var/tmp/** directory on the switch:

```
user@satellite-device> request chassis device-mode satellite  
/var/tmp/satellite-3.1R1.6-signed.tgz
```

- To install satellite software onto a QFX5100 switch, use the **satellite-3.1R1.6-signed.tgz** satellite software package.

- To install satellite software onto a EX4300 switch, use the **satellite-ppc-3.1R1.6-signed.tgz** satellite software package.
2. 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** statement 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-3.1R1.6-signed.tgz
```

NOTE: Before issuing **request system software add upgrade-group group1**, you must issue a one-time command to expand the storage capacity. Use the **request system storage user-disk expand** command to increase the size of /user partition.

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

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

Converting a Satellite Device to a Standalone Device

In the event that you need to convert a satellite device to a standalone device, you will need to install a new Junos OS software package on the satellite device and remove it from the Junos Fusion topology.

NOTE: The QFX5100-48SH and QFX5100-48TH satellite device models cannot be converted to a standalone switch.

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 the software that includes pxe in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D43 is named install-media-pxe-qfx-5-14.1X53-D43.7-domestic-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://support.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 14.1X53D43 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. You can check the automatic satellite conversion configuration by entering the **show** command at the **[edit chassis satellite-management auto-satellite-conversion]** hierarchy level.

```
[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
```

9. Commit the configuration.

```
[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-D43.7-domestic-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-D43.7-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 into your device, uncable the device from the Junos Fusion topology. See *Removing a Transceiver from a QFX Series Device* or *Remove a Transceiver*, as needed. Your device has been removed from Junos Fusion.

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

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. 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://support.juniper.net/support/eol/junos.html>.

Downgrading from Release 17.3

To downgrade from Release 17.3 to another supported release, follow the procedure for upgrading, but replace the 17.3 **jinstall** package with one that corresponds to the appropriate downgrade release.

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

New and Changed Features 43
Changes in Behavior and Syntax 44
Known Behavior 44
Known Issues 45
Resolved Issues 46
Documentation Updates 47
Product Compatibility 63

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 63

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 43
Changes in Behavior and Syntax 44
Known Behavior 44
Known Issues 45
Resolved Issues 46

Documentation Updates | 47

Migration, Upgrade, and Downgrade Instructions | 47

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- New and Changed Features | 65
- Changes in Behavior and Syntax | 67
- Known Behavior | 67
- Known Issues | 68
- Resolved Issues | 69
- Documentation Updates | 71
- Migration, Upgrade, and Downgrade Instructions | 71
- Product Compatibility | 79

These release notes accompany Junos OS Release 17.3R2 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

New and Changed Features

IN THIS SECTION

- Junos Fusion Enterprise | 66

This section describes the new features and enhancements to existing features in Junos OS Release 17.3R2 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- **Satellite device support (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.3R1, you can configure QFX5100-48T and QFX5100-48S switches as satellite devices in a Junos Fusion Enterprise topology. The satellite device in a Junos Fusion topology is managed and configured by the aggregation device. Junos Fusion Enterprise uses EX9200 switches in the aggregation device role.
[See [Junos Fusion Enterprise Overview](#).]
- **LAG to single satellite device (Junos Fusion Enterprise)**—Starting in Junos OS Release 17.3R1, you can configure LAGs in a Junos Fusion Enterprise using extended port member links to increase uplink bandwidth and high availability for endpoint devices connected to a satellite device. The member links of the LAG must be on the same satellite device. The LAG can be configured to use LACP, which automates the addition and deletion of individual links to the LAG and can also prevent communication failures by detecting misconfigurations within a LAG.
[See [Configuring Link Aggregation on Satellite Devices in a Junos Fusion Enterprise](#).]

SEE ALSO

Changes in Behavior and Syntax 67
Known Behavior 67
Known Issues 68
Resolved Issues 69
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71
Product Compatibility 79

Changes in Behavior and Syntax

IN THIS SECTION

- [Junos Fusion Enterprise | 67](#)

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 17.3R2 for Junos Fusion Enterprise.

Junos Fusion Enterprise

- For the **request chassis satellite beacon** operational command, the **slot-id** option has been changed to **fpc-slot**. This change was made to support enabling beacon functionality for individual FPCs. [PR1272956](#)

SEE ALSO

New and Changed Features 65
Known Behavior 67
Known Issues 68
Resolved Issues 69
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71
Product Compatibility 79

Known Behavior

IN THIS SECTION

- [Junos Fusion Enterprise | 68](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 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

- In a Junos Fusion Enterprise topology with dual aggregation devices, firewall statistics are not synced across the aggregation devices. [PR1105612](#)
- in a Junos Fusion Enterprise, **show ethernet-switching table** takes a few minutes to show entries when an received on an extended port with MAC count set to 150K. [PR1117567](#)
- In a Junos Fusion Enterprise, in order to use a non-default port as a clustering port in a clustering port policy, the policy must include at least one port that is a default uplink/clustering port for that platform. [PR1241808](#)
- In a Junos Fusion Enterprise, the satellite device link goes down with auto-negotiation enabled when the link partner speed is 100m. [PR1272107](#)
- In a Junos Fusion Enterprise, Auto_MDIX and EEE are not supported on QFX5100 as a satellite device. The configuration commit succeeds on ge- ports but the features are not enabled. [PR1279928](#)

SEE ALSO

New and Changed Features 65
Changes in Behavior and Syntax 67
Known Issues 68
Resolved Issues 69
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71
Product Compatibility 79

Known Issues

IN THIS SECTION

- [Junos Fusion Enterprise | 69](#)

This section lists the known issues in hardware and software in Junos OS Release 17.3R2 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- On a Junos Fusion, when using LLDP, the "Power via MDI" and "Extended Power via MDI" TLVs are not transmitted. [PR1105217](#)
- On a Junos Fusion, the TCPDUMP command does not capture packets on satellite devices. [PR1125568](#)
- 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, 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](#)

SEE ALSO

New and Changed Features 65
Changes in Behavior and Syntax 67
Known Behavior 67
Resolved Issues 69
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71
Product Compatibility 79

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2 | 70](#)
- [Resolved Issues: 17.3R1 | 70](#)

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

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

Resolved Issues: 17.3R2

Junos Fusion Enterprise

- 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.

[PR1183680](#)

- Request chassis satellite beacon functionality to specific SD is not working, causing all the SDs to enable the beacon LED. [PR1272956](#)
- VRRP has a split brain in dual autodiscovery Junos Fusion. [PR1293030](#)
- POE not working on one of the satellite in cluster. Dual AD Fusion setup. [PR1295556](#)
- On Dual-AD JFE setup, while applying Routing Engine lo0 filters and setting the cascade port down on AD2, the SD goes to "ProvSessionDown" on that AD2 while it stays online on AD1. [PR1275290](#)
- Aggregation device without cascade port cannot reach hosts over ICL link if they are authenticated by 802.1X in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)
- 802.1X authentication fails on a Junos Fusion setup. [PR1299532](#)
- 802.1X authentication might crash in a Junos Fusion setup with dual aggregation devices. [PR1303909](#)
- All the 802.1X authentication 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](#)

Resolved Issues: 17.3R1

Junos Fusion Enterprise

- On a Junos Fusion Enterprise, an upgrade group's association with a satellite software version is removed if the chassis satellite-management redundancy-groups configuration is deleted. [PR1267370](#)
- On Junos Fusion Enterprise, traffic shaping is not supported on the extended ports. [PR1268084](#)
- When a race condition results in a dynamic VLAN assignment, the MAC-based VLAN (MBV) entry might not get created on the peer AD. This situation can result in traffic loss when it flows through the peer AD. [PR1282828](#)

SEE ALSO

New and Changed Features	 65
Changes in Behavior and Syntax	 67
Known Behavior	 67
Known Issues	 68
Documentation Updates	 71
Migration, Upgrade, and Downgrade Instructions	 71
Product Compatibility	 79

Documentation Updates

There are no errata or changes in Junos OS Release 17.3R2 for Junos Fusion Enterprise documentation.

SEE ALSO

New and Changed Features	 65
Changes in Behavior and Syntax	 67
Known Behavior	 67
Known Issues	 68
Resolved Issues	 69
Migration, Upgrade, and Downgrade Instructions	 71
Product Compatibility	 79

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | 72
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | 74
- [Preparing the Switch for Satellite Device Conversion](#) | 74
- [Converting a Satellite Device to a Standalone Switch](#) | 76

- Upgrade and Downgrade Support Policy for Junos OS Releases | 78
- Downgrading from Release 17.3 | 79

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 17.3R1:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list on the right of the page.
5. Select the **Software** tab.
6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **junos-install** package on the aggregation device.

NOTE: We recommend that you upgrade all software packages out of band using the console because in-band connections are lost during the upgrade process.

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

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-17.3R1.n.tgz
```

All other customers, use the following commands, replacing *n* with the spin number.

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-17.3R1.n-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: To upgrade from Junos OS 17.1 to 17.3R1, use the procedure for upgrading from Junos OS 17.1 to 17.2, as documented in the Release Notes for Junos Fusion Enterprise, Junos OS Release 17.2.

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: For EX4300 and QFX5100 switches running Junos OS Release 14.1X53-D43 or EX2300 and EX3400 switches running Junos OS Release 15.1X53-D55.5, the following conditions must be met before the Junos switch 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:

```
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. (EX3400 and 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 EX3400 and 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.

The following steps explain how to convert a satellite device that is participating in a Junos Fusion to a standalone device running Junos OS. If you have a standalone switch that is not part of a Junos Fusion but is running satellite software, and you want the switch to run Junos OS software, see [Installing Junos OS Software on a Standalone Device Running Satellite Software](#).

NOTE: Conversion of EX2300 and EX3400 switches from satellite devices to standalone devices cannot be initiated from the aggregation device. To install Junos OS software on an EX2300 or EX3400 switch acting as a satellite device, see [Installing Junos OS Software on a Standalone Device Running Satellite Software](#).

The following steps explain how to download software, remove the satellite device from the 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://support.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 menu and select the switch platform series and model for your satellite device.
4. Select the software image for your platform, using the following guidelines:
 - If the satellite device is a EX4300 switch, you install a standard, signed **jinstall** version of Junos OS.

- If the satellite device is a QFX5100 switch that can be converted to a standalone device, you must install a Preboot eXecution Environment (PXE) version of Junos OS. The PXE version of Junos OS software supports the same feature set as the other Junos OS software packages for a release, but is specially engineered to install Junos OS onto a device running satellite software. The PXE Junos OS package name uses the format **install-media-pxe-qfx-5-version-domestic-signed.tgz**.

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

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

7. 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 the 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.

8. Commit the configuration.

To commit the configuration to both Routing Engines:

[edit]

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

To commit the configuration to a single Routing Engine:

[edit]

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

9. Install Junos OS 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 software package stored in the `/var/tmp` directory on the aggregation device onto a switch acting as the satellite device using FPC slot 102:

```
[edit]
```

```
user@aggregation-device> request chassis satellite install /var/tmp/package-name fpc-slot 102
```

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

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

11. When you are prompted to log back in to your device, uncable the device from the Junos Fusion topology. See *Remove a Transceiver*. Your device is removed from the Junos Fusion.

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

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. 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://support.juniper.net/support/eol/junos.html>

Downgrading from Release 17.3

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

NOTE: You cannot downgrade more than three releases.

To downgrade a Junos Fusion Enterprise, follow the procedure for upgrading, but replace the 17.2 **junos-install** package with one that corresponds to the appropriate release.

NOTE: We recommend that you do not downgrade the aggregation device from 17.3R1 to 17.2 if there are cluster satellite devices in the setup.

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

SEE ALSO

New and Changed Features	 65
Changes in Behavior and Syntax	 67
Known Behavior	 67
Known Issues	 68
Resolved Issues	 69
Documentation Updates	 71
Product Compatibility	 79

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility](#) | 80
- [Hardware Compatibility Tool](#) | 80

Hardware and Software Compatibility

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

To determine the features supported 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/>

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	 65
Changes in Behavior and Syntax	 67
Known Behavior	 67
Known Issues	 68
Resolved Issues	 69
Documentation Updates	 71
Migration, Upgrade, and Downgrade Instructions	 71

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- New and Changed Features | 81
- Changes in Behavior and Syntax | 83
- Known Behavior | 83
- Known Issues | 84
- Resolved Issues | 84
- Documentation Updates | 85
- Migration, Upgrade, and Downgrade Instructions | 86
- Product Compatibility | 94

These release notes accompany Junos OS Release 17.3R2 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.3R2 New and Changed Features | 82
- Release 17.3R1 New and Changed Features | 82

This section describes the new features and enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 17.3R2.

Release 17.3R2 New and Changed Features

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 17.3R2.

Release 17.3R1 New and Changed Features

Junos Fusion

- **Power over Ethernet (PoE) for Junos Fusion Provider Edge**—Starting in Junos OS Release 17.3R1, PoE is supported on Junos Fusion Provider Edge. PoE enables electric power, along with data, to be passed over a copper Ethernet LAN cable. Powered devices—such as VoIP telephones, wireless access points, video cameras, and point-of-sale devices—that support PoE can receive power safely from the same access ports that are used to connect personal computers to the network. This reduces the amount of wiring in a network, and also eliminates the need to position a powered device near an AC power outlet, making network design more flexible and efficient.

In a Junos Fusion system, PoE is used to carry electric power from an extended port on a satellite device to a connected device. An extended port is any network-facing port on a satellite device in a Junos Fusion Provider Edge. All extended ports that support PoE on satellite devices in a Junos Fusion Provider Edge support the IEEE 802.3at PoE+ standard. The aggregation device on Junos Fusion Provider Edge manages PoE support of PoE-capable interfaces on satellite device. Junos Fusion Provider Edge only supports PoE with EX series switches as satellite devices and a MX Series 3D Universal Router as the aggregation device.

[See [Understanding Power over Ethernet in a Junos Fusion](#).]

- **Port-based network access control**—Starting in Junos OS Release 17.3R1, Junos Fusion Provide Edge supports port-based authentication as defined by the IEEE 802.1X standard and central Web authentication to prevent unauthorized network access on extended ports of the satellite devices. This feature allows you to configure satellite devices to block access to the network until the client is authenticated. This feature allows you to configure satellite devices to block access to the network until the client is authenticated.

[See [Understanding port-based authentication in a Junos Fusion Provider Edge](#).]

- **Metro Ethernet Forum (MEF) Carrier Ethernet 2.0 Certification**—Starting in Junos OS Release 17.3R1, Junos Fusion Provider Edge qualifies for Carrier Ethernet 2.0 (CE2.0) certification. This ensures that the routers and switches in a Junos Fusion Provider Edge system comply with the Carrier Ethernet specification set by the MEF.

SEE ALSO

[Changes in Behavior and Syntax | 83](#)

[Known Behavior | 83](#)

Known Issues	84
Resolved Issues	84
Documentation Updates	85
Migration, Upgrade, and Downgrade Instructions	86
Product Compatibility	94

Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 17.3R2.

SEE ALSO

New and Changed Features	81
Known Behavior	83
Known Issues	84
Resolved Issues	84
Documentation Updates	85
Migration, Upgrade, and Downgrade Instructions	86
Product Compatibility	94

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software for Junos Fusion Provider Edge in Junos OS Release 17.3R2.

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

SEE ALSO

New and Changed Features	81
Changes in Behavior and Syntax	83
Known Issues	84

Resolved Issues	84
Documentation Updates	85
Migration, Upgrade, and Downgrade Instructions	86
Product Compatibility	94

Known Issues

There are no known issues in the Junos OS Release 17.3R2 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

New and Changed Features	81
Changes in Behavior and Syntax	83
Known Behavior	83
Resolved Issues	84
Documentation Updates	85
Migration, Upgrade, and Downgrade Instructions	86
Product Compatibility	94

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2](#) | [85](#)
- [Resolved Issues: 17.3R1](#) | [85](#)

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: 17.3R2

Junos Fusion Provider Edge

- Some vlan bridges and VTEP bindings might be lost after deleting or deactivating a vlan and then committing configuration. [PR1298659](#)
- The LAG interface might flap if rebooting aggregation device.[PR1315879](#)
- Duplicated packets might be received on the multicast downstream devices and multicast receivers. [PR1316499](#)

Junos Fusion Satellite Software

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

Resolved Issues: 17.3R1

There are no fixed issues in the Junos OS Release 17.3R1 for Junos Fusion Provider Edge.

SEE ALSO

New and Changed Features 81
Changes in Behavior and Syntax 83
Known Behavior 83
Known Issues 84
Documentation Updates 85
Migration, Upgrade, and Downgrade Instructions 86
Product Compatibility 94

Documentation Updates

There are no errata or changes in Junos OS Release 17.3R2 for Junos Fusion Provider Edge documentation.

SEE ALSO

New and Changed Features	81
Changes in Behavior and Syntax	83
Known Behavior	83
Known Issues	84
Resolved Issues	84
Migration, Upgrade, and Downgrade Instructions	86
Product Compatibility	94

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 86
- Upgrading an Aggregation Device with Redundant Routing Engines | 89
- Preparing the Switch for Satellite Device Conversion | 89
- Converting a Satellite Device to a Standalone Device | 91
- Upgrading an Aggregation Device | 93
- Upgrade and Downgrade Support Policy for Junos OS Releases | 93
- Downgrading from Release 17.3 | 93

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 17.3R1 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://support.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the page.
5. Select the **Software** tab.
6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the aggregation device.

NOTE: We recommend that you upgrade all software packages out-of-band using the console, because in-band connections are lost during the upgrade process.

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

For upgrades from Junos Release 14.2 and earlier:

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

All other upgrades:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.3R1 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

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

NOTE: The following conditions must be met before a Junos switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The Junos switch can only be converted to SNOS 3.1 and higher.
- The Junos switch must be either set to factory-default configuration using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

```
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-D43 is named `install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz`. If the satellite device is an EX4300 switch, you install a standard `jinstall-ex-4300` version of Junos OS.

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

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:
<https://support.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-D43 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 into your device, uncable the device from the Junos Fusion topology. See *Removing a Transceiver from a QFX Series Device* or *Remove a Transceiver*, as needed. Your device has been removed from Junos Fusion.

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

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 17.3R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. 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://support.juniper.net/support/eol/junos.html>.

Downgrading from Release 17.3

To downgrade from Release 17.3 to another supported release, follow the procedure for upgrading, but replace the 17.3 **jinstall** package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

New and Changed Features	 81
Changes in Behavior and Syntax	 83
Known Behavior	 83
Known Issues	 84
Resolved Issues	 84
Documentation Updates	 85
Product Compatibility	 94

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility](#) | 94

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

New and Changed Features 81
Changes in Behavior and Syntax 83
Known Behavior 83
Known Issues 84
Resolved Issues 84
Documentation Updates 85
Migration, Upgrade, and Downgrade Instructions 86

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

IN THIS SECTION

- [New and Changed Features | 96](#)
- [Changes in Behavior and Syntax | 120](#)
- [Known Behavior | 125](#)
- [Known Issues | 131](#)
- [Resolved Issues | 144](#)
- [Documentation Updates | 163](#)
- [Migration, Upgrade, and Downgrade Instructions | 164](#)
- [Product Compatibility | 171](#)

These release notes accompany Junos OS Release 17.3R2 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.3R2 New and Changed Features | 96](#)
- [Release 17.3R1 New and Changed Features | 96](#)

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

Release 17.3R2 New and Changed Features

Subscriber Management and Services

- **Preventing validation of magic numbers in PPP peer-originated keepalive messages (MX Series)**—Starting in Junos OS Release 17.3R2-S3, you can include the **ignore-magic-number-mismatch** statement to disable the Packet Forwarding Engine from validating PPP magic numbers received during PPP keepalive (Echo-Request/Echo-Reply) exchanges. Because validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends a magic number that does not match the number agreed upon during LCP negotiation. This prevents PPP from tearing down the session in the event of a mismatch. This capability is useful when the remote PPP peers include arbitrary magic numbers in the keepalive packets. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.

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

Release 17.3R1 New and Changed Features

Class of Service (CoS)

- **Support for efficient use of CoS resources on targeted interfaces (MX Series)**—Starting in Junos OS Release 17.3R1, when you configure Junos OS to target the egress traffic for a subscriber on a single member link, Junos OS applies CoS resources only to the active link, optimizing the use of available scheduling nodes. If the assigned primary link goes down, CoS scheduling resources are switched to the backup link.

[See [targeted-distribution \(Dynamic Demux Interfaces over Aggregated Ethernet\)](#).]

- **Support for setting the DSCP code point for host-originating IS-IS traffic sent over a GRE tunnel (MX Series)**—Starting in Junos OS Release 17.3R1, you can determine traffic prioritization for IS-IS traffic originating on a host and being sent over a GRE tunnel by assigning a DSCP code point to the IS-IS packets. You can set the DSCP code point by including the `isis-over-gre dscp-code-point value` statement at the `[edit class-of-service host-outbound-traffic protocol]` hierarchy level.

[See [protocol \(Host Outbound Traffic\)](#).]

Dynamic Host Configuration Protocol (DHCP)

- **Support for single-session DHCP dual-stack subscriber for S-VLAN model server and relay (MX Series)**—Starting in Junos OS Release 17.3R1, DHCP dual-stack subscriber for N:1 (IP demux) access models support multiple household share the same S-VLAN.

A dual-stack DHCP subscriber is represented as a single subscriber with a single session database (SDB) session.

The benefits of a single-session dual-stack model are as follows:

- Simplifies router configuration.
- Reduces RADIUS message load.
- Reduces the backend correlation of multiple accounting sessions for the same household.
- Is compatible with existing RADIUS messaging.

[See [Single-Session DHCP Local Server Dual-Stack Overview](#) and [Single-Session DHCP Dual-Stack Overview](#).]

- **Support for single-session DHCP dual-stack subscriber single BNG connect (MX Series)**—Starting in Junos OS Release 17.3R1, DHCP single-session dual-stack subscribers connect to a single broadband network gateway (BNG) in a load sharing access model.

For a DHCP dual-stack subscriber, the DHCPv4 and DHCPv6 protocol handshakes are generally completely independent of each other. So it is theoretically possible that each arm of a given dual-stack subscriber could connect to a different BNG. A configured mode of operation is supported to avoid this scenario

A given address family is designated as the protocol master for a dual-stack subscriber. Any binding attempt from the secondary address family client for a given dual-stack subscriber is rejected if a binding from the protocol master family of the same dual-stack subscriber is not currently active.

In case bindings for both arms of a DHCP dual-stack subscriber are currently active when the **protocol-master** family binding is released (or otherwise deleted for any reason), then the secondary address family binding for that subscriber will be automatically torn down.

[See [Single-Session DHCP Local Server Dual-Stack Overview](#) and [Single-Session DHCP Dual-Stack Overview](#).]

- **Support for DHCP local server dual-stack single-session (MX Series)**—Starting in Junos OS Release 17.3R1, DHCP local server dual-stack subscribers are supported on a single VLAN session. This reduces the required number of session database (SDB) entries utilized and simplifies RADIUS authentication and accounting operations.

The benefits of a single-session dual-stack model are as follows

- Simplifies router configuration.
- Reduces RADIUS message load.
- Reduces the backend correlation of multiple accounting sessions for the same household.
- Is Compatible with existing RADIUS messaging.

[See [Single-Session DHCP Local Server Dual-Stack Overview](#).]

- **Support for DHCPv6 prefix exclude option(MX Series)**—Starting in Junos OS Release 17.3R1, you can exclude one specific prefix that is bigger than the prefix length from a delegated prefix set while using DHCPv6 based prefix delegation. This specific prefix is used as the link between the delegating router and the requesting router, where the delegating router exchanges DHCPv6 messages with the requesting router. Configure the **exclude-prefix-len** statement at the **[edit access address-assignment pool delegated-address-pool family inet6 dhcp-attributes]** hierarchy level to exclude the prefix from the delegated prefix set. You can configure the **support-option-pd-exclude** statement at either the **[edit system services dhcp-local-server dhcpv6 reconfigure]** or the **[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]** hierarchy level to exclude prefix support in the reconfigure message.

[See [Understanding Support for DHCPv6 Prefix Exclude Option](#)]

EVPNs

- **EVPN-VXLAN support for VXLAN gateways using an IPv6 underlay (MX Series with MPC and MIC)**—Starting in Junos OS Release 17.3R1, MX Series routers with MPC and MIC interfaces extend support for Virtual Extensible LAN (VXLAN) gateways from IPv4 to IPv6 underlays. With this feature enhancement, each VXLAN gateway supports the following functionalities in addition to the IPv4 functionalities already supported:
 - VLAN-based service
 - VLAN-bundle service
 - Port-based service
 - VLAN-aware service

Similar to IPv4 underlay support, the IPv6 EVPN-VXLAN underlay supports the Type 2 MAC address with IP address advertisement and the proxy MAC address with IP address advertisement.

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#).]

- **Preference-based DF election for EVPN and PBB-EVPN (MX Series with MPC and MIC interfaces)**—Starting in Junos OS Release 17.3, the designated forwarder (DF) election in a multihomed Ethernet VPN (EVPN) environment can be controlled using an administrative preference value for an Ethernet segment identifier (ESI). Currently, the DF election (as specified in RFC 7432) is performed randomly by all the multihoming devices using the modulo operation. With the preference-based DF election, the DF is elected manually using interface configuration options, such as the preference value and the router ID or loopback address. This method of DF election is useful when there is a need to choose the DF based on interface attributes like bandwidth associated with the interface.

To enable preference-based DF election, include the **df-election-type preference value value** statements at the **[edit interfaces interface-name esi]** hierarchy level.

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

- **Support for seamless migration from LDP-VPLS to EVPN (MX Series)**—Currently, a virtual private LAN service (VPLS) network can be connected to an Ethernet VPN (EVPN) network using logical tunnel interfaces on the interconnection point of the VPLS and EVPN routing instances. In this case, the provider edge (PE) devices in each network are unaware of the PE devices in the other technology network. Starting in Junos OS Release 17.3R1, a solution is introduced for enabling staged migration from FEC128 LDP-VPLS toward EVPN on a site-by-site basis for every VPN routing instance. In this solution, the PE devices running EVPN and VPLS for the same VPN routing instance and single-homed segments can coexist. During the migration, there is minimal impact to the customer edge (CE) device-to-CE device traffic forwarding for affected customers.

[See [Migrating From FEC128 LDP-VPLS to EVPN Overview](#).]

General Routing

- **Commit process split into two steps (MX Series)**—Starting in Junos OS Release 17.3R1, new configuration statements are introduced for **commit** to split the commit process into two steps. These configuration statements are **prepare** and **activate**.

In the first step, known as the preparation stage, **commit prepare** validates the configurations and then creates the necessary files and database entries so that the validated configurations can be activated at a later stage.

In the second step, referred to as the activation stage, **commit activate** activates the previously prepared commit. A new configuration statement, **prepared**, is added to **clear system commit**, which clears the prepared commit cache

This feature enables you to configure a number of Junos OS devices and simultaneously activate the configurations. This approach is helpful in time-critical scenarios.

[See [Commit Preparation and Activation Overview](#).]

High Availability (HA) and Resiliency

- **Mandatory action before initiating GRES in the presence of PIC bounce alarms (MX10003 router)**—In Junos OS Release 17.3R1, before initiating graceful Routing Engine switchover (GRES) on an MX10003,

you must bounce the PIC (by issuing offline/online of the PIC) using `request chassis pic` command before performing switchover operation. Otherwise, it will provide negative results as the alarms are not preserved on GRES currently. It may also result in unstable behavior of MPC.

Consider the example of PIC bounce alarm shown below. In this case, you must bounce the PIC before initiating a switchover.

```
user@host# run show chassis alarms
Apr 17 01:50:13
4 alarms currently active
Alarm time          Class  Description
2017-04-17 01:48:57 PDT  Minor  FPC 0 PIC 1 Need bounce
2017-04-14 09:14:03 PDT  Major  PEM 4 Not Present
2017-04-14 09:14:03 PDT  Major  PEM 3 Not Present
2017-04-14 09:14:03 PDT  Major  PEM 1 Not Present
```

- **VRRP scale improvements per aggregated Ethernet bundle(MX Series)**—Starting in Junos OS Release 17.3R1, you can configure up to 4000 active VRRP sessions per aggregated Ethernet bundle on MX Series routers. To configure VRRP support, include 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 and Chassis

- **Support for new MX150 Universal Routing Platform**—Starting in Junos OS Release 17.3R1, Junos OS supports a new MX Series edge router—the MX150—which is a compact, high-performance edge router that is ideally suited for lower bandwidth service provider applications and distributed service architectures, and for enterprise WAN use-cases. The MX150 is 1 rack unit (RU) tall and supports bandwidth that can be upgraded from 100 Mbps to 20 Gbps.
- **Support for FRU control, power management, and environmental monitoring in MX10003 routers**—Starting with Junos OS Release 17.3R1, Junos OS chassis management software for the MX10003 routers provides enhanced environmental monitoring and FRU control. MX10003 has a pair of Routing Engines, which support virtualization. Each Routing Engine board is a single FRU. The MX10003 router has two MPCs, each supporting a bandwidth up to 1.2 Tbps. Each MPC has three Packet Forwarding Engines, each providing a maximum bandwidth of 400 Gbps. Each MPC supports a fixed PIC comprising six QSFP ports and a modular interface card (MIC) comprising 12 QSFP28 ports. All FRUs are upgradable. The MX10003 chassis has two power supply modules (PSM)—a DC PSM and an AC PSM. The MX10003 cooling system contains four fan assemblies, with two fans in each. MX10003 supports temperature thresholds for each temperature sensor, which enables the router to precisely control the cooling, raise alarms, and shut down an FRU. The router also supports preserving power-on sequence for the FPCs, and power management using ambient-temperature.

[See [Understanding How Dynamic Power Management Enables Better Utilization of Power](#).]

- **Fabric management in MX10003 routers**—Starting with Junos OS Release 17.3R1, Junos OS supports management and control of fabric operations on MX10003 routers. On the MX10003 router, the switching fabric is located on the MPC. The router has two MPCs, each supporting a bandwidth up to 1.2 Tbps. The switching fabric has 22 planes and each plane supports a maximum link speed of 24.883 Gbps. MX10003 routers do not have a dedicated fabric card. The router supports features such as fabric hardening and forward error correction.

[See [MX Series Routers Fabric Resiliency](#).]

- **MPCs, PICs, and MICs supported on MX10003 routers**—Starting with Junos OS Release 17.3R1, the MX10003 router supports a new MPC, MX10003 MPC. The MX10003 MPC supports three Packet Forwarding Engines. The forwarding capacity of each Packet Forwarding Engine is 400Gbps which cannot be oversubscribed. Each MPC supports a fixed-port PIC and modular MICs, JNP-MIC1 (MIC without MACsec support) and JNP-MIC1-MACSEC (MIC with MACsec support). The fixed port PIC is mapped to PIC 0 and each PFE is mapped to 2 ports in PIC 0. The MIC is mapped to PIC 1 and each PFE is mapped to 4 ports in PIC 1. The PIC/MIC ports on MX10003 router MPCs support multiple port speeds (10/40/100GE). Hence, these ports are classified as multi-rate ports. However, all the PIC/MIC ports do not support all the port speeds. On MPC all the 12 ports are active and are capable of running in 40-Gigabit Ethernet, 100-Gigabit Ethernet, and 4x10-Gigabit Ethernet mode. [See [MX10003 MPC on MX10003 Router Overview](#) for more details.]
- **Support for inline flow monitoring on MPCs on MX10003 routers**—Starting with Junos OS Release 17.3R1, MPCs on MX10003 router support inline flow monitoring. Inline flow monitoring results in higher scalability and performance, as the scaling and performance are not dependent on the capacity of the services interface. MX10003 router contains two MPCs, each supporting a bandwidth up to 1.2 Tbps.
- **Broadband edge (BBE) telemetry sensors(MX Series)**—Starting in Junos OS Release 17.3R1, support is added for BBE telemetry sensors. These sensors are used to proactively manage a broadband network gateway (BNG) and are configured using both Junos Telemetry Interface (JTI) and gRPC streaming.

The new sensors are grouped into the following functional areas:

- Chassis and system extensions
- AAA
- DHCP
- PPP
- L2TP
- MX Series routers Virtual Chassis
- ERA
- BBE infrastructure
- Packet Forwarding Engine resource and monitoring

- **Support for inline NAT services on MX10003**—Starting with Junos OS Release 17.3R1, MX10003 routers support inline Network Address Translation (NAT) services on Modular Port Concentrators (MPCs). This enables you to achieve line-rate, low-latency address translations (up to 120 Gbps per slot) without having to use a dedicated MS-MPC for NAT.
- **MAC address persistence after a Routing Engine switchover**—In Junos OS Release 17.3R1 and later, if you configure multiple aggregated Ethernet interfaces, the MAC addresses of the aggregated Ethernet interfaces are saved on a file that is stored on the master Routing Engine and is synchronized with the backup Routing Engine. The file is updated after each successful commit that required changes to the MAC addresses table.

In earlier releases, 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.

IPsec

- **Support for configuring IPsec (site-to-site) VPN tunnels (MX150)**—Starting in Junos OS Release 17.3R1, the MX150 supports IPsec VPN connections or tunnels. You can configure a route-based VPN or a policy-based VPN. You implement a policy-based VPN if the remote VPN device is a non-Juniper Networks device and only one subnet or network at the remote site across the VPN needs to be accessed.

IPv6

- **IPv6 support (MX150)**—Starting in Junos OS Release 17.3R1, Junos OS supports IPv6 features on the MX150. The following is a list of some of the IPv6 features supported:
 - IPv6 forwarding
 - IPv6 path maximum transmission unit (MTU) discovery
 - Neighbor discovery
 - Static routes for IPv6
 - Internet Control Message Protocol (ICMP) version 6

Layer 2 Features

- **Support for Junos Fusion Provider Edge (MX10003 routers)**—Starting in Junos OS Release 17.3R1, you can configure MX10003 Universal Routing Platforms as aggregation devices in a Junos Fusion Provider Edge topology. Junos Fusion Provider Edge brings the Junos Fusion technology to the service provider edge. In a Junos Fusion Provider Edge, MX Series routers act as aggregation devices, while EX4300 and QFX5100 switches act as satellite devices.

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

- **Support for Layer 2 protocols on MX10003 routers**—Starting in Junos OS Release 17.3R1, all Layer 2 bridging features are supported on MX10003 routers.

- **Support for Layer 2 and Layer 3 features (MX150)**—Starting in Junos OS Release 17.3R1, the MX150 supports the following Layer 2 and Layer 3 features:
 - Layer 2 protocols and including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
 - VLAN support—VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
 - Link Layer Discovery Protocol (LLDP)—Enables advertising the identity and capabilities on a LAN, and receive information about other network devices.
 - Layer 3 routing protocols and MPLS

Layer 2 VPN

- **Support of ping utility for testing CE device connectivity (MX Series with MPC and MIC)**—Starting in Junos OS Release 17.3R1, reachability to the customer endpoint can be achieved from the service endpoint in a network. This feature is supported in a virtual private LAN service (VPLS), hierarchical VPLS (H-VPLS), and Ethernet VPN (EVPN) network. It is based on the LSP ping infrastructure, where the **ping** utility is extended to use the CE device IP address as the target host and the PE device loopback address as the source for a specific VPLS or EVPN routing instance.

To implement this feature, issue the **ping ce-ip destination-ip-address instance routing-instance-name source-ip source-ip-address** command on a PE device. Based on the configured routing instance type, the command output displays the connectivity information of the CE device.

[See [Pinging Customer Edge Device IP Address](#).]

- **Support for Group VPN (MX150)**—Starting in Junos OS Release 17.3R1, Junos OS supports Group VPN on the MX150. Group VPN extends existing IPsec architecture to support group-shared security associations. The group server manages group keys and policies and distributes them to group members. Group VPN provides the following benefits:
 - Data security and transport authentication.
 - High-scale network meshes, eliminating complex peer-to-peer key management with group encryption keys.
 - Full-time, direct communications between sites, without requiring transport through a central hub.

[See [Group VPN Overview](#).]

- **Support for connectivity fault management**—Starting in Junos OS Release 17.3R1, Junos OS supports multiple up maintenance association end points (MEPs) for a single combination of maintenance association ID and maintenance domain ID for Layer 2 VPN local switching.

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

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

- **Support for chained composite next hops**—Starting in Junos OS Release 17.3R1, you can enable composite chained next hops on MPCs on MX Series routers to manage ingress traffic for Layer 2 circuits and Layer 2 VPNs. A chained composite next hop allows the router to direct sets of routes sharing the same destination to a common forwarding next hop, rather than having each route also include the destination. This helps facilitate large volumes of traffic.

To enable composite chained next hop for ingress traffic, include the `l2ckt` or `l2vpn` statement at the `[edit routing-options forwarding-table chained-composite-next-hop ingress]` hierarchy level.

[See [Chained Composite Next Hops for Layer 2 VPNs and Layer 2 Circuits](#).]

Layer 3 Features

- **Junos Fusion support (MX2008)**—Starting in Junos OS Release 17.3R1, the Junos OS supports a network system named Junos Fusion. Based on the 802.1BR standard, Junos Fusion is a combination of aggregation devices and satellite devices that appear to the rest of the network as a single device. Junos Fusion expands the port density of the aggregation device and allows it to send and receive traffic using the customer-facing ports of the directly connected satellite devices. The composite of the aggregation device and satellite devices—the Junos Fusion—is configured and managed through the aggregation device. You can configure MX2008 Universal Routing Platforms as an aggregation device.

[See [Junos Fusion Provider Edge Overview](#).]

- **Support for Layer 3 protocols (MX10003)**—Starting in Junos OS Release 17.3R1, Layer 3 protocols are supported on MX10003 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 Junos Telemetry Interface (MX150)**—Starting with Junos OS Release 17.3R1, the Junos Telemetry Interface is supported on the MX150 router. Junos Telemetry Interface enables you to provision sensors to stream telemetry data for network elements without involving polling.

On the MX150 router, only the following sensors are supported:

- Physical interfaces (UDP and gRPC streaming)
- Network Discovery Protocol table state (gRPC streaming only)
- Address Resolution Protocol table state (gRPC streaming only)
- IPFIX inline flow aggregation (UDP streaming only)
- Chassis components (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 commands 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 to configure YANG files for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.3R1, you can add user-defined YANG files that provide mappings between the XML path and the OpenConfig path for data streamed through the Junos Telemetry Interface. Previously, only the Junos OpenConfig package was available for providing these mappings to the XML proxy for data streamed through gRPC. To add YANG files, include the **request system yang add package *package-name* proxy-xml module *yang-file-path*** operational command. You can validate the YANG module by using the **request system yang validate proxy-xml module *yang-file-path*** command. To delete a YANG file, use the **request system yang delete package *package-name* proxy-xml *yang-file-path*** operational command.

[See [Creating YANG Files for XML Proxy for Junos Telemetry Interface](#).]

- **Enhancements to BGP peer sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.3R1, telemetry data streamed through gRPC for BGP peers is reported separately for each routing instance. To export data for BGP peers, you must now include the following path in front of all supported paths: **/network-instances/network-instance/[name_ 'instance-name']/protocols/protocol/**

Additionally, the following paths are also now supported:

- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/accepted**
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/snmp-peer-index**
- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/afi-safis/afi-safi/state/queues/output**
- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/afi-safis/afi-safi/state/queues/input**
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/ImportEval**
- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/state/ImportEvalPending**

Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Support for packet loss priority for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.3R1, you can specify loss priority for telemetry packets streamed through UDP only. Loss priority settings help determine which packets are dropped from the network during periods of congestion.

To configure, include the **loss-priority** (**high** | **low** | **medium-high** | **medium-low**) statement at the **[edit services analytics export-profile *profile-name*]** hierarchy level. To apply an export profile to a sensor, include the **export-name *profile-name*** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level. The **show agent sensors** command includes a new **loss-priority** field that is displayed for each sensor when this new option is configured.

[See [Configuring a Junos Telemetry Interface Sensor.](#)]

- **Junos Telemetry Interface support (MX10003 and MX204)**—Starting with Junos OS Release 17.3R1, MX10003 and MX204 routers support the Junos Telemetry Interface, which enables you to provision sensors to export telemetry data for various network elements. 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.

MPLS

- **Anchor point redundancy support for pseudowire subscriber logical Interfaces (MX Series)**—Starting in Junos OS Release 17.3R1, stateful anchor point redundancy support is provided for pseudowire subscriber logical interfaces by the underlying redundant logical tunnel interface in active-backup mode. This redundancy protects the access and the core facing link against anchor Packet Forwarding Engine failure.

Both transport and services logical interfaces created for the pseudowire subscriber logical interface are stacked on the underlying redundant logical tunnel control logical interface. This logical interface stacking model is used for both redundant and non-redundant pseudowire subscriber logical interfaces.

[See [Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview.](#)]

- **Support for features on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.3R1, MPC7E, MPC8E, and MPC9E support the following features:
 - LDP uses the longest match to learn the routes aggregated or summarized across OSPF areas or IS-IS levels in the interdomain.
 - Support for notifications on the service node when the access pseudowire goes down, and efficient termination capabilities when Layer 2 and Layer 3 segments are interconnected.

[See [Pseudowire Termination: Explicit Notifications for Pseudowire Down.](#)]

- BGP PIC Edge for RSVP enables you to implement a solution where a protection path is calculated in advance to provide an alternative forwarding path in case of path failure.

[See [show rsvp version.](#)]

- Circuit cross-connect (CCC) encapsulation is supported on the transport side of an MPLS pseudowire subscriber logical interface. This feature helps in migrating or deploying seamless MPLS architectures in access networks.

[See [Pseudowire Subscriber Logical Interfaces Overview.](#)]

- `inet` and `inet6` families are supported on the services side of an MPLS pseudowire subscriber as well as non subscriber logical interfaces.
- Distributed denial-of-service (DDoS) protection is supported on the services side of an MPLS pseudowire subscriber logical interface.
- Policer and filter are supported on the services side of an MPLS pseudowire subscriber logical interface.
- Accurate transmit logical interface statistics are supported on the services side of an MPLS pseudowire subscriber logical interface.
- Inline IPFIX is supported on the services side of an MPLS pseudowire subscriber logical interface.
- Port mirroring is supported on the services side of an MPLS pseudowire subscriber logical interface.

Multicast

- **PIM resolve type-length-value (TLV) for multicast in seamless MPLS (MX Series)**—Starting in Junos OS Release 17.3R1, Junos OS adds support for RFC 5496, Reverse Path Forwarding (RPF) Vector TLV . With this support, Protocol Independent Multicast (PIM) can be used in environments where the core routers do not maintain external routes, for example in a seamlessMPLS network.

[See [rpf-vector](#).]

- **Support for IPv6 multicast Rosen version 7 (MX Series)**—Starting in Junos OS Release 17.3R1, Junos OS multicast support extends to the default multicast distribution tree (MDT) for Rosen 7 multicast virtual private networks (MVPN) and data MDT for both Rosen 6 (PIM-ASM) and Rosen 7 (PIM-SSM). The IPv6 support applies to the customer space only.

[See [Draft-Rosen Multicast VPNs Overview](#) .]

Network Management and Monitoring

- **mLDP MIB extends support to LDP point-to-multipoint (P2MP) LSPs (MX Series)**—Starting in Junos OS Release 17.3R1, the mLDP MIB builds on the objects and tables that are defined in RFC 3815, which only support LDP point-to-point label switched paths (LSPs). This mLDP MIB provides support for managing multicast LDP point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs. The mLDP MIB tables are directly accessible through SNMP. All objects in the mLDP MIB are read-only and cannot be created or set through SNMP. This implementation of mLDP MIB is specified in `draft-ietf-mpls-mldp-mib`.
- **Support for automatic targeted distribution of logical interface sets of static VLANs over aggregated ethernet logical interfaces (MX Series)**—Starting in Junos OS Release 17.3R1, automatic targeted distribution of logical interface sets of static VLANs over aggregated Ethernet logical interfaces is supported. When targeted distribution is set for a logical interface sets then the logical interface set participates in targeting and the link selected for the logical interface set is propagated to the underlying logical interfaces. You can assign weight for all the targeted subscribers like PPPoE, demux, and conventional VLANs based on the business, CoS, or bandwidth requirement. To configure the **weight** statement at either the `[edit interfaces interface-set interface-set-name targeted-options]` or the `[edit`

interfaces *interface-name* unit *unit-number* targeted-options] hierarchy level to assign the member links for the logical interface set or logical interface based on the weight value.

[See [Understanding Support for Targeted Distribution of Logical Interface Sets of Static VLANs over Aggregated Ethernet Logical Interfaces.](#)]

Operation, Administration, and Maintenance (OAM)

- **Junos OS daemons to natively emit JSON output (MX Series)**—Starting with Junos OS Release 17.3R1, the operational state emitted by daemons is supported in JSON format as well as XML format. To configure JSON format, specify the following CLI command: **set system export-format state-data json compact**. To specify JSON format for specific command output, include **display json** in specific CLI commands.
- **Support for Ethernet OAM Rx statistics for CCM (MX Series)**—Starting in Junos OS Release 17.3R1, the **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *md-name* maintenance-association *ma-id* local-mep *mep-id* remote-mep *mep-id*** command displays Ethernet OAM Rx statistics. The Ethernet OAM Rx statistics displays the number of CCM PDUs received for a particular maintenance association and remote MEP and does not include error packets received.

NOTE: The Ethernet OAM Rx statistics are not displayed for UP MEP on trunk modes if the network-services mode is configured as IP.

If you perform unified ISSU, the counter is reset to zero. The counter is also reset to zero when the session flaps or if the session is down.

NOTE: If you do not provide the local MEP and remote MEP IDs, the **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *md-name* maintenance-association *ma-id* local-mep *mep-id* remote-mep *mep-id*** command does not display latest statistics. Also, if you do not provide the remote MEP ID, then actual received statistics display zero.

- **Support for connectivity fault management (CFM) monitoring between customer-edge (CE) and provider-edge (PE) devices (MX Series)**—Starting in Junos OS Release 17.3R1, you can enable CFM monitoring between PE devices and CE devices when the CE device is not a Juniper Networks device by using the remote defect indication (RDI) bit. When the status of the EVPN provider edge device is standby, the EVPN VPWS service is notified and it sets the interface status to CCC-down. When the interface status is CCC-down, it indicates that the PE service is down. When you enable CFM monitoring, CFM propagates the status of the PE device via the RDI bit in the CC messages. Thus, the CE device is aware that the PE device is down. The RDI bit is cleared when the service is back up.

To enable CFM monitoring by using the RDI bit, use the **interface-status-send-rdi** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name continuity-check]** hierarchy level.

Alternately, you can enable CFM monitoring by using the **interface-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name continuity-check]** hierarchy level.

- **Nonstop active routing support for link fault management (LFM) (MX Series)**—Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when GRES is configured. When the lfmd daemon runs on the backup Routing Engine as well, the LFM states are kept in sync and so minimal work is required by the lfmd daemon after switching over. To verify if the LFM states are in sync, use the **show oam ethernet link-fault-management** command on both master and backup Routing Engines. In Junos OS Release 17.2R1 and earlier, the lfmd daemon runs only on the master Routing Engine when GRES is configured.
- **Junos OpenConfig to support adjacent RIB operational state model (MX Series)**—Starting with Junos OS Release 17.3R1, **adj-rib-in-pre** and **adj-rib-out-post** tables have been added for the OpenConfig RIB operational state mode. The BGP RIB consists of several tables per address family, consisting of **loc-rib** and **per-neighbor** tables.
- **Support for inline CCM and BFD on MX10003 routers**—MX10003 routers support inline transmission of continuity check messages (CCMs) to achieve maximum scaling of CCMs. By enabling inline transmission of CCMs, you can delegate transmission of CCMs to the forwarding ASIC (that is, to the hardware). Inline transmission enables the system to handle more connectivity fault management (CFM) sessions per line card. MX10003 routers also support the Bidirectional Forwarding Detection (BFD) protocol, which is a mechanism that detects failures in a network.

Port Security

- **Media Access Control Security (MACsec) support on Terabit Interface card (MX10003)**—Starting in Junos OS Release 17.3R1, JunosOS supports MACsec on the 12x QSFP28 Terabit Interface card (TIC) in MX10003 routers. MACsec is an 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. MACsec is standardized in IEEE 802.1AE.

Routing Policy and Firewall Filters

- **Support for packet forwarding features (MX150)**—Starting in Junos OS Release 17.3R1, the MX150 supports the following key packet forwarding features:

- **Basic Layer 2 features and protocols**—You can configure layer 2 features that can vary from the very simple (aggregated Ethernet trunk interfaces, spanning trees), to the more complex (inner and outer VLAN tags, broadcast domains), to the very complicated (integrated bridging and routing, layer 2 filtering).
- **Class of service (CoS)**—You can configure CoS features to provide multiple classes of service for different applications. CoS enables you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. It enables you to provide differentiated services when best-effort traffic delivery is insufficient.
- **Firewall filters and policers**—You can configure firewall filters 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. You can use policing to apply limits to traffic flow and specify the action to be taken for packets that exceed those limits.
- **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 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.
- **Bypass loopback with firewall filter tunnel encapsulation (MX Series)**—Starting in Junos OS Release 17.3R1, static filter based generic routing encapsulation (GRE) tunnels no longer use a loopback stream for transit traffic. The new default, which allows for increased bandwidth utilization on MPCs using the MX Series chipset, is to skip the loopback. In addition, support for IPv4 as the outer IP is available (the inner payload supports both IPv4 and IPv6). Egress sampling on the outer header is not affected. This change does not apply to GRE in UDP or to dynamic tunnels.

This change applies to the following filter-based tunneling commands in the CLI:

set firewall family inet6 filter *filter* term *term* then encapsulate *tunnel*

set firewall tunnel-end-point *tunnel* ipv4 source-address *ipv4 address*

set firewall tunnel-end-point *tunnel* ipv4 destination-address *ipv4 address*

set firewall tunnel-end-point *tunnel* gre

[See [Filter-Based Tunneling Across IPv4 Networks](#).]

Routing Protocols

- **Support for timing and synchronization on Terabit Interface card (MX10003)**—Starting in Junos OS Release 17.3R1, 12x QSFP28 Terabit Interface card (TIC) in MX10003 routers support the following timing and synchronization features:

- **SyncE support with ESMC**—Synchronized Ethernet with Ethernet synchronization Message Channel (ESMC) is supported as per the ITU G.8264 specification. ESMC is a logical communication channel. It transmits synchronization status message information, which is the quality level of the transmitting synchronous Ethernet equipment clock, by using ESMC protocol data units.
- **PTP support**—Precision Time Protocol (PTP), also known as IEEE 1588v2, is a packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks. IEEE 1588 PTP (Version 2) clock synchronization standard is a highly precise protocol for time synchronization that synchronizes clocks in a distributed system. The time synchronization is achieved through packets that are transmitted and received in a session between a master clock and a slave clock. One step clock mode operation for the master clock is supported.
- **BITS (T1/E1) Interface support**—BITS support for input and output on T1/E1 framed and 2.048MHz unframed clock input.
- **GPS external clock interface and TOD support**—GPS input and output support for 1 MHz/5 MHz/10 MHz and PPS signal.

[See [Ethernet Synchronization Message Channel Overview](#)].

- **Routing protocol process (rpd) recursive resolution over multipath (MX Series)**—Starting in Junos OS Release 17.3R1, when a BGP prefix that has a single protocol next hop is resolved over another BGP prefix that has multiple resolved paths (unilist), all the paths are selected for protocol next-hop resolution. In prior Junos OS releases, only one of the paths is picked for protocol next-hop resolution. This new feature benefits densely connected networks where BGP is used to establish infrastructure connectivity such as WAN networks with high equal-cost multipath and seamless MPLS topology.

To configure recursive resolution over multipath, define a policy that includes the **multipath-resolve** action at the **[edit policy-options policy-statement *policy-name* then]** hierarchy level and import the policy at the **[edit routing-options resolution rib *rib-name*]** hierarchy level.

Currently, if you apply the policy on **bgp.l2vpn.0** only, the RIB, also known as the routing table reflects recursively resolved multiple paths only in the control plane, you need to explicitly apply the policy on **mpls.0** to reflect recursively resolved multiple paths on the data plane also.

[See [Configuring Recursive Resolution over BGP Multipath](#).]

- **Redistribution of IPv4 routes over IPv6 routes into BGP through tunnels (MX Series)**—Starting in Release 17.3R1, Junos OS devices can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, IPv4 traffic is tunneled from CPE devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers, which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop. Currently the dynamic IPv4-over-IPv6 tunnel feature does not support unified ISSU.

To configure a dynamic IPv4-over-IPv6 tunnel, include the **dynamic-tunnels** statement at the **[edit routing-options]** hierarchy level.

[See [Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.](#)]

- **Support for IS-IS SPRING and RSVP coexistence (MX Series)**—Starting in Junos OS Release 17.3R1, the routing protocol process (rpd) takes into account the bandwidth used by SPRING traffic to calculate the balance bandwidth available for RSVP-TE. The allocated bandwidth for RSVP is periodically modified based on the traffic on the SPRING interface and its bandwidth utilization. To configure automatic bandwidth calculation, include the **auto-bandwidth template** statement at the **[edit routing-options]** hierarchy level. You can apply the **auto-bandwidth template** configuration either globally at the **[edit protocols isis source-packet-routing traffic-statistics]** hierarchy level or at the **[edit protocols isis interface *interface-name*]** hierarchy level. This feature is useful for networks that are moving to SPRING but also have RSVP deployed, and continue to use both SPRING and RSVP.

[See [auto-bandwidth.](#)]

- **Support for BGP large communities (MX Series)**—Starting in Junos OS Release 17.3R1, BGP community is enhanced to support a BGP large community, which uses 12-byte encoding. The most significant 4 bytes encode an autonomous system number or global administrator and the remaining two 4 bytes encode operator defined local values. Currently, BGP normal community (4 byte) and BGP extended community (6 byte) provide limited support for BGP community attributes after the introduction of a 4 byte autonomous system number. Configure the large BGP community attributes at the **[edit policy-options community *community-name* members]** hierarchy level and at the **[edit routing-options static route *route* community]** hierarchy level with keyword **large** followed by three 4-byte unsigned integers separated by colons. The attributes are represented as large:autonomous system number:local value 1:local value2.

[See [Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions](#)]

- **Support for inline Two-Way Active Measurement Protocol (TWAMP) server and client on MX10003 routers**—Starting in Junos OS Release 17.3R1, supports the inline Two-Way Active Measurement Protocol (TWAMP) control-client and server for transmission of TWAMP IPv4 UDP probes between the session-sender (control-client) and the session-reflector(server). The TWAMP control-client and server can also work with a third-party server and control-client implementation. TWAMP is an open protocol for measuring network performance between any two devices that support TWAMP.

Security

- **Secure boot (MX10003)**—Starting in Junos OS Release 17.3R1, 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. secure boot is enabled by default on supported platforms.

Services Applications

- **ECDSA authentication for IKE SA and AES-GCM encryption for IPsec SA (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.3R1, you can configure the Elliptic Curve Digital Signature Algorithm (ECDSA) authentication method for an IKE security association (SA) and the Advanced Encryption Standard in Galois/Counter Mode (AES-GCM) encryption algorithm for an IPsec SA for MS-MPCs and MS-MICs. Junos OS supports the ECDSA 256-bit and 384-bit moduli methods and the AES-GCM 128-bit, 192-bit, and 256-bit encryption algorithms.

[See [Configuring IKE Proposals](#) and [Configuring IPsec Proposals](#).]

- **Support for IPv6 GRE tunnels (MX Series)**—Starting in Junos OS Release 17.3R1, you can configure IPv6 generic routing encapsulation (GRE) tunnel interfaces on MX Series routers. This lets you run a GRE tunnel over an IPv6 network. Packet payload families that can be encapsulated within the IPv6 GRE tunnels include IPv4, IPv6, MPLS, and ISO. Fragmentation and reassembly of the IPv6 delivery packets is not supported.

To configure an IPv6 GRE tunnel interface, specify IPv6 addresses for **source** and **destination** at the **[interfaces gr-0/0/0 unit 0 tunnel]** hierarchy level.

[See [GRE Keepalive Time Overview](#).]

- **Increased number of IPv4 RPM probes (MX Series with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.3R1, you can increase the number of IPv4 **icmp-ping** and **icmp-ping-timestamp** real-time performance monitoring (RPM) probes that can run simultaneously. Use the **delegate-probes** statement to configure an MS-MPC or MS-MIC services interface to perform the RPM processing for the probes, enabling more probes to run simultaneously.

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

- **Inline TWAMP requester support (MX2010 and MX2020 routers)**—Starting in Junos OS Release 17.3R1, MX2010 and MX2020 routers support the inline Two-Way Active Measurement Protocol (TWAMP) control-client and session-sender for transmission of TWAMP probes using IPv4 between the sender (control-client or session-sender) and the receiver (server or session-reflector). The control-client and session-sender reside on the same router. The TWAMP control-client can also work with a third-party server implementation.
- **Support for enhancing the current Inline JFlow scale limits for XL-based and EA-based linecards for MX routers**—Starting in Junos OS Release 17.3R1, the **ipv4-flow-table-size**, **ipv6-flow-table-size**, **vppls-flow-table-size**, and **mpls-flow-table-size** allow upto 245 **flow-table-size** to support 64M flows at the **[edit chassis fpc slot-number inline-services flow-table-size]** hierarchy level. The existing limit on

flow-export-rate under **inline-jflow** for each family in the sampling instance is increased to 3200 from 400.

- **Support for Inline services (MX150)**—Starting in Junos OS Release 17.3R1, the MX150 supports inline active flow monitoring services. Inline active flow monitoring provides for higher scalability and performance and is implemented on the Packet Forwarding Engine. Version 9 template and IP Flow Information Export (IPFIX) template are supported to define a flow record template suitable for IPv4 or IPv6 traffic.

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

- **RPM support for IPsec and GRE tunnels (MX Series router with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.3R1, you can apply real-time performance monitoring (RPM) to IPsec tunnels and GRE tunnels for PIC-based and Routing Engine based RPM clients and servers if you are using MS-MPCs or MS-MICs. Packet Forwarding Engine based RPM is not supported for IPsec tunnels. Support of RPM on IPsec tunnels enables service-level agreement (SLA) monitoring for traffic transported in IPsec tunnels.

[See [Real-Time Performance Monitoring Services Overview](#).]

- **NAT with deterministic IP address and port mapping (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.3R1, support for deterministic NAT mapping for NAPT44 is extended to the MS-MPC and MS-MIC. Deterministic NAT mapping ensures that a given internal IP address and port are always mapped to the same external IP address and port range, and the reverse mapping of a given translated external IP address and port are always mapped to the same internal IP address. Deterministic NAT mapping eliminates the need for logging address translations.

[See [Configuring Deterministic NAPT](#).]

- **Support for TWAMP server and client (MX150)**—Starting in Junos OS Release 17.3R1, the MX150 supports the inline Two-Way Active Measurement Protocol (TWAMP) control-client and server for transmission of TWAMP IPv4 UDP probes between the session-sender (control-client) and the session-reflector (server). The TWAMP control-client and server can also work with a third-party server and control-client implementation. TWAMP is an open protocol for measuring network performance between any two devices that support TWAMP.

[See [Two-Way Active Measurement Protocol Overview](#).]

- **Increase in IKE tunnel setup rate (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.3R1, the IKE tunnel setup rate has increased if you are using MS-MPCs or MS-MICs. This increase is the result of moving the public key cryptographic operations to the MS-MPC or MS-MIC.

[See [Understanding Junos VPN Site Secure](#).]

- **Maximum number of RPM probes increased (MX Series routers)**—Starting in Junos OS Release 17.3R1 and 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](#).]

Software Defined Networking (SDN)

- **Support for Junos Node Slicing on MX480 routers**—Starting with Junos OS Release 17.3R1, MX480 routers support Junos Node Slicing. Junos node slicing is the capability to partition an MX Series router to make it appear as multiple, independent routers. Each partition has its own independent Junos OS control plane and dataplane, which run as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF). In the node slicing setup, the MX Series router functions as the base system (BSYS). Junos node slicing enables the convergence of multiple services on a single physical infrastructure while avoiding the operational complexity involved.

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

- **Support for OpenDaylight (ODL) controller on MX Series routers**—Starting with Junos OS Release 17.3R1, MX Series router supports OpenDaylight (ODL) controller (Boron-SR1 release), which provides an open source platform for network programmability aimed at enhancing software-defined networking (SDN). The ODL controller provides a southbound Network Configuration Protocol (NETCONF) connector API, which uses NETCONF and YANG models to interact with a network device. You can use the ODL controller to orchestrate and provision MX Series routers, and execute remote procedure calls (RPCs) to the routers to get state information. Also, the ODL controller enables you to carry out configuration changes in the routers. To configure the ODL controller to interoperate with MX Series routers, include the **netconf rfc-compliant** and **netconf yang-compliant** statements at the **[edit system services]** hierarchy level.

[See [Configuring Interoperability Between MX Series Routers and OpenDaylight](#)]

- **Advanced Forwarding Interface (AFI) API is available for vMX routers**—Starting in Junos OS Release 17.3R1, the Advanced Forwarding Interface (AFI) version 1.0 is available for vMX routers. AFI APIs are provided as C++ APIs only. The APIs allow developers to interact with the Packet Forwarding Engine by accessing a section of the forwarding path from within a sandbox to affect the traffic that enters that part of the path. The sandbox is provided by Junos OS after CLI-based configuration and has one or more pairs of input and output ports that represent the points along the forwarding path at which the AFI clients enter and exit the path to do their work.

Subscriber Management and Services

- **Support for excluding tunnel attributes from RADIUS Access-Request messages (MX Series)**—Starting in Junos OS Release 17.3R1, 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

[See [Configuring How RADIUS Attributes Are Used for Subscriber Access.](#)]

- **Clearing accounting option statistics from the Packet Forwarding Engine (MX Series)**—Starting in Junos OS Release 17.3R1, you can issue the **clear interfaces statistics *interface-name*** command to clear counters for accounting statistics received on the logical interface from the Packet Forwarding Engine. The existing statistics are stored as the new current baseline statistics and the counters are reset to zero. This applies to interfaces for which accounting statistics are collected as specified by the **interface-profile** statement at the **[edit accounting-options]** hierarchy level.

Include the **allow-clear** statement in the interface profile to enable reporting of the cleared (new current baseline) statistics to the accounting flat file. Reporting is disabled by default. When you clear statistics for an interface that does not have this statement in its interface profile, the CLI displays the statistics as cleared, but this is not reported to the flat file.

[See [Configuring the Interface Profile.](#)]

- **Filter actions extended to dynamic filters (MX Series)**—Starting in Junos OS Release 17.3R1, you can include the **dscp *value*** action for the inet address family and the **traffic-class *value*** action for the inet6 address family in dynamic, parameterized filters. This means that you can configure a user-defined dynamic variable or a static value for the action value. In earlier releases, these actions are supported only for static (nonparameterized) filters.

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

- **Support for inline IP reassembly on GRE tunnel interfaces (MX Series routers with MPCs)**—Starting in Junos OS Release 17.3R1, you can configure fragmentation and inline reassembly of generic routing encapsulation (GRE) packets on GRE tunnel interfaces on MX Series routers with the following Modular Port Concentrators: MPC7E, MPC8E, and MPC9E.

[See [Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation](#)]

- **Limiting subscribers based on client type for different hardware elements (MX Series)**—Starting in Junos OS Release 17.3R1, use the **subscribers-limit** stanza at the **[edit system services resource-monitor]** hierarchy level to configure the maximum number of subscribers by client type (DHCP, L2TP, PPPoE, or the sum of all three) that are allowed per chassis, MPC, MIC, and port. Subscriber login is denied when the number of subscribers having that type exceeds the configured limit. This feature ensures that the number of subscribers per hardware element does not exceed the number that your network can serve with stability at the desired bandwidth. When the limit is reached for a hardware element, new subscribers can connect to another hardware element in the same broadcast domain. When you configure the limit on one or more legs of an aggregated Ethernet interface, login is denied if the subscriber count exceeds the value on any of the legs.

Use the **show system resource-monitor subscribers-limit** command to display information about subscriber limits.

[See [Limiting Subscribers by Client Type and Hardware Element with Resource Monitor.](#)]

- **Support for sending LAC NAS-port and LAC IP-address attributes to RADIUS for MX Routers**—Starting in Junos OS Release 17.3R1, you can override the following at the `[edit access profile set radius options override]` hierarchy level:

- **nas-port** with the LAC side **nas-port** information.
- **nas-ip-address** with the l2tp LAC endpoint IP address information.

- **Support for load-based throttling of subscribers (MX Series)**—Starting in Junos OS Release 17.3R1, the **no-load-throttling** statement disables line card load-based throttling when configured at the `[edit system services resource-monitor]` hierarchy level. Load-based throttling is also disabled when the **no-throttle** statement is configured at the `[edit system services resource-monitor]` hierarchy level.

- **DDoS protection flow detection for enhanced subscriber management (MX Series Routers)**—Starting in Junos OS Release 17.3R1, enhanced subscriber management supports flow detection for DDoS protection. Enable flow detection by including the **flow-detection** statement at the `[edit system ddos-protection global]` hierarchy level. Flows that violate a DDoS protection policer are tracked as suspicious flows; they become culprit flows when they violate the policer bandwidth for the duration of a configurable detection period. Culprit flows are dropped, kept, or policed to below the allowed bandwidth level. Suspicious flow tracking stops if the violation stops before the detection period expires.

Most flow detection attributes are configured at the packet level or flow aggregation level of the CLI hierarchy (`[edit system ddos-protection protocols protocol-group packet-type]`). By default, flow detection automatically generates reports for events associated with the identification and tracking of culprit flows and bandwidth violations. Use commands at the **show ddos-protection** hierarchy level and **culprit-flows** or **culprit-flows detail** to display flow detection information and statistics on the basis of protocol, packet type, or subscriber management.

[See [DDoS Protection Flow Detection Overview](#)]

- **Excluding channel information from interface descriptions (MX Series)**—Starting in Junos OS Release 17.3R1, you can exclude channel information from being reported by default in the description for channelized interfaces that are included in RADIUS attributes such as NAS-Port-ID (87) and Calling-Station-ID (31). In earlier releases, you can exclude only adapter (PIC) and subinterface (logical interface number) information from an interface description.

[See [Interface Text Descriptions for Inclusion in RADIUS Attributes](#).]

- **BPCEF phase 2 enhancements (MX Series)**—Starting in Junos OS Release 17.3R1, support for additional OCS and PCRF features are added using Gy and Gx protocols. The new statements:
 - **accept-sdr** is added for PCRF partition at the `[edit access pcrf partition partition-name]` hierarchy level.
 - **alternative-diameter-partition** is added for OCS partition at the `[edit access ocs partition partition-name]` hierarchy level.

[See [Understanding Gx Interactions Between the Router and the PCRF](#) and [Configuring the Diameter Transport](#).]

- **System logs and traps added for Diameter peer connect/disconnect state changes (MX Series)**—Starting in Junos OS Release 17.3R1, the following event options related to Diameter peer connect and disconnect events are available to raise a trap when the corresponding state change occurs:
 - `jdiameterd_dne_state_connected`—Diameter network element (DNE) connected over a single peer.
 - `jdiameterd_dne_state_fully_connected`—DNE connected through at least two peers.
 - `jdiameterd_dne_state_disconnected`—DNE lost its connection.
 - `jdiameterd_peer_premiership_acquired`—Peer became primary for DNE.
 - `jdiameterd_peer_premiership_released`—Peer stopped being primary for DNE.
 - `jdiameterd_peer_state_down`—Peer is closing.
 - `jdiameterd_peer_state_open`—Peer reached i-open state.
 - `jdiameterd_peer_state_suspected`—Peer is downgraded to suspected state.

You can configure these at the `[edit event-options policy policy-name]` hierarchy level. Each of the event traps generates a corresponding ERRMSG system log.

[See [System Log Explorer](#).]

- **Diameter peers and transports support IPv6 addresses (MX Series)**—Starting in Junos OS Release 17.3R1, you can use IPv6 addresses for Diameter peers and transport connections. You must configure the same address family type for corresponding peers and transport connections. In earlier releases, only IPv4 addresses are supported, requiring the use of NAT to enable peering between IPv4 and IPv6 Diameter nodes.

[See [Configuring Diameter Peers](#) and [Configuring the Diameter Transport](#).]

- **Support for concurrent subscriber secure policy and FlowTapLite (MX Series)**—Starting in Junos OS Release 17.3R1, you can enable both DTCP-based flow-tap services on tunnel interfaces (FlowTapLite) and DTCP-initiated and RADIUS-initiated subscriber secure policies concurrently on the same router. Concurrent support enables using DTCP for monitoring both dynamic subscribers and static logical interfaces for business subscribers, as in a Layer 2-based wholesale topology that uses Extensible Subscriber Services Manager (ESSM). In earlier releases, concurrent use of subscriber secure policies and FlowTapLite is not supported.

[See [Guidelines for Configuring Subscriber Secure Policy Mirroring](#).]

- **Disabling RADIUS-initiated subscriber secure policy mirroring (MX Series)**—Starting in Junos OS Release 17.3R1, you can use the `dtcp-only` statement to prevent RADIUS-initiated subscriber secure policy mirroring from being enabled, while allowing both DTCP-initiated mirroring and DTCP-based flow-tap services (FlowTapLite) to be enabled. Requests from RADIUS to attach a subscriber secure policy (mirroring service) to a subscriber are rejected. This statement has no effect on existing RADIUS-initiated mirroring services. You must issue the statement before such services are activated for a subscriber. Subscriber login and session establishment are not affected.

[See [Disabling RADIUS-Initiated Subscriber Secure Policy Mirroring](#).]

- **Appending subscriber information to redirect URLs (MX Series)**—Starting in Junos OS Release 17.3R1, you can append information about the subscriber retrieved from the subscriber session database when the redirect URL is returned to the HTTP client. You specify the attributes in the redirect URL format in the Activate-Service VSA (26–65) or Deactivate-Service VSA (26–66) included in the RADIUS Access-Accept message when the subscriber is authenticated or in a Change of Authorization (CoA) message. Only the following attributes are supported: subscriber IP or IPv6 address, NAS IP address, requested URL, NAS port ID, MAC address, subscriber session ID, and username.

[See [Adding Subscriber Information to HTTP Redirect URL](#).]

- **HTTP status code 307 support (MX Series)**—Starting in Junos OS Release 17.3R1, the HTTP status code returned with the redirect URL by the redirect server depends on the HTTP version used by the HTTP client that sent the GET message. When the version is later than 1.0, the 307 (Temporary Redirect) status code is returned. When the version is 1.0, the 302 (Found) status code is returned. In earlier releases, only the 302 status code is returned with the redirect URL. Both codes inform the HTTP client to use the original URL for subsequent GET requests.

[See [HTTP Redirect Service Overview](#).]

- **Subscriber management support for Junos Node Slicing**—Starting with Junos OS Release 17.3R1, the MX Series routers that have Junos Node Slicing configured support all subscriber management features and services. Subscriber management provides capabilities such as subscriber access, authentication, and service creation, activation, and deactivation. The subscriber management services include DHCP, PPP, L2TP, VLAN, and pseudowire. However, in this release, the subscriber management services for Junos Node Slicing do not include advanced services and do not support unified in-service software upgrade (unified ISSU).
- **Support for Broadband Edge on MX10003 routers**—Starting in Junos OS Release 17.3R1, MX10003 supports the next-generation broadband edge software architecture for wireline subscriber management. With enhanced subscriber management, you can take advantage of optimized scaling and performance for configuration and management of dynamic interfaces and services for subscriber management.

Virtual Chassis

- **Support for host infrastructure(MX10003)**—Starting in Junos OS Release 17.3R1, MX10003 supports host infrastructure that can launch Junos OS virtual machine (VM) based on configuration data, monitor and manage the VM and the host-networking infrastructure, support Junos OS and host software upgrade, collect hardware errors for Junos OS error reporting and act as a proxy to Junos OS for executing host operations. Only one VM is supported per Routing Engine.

SEE ALSO

[Changes in Behavior and Syntax](#) | 120

[Known Behavior](#) | 125

[Known Issues](#) | 131

[Resolved Issues | 144](#)

[Documentation Updates | 163](#)

[Migration, Upgrade, and Downgrade Instructions | 164](#)

[Product Compatibility | 171](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [EVPNs | 120](#)
- [General Routing | 121](#)
- [Interfaces and Chassis | 121](#)
- [Management | 121](#)
- [MPLS | 122](#)
- [Network Management and Monitoring | 122](#)
- [Routing Protocols | 123](#)
- [Security | 124](#)
- [Services Application | 124](#)
- [Subscriber Management and Services | 124](#)
- [VLAN Infrastructure | 125](#)

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 17.3R2 for MX Series routers.

EVPNs

- **commit check command successful with trunk port and EVPN-MPLS/EVPN-VXLAN EVI configured**—Starting in Junos OS Release 17.3R1, when adding a trunk port with dual tags to an EVPN and MPLS routing instance, or an EVPN and VXLAN routing instance, the CLI commit check configuration considers the **inner-vlan-id-list** statement and is successful.
- **Changes in the output of show route table command**—Starting in Junos OS Release 17.3R2, 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.

General Routing

- **Change in boot up behavior(MX10003)**—Starting in Junos OS Release 17.3R1, when the MPC is removed and plugged into the slot, the MPC is brought online automatically. In Junos OS 17.3R1 prior releases, the MPC could be brought online only after issuing the **request chassis fpc slot number online** command.
- **Commit preparation on MX-VC setup**—On MX Series virtual chassis setup, you see the following:
 - When you issue **commit prepare** on one Routing Engine followed by switchover, the Routing Engine where the switchover command is issued reboots. Therefore, the prepared cache gets cleared in that Routing Engine.
 - **clear system commit prepared** clears the plus files and prepared cache only in the device where the command is issued.

Interfaces and Chassis

- **show chassis environment cb command not supported on MX10003 backup Routing Engine**—In Junos OS Release 17.3R1, you cannot get the environmental information about the Control Boards (CBs) installed in an MX10003 because the router does not support the [show chassis environment cb](#) CLI command on a backup Routing Engine. No output is displayed if you execute this command on an MX10003 backup Routing Engine.

Management

- **Changes to custom YANG RPC syntax (MX Series)**—Starting in Junos OS Release 17.3, custom YANG RPCs have the following changes in syntax:
 - The **junos:action-execute** statement is a substatement to **junos:command**. In earlier releases, the **action-execute** and **command** statements are placed at the same level, and the **command** statement is optional.
 - The CLI formatting for a custom RPC is defined within the **junos-odl:format** statement, which takes an identifier as an argument. In earlier releases, the CLI formatting is defined using a container that includes the **junos-odl:cli-format** statement with no identifier.
 - The **junos-odl:style** statement defines the formatting for different styles within the statement. In earlier releases, the CLI formatting for different styles is defined using a container that includes the **junos-odl:cli-format** and **junos-odl:style** statements.
- **Enhancement to show agent sensors command (MX Series)** —Starting with Junos OS Release 17.3R1, the **show agent sensors** command, which displays information about Junos Telemetry Interface sensors, displays the default value of **0** for the **DSCP** and **Forwarding-class** values. Previously, the displayed default value for these fields was **255**. The default value is displayed when you do not configure a DSCP

or forwarding-class value for a sensor at the `[edit services analytics export-profile profile-name]` hierarchy level.

[See [export-profile](#) and [show agent sensors](#).]

MPLS

- Starting in Junos OS Release 17.3R1, 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](#).]

- Starting in Junos OS Release 17.3R2-S2, the * (asterisk) wildcard character is supported for the interface name of the **show ppp interfaces** command for debugging purpose. With this support, you can match any string of characters in that position in the interface name. For example, `so*` matches all SONET/SDH interfaces.

[See [show ppp interface](#).]

Network Management and Monitoring

- Enhancement to SNMPv3 traps for contextName field (MX Series)**—Starting in Junos OS Release 17.3R1, 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.
- Enhancement to about-to-expire logic for license expiry syslog messages (MX Series)**—Starting in Junos OS Release 17.3R1, the logic for multiple capacity type licenses and when their expiry raises alarms was changed. Previously, the behavior had alarms and syslog messages for expiring licenses raised based on the highest validity, which would mislead users in the case of a license expiring earlier than the highest validity license. The new behavior has the about-to-expire logic based on the first expiring license.
- SNMP syslog messages changed (MX Series)**—In Junos OS Release 17.3R1, 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 17.3R2, 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](#).]

- **Customer-visible SNMP trap name changes (MX Series)**—In Junos OS Release 17.3R2, on the Enhanced Switch Control Board (SCBE), name changes include the control board slot when jnxTimingFaultLOSset and jnxTimingFaultLOSClear traps are generated in the case of BITS interfaces (T1 or E1). SNMP traps for the backup Routing Engine clock failure event have been added, and the control board name is included in the SNMP trap interface name (jnxClksyncIntfName), for example, value: "external(cb-0)".

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

Routing Protocols

- **Change in output of show configuration routing-options flow operational command**—Starting in Junos OS Release 17.3R1, the sequence of statements in the output of **show configuration routing-options flow operational** command has changed to improve readability. The **then** statements are now displayed after the **match** conditions in a logical sequence.
- **BGP GR stale routes are not removed when BFD goes down**—Starting in Junos OS Release 17.3R1, 17.2R2, 17.1R3, 16.2R3, 16.1R5, and 15.1R7, when a BGP session that has BFD configured without the **hold-down-interval** fails, the BFD session remains active. The BFD session is not impacted even when graceful restart is enabled. BGP deletes the BFD session when user explicitly disables BFD on a BGP peer. Note that BFD session is created only when a BGP session is **Established**. In earlier Junos OS releases, BFD sessions are deleted when the BGP session fails and the **hold-down-interval** option is not configured.

Security

- **Support for SSH protocol version 2**—Starting in Junos OS Release 17.3R2, SSH protocol version 1 (SSHv1) is not supported. SSH protocol version 2 (SSHv2) is the default protocol-version option available under the `[edit system services ssh]` hierarchy level.

[See [protocol-version](#)]

Services Application

- **Changes to the show services rpm history-results command (MX Series)**—Starting in Junos OS Release 17.3R1, you must include the **owner** *owner* and **test** *name* options when using the **show services rpm history-results** command.

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

- In Junos OS Release 17.3R1 and later, for PIC-based J-Flow on MX Series routers and inline J-Flow on PTX Series routers, the Options template and Options data records include the **Sampling Interval** field as part of the **ScopeTemplate** field instead of the **ScopeSystem** field.

Subscriber Management and Services

- **Source-specific multicast (SSM) CLI changes for dynamic IGMP and dynamic MLD (MX Series)**—Starting in Junos OS Release 17.3R1, 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 does not appear in the CLI. 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.

Before you upgrade from an earlier release with a configuration that includes **ssm-map**, delete the **ssm-map** statement. If you do not, the upgrade fails. If you perform the upgrade without validation (**no-validate**), the upgrade passes and the **ssm-map** configuration is accepted, but it has no effect.

[See [ssm-map-policy \(Dynamic IGMP Interface\)](#) and [ssm-map-policy \(Dynamic MLD Interface\)](#).]

- **Memory mapping statement removed for Enhanced Subscriber Management (MX Series)**—Starting in Junos OS Release 17.3R1, 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 17.3R2, 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.

VLAN Infrastructure

- **LAG interface flaps while adding/removing a VLAN**—From Junos OS Release 17.3 or later, the LAG interface flaps while adding or removing a VLAN. The flapping happens when a low-speed SFP is plugged into a relatively high-speed port. To avoid flapping, configure the port speed to match the speed of the SFP.

SEE ALSO

New and Changed Features	 96
Known Behavior	 125
Known Issues	 131
Resolved Issues	 144
Documentation Updates	 163
Migration, Upgrade, and Downgrade Instructions	 164
Product Compatibility	 171

Known Behavior

IN THIS SECTION

- [Class of Service \(CoS\)](#) | [126](#)
- [EVPN](#) | [126](#)
- [General Routing](#) | [127](#)
- [High Availability \(HA\) and Resiliency](#) | [129](#)

- Interfaces and Chassis | 129
- MPLS | 129
- Routing Protocols | 130
- Subscriber Management and Services | 130

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 for MX Series routers.

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

Class of Service (CoS)

- In Junos OS Release 17.2, egress rate limit at the extended port does not work properly when you have a rate-limit configuration applied at the extended port physical interface (IFD) level by **traffic-control-profile-remaining** and also at some of the extended port logical interfaces (IFL) by **explicit traffic-control-profile** in hierarchical-scheduler mode. [PR1271719](#)

EVPN

- Routing instances of type EVPN configured with a VLAN ID will advertise MAC (type 2) routes with the VLAN value in the Ethernet tag field of the MAC route. As a workaround, use **vlan-id-none** to claim the RFC compliance. [PR945247](#)
- A provider edge (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 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](#)
- Configurable udp-port for VXLAN in EVPN-VXLAN scenario is currently not supported. [PR1249310](#)
- On an MX Series router running Junos OS, while migrating a routing instance from VPLS to EVPN, if an EVPN command (for example, **control-word**) is handled in a catastrophic manner by daemons (processes), traffic loss can occur while the control plane state is cleaned up and reconstructed. [PR1268428](#)
- In a hub-and-spoke VPLS environment running Junos OS, if local switching is enabled on the hub-and-spoke PE devices migrated to EVPN (for which the hub remains VPLS-only), the following issue could occur: (1) Two copies of BUM traffic could be received at the spoke PE device (one copy through the EVPN next hop from the ingress spoke and the other copy through the VPLS pseudowire from the hub) and (2) MACs behind a spoke PE device would use the VPLS pseudowire to the hub as the next

hop on the remote spoke PE devices (instead of the EVPN next hop). This issue occurs because the VPLS-only hub continues to provide an alternative forwarding path between the spoke PE devices (migrated to EVPN). [PR1272449](#)

- An IPv6 underlay with an IPv6 overlay with IRB is not supported in a bridge domain, because having two IPv6 headers exceeds the 128-byte parcel size for the line card. [PR1274709](#)
- When changing encapsulation from VXLAN to MPLS or vice versa, you need to deactivate and reactivate the instance. [PR1326430](#)

General Routing

- On MX Series routers, parity memory errors occur in the pre-classifier engines within an MPC. Packets silently discarded earlier are reported in syslogs and alarms when parity memory errors occur.
- On an MX10003 router, when the management interface (fxp0 or em0) is down on the master Routing Engine, in addition to the **Ethernet Link Down** alarm, an additional **Management Ethernet Link Down** alarm is also raised.
- A provider edge (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 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 with MS-MPC or MS-MIC, memory leaks will be seen with `jnx_msp_jbuf_small_oc` object, upon sending millions of point-to-point tunneling protocol control connections (3-5M) alone at higher cells per second (cps) (greater than 150K cps). This issue is not seen with up to 50,000 control connections at 10,000-30,000 cps. [PR1087561](#)
- NAT64: Source-prefix filtering and protocol filtering of the CGNAT sessions are incorrect. For example, **show services sessions extensive protocol udp source-prefix <0:7000::2>** displays incorrect filtering of the sessions. [PR1179922](#)
- Chef for Junos OS supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev-resources. The netdev-resource developed for interface configuration has a limitation to configure the XE interface. Netdev-interface resource assumes that speed is a configurable parameter that is supported on a GE interface but not on an XE interface. Hence, netdev-interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages `chef-11.10.4_1.1.*.tgz` `chef-11.10.4_2.0.*.tgz` in all platforms {i386/x86-32/powerpc}. [PR1181475](#)
- As described in RFC 7130, 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](#)
- In certain interface scaling scenarios, during configuration commit/rollback, you might see an `fpcx` error message. You can safely ignore this message because of the FPGA monitor mechanism on DPC cards

for logical interface mapping (ifl_map). Between the deletion of a physical interface and the monitoring event, this 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 the ifl_map, harmless messages might populate the log file. [PR1210877](#)

- The ptp master streams on IP and Ethernet are not supported simultaneously. [PR1217427](#)
- There is no unified ISSU from Junos OS Release 15.1 and earlier releases to Junos OS Release 16.2R1. [PR1222540](#)
- The following MICs in MPC2E-NG/MPC3E-NG are non PHY-Timestamping capable: MIC-3D-4XGE-XFP MIC3-3D-10XGE-SFPP, MIC-3D-2XGE-XFP, MIC-3D-20GE-SFP. The 2Way/T1/T4 time error can be up to +/-450 nsec in these MICs. [PR1243646](#)
- 1PPS TE/cTE performance metric can be as high as +/-550 nsec in MPC2E/3E NG QoS/3D 20x 1GE(LAN)-E,SFP with no PHY-timestamp and non-hybrid mode. [PR1263235](#)
- This issue occurs when an interface comes online and both the OAM protocol and MKA protocol try to establish their respective sessions. Because of contention between these two protocols OAM takes down the interface and MKA fails to establishes a connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
- PCC controlled LSP metric not getting updated on the controller, PCE-delegated LSPs do not come up. [PR1265864](#)
- On an MX Series router running Junos OS, while migrating a routing instance from VPLS to EVPN, if an EVPN command (for example, control-word) is handled in a catastrophic manner by daemons (processes), traffic loss can occur while the control plane state is cleaned up and reconstructed. [PR1268428](#)
- In a hub-and-spoke VPLS environment running Junos OS, if local switching is enabled on the hub-and-spoke PE devices migrated to EVPN (for which the hub remains VPLS-only), the following issue could occur: (1) Two copies of BUM traffic could be received at the spoke PE device (one copy through the EVPN next hop from the ingress spoke and the other copy through the VPLS pseudowire from the hub) and (2) MACs behind a spoke PE device would use the VPLS pseudowire to the hub as the next hop on the remote spoke PE devices (instead of the EVPN next hop). This issue occurs because the VPLS-only hub continues to provide an alternative forwarding path between the spoke PE devices (migrated to EVPN). [PR1272449](#)
- The device might not power up when crossover cables are used. We recommend using straight cables. [PR1274613](#)
- An IPv6 underlay with an IPv6 overlay with IRB is not supported in a bridge domain, because having two IPv6 headers exceeds the 128-byte parcel size for the line card. [PR1274709](#)
- On QFX10000 line switches implementing EVPN-VXLAN, if the Routing Engine is repeatedly restarted on redundant gateways, then inter-VRF traffic will be dropped without notification. [PR1289091](#)
- On MX150 routers, if you connect an even-numbered port to another even-numbered port using external loopback, they cannot communicate with each other. On MX150 routers, **ge-0/0/0,2,4,6,8,10** and **xe-0/0/12** are identified as even-numbered ports. Also, if you connect an odd-numbered port to another

odd-numbered port using external loopback, they cannot communicate with each other. On MX150 routers, **ge-0/0/1,3,5,7,9,11** and **xe-0/0/13** are identified as odd-numbered ports.

For instance, if you connect port (**ge-0/0/0**) to port (**ge-0/0/6**) using external loopback, the two ports cannot communicate with each other. Also, if you connect port (**ge-0/0/3**) to port (**ge-0/0/9**) using external loopback, the two ports cannot communicate with each other. To configure external loopback, connect an even-numbered port (for instance, **xe-0/0/12**) to an odd-numbered port (for instance, **xe-0/0/13**).

High Availability (HA) and Resiliency

- **MPC7E MPC8E and MPC9E line card restrictions for MX Series Virtual Chassis unified ISSU (MX Series)**—MPC7E, MPC8E, and MPC9E line cards do not support unified ISSU in Junos OS Release 17.3R1 for MX Series Virtual Chassis configurations. These line cards must be removed or configured to power off during the MX-VC ISSU process. ISSU in Junos OS Release 17.3R1 is supported for MX Series standalone chassis configurations.

[See [Preparing for a Unified ISSU in an MX Series Virtual Chassis](#).]

Interfaces and Chassis

- Convergence time for VRRP traffic is higher when the router or Routing Engine is rebooted in a single Routing Engine system. We recommend having a dual Routing Engine system with redundancy enabled. In this case, if the master Routing Engine is rebooted, the backup Routing Engine will take over mastership. There will not be any disruption in VRRP traffic. [PR1270168](#)
- When an FPC with both core link and member link of an aggregated Ethernet interface (running VRRP) is restarted or offlined, the convergence time will be higher. [PR1270811](#)
- Higher MTU configuration on an IRB than on the member link of its VLAN might bring down a VRRP session configured on the IRB. As a workaround, always have the MTU configured on the IRB of the VLAN be less than or equal to the MTU configured on its member links of the same VLAN because MX Series devices do not throw error or warning messages during configuration commit. [PR1295763](#)

MPLS

- When NG-MVPN is configured with RSVP provider tunnels and NSR is used, then the egress router for the tunnel might not correctly replicate some of the tunnel state to the backup routing engine, leading to temporary traffic loss during NSR failover for the affected tunnels. [PR1293014](#)
- In Junos OS Release 17.1R1 or earlier releases, labels from within the following ranges can be used as incoming labels for static VPLS LSI-based services by default: R1. [29696 - 41983]; R2. [1000000 - 1048575]. In Junos OS Release 17.1R1 and later releases on a system operating in enhanced-IP mode, range R1 cannot be used any longer for static VPLS LSI-based services incoming label assignment by

default. This limitation is applicable only for range R1 and is not applicable for range R2. The latter works on Junos OS Release 17.1R1 and later releases just as it does on previous Junos OS releases. [PR1307402](#)

Routing Protocols

- When a Junos OS aggregation gateway uses a IPv6 address as next hop for IPv4 aggregates announced to downstream, 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 BGP inet6-labeled-unicast family. [PR1220235](#)
- PIM is not supported on a tunnel interface configured with an inet6 address. Configuring PIM over a tunnel interface with an inet6 address might cause the routing protocol process (rpd) to crash and generate a core file. [PR1267570](#)

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.

SEE ALSO

New and Changed Features 96
Changes in Behavior and Syntax 120
Known Issues 131
Resolved Issues 144
Documentation Updates 163
Migration, Upgrade, and Downgrade Instructions 164
Product Compatibility 171

Known Issues

IN THIS SECTION

- [Application Layer Gateways \(ALGs\) | 131](#)
- [Class of Service \(CoS\) | 131](#)
- [EVPN | 132](#)
- [Forwarding and Sampling | 133](#)
- [General Routing | 133](#)
- [Infrastructure | 137](#)
- [Interfaces and Chassis | 138](#)
- [Layer 2 Ethernet Services | 138](#)
- [Layer 2 Features | 139](#)
- [MPLS | 139](#)
- [Platform and Infrastructure | 140](#)
- [Routing Protocols | 141](#)
- [Services Applications | 143](#)
- [Subscriber Access Management | 143](#)
- [VPNs | 144](#)

This section lists the known issues in hardware and software in Junos OS Release 17.3R2 for MX Series routers.

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

Application Layer Gateways (ALGs)

- An IPsec VPN connection cannot be established successfully because the Internet Key Exchange (IKE) ALG drops the first response message during the IPsec IKEv2 negotiation. [PR1300448](#)

Class of Service (CoS)

- A CoS scheduler update can fail when all of the following conditions are met:

- Dynamic subscribers exist on an aggregated Ethernet bundle.
- The CoS traffic-control-profile and/or scheduler-map applied to these dynamic subscribers is from a static configuration.
- The relevant static CoS is modified in the same configuration commit as a modification to the aggregated Ethernet bundle (either a leg add or leg remove) containing the subscribers.
- The leg add or leg remove in the commit is the first or last leg to be added or removed from a line card.

To avoid this issue, do not commit a bundle change in the same commit as a static CoS change. In this event, one of the following logs is displayed in the "message" system log:

- **subscriber cos update not applied to interface <interface-name> status <id>**
- **subscriber cos update not applied to interface-set <interface-set-name> status <id>**

This indicates that the last update to the subscriber or interface set was not applied.

[PR1276459](#)

EVPN

- 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](#)
- In an EVPN scenario with static MAC configured in the EVPN instance, MAC route information is visible to the remote EVPN instance. However, after deactivating and reactivating the static MAC in the EVPN instance, there is no such MAC route in the EVPN routing table. [PR1193754](#)
- On MX Series routers with EVPN, the routing protocol process might crash when MAC moves between multihomed PE routers, resulting in traffic loss. [PR1216144](#)
- In an EVPN and VXLAN scenario with a dual-homing Layer 3 gateway configuration, if one Layer 3 gateway receives the changed ARP entry from another Layer 3 gateway, it might delete the related ARP entry (it should update this ARP entry) and host route. In this error state, traffic might be dropped. [PR1306024](#)
- The issue is applicable to mac-in-mac PNN-EVPN and does not affect any other scenario. When PBB EVPN configuration is reloaded on MX Series routers, error logs are seen while deleting interfaces related to backbone bridge component. These errors does not result in any functional issues. [PR1323275](#)
- PBB EVPN will not be able to flood traffic towards core. Traffic recovers by performing **restart l2-learning**. In addition to this, there is a limitation in PBB EVPN A/A Unicast traffic forwarding. If entropy in the traffic not sufficient, then uneven load balancing causes a problem on MH peer A/A routers. This will cause a drop for return traffic. These issues are applicable to mac-in-mac PNN-EVPN and does not affect any other scenario. [PR1323503](#)

- On an EVPN VXLAN enabled MX Series routers, if the underlying interface for the VXLAN tunnel is a LACP enabled aggregated Ethernet interface with multiple members, and one of the member is flapped. There might be a momentary IPv4 or IPv6 inter-VNI traffic loss. [PR1326572](#)
- In EVPN-Etree, traffic loss is seen on deactivating CE facing interface both with NSR enabled and normal scenario. CE interface which is leaf interface is deleted completely and added back to restore same old state logical interface being part same EVPN. The leaf-to-leaf traffic might not get blocked. [PR1330134](#)
- On deactivating ESI for a physical interface where its logical interface is used for the EVPN VPWS and then deactivating that EVPN VPWS routing-instance, the rpd might crash and generate a core file. [PR1332652](#)
- In EVPN with VXLAN Data Plane Encapsulation configuration, RPD core got generated in provider edge (PE) router when 'restart routing' is applied. This provider edge (PE) router has EVPN-VXLAN configuration and VXLAN tunnel endpoint (VTEP) interface created for VXLAN acts as access link. The loopback interface that is configured as vtep-source-interface is also configured with ESI value and ESI mode as single-active. [PR1333331](#)
- In some scenario the same MAC message is added twice (without deleting in between) from L2-learning module to routing protocol process (rpd). The backup of rpd might generate a core file. [PR1336881](#)

Forwarding and Sampling

- When a policing filter is applied to an active LSP carrying traffic, the LSP resignals and drops traffic for approximately 2 seconds. It might take up to 30 seconds for the LSP to come up under the following conditions:
 - Creation of the policing filter and application of the same 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](#)
- The sampled process crashes aggressively and generates a core file when connecting to L2BSA and EVPN subscribers. [PR1293237](#)

General Routing

- On chassis-based line cards, the **FI: Protect: Parity error for CP freepool SRAM** SRAM parity error might be seen. It is harmless and can be ignored. [PR1079726](#)
- Deleting the whole maintenance-domain section under the [edit protocols oam ethernet connectivity-fault-management] hierarchy level and committing the changes, MX Series BNG still persists to generate CCM frames for the already deleted **maintenance-domain domain-name maintenance-association** statement. [PR1107542](#)

- In scaled up EVPN VPWS configurations (approximately 8000 EVPN VPWS), during Routing Engine switchover, rpd scheduler slip messages might be seen. [PR1225153](#)
 - An incorrect PE router is attached to an ESI when the router receives two copies of the same AD or ESI route (for example, one through eBGP and another one received from an iBGP neighbor). This causes a partial traffic black hole and stale MAC entries. You can confirm the issue by checking the members of the ESI: `user@host> show evpn instance extensive ...` **Number of ethernet segments: 5 ESI: 00:13:78:00:00:00:00:00:01 Status: Resolved Number of remote PEs connected: 3 Remote PE MAC label Aliasing label Mode 87.233.39.102 0 0 all-active 87.233.39.1 200 0 all-active <<<< this PE is not part of the ESI 87.233.39.101 200 0 all-active.** [PR1231402](#)
 - When 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](#)
 - In a BGP and MPLS scenario, if the next-hop type of label route is indirect, then the following changing events about "family mpls" of the next-hop interface might cause the route to be in the dead state, and the route will remain dead even though the family mpls is activated again:
 - Deactivating and activating the interface family mpls.
 - Deleting and adding back the interface family mpls.
 - Changing maximum labels for the interface.
- [PR1242589](#)
- On MX Series routers with XM chipset (for example, MPC3E, MPC4E, MPC5E, MPC6E, MPC2E-NG, and MPC3E-NG), the MPC might reboot after a unified ISSU completion. [PR1256145](#)
 - When nonstop active routing (NSR) is configured at high scale in a node virtualization setup, the guest network function (GNF) might restart MPC9 line cards rarely during a Routing Engine switchover. [PR1259910](#)
 - On MX Series platforms, specific MPC card traffic drops when late_cell counter reaches 65,000 or a very high value at each polling interval after the link sanity is checked. [PR1262868](#)
 - This issue occurs an interface comes online and both the OAM protocol and the MKA protocol try to establish their respective sessions. Because of contention between these two protocols, OAM takes down the interface and MKA fails to establish a connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
 - When a unified ISSU is performed on an MX Series Virtual Chassis system in a scaled subscriber management scenario, the BGP protocol sessions are active. When the sessions are clients of BFD, then the BGP sessions might go down and come up again, which might cause traffic loss. [PR1265407](#)
 - The issue occurs when the Packet Forwarding Engine is oversubscribed with unknown unicast flood with no MAC learning, which is not a common configuration. During unified ISSU, only the Packet Forwarding Engine gets wedged. However, this issue is not seen when the Packet Forwarding Engine is oversubscribed with L3 traffic or with L2 traffic with MAC learning. [PR1265898](#)

- Currently, BBE advanced services is not supported on the node virtualization platform. Hence, mobility is disabled on the node virtualization platform BSYS and GNF Routing Engines. For legacy purposes, BBE functionality needs to work properly on the node virtualization platform. Reboot is required when the BSYS Routing Engine is changed to standalone Routing Engine mode (normal) and vice versa. [PR1266615](#)
- DEP does not support dh group group19, encryption algorithm aes-256-cbc, and hash sha-384 in its list of default proposals. These must be configured explicitly in the configuration. [PR1269160](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- The load based throttling feature should be enabled by default. Due to its high risk, the feature was disabled by default when it was released. At the time of the release, a workaround was provided for the customers to enable the feature by configuration, if required. Now, the feature is enabled by default. [PR1271739](#)
- In an L2BSA scaling senario, after bringing up about 12,000 subscribers, one or more FPCs will reboot. [PR1273353](#)
- Incorrect counters for output packets on child links ae0 interface when configured with new feature 'revertive'. [PR1273983](#)
- Interfaces might flap on the 20x1GE SFP MIC while performing a unified ISSU from Junos OS Release 17.3R1. [PR1276816](#)
- If a vmhost snapshot is taken on an alternate disk and there is no further vmhost software image upgrade, the expectation is that if the current vmhost image gets corrupted, the system will boot with an alternate disk so that the user can recover the primary disk to restore the state. However, under conditions where corruption is with the host root file system, the node boots with the previous vmhost software instead of booting from an alternate disk space. [PR1281554](#)
- When using a channelized configuration on MX Series Virtual Chassis MPC7/8/9 MRATE PIC QSFP interfaces for VCP connections between members, a VCP interface needs to be configured on channel 0 of each QSFP to activate the port. [PR1283283](#)
- Due to vendor code limitation, ungraceful removing of summit MACsec TIC from chassis might cause a crash or unpredictable result. [PR1284040](#)
- TVP platforms do not support the **chassisd hard restart** command, because of an infrastructure limitation. FPC power off does not happen properly, because the old chassisd process initiates the **FPC power off** command and exits. The CLI command **restart chassisd hard** with GRES on MX10003 causes a new chassisd process to open a reconnect window and wait for a connection. The Routing Engine and FPC get out of synchronization when the FPC restarts multiple times to connect and synchronize. Finally, the FPC comes online. [PR1293314](#)
- During PPPoE subscriber login errors like [**vbf_flow_src_lookup_enabled**] and [**failed to find iff structure, ifl**] were seen on FPC. [PR1294710](#)
- The Routing Engine gets stuck and boots from the other SSD after vmhost reboot. [PR1295219](#)

- If a JET user uses TACACS to authenticate the remote user, a bug in the code results in a mismatch of local and remote users, resulting in authentication failure. [PR1296237](#)
- PTP slave is taking longer time (more than 1 hour) to lock master in T-BC scenario test. [PR1298792](#)
- This issue is applicable when using MPLS Label Switched Paths (LSP) and RSVP - Traffic Engineering (RSVP-TE) self-ping. When rpd sends out a self-ping packet and a RSVP packet at the same time these packets might overwrite kernel's packet buffers, causing memory corruption. As a result, the kernel panics. [PR1303798](#)
- MX2000 platforms with MPC9E or MPC8E (bandwidth 1.6T mode) and SBF2 fabric with certain high amount of traffic volume might cause transient traffic drops with cell underflow messages at the fabric input block. The fabric parameters are tuned to avoid such traffic overflow conditions and traffic drops. If the traffic drop conditions occur at sustained rate, it could lead to permanent impact of traffic forwarding. [PR1304801](#)
- Inline J-Flow VMX, OIF field of VPLS data records sometimes reports the SNMP index value of the LSI interface instead of egress physical interface. [PR1305411](#)
- When syslog errors such as **pfeman_inline_ka_steering_gencfg_handler: nh not found for nh=<pfh-nhid>** are seen on the FPC after it reboots, it is likely that the steering rules used for BFD packet redirection are not installed correctly. This is because of the unexpected replay order of IPC messages from the kernel when the FPC reboots. [PR1308884](#)
- In streaming telemetry, when a user logs in and logs out quickly from TACACS, the following messages are displayed: **Oct 2 06:26:20 jsd[11821]: early: bad stored heap: heap-ptr=0x0 data-ptr=0x1481cbf8 Oct 2 06:26:29 jsd[11821]: early: bad stored heap: heap-ptr=0x0 data-ptr=0x1481ccb8 Oct 2 06:27:12 jsd[11821]: early: bad stored heap: heap-ptr=0x0 data-ptr=0x1481cd78 Oct 2 06:27:23 jsd[11821]: early: bad stored heap: heap-ptr=0x0 data-ptr=0x1481ce38 Oct 2 06:27:58 jsd[11821]: early: bad stored heap: heap-ptr=0x0 data-ptr=0x1481cef8 Oct 2 06:28:00 jsd[11821]: early: bad stored heap: heap-ptr=0x0 data-ptr=0x1481cfb8 .** [PR1311482](#)
- VLAN-CCC logical interface for Layer 2 circuit remains in CCC-Down state upon Layer 2 circuit to EVPN-VPWS service change unless it is deactivated and reactivated manually. [PR1312043](#)
- MPC type 5/6/7/8/9 supports FPC specific sensors for intake or exhaust temperature and software has corresponding threshold definitions. During SNMP periodic temperature check, the software might send false trap over temperature SNMP trap because of a software bug that it checks default threshold instead corresponding threshold. **Oct 11 15:25:07 CHASSISD_SNMP_TRAP6: SNMP trap generated: Over Temperature! (jnxContentsContainerIndex 7, jnxContentsL1Index 9, jnxContentsL2Index 0, jnxContentsL3Index 0, jnxContentsDescr FPC: MPC9E 3D @ 8/*/*, jnxOperatingState/Temp 76).** [PR1313391](#)
- While performing multiple switchovers with thousands of subscribers, backup on the upcoming master Routing Engine might generate core files repeatedly, if it is unable to set up distributed multicast for a few subscribers. [PR1314651](#)

- On MX Series platforms, the router might run into a KRT stuck issue. The KRT asynchronous queue is stuck because of the kernel KRT_STATE_BLOCKED state (KRT_BLOCK_REASON_MEM_BLOCK) with an unknown trigger. [PR1315212](#)
- The identical logs are generated and the severity of the logs are different between the two releases. The precise severity is observed in later release. The reason to find dissimilar severity in the earlier release is not identified. [PR1318884](#)
- When BGP is brought over many subscriber L2TP logical interfaces, it does not establish on a few of the subscriber logical interfaces. [PR1322001](#)
- Commit operation gets stuck when commit check is performed when the **fast-synchronize** option is enabled. [PR1322431](#)
- On MX10003 platform, MACsec might fail to establish the session because MACsec is wedged during the key rollover. This is a timing issue and the interface associated with the wedged MACsec might also go to a down state. [PR1325331](#)
- Any-Any in traffic selector for both v4 and v6 cannot be configured. [PR1334966](#)

Infrastructure

- The configuration statement **set system ports console log-out-on-disconnect** logs the user out from the console and closes the console connection. If the configuration statement **set system syslog console any warning** is used along with the earlier configuration and if there is no active telnet connection to the console, the processes (daemons) try to open the console and hang as they wait for a "serial connect" that is received only by doing a telnet to the console. [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 "any any", then you would need to skip these logs with an appropriate match condition. [PR1239651](#)
- When the configuration statement **set system log-out-on-disconnect** is enabled, the Junos OS eventd processes (daemon) blocks the console-open(). But during this stage with syslog console configured (always logs on console), logging continues even if the console session ends. While console logging is in the wait state by eventd, syslog rotation freezes and some processes directly attached to logging in the system would also get into the wait state, causing an undesirable behavior. [PR1253544](#)

Interfaces and Chassis

- During configuration changes and reuse of virtual IP on an interface as an interface address, it is required to delete the configuration and do a commit and then add the interface address configuration in the following commit. [PR1191371](#)
- MTU on the BNG and CPE sides has different values. In a rare situation, MX Series routers might calculate the MTU value for the corresponding pp0 logical interface incorrectly. [PR1240257](#)
- Rate-limit dropped packets are not displayed by `[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](#)
- In a VPLS multihoming scenario, the CFM packets are forwarded over the standby PE link, resulting in duplicate packets or a loop between the active and standby link. [PR1253542](#)
- Junos OS upgrade involving Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later main releases with CFM configuration can cause cfmd to generate a core file after the upgrade. This is because of the old version of `/var/db/cfm.db`. [PR1281073](#)
- 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](#)

Layer 2 Ethernet Services

- After changing the underlying physical interface (IFD) for a static VLAN demux interface, the NAS-Port-ID formed is based on the previous physical interface. [PR1255377](#)
- Whenever an MC-aggregated Ethernet interface is deactivated or activated on an MC-LAG node, once the MC-aggregated Ethernet interfaces are up again, the system clears neighbor discovery entries on the ICL which triggers a neighbor discovery solicit and thereby neighbor discovery entries are learned on the MC-aggregated Ethernet interface. As a workaround, clear neighbor discovery entries on the ICL whenever MC-aggregated Ethernet interfaces have been deactivated or activated on MC-LAG nodes. [PR1294958](#)
- On RE-2000 running FreeBSD 10.x-based Junos OS release, the following false positive CB alarms might be seen: **Aug 19 16:56:18.119 acb_pmbus_read: unable to convert voltage for CB 2 pmbus device XF ASIC A** **Aug 19 16:56:19.087 send: yellow alarm set, device CB 2, reason CB 2 PMBus Device Fail** **Aug 19 16:57:18.663 send: yellow alarm clear, device CB 2, reason CB 2 PMBus Device Fail** **Aug 21 22:45:04.145 acb_pmbus_read: pmbus command READ_VOUT_CMD to CB 2 XF ASIC B failed** **Aug 21 22:45:04.219 send: yellow alarm set, device CB 0, reason CB 2 PMBus Device Fail** **Aug 21 22:46:04.147 send: yellow alarm clear, device CB 0, reason CB 2 PMBus Device Fail.** [PR1298612](#)
- After rebooting the router or after smg-service is restarted, DHCPv6 packets get dropped when a **no-snoop** configuration is used. The issue is observed in a setting where subscribers connect over a static VLAN demux interface. [PR1316274](#)

Layer 2 Features

- In EVPN-VLAN, when VXLANs are interconnected via PIM over IRB where L3 multicast has only IRB in multicast output interface list and IRB traffic needs to be forwarded over VTEP, then such packets get dropped. [PR1318706](#)

MPLS

- When using **mpls traffic-engineering bgp-igp-both-ribs** with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit ABR when there are two or more such area border routers (ABRs). This causes interdomain RSVP LSPs to break. RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. [PR1048560](#)
- This issue occurs when there is a GRES between the master and backup Routing Engines of different memory capabilities. For example, one Routing Engine has only enough memory to run routing protocol process (rpd) in 32-bit mode while the other is capable of 64-bit mode. The situation 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 might crash on the new master Routing Engine. As a workaround, this issue can be avoided by using the CLI command **set system processes routing force-32-bit**. [PR1141728](#)
- Because of the current way of calculating bandwidth, you see a minimal discrepancy between MPLS statistics and adjusted bandwidth reported. The algorithm will be enhanced so that both values match 100 percent. [PR1259500](#)
- In an L2 circuit scenario, while processing an advertisement of LDP signaled L2 circuit, it gets stale binded because of the corrupted LDP structure. As a result, the rpd crashes. The reason for this corruption is not found and this issue is not reproduced. [PR1275766](#)
- With nonstop active routing (NSR), when a routing protocol process (rpd) restarts on the master Routing Engine, rpd might also restart on the backup Routing Engine. [PR1282369](#)
- The **show mpls container-lsp** output will not show any egress LSP until the enhanced FRR is enabled for these egress LSPs. [PR1314960](#)
- In an LDP over RSVP setup, when the RSVP label-switched paths (LSPs) have protection and a route can be reached through both LDP direct neighbor (IP next hop) and LDP remote neighbor over RSVP LSPs (RSVP next hop), the LDP route next hop is transitioned between the IP next hop and the RSVP LSP next hop. Then RSVP LSP make-before-break (MBB) can happen, and the LDP route might use stale RSVP LSP next hop because of a timing issue. This might cause the rpd process to crash. [PR1318480](#)
- The routing protocol process (rpd) generates a core file in jemalloc_block_mallocx because of a memory leak. [PR1321952](#)

Platform and Infrastructure

- When using the **show | compare** method to commit, part of the configuration might be treated as noise and return a syntax error. [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](#)
- 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](#)
- Scale used: 120 bridge domains among 1000 bridge domains, with XE/GE links toward downstream switch and LAG bundles as uplinks toward upstream routers. The XE/GE link is part of the physical loop in the topology. Spanning-tree protocols such as VSTP, RSTP, and MSTP are used for loop avoidance. Some MAC addresses are not learned on DUT when LAG bundles part of such bridge domains are flapped, along with other events such as spanning-tree root bridge change. [PR1275544](#)
- With unified ISSU, momentary traffic loss is expected. In an EVPN E-Tree, the known unicast frames can be flooded for around 30 seconds during unified ISSU before all the forwarding states are restored. This issue does not affect BUM traffic. As a workaround, nonstop bridging (NSB) can be configured at **set protocols layer2-control nonstop-bridging**. This reduces traffic flood to around 10 seconds in a moderate setup. [PR1275621](#)
- Due to a transient hardware error condition, **CPQ Sram parity error** and **CPQ RLDRAM double bit ECC error** syslog errors on MQCHIP need to raise a major CM alarm. [PR1276132](#)
- MX-MPC1-3D, MX-MPC2-3D, and MPC-3D-16XGE do not raise major CMERROR alarm upon a high rate of cell underflow events and link sanity interrupts. This might have a permanent impact on packet forwarding because of the transient hardware failure. [PR1276144](#)
- In the MoFRR scenario, if the core-facing links flap, multicast traffic is forwarded from both links. This might lead to doubling of multicast traffic. [PR1318129](#)
- If RPM delegate probes are configured, every time after any of the following configurations have been added, deleted, or changed, and committed–
 - **protocols bgp group <group> ...**
 - **protocols bgp local-address ...**

The user must restart RMOPs by issuing the **restart remote-operations** CLI command. Otherwise, RMOPd clears its configuration and all configured RPM delegate probes stop working. [PR1322097](#)

- When the label-switching router (LSR) works on MX Series with MPCs/MICs platforms or vMX and LSR carries MPLS pseudowire (such as I2circuit(LDP based)/I2vpn(BGP based)/VPLS) traffic, the packet might

get dropped if the MPLS pseudowire payload does not have control word and its destination MAC starts with '4' or '6'. [PR1327724](#)

- Traffic statistics might not match on pseudowire-service (PS) after clearing interface statistics. [PR1328252](#)

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 when the import policy has a non-default setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- Continuous soft core files might be generated because of bgp-path-selection code. The routing protocol process (rpd) forks a child and the child asserts to produce a core file. The problem is with route ordering and it is auto-corrected after collecting the **soft-assert-core** file, without any impact to the traffic or service. [PR815146](#)
- The rpd might crash because the IS-IS database tree gets corrupted with Shared Risk Link Group (SRLG) enabled under corner conditions, after executing the CLI command **clear isis database**. [PR1152940](#)
- VRF routes present in the access routers leaked into inet.0 and are not getting advertised into global inet.0 table on the core. [PR1200883](#)
- In a rare case, when LDP is deactivated, the result of a remote loop free alternate (RLFA) might be computed to go through the deactivated LDP node. The situation is self-recovered in the next SPF calculation. [PR1202392](#)
- Certain BGP traceoption flags (for example, "open", "update", "keepalive") might result in logging trace of debugging messages that do not fall within the specified traceoption category. This results in some unwanted BGP debug messages logged to the BGP traceoption file. [PR1252294](#)
- LDP and OSPF are in "sync" state and the reason observed for this is "IGP interface down" with **ldp-synchronization** enabled for OSPF. **user@host> show ospf interface ae100.0 extensive** Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. As per the current analysis, "IGP interface down" is observed as the reason because although LDP notified OSPF that LDP synchronization was achieved, OSPF was not able to take note of the LDP synchronization notification, because the OSPF neighbor was not up yet. The issue is under investigation. [PR1256434](#)
- When switchover and zeroize are done in succession quickly, "zeroize" deletes the databases. If dfwd is going to start SIHUP processing after the zeroize, it might generate a core file because the database is not present. Zeroize should be done when the system is in a stable state; that is, signup processing by daemons is over. [PR1262385](#)

- Performance degradation in computing LFA and remote LFAs is observed. [PR1264564](#)
- The BMP session sends a peer-down event and route withdrawals when peer monitoring is disabled through a configuration event. After the 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 commit is followed by an rpd process crash. [PR1278582](#)
- Currently, two multicast tunnel (mt) interfaces are seen for each of the PIM neighbors after the VPN-Tunnel-Source activation or deactivation. However, ideally, you should have the same tunnel source for both IPv4 and IPv6 address families, if both are using the same PIM tunnel. [PR1281481](#)
- The mcsnoopd process generating a core file will be seen. [PR1285727](#)
- When BGP monitoring protocol (BMP) sends out route monitoring messages for BGP routes that have unusable or unresolved next hops, the route monitoring messages might contain a BGP update with an MP_REACH_NLRI path attribute that specifies an incorrect path attribute length. This might occur for any of the address families, except IPv4 and flowspec (for example, IPv6 can be impacted). This issue could result in unexpected behavior or failures in BMP station applications. [PR1292848](#)
- If a router works as the GR helper during a peering establishment, the newly established peer might lose some of the negotiated capabilities and interpret the updates incorrectly. This can cause peer drops or invalid routes. [PR1293174](#)
- Multicast flow interruption might be observed on a transit router in a Protocol Independent Multicast (PIM) scenario. If (*,G) join is received on one interface, and (*,G) join and (S,G,RPT) prune are received on another interface, which then receives (*,G) prune. Multicast flows on the first interface get reset and interrupted for a short time (for example, 1 second). [PR1293900](#)
- The routing protocol process (rpd) might restart unexpectedly when configuring rib-groups and routing instances with static routes in a certain order. [PR1298262](#)
- When the mcsnoopd process tries to terminate gracefully, it tries to clean up all the resources it has used. For this cleanup to happen, the task infrastructure waits for 10 minutes. In these 10 minutes, if the KRT task does not clean up properly, it generates a core file. [PR1305239](#)
- The trigger to the problem is the configuration of a policy with an **llgr-stale community** and a long-lived graceful restart (LLGR) **LLGR** configuration. With this configuration, initialization does not happen correctly and a core file might be generated. Workaround would be to avoid configurations with llgr-stale community. [PR1310751](#)
- An MX104 is connected to SRX1500. IS-IS is running between these device and BFD has been configured between the IS-IS peers. Unfortunately, BFD is not coming up between these devices successfully. [PR1312298](#)
- In some scenarios, the routing protocol process (rpd) might generate a core file on the router while importing IS-IS routes because of a configuration change or a network event. [PR1312325](#)
- Routing protocol process (rpd) might crash and generate core files in distributed IGMP environment. [PR1314679](#)

- On a chassis with BMP configured, the rpd might crash when the rpd process is gracefully terminated. [PR1315798](#)
- Layer 2 VPN, nonstop active routing (NSR): The rpd process generates core files in the backup Routing Engine because of a route distinguisher clash between the new RT instance updated by the master Routing Engine and the deleted RT instance in the backup Routing Engine. [PR1319587](#)

Services Applications

- Session counters for cleartext traffic are not updated after decryption. The decrypted packet count can be obtained by running the following command: **show security group-vpn member ipsec statistics**. [PR1068094](#)
- We recommend that you do not configure **ms- interface** when an AMS bundle in one-to-one mode has the same member interface. [PR1209660](#)
- When configuring a NAT pool that is shared between PCP and standard NAT the PCP mappings cannot be manually cleared. Some operators will take their NAT pools and have the lower port ranges set for their PCP pool and then take the same NAT'd IPs and use them in another pool with the higher ranges.

```

services { nat { pool NAT_NON_PCP { address-range low 100.100.100.0 high 100.100.100.255; port {
range low 2048 high 65535; secured-port-block-allocation block-size 512 max-blocks-per-address 10
active-block-timeout 0; } address-allocation round-robin; mapping-timeout 120; ei-mapping-timeout
120; } pool PCP_Pool { address-range low 100.100.100.0 high 100.100.100.255; port { range low 1025
high 2047 random-allocation; } root@MX-RE0> show services nat mappings pcp Interface: sp-1/1/0,
Service set: SS_PCP_NAT NAT pool: PCP_Pool NAT pool: NAT_NON_PCP . PCP Client : 1001::1 PCP
lifetime : 1514 Mapping : 192.168.100.3 : 8080 --> 100.100.100.10 : 1500 Session Count : 0 Mapping
State : Active B4 Address : 1001::1 root@MX-RE0> clear services nat mappings pcp b4address 1001::1
internal-host 192.168.100.3 port 8080 service-set SS_PCP_NAT Interface Service set Mappings removed
Flows removed sp-1/1/0 SS_PCP_NAT 0 0 . PR1284261

```
- One of the internal high availability (HA) queues gets corrupted, which eventually generates a mspmand core file on the backup SDG because sometimes different threads of mspmand might have different timestamps. [PR1291664](#)
- Layer 2 Tunneling Protocol (L2TP) and L2TP access concentrator (LAC) subscribers might get stuck in terminating state because of the race condition during login. [PR1298175](#)
- Remove nonzero check for dsl-type in ICRQ transmission. Dsl-type 0 is valid and should be transmitted in ICRQ. [PR1313093](#)

Subscriber Access Management

- Subscribers get stuck in terminated state during PPPoE login or logout test. [PR1262219](#)
- After Virtual Chassis switchover, RADIUS assigned addresses that do not belong to any configured pool are incorrectly added to a pool. [PR1286609](#)

VPNs

- The routing protocol process (rpd) might eventually become exhausted and crash when Layer 2 Circuit, Layer 2 VPN, or virtual private LAN service (VPLS) configurations are committed. These commit activities might create a small memory leak of 84 bytes in the rpd. If the rpd memory is exhausted, recovery can be accomplished by restarting the rpd. If nonstop active routing (NSR) is configured, the master Routing Engine can be switched over to the standby Routing Engine, causing the master rpd to exit and restart and free the leaked memory. [PR1220363](#)
- In a multicast virtual private network based with Border Gateway Protocol (next-generation-MVPN) scenario with only shortest path tree (SPT) mode configuration, in certain condition the Protocol Independent Multicast (PIM) register-stop packet might be sent before the source tree join (Type-7) packet, which causes some multicast packets to be dropped. [PR1238916](#)
- In a next-generation MVPN scenario with only SPT mode, if the rendezvous point (RP) on the receiver site receives the first register packet, it sends a register stop packet back to the source before sending a type-7 routes packet. Then, the first-hop router might receive the register stop packet, but not receive the join packet. This causes multiple packets to be dropped because there is a multicast route item. [PR1269234](#)

SEE ALSO

New and Changed Features 96
Changes in Behavior and Syntax 120
Known Behavior 125
Resolved Issues 144
Documentation Updates 163
Migration, Upgrade, and Downgrade Instructions 164
Product Compatibility 171

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2 | 145](#)
- [Resolved Issues: 17.3R1 | 158](#)

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: 17.3R2

Application Layer Gateways (ALGs)

- IPsec IKEv2 negotiation fails with IKE ALG enabled. [PR1300448](#)

EVPN

- The traffic might drop after receiving an updated ARP route packet from the peer Layer 3 gateway in an EVPN and VXLAN scenario. [PR1306024](#)
- Split horizon label is not allocated when the ESI configuration switches from **single-active** to **all-active**. [PR1307056](#)
- Core link flap might result in an inconsistent global MAC count. [PR1328956](#)

Forwarding and Sampling

- Some account 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](#)
- ACCT_FORK_LIMIT_EXCEEDED log level is ERROR even when the 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 fail when the **nexthop-learning** configuration statement is enabled for J-Flow v9. [PR1316349](#)
- Some firewall filter counters might not be created in SNMP. [PR1335828](#)

General Routing

- On MX Series platforms, the configuration of enhanced-IP and enhanced-Ethernet network mode is not compatible with MS-DPC card. Hence, the MS-DPC might not work correctly. [PR1035484](#)
- Ksyncd might crash because of the transient replication errors between Routing Engines. [PR1161487](#)
- Stale VBF states occur without sdb sessions. [PR1204369](#)
- 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 Routing Engines. [PR1240524](#)

- Disabling and enabling the "family mpls" of the next-hop interface might cause the route to be in a dead state in a BGP and MPLS scenario with a route of indirect next hop type. [PR1242589](#)
- The **validation-state:unverified** routing entry might not be shown with proper location when users run **show route**. [PR1254675](#)
- The rpd might crash during the next-hop change if unicast reverse-path- forwarding (uRPF) is used. [PR1258472](#)
- Status LED for the ge-0/0/0 interface does not glow. [PR1259112](#)
- PCC controlled LSP metric is not getting updated on the controller, PCE delegated LSPs do not come up. [PR1265864](#)
- MPC might report a parity error with the **fast-lookup-filter** configuration statement. [PR1266879](#)
- On MX Series routers, **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](#)
- On MX Series platforms, if a large number of routes are processed, then the Packet Forwarding Engine of the MS-MPC might crash. [PR1277264](#)
- I2C BUS stuck causes SFP+ thread hogging and restarting of MPC. [PR1277467](#)
- The bbe-smgd process might generate a core file in certain cases when using iflsets in universal call admission control policy mode. [PR1278543](#)
- The chassis network services does not get set as "Enhanced-IP". [PR1279339](#)
- After an MS-MPC PIC goes offline or online or gets bounced (because of an AMS configuration change), sometimes the PIC can take approximately 400 seconds to come up. [PR1280336](#)
- Syslog messages **CM_FPC: Error requesting SET BOOLEAN, illegal setting 132,111** are seen after a unified ISSU from Junos OS Release 16.2R2 to Junos OS Release 17.1R2. [PR1280878](#)
- BIOS firmware upgrade or downgrade support is not available with Junos OS Release 17.3R1. [PR1281050](#)
- The **ingress service-accounting-deferred** command is not providing the correct IP traffic statistics for L2BSA subscribers. [PR1281201](#)
- Subscribers might not be able to connect to MX Series BNG in certain scenarios. [PR1281896](#)
- The kernel might crash in a rare corner case. [PR1282573](#)
- Layer 2 circuit will flap repeatedly, after the link up between PE and CE devices in "asynchronous-notification" and a specific CE device environment. [PR1282875](#)
- Error messages such as **IFRT: 'IFL**, **IFRT: 'Aggregate interface** and **IFRT: 'IFD** are seen when there is a change in configuration. [PR1282938](#)
- On MX Series routers, the CLI command **show interfaces** does not display the reason for bringing down the interfaces when the Packet Forwarding Engine is disabled. [PR1283323](#)

- The log message **VTAG not found in uflow** might be seen when a PPPoE subscriber logs on to a static VLAN logical interface. [PR1284966](#)
- 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 flaps, traffic carried by the other next-hop LSP could get load-balanced across the primary and secondary paths. [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](#)
- Unified ISSU is not supported in Junos OS Release 15.1 or later releases, because the source release includes one or more BBE features such as logical interface (IFL) options, CoS fragmentation map, MLPPP, advisory options, advanced services, and multicast distribution. [PR1286507](#)
- DDoS culprit flows are not reported by CLI or logs in a single Packet Forwarding Engine MX Series router. [PR1286521](#)
- Framed routes might get struck in the KRT queue. [PR1286849](#)
- The A10NSP interface does not get attached to the Layer 2 routing instance after renaming the routing instance. [PR1287070](#)
- SNMP query for 'IF-MIB::ifOutQLen' reports the wrong type. It should be Gauge32 or Unsigned32 for a dynamic VLAN DEMUX0 interface. [PR1287852](#)
- During unified ISSU (FRU upgrade) 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, smid was declared thrashing and was not restarted after a fatal SDB error. [PR1288871](#)
- Kernel "rtdata" memory leak is found on an MX Series Virtual Chassis with the configuration statement heartbeat enabled. [PR1289363](#)
- The FPC memory leak might happen in a BBE subscriber environment. [PR1289365](#)
- The interfaces might go down state after performing GRES. [PR1289493](#)
- The **request system zeroize** command deletes the **/var/db/scripts** directory which does not get re-created until the next USB or Netboot recovery. [PR1289692](#)
- The output **jnxContainersType** is not displayed for PIC and MIC as correctly as it is displayed on other Juniper Networks platforms. [PR1289778](#)
- If any of the vmhost applications are not running, then the alarm string will have "Application" name embedded in it. [PR1290150](#)
- The NAT-T and DPD functionality do not work for aggressive mode. [PR1290689](#)
- Incorrect temperature is displayed for MPCP5 and MPC7 in **show chassis fpc** output. [PR1290771](#)

- Memory leak occurs in the bbe-smgd daemon on subscriber logout for subscribers who have joined any multicast group. [PR1290918](#)
- LSP traffic might silently drop and get discarded after a link goes down in the bypass path. [PR1291036](#)
- The routing protocol process might generate a core file when restarting the process using a CLI command. [PR1291110](#)
- The switch might incorrectly learn its own IRB MAC address. [PR1291184](#)
- Device might lead to the DB prompt `db@jsr_jsm_send_ka_after_merge,send_proto_keepalive`. This is observed on master Routing Engine. [PR1291247](#)
- The **Rescue configuration is not set** minor alarm getting set for MX10003. [PR1291525](#)
- l2tp incoming-call-connected messages retransmit fast and declare that the tunnel is down. [PR1291557](#)
- An error in `vbf_filter_add_orphan_check` might be seen when the subscribers use filter log out or log in. [PR1292582](#)
- An error message might be seen while bringing up the subscriber in a subscriber management environment. [PR1293057](#)
- **DDR3 TEMP ALARM** messages are logged in the chassisd log. [PR1293543](#)
- The `show extensible-subscriber-services sessions` command displays an incorrect timestamp after a unified ISSU. [PR1293800](#)
- On MPC6E linecard inline sampling, the flow export rate remains lower than the configured export rate. [PR1294296](#)
- Loss of DHCP and PPPoE subscribers is observed during unified ISSU from Junos OS Release 16.1-20170718_161_r4_s5.0 to Junos OS Release 16.1-20170718_161_r4_s5.0. [PR1294709](#)
- An rpd core file is generated after interface or BGP flapping. [PR1294957](#)
- The KRT queue might get stuck with the error of `RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0`. [PR1295756](#)
- The bbe-smgd process might generate a core file at `bbe_mcast_ifl_vbf_encoder` on service activation or deactivation along with smg-service daemon restart. [PR1295938](#)
- 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](#)
- TACACS remote user is unable to run JET applications because of a bad stored heap. [PR1296237](#)
- The mspmand process might crash when using TDF gateway services on MS-MPC and MS-MIC. [PR1296422](#)
- The jdhcpd might crash when using 'dhcp-security' related command in enhanced subscriber management. [PR1296461](#)
- LLDP sensor on telemetry uses a lot of bandwidth. [PR1296869](#)

- The kernel might crash continuously when a lot of terms are configured for firewall filters. [PR1296884](#)
- In ECMP fast reroute scenario, traffic might get silently dropped and discarded because next hop is in "hold" state. [PR1297251](#)
- The bbe-smgd memory leak occur in multicast through dax/ddl. [PR1297454](#)
- When a service multicast profile uses variables for group policy or optical internetworking forum (OIF) or SSM-MAP-POLICY and if nonexistent policy names are sent down from the external system during service activation, core files are generated. [PR1297612](#)
- The routing protocol process crashes and generates a core file. [PR1298587](#)
- The commit error [**First_Net**] is thrown when trying to commit a configuration with applied groups. [PR1298649](#)
- The bbe-smgd process might crash when traceoption is enabled because of an invalid username character. [PR1298667](#)
- The bbe-smgd core files are constantly generated while running ESSM and PPPoE stress test with concurrent GRES. [PR1298742](#)
- MX Series BNG does not respond to PADI after GRES on some ports and VLANs. [PR1298890](#)
- When the subscriber limit feature is configured, any new login request after the maximum number of subscribers is denied. [PR1298924](#)
- The "asynchronous notification" feature cannot be implemented properly in a circuit that has MIC-3D-20GE-SFP-E and Tri Rate Copper SFP(740-013111). [PR1299574](#)
- Flat accounting files are not generated according to the configured timers. [PR1299597](#)
- Subscriber database is stuck in "not ready" state after GRES. [PR1299940](#)
- 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 are enabled for service analytics. [PR1300829](#)
- ICMP and ICMPv6 error messages might be discarded while forwarding through an AMS interface. [PR1301188](#)
- Configured sub-interface might not be created correctly after commit. [PR1301823](#)
- Continuous interface flapping might lead to unwanted MIC reset. [PR1302246](#)
- The rpd might crash when toggling **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statements. [PR1302504](#)
- Chassisd.core-tarball.0.tgz is found during unified ISSU aborted in FRU upgrade phase. [PR1303086](#)
- 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 the **show subscribers routing-instance** command. [PR1303199](#)
- The inline-ka PPP echo requests are not generated for aggregated Ethernet interfaces. [PR1303249](#)
- Blocking PPPoE or DHCP to initiate VLAN auto-sensing, if VLAN-OOB connected is in pending state. [PR1303338](#)
- Fan speed changes frequently on MX Series Virtual Chassis. [PR1303459](#)
- MX Series router with MIB polling returns a value that has "sdg". Polling result should include svc generic value. [PR1303848](#)
- Truncated output is shown for the **show pppoe lockout** CLI command. [PR1304016](#)
- Effective rate of E3 in framed mode is limited to 30 Mbps on certain channelized MICs. [PR1304344](#)
- RPF-check strict causes traffic drop in next-generation subscriber management release. [PR1304696](#)
- 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](#)
- Inline J-Flow vMX: OIF field of VPLS data records sometimes report SNMP index value of LSI interface instead of egress physical interface. [PR1305411](#)
- MX Series router sends immediate-interim for the services pushed by SRC. [PR1305425](#)
- The routing protocol process (rpd) crashes on loading EVPN configurations. [PR1305440](#)
- JET **daemonize** application restarts even on normal exit. [PR1305615](#)
- L2BSA subscriber connection attempts failed with VLAN profile-request-error. [PR1305962](#)
- L2BSA subscribers came up, while 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](#)
- IPsec key management process (kmd) stops key exchange process after sending out **UI_DBASE_OPEN_FAILED Too many open files** error message. [PR1308380](#)
- License is lost during Routing Engine switchover in scale-subscriber scenario. [PR1308620](#)
- CoS applied to a subscriber demux logical interface (IFL) is not working. [PR1308671](#)
- All the MICs on FPC, with PS interfaces configured, went offline during the restart of the FPC in another slot. [PR1308995](#)
- Error messages **%PFE-3: fpc0 vbf_var_iflset_add:633: vbf container 11 not found in the msg for ifl .demux.6514** are often seen after MPC restart. [PR1309013](#)
- Incorrect values are found in the event-timestamp of RADIUS accounting-stop packets for L2BSA subscribers. [PR1309212](#)

- On MX2020 and MX2010, after smooth SFB to SFB2 upgrade, if one plane is restarted, link training fails between that plane and the MPC6 cards. [PR1309309](#)
- First access-request fails for L2BSA subscribers when changing the MTU of LACP aggregated Ethernet A10NSP interface. [PR1309599](#)
- DHCP client gets stuck in selecting state while verifying untagged DHCP subscribers after modifying router configuration. [PR1309730](#)
- DT_BNG : 9000 out of 10000 terminated subscribers go down during the unified ISSU from Junos OS Release 16.1 through Junos OS Release 17.3. [PR1309983](#)
- The bbe-smgd process memory leak might be seen after deleting or adding the address pool in next-generation subscriber management release. [PR1310038](#)
- The MS-MIC and 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 logs for SI interfaces when running SNMP walk on service PIC NAT OIDs. [PR1310081](#)
- **krt_junos_sanity_check_ctrl_resp: rtsock** request finally succeeded after error 16 syslog message in Junos OS Release 17.1R1.8. [PR1310678](#)
- Local IPv6 interface from NDRA prefix is not removed from service interface, while subscriber dual-stack session is removed. [PR1310752](#)
- After bsys reboot sometimes rpd is unresponsive on one or more GNFs. [PR1310765](#)
- Bad stored heap: heap-ptr=0x0 data-ptr=0x1481cbf8. [PR1311482](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Counter at PPPoE session logical interface increments incorrectly, causing the accounting packet to contain incorrect acct-input-packets value and incorrect acct-input-octets value. [PR1312998](#)
- The CLI command **show version detail | no-more** hangs for more than 120 seconds in the master Routing Engine and more than 60 seconds in the backup Routing Engine. [PR1314242](#)
- The smgd process generates a core file with reference to bbe_cos_ifl_publish() bbe_cos_if.c:6543. [PR1314651](#)
- The rpd might crash in MoFRR scenario. [PR1314711](#)
- The RIB and FIB might get out of synchronization because the KRT asynchronous queue might get stuck. [PR1315212](#)
- The CLI command **show version detail** gives severity error log **main: name: SRD ret: 0**. [PR1315436](#)
- Transit traffic over GRE tunnel might hit the CPU and trigger a DDoS violation on the L3 next hop. [PR1315773](#)
- The **show auto-configuration out-of-band** CLI command used with a different configuration statement shows 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 rpd might crash when the link flaps on an adjacent router. [PR1318476](#)
- MS-MPC and MS-MIC might crash after a new IPsec tunnel is added. [PR1318932](#)
- MX Series routers sends the IPv6 router advertisements and the DHCPv6 advertisements before sending IPCPv6 ACK from CPE. [PR1321064](#)
- In commit fast-synchronize mode, the commit operation might get stuck after the commit check is performed. [PR1322431](#)
- An incorrect output is observed while verifying the command **show subscribers client-type vlan subscriber-state active logical-system default routing-instance default**. [PR1322907](#)
- Subscribers might fail to login after the interface is deactivated or activated. [PR1324446](#)
- Approximately three percent of Packet Forwarding Engine forwarding capacity might be seen on XM-chip when its temperature is higher than 67 degrees Celsius. [PR1325271](#)
- Minor alarm **LCM Peer Connection un-stable** on the MX150. [PR1328119](#)
- When using **show subscribers** and FPC number has two digits, the interface and IPv6 address get connected together for DHCPv6 PD. [PR1334904](#)

High Availability and Resiliency

- Insufficient available space on hard disk lead by the crashinfo files is generated by ksyncd when GRES is configured in large-scale configuration scenario. [PR1332791](#)

Infrastructure

- 3RU: "Last flapped" time stamp is not getting updated for fxp0 interface as per the expectation. [PR1244502](#)
- The **show system users** CLI command output displays more users than are actually using the router. [PR1247546](#)
- The MX Series router might fail to upgrade Junos OS Release 14.2R6-S4 to Junos OS Release 16.1R4-S4. [PR1298749](#)
- The syscalltrace.sh might create a huge output file, which could cause the router to run out of storage space. [PR1306986](#)

Interfaces and Chassis

- The output value is incorrect when querying the optical power of OTN interfaces on the router. [PR1216153](#)
- [SIRT] MX Series Packet Forwarding Engine and MX Series MPC7E, MPC8E, and MPC9E Packet Forwarding Engine crash when fetching interface statistics with extended-statistics enabled (CVE-2017-10611). [PR1247026](#)

- At a high logical interface scale, an ifinfo process (daemon) generates a core file on executing the command **show interfaces**. [PR1254189](#)
- Monitor interface on aggregated Ethernet logical interfaces displays incorrect bps value compared to **show interface** output. [PR1283831](#)
- The family inet shows as not configured after adding or deleting the loopback address. [PR1294267](#)
- A VRRP track interface-down does not trigger a mastership election immediately. [PR1294417](#)
- IRB interface is showing incorrect bandwidth value. [PR1302202](#)
- AFEB might not come up when LFM is deactivated. [PR1306707](#)
- After executing **request system reboot both** CLI command, the Juniper PPP daemon might become unresponsive. [PR1310909](#)
- The PPPoE subscriber might not log in correctly after authentication failure. [PR1311113](#)
- MX Series Virtual Chassis unified ISSU emits benign error message if unsupported FRUs are present. [PR1316374](#)
- IPv6 Framed Interface Id field is not showing correctly in **show subscribers extensive** output. [PR1321392](#)
- The interface might not work properly after FPC restarts. [PR1329896](#)

Layer 2 Features

- A misconfiguration adds an aggregated Ethernet interface bundle, and its member links to a VPLS instance might cause 100 percent routing protocol process (rpd) utilization. [PR1280979](#)
- On MX Series routers with MPCs or MICs based platforms, packets received on the IRB interface in virtual private LAN services (VPLS) get double tagged. [PR1295991](#)

Layer 2 Ethernet Services

- DHCPV6 client bound to IA_PD prefix on reception of DHCPV6 request for IA_NA, MX Series deletes the existing binding. [PR1286359](#)
- ARP requests are not generated for IRB configured in VPLS over GRE tunnel. [PR1295519](#)
- In a PPPoE and DHCP dual-stack subscriber scenario with an incorrect DHCP configuration, MX Series router might eventually stop logging in PPPoE and DHCP clients. [PR1298976](#)
- Multiple jdncpd core files are observed in jdncpd_update_groups at `../../../../src/junos/usr/sbin/jdncpd/jdncpd_config.c:2290`. [PR1311569](#)

MPLS

- RSVP p2mp sub-LSPs having more than one sub-LSP in down state might not get re-optimized after transit path goes down. [PR1174679](#)
- The rpd might crash when moving static LSP from one routing instance to another. [PR1238698](#)

- The created time value in **show mpls lsp extensive** might drift by a second when the **show** command is issued multiple times. [PR1274612](#)
- MPLS layer 2 circuit ping packet is incorrectly parsed by the output loopback filter. [PR1288829](#)
- Received MTU might not get updated in RSVP MTU signaling. [PR1291533](#)
- Stale RSVP LSP entry occurs after NSR switchover and session is not refreshed. [PR1292526](#)
- The rpd might crash if MPLS LSP path change occurs. [PR1295817](#)
- When using IS-IS traffic engineering (TED), if an LSP's state changed, routing protocol process might lose track of memory. [PR1303239](#)
- BGP multipath might not work when interface flaps. [PR1305228](#)
- Feature "explicit-null" might block host-bound traffic incoming from LSP. [PR1305523](#)
- The rdp process might crash during interface-down events when UHP-based LSPs are configured. [PR1309397](#)

Network Management and Monitoring

- Mib2d-related syslog messages **MIB2D_RTSLIB_READ_FAILURE: rtslib_iflm_snmp_pointchange** are seen during remove and restore configurations. [PR1279488](#)
- The mib2d process might crash when polling the OID ifStackStatus.0 after a logical interface of lo0 is deleted. [PR1286351](#)
- The **show arp no-resolve interface X** command output for nonexistent interface X is showing all unrelated static ARP entries. [PR1299619](#)
- After SNMP configuration activation, the snmpd process started to consume more CPU time. [PR1300016](#)
- The syslog duplicate entries of hostname and timestamp are breaking the standard logging format. [PR1304160](#)

Platform and Infrastructure

- Traffic drop might occur under a large-scale firewall filter configuration. [PR1093275](#)
- The "forwarding-class-accounting enhanced" feature is not supported in combination with "forwarding-options hyper-mode". Using both features together results in traffic getting silently dropped and discarded. [PR1198021](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- Generate-event time-interval usage now triggers the event only on the actual expiry of time interval. [PR1286803](#)
- Incorrect load-balance on ae interface might occur if traffic transits from MS-DPC to MPC card in enhanced-IP mode. [PR1287086](#)
- Packet Forwarding Engine heap memory leak was found in three routers with PPPoE subscribers. [PR1287870](#)

- While adding a new package to the router, you might see the following message: **mgd: error: Could not open library: /usr/lib/render/libvccpd-render.tlv**. [PR1289158](#)
- The syslog error **not a proper library: /usr/lib/render/libdcd-render.so: Cannot open "/usr/lib/render/libdcd-render.so** appears when any non-superuser/non-root user tries to log in to the router.. [PR1289974](#)
- Dynamic MAC learning might fail on GRE tunnel interface. [PR1291015](#)
- The scale-subscriber license might leak on the backup Routing Engine during bulk subscriber logout. [PR1294104](#)
- The management daemon might crash and generate a core file after GRES in a subscriber environment. [PR1298205](#)
- **RMPD_HW_TIMESTAMP_INVALID** is reported two to four times a day, which raises an alarm when polled through **jnxRpmResSumPercentLost** MIB. [PR1300049](#)
- On MX Series platforms with firewall filter configuration, the MPC might reset while loading the configuration. [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 ae member links on DPC (I-chip) based platforms with CoS configured. [PR1301723](#)
- MX Series FPC wedges when creating more than 4000 logical-tunnel interfaces per Packet Forwarding Engine. [PR1302075](#)
- The CLI command **mk destroy-all** is displaying the error **Could not find jnx.wrlsb.mk**. [PR1302974](#)
- The interface-mac-limit might fail for ae interface. [PR1303293](#)
- The **TWAMP Request-TW-Session** message Type-P descriptor format is not RFC-compliant. [PR1305752](#)
- jlaunchd: System reaching processes ceiling <low or high or critical> watermark because of auditd. [PR1305964](#)
- On MX Series routers with MPCs or MICs, the resource monitor (RSMON) thread might get stuck in a loop, consuming 100 percent of FPC CPU. [PR1305994](#)
- The **show system resource-monitor fpc slot <>** command reported memory free percentages that were not accurate. [PR1287592](#)
- The source MACs might leak (or not learn) between different VPLS instances at the receiving end VPLS PE devices. [PR1306293](#)
- This PR addresses the ICMP error messages in Packet Forwarding Engine and is forwarded to the correct pic in the AMS bundle. [PR1313668](#)
- Multicast traffic is not forwarded on the newly added p2mp branch and receiver. [PR1317542](#)
- Multicast traffic might get duplicated when MoFRR is configured. [PR1318129](#)

- Errors might be observed when the **fabric-header-crc-enable** statement is enabled. [PR1320874](#)
- RPM probes delegated to MS-MIC get stuck when any change is made to the BGP group stanza. [PR1322097](#)

Routing Policy and Firewall Filters

- The rpd might crash when **vrf-target auto** is configured under routing-instance. [PR1301721](#)

Routing Protocols

- No multicast forwarding in ASM mode occurs after unified ISSU. [PR1146621](#)
- MPLS over UDP tunnel creation fails in absence of a routing instance table. [PR1270955](#)
- The rpd might crash after deactivating or activating BGP. [PR1272202](#)
- A few bfd sessions flap while coming up after FPC reboots. [PR1274941](#)
- BGP updates might not be advertised to peers completely under certain conditions. [PR1282531](#)
- Some BGP-related traceoptions flag settings might not take effect 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 bounces because of any fiber cut or manual events causing high CPU spike. [PR1291079](#)
- BGP Monitoring Protocol (BMP) might send malformed route monitoring messages. [PR1292848](#)
- The rpd might crash if BGP flap occurs. [PR1295062](#)
- The backup Routing Engine scheduler slips when the import policy is configured improperly. [PR1295712](#)
- Unified ISSU might take more time to complete and the MPC card 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 subsequent anchor FPC goes offline. [PR1298369](#)
- BGP might send an incorrect AS path when an alias is enabled and multiple peers are under the BGP group. [PR1300333](#)
- The rpd process might crash and generate a core file while deleting a multipath route. [PR1302395](#)
- The mcsnoopd process generates a core file during task cleanup. [PR1305239](#)
- Junos OS Release 16.2 and later releases might give the following error: **Request failed: OID not increasing: ospflfpAddress.0.0.0.0.0** . [PR1307753](#)

- The route's next-hop resolution might fail if the static route is configured with **qualified-next-hop** and **resolve** options over a numbered interface. [PR1308800](#)
- BGP labeled-unicast protection might break multicast RPF. [PR1310036](#)
- CST: rpd generates a core file in **bgp_rt_send_message** at `../../../../src/junos/usr/sbin/rpd/bgp/bgp_io.c:1460`. [PR1310751](#)
- 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](#)
- IS-IS SPF gets triggered by LSP updates containing changes in the reservable bandwidth in traffic engineering extensions. [PR1313147](#)
- 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](#)
- On a chassis with BMP configured, the rpd might crash when the rpd process is gracefully terminated. [PR1315798](#)
- BGP-LU update oscillation occurs with BGP-PIC. [PR1318093](#)
- Need to remove the syslog message that got added to code unintentionally. [PR1318458](#)

Services Applications

- TLVs in ICRQ for **actual-rate-downstream** and **actual-data-rate-upstream** do not reflect PPPoE-IA value. [PR1286583](#)
- Mspmand core file "@_arena_mALLOc" is seen in backup SDG's MS70. [PR1291664](#)
- L2TP subscribers are down after a GRES while verifying the framed IPv6 route support for L2TP network server (LNS) at a higher scale with a maximum number of Framed-IPv6-Route. [PR1293783](#)
- The jl2tpd process might crash shortly after a GRES switchover. [PR1295248](#)
- L2TP subscribers might get stuck in terminating state during login. [PR1298175](#)
- The "jl2tpd_era_lns" log files are continuously generated even when L2TP is not configured. [PR1302270](#)
- LTS clients experience packet drop in large packets because of fragmentation in LTS. [PR1312691](#)
- AVP 145 is not present in IRQ when ANCP DSL-type = 0. [PR1313093](#)
- IPCP active mode is not enabled for MLPPP on LNS. [PR1319580](#)

Software Installation and Upgrade

- Junos Selective Update (JSU) package is not activated after a reboot. [PR1298935](#)

Subscriber Access Management

- Service interim for DHCP subscriber is not working in JSRC scenario. [PR1303553](#)
- The output of the **show network-access aaa accounting** command might display additional entries. [PR1304594](#)
- Incorrect Acct-Delay-Time in RADIUS Accounting-On message after rebooting the MX Series BNG. [PR1308966](#)
- When the subscriber is removed manually or through a script, memory leak might be seen. [PR1312517](#)
- The delegated prefix from RADIUS is parsed incorrectly when the length is less than 20 bytes. [PR1315557](#)
- Unified ISSU is not allowed when the account is suspended. [PR1320038](#)
- Authd considers RADIUS attribute Framed-IPv6-Prefix = ::/64 or Delegated-IPv6-Prefix = ::/56 as valid parameters. [PR1325576](#)

VPNs

- Next-generation MVPN SG entry and MVPN route persist after data stop. [PR1236733](#)
- Next-generation MVPN IPv6 RP bootstrap type 3 S-PMSI AD route prefix ff02::d persists after BSR data stop. [PR1269234](#)
- Layer 2 circuits stitched through It peer interfaces might get stuck in local site signal down (LD) status. [PR1305873](#)

Resolved Issues: 17.3R1

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

- Unexpected messages might be seen in logs. [PR1270686](#)
- The sampled process stops collecting data on Routing Engine based sampling supported platforms. [PR1270723](#)
- The sampled process might crash if traceoptions are enabled. [PR1289530](#)

General Routing

- On MX240/480/960 platforms, due to I2C bus hardware issue, FPC might reboot and error message might appear. [PR1174001](#)
- In MX Series subscriber management environment, the rpd might crash in the backup Routing Engine after executing Routing Engine switch over. [PR1206804](#)
- On MX Series routers with MPC2E-3D-NG/MPC2E-3D-NG-Q/MPC3E-3D-NG/MPC3E-3D-NG-Q line card, if the FPC-MIC link failure happens, the bridge might keep sending register messages in an infinite loop, which would cause continuous PCI exceptions, the MPC might crash and traffic forwarding might be affected. This is a rare issue, it is hard to reproduce. [PR1231167](#)
- XM chip based line card (MPC3E/4E/5E/6E/2E-NG/3E-NG) might drop traffic under high temperature (67C or higher). [PR1244375](#)
- On MX2000 with MPC6E, EOAM LFM adjacency flaps when an unrelated MIC accommodated in the same MPC6E slot is brought online by configuring OAM pdu-interval 100 ms and pdu-threshold 3. [PR1253102](#)
- When unified ISSU is performed under scaled scenarios where the Packet Forwarding Engine next-hop memory uses more than 4 Million Dwords, PPE traps and traffic loss may be observed during the software-sync phase until the end of the hardware sync. [PR1267680](#)
- The mspmand log messages about memory zone level which should not be generated are generated. It will occur every 49.7 days and will recover by itself. This is a display issue and will not affect the traffic. [PR1273901](#)
- The CLI commands fails for the following commands: **show subscribers detail**, **show subscribers extensive**, **show subscribers count client-type <>**, and other commands. The failure occurs because the subscriber-management database is unavailable. [PR1274464](#)
- Link stays down after a flap on MPC next generation cards with QSFP+-40G direct attach copper (DAC). [PR1275446](#)
- VT interface flaps during unrelated commit operations if MTU is configured on it. [PR1277600](#)
- vlan-oob subscriber session fails in autoconfd due to physical interface down even if the interface is up. [PR1279612](#)
- **MIC Error code: 0x1b0001** alarm was not clear even after the voltage was returned to normal. [PR1280558](#)
- In a subscriber management environment, if authenticated subscriber dynamic VLAN receives idle timeout from the Radius server, due to a rare timing issue such dynamic VLAN interface can be removed immediately after it was successfully created. [PR1280990](#)
- Establishment of IPsec SAs for link type tunnels might fail under certain conditions in case of scaled IPsec link type service set configuration. In such cases the inside IFL corresponding to service set would remain down. This can be resolved by restarting ipsec-key-management daemon by issuing the following command -----8< -----8< ----- restart ipsec-key-management -----8< -----8< ----- Additionally sometimes the traffic may be affected after restarting IPsec

management daemon. Clearing IPsec SAs corresponding to such service set would resolve this issue. This can be achieved by running the following commands -----8< -----8<
 ----- clear services ipsec-vpn ipsec security association <service-set> -----8<
 -----8< ----- [PR1281223](#) [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](#)
- Inline-JFlow unrelated configuration changes related to a routing-instance results in invalid/incomplete JFlow data packets. Commit-full resumes proper functionality. [PR1282580](#)
- Error messages related to "IFRT: 'IFL'", "IFRT: 'Aggregate interface'" and "IFRT: 'IFD'" seen on config change [PR1282938](#)
- VBF flows are not programmed correctly on ae interfaces resulting in 50% traffic loss. [PR1282999](#)
- OAM fails to come up when GRE tunnel source and family inet address are the same. [PR1283646](#)
- PPTP session could not be established on MSMPC when it is bothstateful-firewall and NAT enabled, and the address could not be translated. [PR1285207](#)
- Possible High CPU on MPC4E when interfaces have been disabled by administrator. [PR1285673](#)
- The J-Flow data template sequence number is zero for MPLS flows. [PR1285975](#)
- Process routing protocol daemon might crash while logging in or logging out with multicast service enabled and performing a GRES switchover. [PR1286653](#)
- L2TP tunnel switch functionality is not working on Junos OS Release 16.1R4-S2 if rewrite-rule configuration is applied to the dynamic profile. [PR1287788](#)
- services-oids-ev-policy.slax & services-oids.slax files built in Junos OS images are not using latest versions. [PR1287894](#)
- After offlining and onlining fabric planes, a few planes are stuck in the offline state in the MX480 router. [PR1287973](#)
- Backup bbe-smgd.core with distributed IGMP configuration. [PR1288465](#)
- If any of the vmhost application is not running then the alarm string will have "Application" name embedded in it. [PR1290150](#)
- BBE-SMGD generates a core file following a stress test in bbe_iff_add_ifa. [PR1291969](#)
- CPCDD might generate core files while using Routing Engine-based http-redirect. [PR1293553](#)
- Not able to edit dynamic profiles after scaling up to 400 dynamic profiles. [PR1295446](#)
- bbe-smgd core at bbe_mcast_ifl_vbf_encoder on service activation or deactivation along with smg-service restarts. [PR1295938](#)

Interfaces and Chassis

- L2TP sessions are not coming up on some of si interfaces after an MPC restart followed by a Routing Engine switchover. [PR1290562](#)

Layer 2 Features

- All the XML duplications and unformatted output are addressed. For Example, histogram was just declared as a element inside pfkey container, with this change a new container is defined for histogram. [PR1271648](#)

Layer 2 Ethernet Services

- DHCP is not using the configured IRB MAC as the source MAC because DHCP is offering only unicast replies. [PR1272618](#)

MPLS

- NG-MVPN MLDP at the receivers' PE does not join P2MP LSP on changing the root PE route from IGP/LDP to LBGp. [PR1277911](#)

Network Management and Monitoring

- The command Esc-q does not work to toggle the console log/terminal log. [PR1269274](#)
- The MIB II process (mib2d) logs an "RLIMIT curr 1048576000 max 1048576000" message every time a commit is performed. [PR1286025](#)
- The mib2d process might crash when polling the OID ifStackStatus.0 after an IFL of lo0 is deleted. [PR1286351](#)

Platform and Infrastructure

- Traffic drop might occur under a large scale of firewall filter configuration. [PR1093275](#)
- FPC crashes with MAC accounting feature enabled. [PR1173530](#)
- FPC CPU spikes every 6 minutes on MX Series routers with MICs and MPCs chipsets due to micro code rebalance. [PR1207532](#)
- RPM loss percent values for "overall tests" through SNMP is incorrect. [PR1272566](#)
- The CLI command **request routing-engine login other-routing-engine** might require a password. [PR1283430](#)
- Transit traffic with DMAC starting with "02" will be punted to Routing Engine when mac-learn-enable is configured. [PR1285874](#)
- The source MAC learned over cross-PFE ae might bounce between ae member Packet Forwarding Engines for a long time and which might cause MLP-ADD storm. [PR1290516](#)
- RMOPD might get stuck in the sbwait state upon receiving a specific response from the HTTP agent. [PR1292151](#)

Routing Protocols

- Routing protocol daemon on the backup Routing Engine might restart unexpectedly upon the addition of a new L2VPN routing instance. [PR1233514](#)
- When the **advertise-from-main-vpn-tables** configuration statement is used under BGP and if RR functionality is added, a refresh message is not sent, and as a result, some routes are missed. [PR1254066](#)
- MPLSoUDP tunnel creation failure in the absence of a routing instance table. [PR1270955](#)
- After Routing Engine switchover (GRES+GR) default mdt failed to come up also seen with core facing interface flap. [PR1279459](#)
- Routing protocol daemon might crash due to a certain chain of events in the BGP-LU protection scenario. [PR1282672](#)
- The second multicast packet might be discarded on RP router. [PR1282848](#)
- Routing protocol daemon crashes while deactivating in a routing instance protocols pim static. [PR1284760](#)
- Routing protocol daemon might crash if dynamic RP goes down in ECMP topology when PIM join load balancing automatic is configured. [PR1288316](#)

Services Applications

- DTCP LI filters are very slow to program when using the "X-RM-Circuit-ID" trigger. [PR1269770](#)
- Business service fails to get deactivated post Routing Engine switchover. [PR1280074](#)
- Backup Routing Engine is going to the database prompt with a vmcore if the down ASI interface configuration is deleted. [PR1281882](#)
- Loss of all L2TP subscribers on an LAC router after smg-service restarts on the L2TP tunnel switch.. [PR1284260](#)
- The l2tpd process generates a core file with reference to 0x084166f5 in L2tpTunnel::createSucceeded (this=0xa04ae84, createFlags=...) at ../src/junos/usr.sbin/jl2tpd/l2tpTunnel.cc:1845. [PR1288029](#)
- Each subscriber session is getting its own L2TP tunnel without "Tunnel-Client-Endpoint" from radius. [PR1293927](#)

Subscriber Management and Services

- MX Series router could not filter some RADIUS attributes with the accounting-Off and accounting-On messages. [PR1279533](#)
- Authenticated subscriber dynamic VLAN interface might get disconnected immediately after a successful connection. [PR1280990](#)
- Authd core file is observed while terminating large number of subscribers. [PR1289215](#)

User Interface and Configuration

- The commitd process might generate a core file by certain configuration removal followed by a commit operation. [PR1267433](#)

VPNs

- Routing protocol daemon memory leak is observed in next-generation-MVPN enviroment. [PR1259579](#)

SEE ALSO

New and Changed Features 96
Changes in Behavior and Syntax 120
Known Behavior 125
Known Issues 131
Documentation Updates 163
Migration, Upgrade, and Downgrade Instructions 164
Product Compatibility 171

Documentation Updates

IN THIS SECTION

- [Subscriber Management Provisioning Guide | 164](#)

This section lists the errata and changes in Junos OS Release 17.3R2 documentation for MX Series.

Subscriber Management Provisioning Guide

- 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 Baseline Accounting Statistics Overview](#).]

SEE ALSO

[New and Changed Features | 96](#)

[Changes in Behavior and Syntax | 120](#)

[Known Behavior | 125](#)

[Known Issues | 131](#)

[Resolved Issues | 144](#)

[Migration, Upgrade, and Downgrade Instructions | 164](#)

[Product Compatibility | 171](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.3 | 165](#)
- [Procedure to Upgrade to FreeBSD 10.x based Junos OS | 166](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 168](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 169](#)
- [Upgrading a Router with Redundant Routing Engines | 170](#)
- [Downgrading from Release 17.3 | 170](#)

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

Starting 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](#).

NOTE: In Junos OS Release 15.1, Junos OS (FreeBSD 10.x) is not available to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to use the existing Junos OS (FreeBSD 6.1).

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
MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 17.3

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

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

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files) might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [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://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system 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.

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-32-17.3R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.3R1.9-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**

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 17.3**jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.1) 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://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system 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.

- Customers in the United States and Canada, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-17.3R1.9-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.3R1.9-export-signed.tgz
```

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

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

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 17.3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 16.2, 17.1 and 17.2 are EEOL releases. You can upgrade from Junos OS Release 16.2 to Release 17.1 or even from Junos OS Release 16.2 to Release 17.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

For more information about EEOL releases and to review a list of EEOL releases, see <https://support.juniper.net/support/eol/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 17.3

To downgrade from Release 17.3 to another supported release, follow the procedure for upgrading, but replace the 17.3 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 96
Changes in Behavior and Syntax 120
Known Behavior 125
Known Issues 131
Resolved Issues 144
Documentation Updates 163
Product Compatibility 171

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 171](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 96
Changes in Behavior and Syntax 120
Known Behavior 125
Known Issues 131
Resolved Issues 144
Documentation Updates 163
Migration, Upgrade, and Downgrade Instructions 164

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [New and Changed Features | 172](#)
- [Changes in Behavior and Syntax | 173](#)
- [Known Behavior | 174](#)
- [Known Issues | 174](#)
- [Resolved Issues | 175](#)
- [Documentation Updates | 175](#)
- [Migration, Upgrade, and Downgrade Instructions | 176](#)
- [Product Compatibility | 177](#)

These release notes accompany Junos OS Release 17.3R2 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Juniper Device Manager | 173](#)

This section describes the new features or enhancements to existing features in Junos OS Release 17.3R2 for NFX Series.

Juniper Device Manager

- **Support for Virtual Route Reflector (NFX250-S2)**—Starting in Junos OS Release 17.3R1, you can implement the virtual router reflector capability by creating and deploying a VRR virtual machine as a VNF (Virtual Network Function) on the NFX250-S2 device. Benefits of implementing virtual route reflectors are:
 - Improved scalability
 - Fast and more flexible deployment
 - Savings as a result of elimination of router hardware

SEE ALSO

Changes in Behavior and Syntax 173
Known Behavior 174
Known Issues 174
Resolved Issues 175
Documentation Updates 175
Migration, Upgrade, and Downgrade Instructions 176
Product Compatibility 177

Changes in Behavior and Syntax

There are no changes in behavior and syntax for NFX Series in Junos OS Release 17.3R2.

SEE ALSO

New and Changed Features 172
Known Behavior 174
Known Issues 174
Resolved Issues 175
Documentation Updates 175
Migration, Upgrade, and Downgrade Instructions 176
Product Compatibility 177

Known Behavior

There are no known limitations in Junos OS Release 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.

SEE ALSO

New and Changed Features 172
Changes in Behavior and Syntax 173
Known Issues 174
Resolved Issues 175
Documentation Updates 175
Migration, Upgrade, and Downgrade Instructions 176
Product Compatibility 177

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.3R2 for NFX Series.

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

SEE ALSO

New and Changed Features 172
Changes in Behavior and Syntax 173
Known Behavior 174
Resolved Issues 175
Documentation Updates 175
Migration, Upgrade, and Downgrade Instructions 176
Product Compatibility 177

Resolved Issues

There are no fixed issues in Junos OS Release 17.3R2 for NFX Series.

SEE ALSO

New and Changed Features 172
Changes in Behavior and Syntax 173
Known Behavior 174
Documentation Updates 175
Known Issues 174
Migration, Upgrade, and Downgrade Instructions 176
Product Compatibility 177

Documentation Updates

There are no errata or changes in Junos OS Release 17.3R2 documentation for NFX Series.

SEE ALSO

New and Changed Features 172
Changes in Behavior and Syntax 173
Known Behavior 174
Known Issues 174
Resolved Issues 175
Migration, Upgrade, and Downgrade Instructions 176
Product Compatibility 177

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 176](#)

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. For example, Junos OS Releases 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. 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://support.juniper.net/support/eol/junos.html>.

SEE ALSO

[New and Changed Features | 172](#)

[Changes in Behavior and Syntax | 173](#)

[Known Behavior | 174](#)

[Documentation Updates | 175](#)

[Known Issues | 174](#)

Resolved Issues 175
Product Compatibility 177

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 177

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 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 172
Changes in Behavior and Syntax 173
Known Behavior 174
Documentation Updates 175
Known Issues 174
Resolved Issues 175
Migration, Upgrade, and Downgrade Instructions 176

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- New and Changed Features | 178
- Changes in Behavior and Syntax | 186
- Known Behavior | 189
- Known Issues | 190
- Resolved Issues | 192
- Documentation Updates | 195
- Migration, Upgrade, and Downgrade Instructions | 196
- Product Compatibility | 200

These release notes accompany Junos OS Release 17.3R2 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.3R2 New and Changed Features | 179
- Release 17.3R1 New and Changed Features | 179

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 17.3R2 New and Changed Features

Software Installation and Upgrade

- **Device serial number added to DHCP option 60 (PTX1000)**—Starting in Junos OS Release 17.3R2, DHCP option 60 (Vendor Class Identifier) includes the serial number of the device when you use zero touch provisioning (ZTP) 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 17.3R1 New and Changed Features

Class of Service

- **Support for setting the DSCP code point for host-originating IS-IS traffic sent over a GRE tunnel (PTX Series)**—Starting in Junos OS Release 17.3R1, you can determine traffic prioritization for IS-IS traffic originating on a host and being sent over a GRE tunnel by assigning a DSCP code point to the IS-IS packets. You can set the DSCP code point by including the **isis-over-gre dscp-code-point value** statement at the **[edit class-of-service host-outbound-traffic protocol]** hierarchy level.

[See [protocol \(Host Outbound Traffic\)](#).]

- **Support for shaping of the traffic exiting a physical interface (PTX10008)**—Starting with Junos OS Release 17.3R1, you can shape the output traffic of a physical interface on PTX10008 routers so that the interface transmits less traffic than it is physically capable of carrying. Shaping on a PTX10008 router interface 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 \(Applying to an Interface\)](#).]

General Routing

- **Commit process split into two steps (PTX Series)**—Starting in Junos OS Release 17.3R1, new configuration statements are introduced for **commit** to split the commit process into two steps. These configuration statements are **prepare** and **activate**.

In the first step, known as preparation stage, **commit prepare** validates the configurations and then creates the necessary files and database entries so that the validated configurations can be activated at a later stage.

In the second step, referred to as the activation stage, **commit activate** activates the previously prepared commit. A new configuration statement, **prepared**, is added to **clear system commit**, which clears the prepared commit cache

This feature enables you to configure a number of Junos OS devices and simultaneously activate the configurations. This approach is helpful in time-critical scenarios.

[See [Commit Preparation and Activation Overview](#).]

Interfaces and Chassis

- **Management Ethernet interface (fxp0) is confined in a non-default virtual routing and forwarding table (PTX 10008)**—Starting in Junos OS Release 17.3R1, you can confine the management interface in a dedicated management instance by setting a new CLI configuration statement, **management-instance**, at the **[edit system]** hierarchy level. By doing so, operators will ensure that management traffic no longer has to share a routing table (that is, the default.inet.0 table) with other control or protocol traffic in the system. Instead, there is a **mgmt_junos** routing instance introduced for management traffic.

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

- **Support for confining management Ethernet Interface (fxp0) in a virtual routing and forwarding table (PTX10008)**—Starting in Junos OS Release 17.3R1, Junos OS is able to confine the management interface in a dedicated management instance by setting a new CLI configuration statement, **management-instance**, at the **[edit system]** hierarchy level. By doing so, operators will ensure that management traffic no longer has to share a routing table (that is, default.inet.0 table) with other control or protocol traffic in the system. Instead, there is a **mgmt_junos** routing instance introduced for management traffic.

For more information, see [Configuring the mgmt_junos Routing Instance](#)

Management

- **Support to configure YANG files for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.3R1, you can add user-defined YANG files that provide mappings between the XML path and the OpenConfig path for data streamed through the Junos Telemetry Interface. Previously, only the Junos OpenConfig package was available for providing these mappings to the XML proxy when streaming data through gRPC. To add YANG files, include the **request system yang add package *package-name* proxy-xml module *yang-file-path*** operational command. You can validate the YANG module by using the **request system yang validate proxy-xml module *yang-file-path*** command. To delete a YANG file, use the **request system yang delete package *package-name* proxy-xml *yang-file-path*** operational command.

[See [Creating YANG Files for XML Proxy for Junos Telemetry Interface](#).]

- **Enhancements to BGP peer sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.3R1, telemetry data streamed through gRPC for BGP peers is reported separately for each routing instance. To export data for BGP peers, you must now include the following path in front of all supported paths:

/network-instances/network-instance/[name_ 'instance-name']/protocols/protocol/

Additionally, the following paths are also now supported:

- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/accepted**
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/snmp-peer-index**
- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/afi-safis/afi-safi/state/queues/output**
- **/network-instances/network-instance/protocols/protocol
/bgp/neighbors/neighbor/afi-safis/afi-safi/state/queues/input**
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/ImportEval**
- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/state/ImportEvalPending**

Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Junos Telemetry Interface support for Routing and Control Board RCB-PTX-X6-32G (PTX3000)**—Starting with Junos OS Release 17.3R1, the Routing and Control Board (RCB) on PTX3000 routers supports the Junos Telemetry Interface, which enables you to provision sensors to export telemetry data for various network elements. The RCB combines the functionality of a Routing Engine, Control Board, and Centralized Clock Generator (CCG) in a single FRU. 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](#).]

- **Enhanced support for Junos Telemetry Interface (PTX1000 routers)**—Starting with Junos OS Release 17.3R1, you can also provision sensors through the Junos Telemetry Interface for the following network elements:

- Logical interfaces, including queue statistics (UDP and gRPC streaming)
- BGP Peers (gRPC streaming only)
- Memory utilization for routing protocol tasks (gRPC streaming only)
- RSVP interface events (gRPC streaming only)
- Firewall filters, including traffic-class counter (UDP and gRPC streaming)
- Chassis components (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)
- Routing Engine logical and physical interfaces (UDP and gRPC streaming)
- Optical interfaces (UDP and gRPC streaming)
- Network Discovery Protocol table state (gRPC streaming only)
- Address Resolution Protocol table state (gRPC streaming only)
- IPFIX inline flow aggregation (UDP 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. 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](#).]

Multicast

- **Support for next generation MVPN and Internet multicast (PTX1000)**—Starting in Junos OS Release 17.3R1, 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 (or next generation) MVPN.

NOTE: Next-generation MVPN is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

[See [Multiprotocol BGP MVPNs Overview](#).]

- **Support for next generation MVPN and Internet multicast (PTX10008)**—Starting in Junos OS Release 17.3R1, 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 (or next generation) MVPN. Next generation MVPN is available only for PTX Series routers that have third-generation FPCs installed.

[See [Multiprotocol BGP MVPNs Overview](#).]

Network Management and Monitoring

- **mLDP MIB extends support to LDP point-to-multipoint (P2MP) LSPs (PTX Series)**—Starting in Junos OS Release 17.3R1, the mLDP MIB builds on the objects and tables that are defined in RFC 3815, which only support LDP point-to-point label switched paths (LSPs). This mLDP MIB provides support for managing multicast LDP point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs. The mLDP MIB tables are directly accessible through SNMP. All objects in the mLDP MIB are read-only and cannot be created or set through SNMP. This implementation of mLDP MIB is specified in draft-ietf-mpls-mldp-mib.
- **Support for inline jflow version 9 flow templates (PTX1000)**—Starting in Junos OS Release 17.3R1, you can use inline-JFlow's 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](#).]

Operation, Administration, and Maintenance

- **Junos OS daemons to natively emit JSON output (PTX Series)**—Starting with Junos OS Release 17.3R1, the operational state emitted by the daemons is supported in JSON format as well as XML format. To configure JSON format, specify the following CLI command: **set system export-format state-data json compact**. To specify JSON format for specific command output, include **display json** in specific CLI commands.
- **Junos OS OpenConfig to support adjacent RIB operational state model (PTX Series)**—Starting with Junos OS Release 17.3R1, **adj-rib-in-pre** and **adj-rib-out-post** tables have been added for the OpenConfig RIB operational state mode. The BGP RIB consists of several tables per address family, consisting of **loc-rib** and **per-neighbor** tables.

Routing Policy and Firewall Filters

- **Optimized performance for DSCP and traffic-class firewall filter match conditions (PTX10008)**—Starting in Junos OS Release 17.3R1, the **promote dscp** and **promote traffic-class** indicators are supported in firewall filters for IPv4 and IPv6 traffic. When either of these are applied to a filter, the entire filter is compiled in a way that optimizes its performance for the **dscp** or **traffic-class** match condition. The indicators are configured at the **[edit firewall family (inet | inet6) filter filter-name]** hierarchy level.

NOTE: Enabling the indicators requires that network services is set to **enhanced-mode**. Use of the indicators may impact the performance of the **source-port** match condition.

- **Optimized performance for DSCP and traffic-class firewall filter match conditions (PTX1000)**—Starting in Junos OS Release 17.3R1, the **promote dscp** and **promote traffic-class** indicators are supported in firewall filters for IPv4 and IPv6 traffic. When either are applied to a filter, the entire filter is compiled in a way that optimizes its performance for the **dscp** or **traffic-class** match condition. The indicators are configured at the **[edit firewall family (inet | inet6) filter filter-name]** hierarchy level.

NOTE: Enabling the indicators requires that network services be set to **enhanced-mode**. Use of the indicators might impact the performance of the **source-port** match condition.

[See [Promote DSCP](#) and [Promote traffic-class](#).]

- **Support for Hop-limit firewall filter match condition (PTX10008)**—Starting in Junos OS Release 17.3R1, you can configure a firewall filter using the **hop-limit hop-limit** and **hop-limit except hop-limit** match conditions for Internet Protocol version 6 (IPv6) traffic (family inet6).

NOTE: The **hop-limit hop-limit** and **hop-limit except hop-limit** match conditions are supported on PTX series routers when you configure the network-services mode as **enhanced-mode** on the router.

For more information, see [Firewall Filter Match Conditions for IPv6 Traffic](#).

- **Hop-limit firewall filter match condition supported (PTX1000)**—Starting in Junos OS Release 17.3R1, 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 PTX1000 routers when **enhanced-mode** is configured on the router.

[See [Firewall Filter Match Conditions for IPv6 Traffic](#).]

Routing Protocols

- **Routing protocol process (rpd) recursive resolution over multipath (PTX Series)**—Starting in Junos OS Release 17.3R1, when a BGP prefix that has a single protocol next hop is resolved over another BGP prefix that has multiple resolved paths (unilist), all the paths are selected for protocol next-hop resolution. In prior Junos OS releases, only one of the paths is picked for protocol next-hop resolution. This new feature benefits densely connected networks where BGP is used to establish infrastructure connectivity such as WAN networks with high equal-cost multipath and seamless MPLS topology.

To configure recursive resolution over multipath, define a policy that includes the **multipath-resolve** action at the **[edit policy-options policy-statement *policy-name* then]** hierarchy level and import the policy at the **[edit routing-options resolution rib *rib-name*]** hierarchy level.

[See [Configuring Recursive Resolution over BGP Multipath](#).]

- **Support for IS-IS SPRING and RSVP coexistence (PTX Series)**—Starting in Junos OS Release 17.3R1, the routing protocol process (rpd) takes into account the bandwidth used by SPRING traffic to calculate the balance bandwidth available for RSVP-TE. The allocated bandwidth for RSVP is periodically modified based on the traffic on the SPRING interface and its bandwidth utilization. To configure automatic bandwidth calculation, include the **auto-bandwidth template** statement at the **[edit routing-options]** hierarchy level. You can apply the **auto-bandwidth template** configuration either globally at the **[edit protocols isis source-packet-routing traffic-statistics]** hierarchy level or at the **[edit protocols isis interface *interface-name*]** hierarchy level. This feature is useful for networks that are moving to SPRING but also have RSVP deployed, and continue to use both SPRING and RSVP.

[See [auto-bandwidth](#).]

- **Support for BGP Large Communities (PTX Series)**—Starting with Junos OS Release 17.3R1, BGP community is enhanced to support BGP large community that uses 12-byte encoding where the most significant 4 bytes encode autonomous system number or global administrator and the remaining two 4 bytes encode operator defined local values. Currently, BGP normal community (4 byte) and BGP extended community (6 byte) provide limited support for BGP community attributes after the introduction of 4-byte autonomous system number. Configure the large BGP community attributes at the **[edit policy-options community *community-name* members]** hierarchy level and at the **[edit routing-options static route *route* community]** hierarchy level with keyword **large** followed by three 4-byte unsigned integers separated by colons. The attributes are represented as large:autonomous system number:local value 1:local value2.
- **Support for BGP to carry flow-specification routes (PTX10008)**—Starting in Junos OS Release 17.3R1, BGP can carry flow-specification network layer reachability information (NLRI) messages on a PTX10008 router. Propagating firewall filter information as part of BGP enables you to propagate firewall filters against denial-of-service (DoS) attacks dynamically across autonomous systems.

[See [Example: Enabling BGP to Carry Flow-Specification Routes](#).]

Services Applications

- **Support for inline JFlow version 9 flow templates (PTX 10008 routers)**—Starting in Junos OS Release 17.3R1, you can use inline-JFlow export capabilities with version 9 flow templates to define a flow record template suitable for IPv4 or IPv6 traffic.

[See [Monitoring Network Traffic Flow Using Inline Flow Monitoring on PTX Series Routers.](#)]

SEE ALSO

Changes in Behavior and Syntax 186
Known Behavior 189
Known Issues 190
Resolved Issues 192
Documentation Updates 195
Migration, Upgrade, and Downgrade Instructions 196
Product Compatibility 200

Changes in Behavior and Syntax

IN THIS SECTION

- [Forwarding and Sampling | 187](#)
- [Interfaces and Chassis | 187](#)
- [Management | 187](#)
- [Network Management and Monitoring | 187](#)
- [Services Application | 188](#)
- [VLAN-Infrastructure | 189](#)

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 17.3R2 for the PTX Series.

Forwarding and Sampling

- In Junos OS Release 17.3R1, 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.

Interfaces and Chassis

- **Restart FPC option supported on PTX1000 router**—In Junos OS Release 17.3, 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

- **Changes to custom YANG RPC syntax (PTX Series)**—Starting in Junos OS Release 17.3, custom YANG RPCs have the following changes in syntax:
 - The **junos:action-execute** statement is a substatement to **junos:command**. In earlier releases, the **action-execute** and **command** statements are placed at the same level, and the **command** statement is optional.
 - The CLI formatting for a custom RPC is defined within the **junos-odl:format** statement, which takes an identifier as an argument. In earlier releases, the CLI formatting is defined using a container that includes the **junos-odl:cli-format** statement with no identifier.
 - The **junos-odl:style** statement defines the formatting for different styles within the statement. In earlier releases, the CLI formatting for different styles is defined using a container that includes the **junos-odl:cli-format** and **junos-odl:style** statements.
- **Enhancement to show agent sensors command (PTX Series)** —Starting with Junos OS Release 17.3R1, the **show agent sensors** command, which displays information about Junos Telemetry Interface sensors, displays the default value of **0** for the **DSCP** and **Forwarding-class** values. Previously, the displayed default value for these fields was **255**. The default value is displayed when you do not configure a DSCP or forwarding-class value for a sensor at the **[edit services analytics export-profile profile-name]** hierarchy level.

[See [export-profile](#) and [show agent sensors](#).]

Network Management and Monitoring

- **SNMP syslog messages changed (PTX Series)**—In Junos OS Release 17.3R1, 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](#).]

- **Enhancement to about-to-expire logic for license expiry syslog messages (PTX Series)**—Starting in Junos OS Release 17.3R1, the logic for multiple capacity type licenses and when their expiry raises alarms was changed. Before, the behavior had alarms and syslog messages for expiring licenses raised based on the highest validity, which would mislead users in the case of a license expiring earlier than the highest validity license. The new behavior has the about-to-expire logic based on the first expiring license.
- **Change in default log level setting (PTX Series)**—In Junos OS Release, 17.3R2, 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](#).]

Services Application

- **Changes to the show services rpm history-results command (PTX Series)**—Starting in Junos OS Release 17.3R1, you must include the **owner owner** and **test name** options when using the **show services rpm history-results** command.

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

- In Junos OS Release 17.3R1 and later, for PIC-based J-Flow on MX Series routers and inline J-Flow on PTX Series routers, the Options template and Options data records include the **Sampling Interval** field as part of the **ScopeTemplate** field instead of the **ScopeSystem** field.

VLAN-Infrastructure

- **LAG interface flaps while adding/removing a VLAN**—Starting in Junos OS Release 17.3, the LAG interface flaps while adding or removing a VLAN. The flapping happens when a low-speed SFP is plugged into a relatively high-speed port. To avoid flapping, configure the port speed to match the speed of the SFP.

SEE ALSO

[New and Changed Features | 178](#)

[Known Behavior | 189](#)

[Known Issues | 190](#)

[Resolved Issues | 192](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 196](#)

[Product Compatibility | 200](#)

Known Behavior

IN THIS SECTION

- [General Routing | 189](#)
- [Multiprotocol Label Switching \(MPLS\) | 190](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 for PTX Series.

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

General Routing

- Uneven load balancing of traffic might occur if the traffic stream changes only in bits 0-15 of the layer3 destination IPv6 address. This limitation might not be visible if the other parameters effecting the load

- balance change along with L3_DST, such as L3 source IP address, L4 source, or destination ports and so on. [PR1065515](#)
- On a PTX Series router with a faulty power-supply(PSM). The PSM might generate excessive interrupt requests (IRQs). Since hardware IRQs are processed by chassis process (chassisd), excessive IRQs might cause chassisd to restart when the condition persists more than 200 seconds. [PR1226992](#)
 - With non-enhanced-mode, traffic loss is seen on version 4 static-lsp with stitch operation, which does not work on PTX Series routers. [PR1290942](#)

Multiprotocol Label Switching (MPLS)

- When NG-MVPN is configured with RSVP provider tunnels and NSR is used, then the egress router for the tunnel might not correctly replicate some of the tunnel state to the backup routing engine, leading to temporary traffic loss during NSR failover for the affected tunnels. [PR1293014](#)

SEE ALSO

New and Changed Features 178
Changes in Behavior and Syntax 186
Known Issues 190
Resolved Issues 192
Documentation Updates 195
Migration, Upgrade, and Downgrade Instructions 196
Product Compatibility 200

Known Issues

IN THIS SECTION

- [General Routing | 191](#)
- [Interfaces and Chassis | 192](#)
- [Routing Protocols | 192](#)

This section lists the known issues in hardware and software in Junos OS Release 17.3R2 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

- PTX Series FPC3 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 rare occasions, upon reboot, the kernel cannot create sysfs entries for the SSDs in the system. This might result in the system entering the 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](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- This issue is applicable when using MPLS LSPs and RSVP-TE self-ping. When rpd sends a self-ping packet and an RSVP packet at the same time, these packets might overwrite the kernel's packet buffers, causing memory corruption and kernel panic. [PR1303798](#)
- On a PTX10000 platform with FPC "LC1101 - 30C / 30Q / 96X" installed, the 10G interface might flap when the interface is active and it is set to 100 Gbps speed. [PR1315079](#)

Interfaces and Chassis

- A Junos OS upgrade involving Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later main releases with a connectivity fault management (CFM) configuration can cause cfmd to generate a core file after the upgrade. This is because of the old version of `/var/db/cfm.db`. [PR1281073](#)

Routing Protocols

- With Shared Risk Link Group (SRLG) enabled under corner conditions, after executing `clear isis database`, the rpd might crash because the IS-IS database tree gets corrupted. [PR1152940](#)

SEE ALSO

New and Changed Features 178
Changes in Behavior and Syntax 186
Known Behavior 189
Resolved Issues 192
Documentation Updates 195
Migration, Upgrade, and Downgrade Instructions 196
Product Compatibility 200

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2 | 193](#)
- [Resolved Issues: 17.3R1 | 195](#)

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: 17.3R2

General Routing

- On PTX1000 routers, the error message `ch_get_product_attribute.324: Cannot find chassisd` is displayed when loading images. [PR1217505](#)
- On PTX Series routers, a faulty power supply module (PSM) might generate excessive interrupt requests. These hardware interrupt requests, processed by chassisd, might restart the chassisd process when the condition persists more than 200 seconds. [PR1226992](#)
- The **validation-state:unverified** routing entry might not be shown with the proper location when users run **show route**. [PR1254675](#)
- The rpd process might crash after BGP sessions and routes flap. [PR1269327](#)
- 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 with a SIB-SFF-PTX-240-S platform. [PR1282983](#)
- The MPLS TTL might reset to 255 on third-generation PTX Series FPCs if the **protocols mpls no-propagate-ttl** statement is configured. [PR1287473](#)
- LSP traffic might silently drop and get discarded after a link goes down in bypass path. [PR1291036](#)
- The routing protocol process (rpd) might generate a core file while restarting the process from the CLI. [PR1291110](#)
- Incorrect SNMP OID values are sent in SNMP traps for removal or insertion of a front panel display on PTX Series routers. [PR1294741](#)
- LINK LED is RED when the port is disabled on PTX Series routers. [PR1294871](#)
- The rpd core file is generated after interface or BGP flapping. [PR1294957](#)
- The chassisd process might run out of memory and restart on a PTX1000 platform. [PR1295691](#)
- On a PTX5000 or 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](#)
- The mgd core file is generated when downgrading from Junos OS Release 17.3-20170721 to Junos OS Release 16.1X65D40.2. The mgd core file is overwritten if downgrading is attempted multiple times. [PR1296504](#)
- On a PTX1000, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3-20170721 fails frequently with sampling enabled. [PR1296533](#)
- Alarms and syslog errors are seen with priority strict-high on an AF4 queue, on the oversubscription cases (1X100G egress to 1X10G egress setup). [PR1297343](#)
- The disable-pfe action upon Hybrid Memory Cube (HMC) fatal errors might have a system-wide impact on PTX Series platforms. [PR1300180](#)

- PTX Series router 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 router FPCs during a multicast route installation into the Packet Forwarding Engine. [PR1302303](#)
- On a PTX3000, powering on an FPC (OPT-3-SFF-PTX/IPLC) card reboots the other FPC cards. [PR1302304](#)
- The third-generation FPC (FPC3-SFF-PTX) might not boot on a PTX3000 with the Control Board or Routing Engine. [PR1303295](#)
- The 100G interfaces might not come up on a PTX3000 and a PTX5000. [PR1303324](#)
- This issue occurs when using MPLS LSPs and RSVP-TE self-ping. When rpd sends out a self-ping packet and an RSVP packet at the same time, these packets might overwrite the kernel's packet buffers causing memory corruption and kernel panic. [PR1303798](#)
- PTX3000 with RCB-PTX Routing Engine might be unable to come online or recognize integrated photonic line cards (IPLCs). [PR1304124](#)
- The routing information base (RIB - also known as routing table) and forwarding information base (FIB - also known as forwarding table) might not synchronize in a large-scale network, because of a timing issue. The root cause is that when the rpd sends route update messages to the kernel, the KRT queue that is used to send the messages can get into a state in which no more messages can be sent to the kernel. [PR1315212](#)
- The physical interfaces might generate framing errors when ports are connected to an odd interface. [PR1317827](#)

Infrastructure

- The **show system users** CLI command output displays more users than that are actually using the router. [PR1247546](#)

Interfaces and Chassis

- 100G interfaces might not come up when **otn-options laser-enable** is configured on PTX Series platforms. [PR1297164](#)
- LFM discovery state might show up as a fault for an aggregated interface after a GRES switchover. [PR1299534](#)

MPLS

- In an RSVP environment, a stale LSP might get created after a Routing Engine switchover with nonstop routing (NSR) enabled. [PR1292526](#)
- The rpd might crash when the MPLS LSP path change occurs. [PR1295817](#)

Platform and Infrastructure

- Continuous log messages occur. For example: `tftpd[23724]: Timeout #35593 on DATA block 85`. [PR1315682](#)

Routing Protocols

- A few BFD sessions flap while coming up after FPC restarts or reboots. [PR1274941](#)
- Multihop BFD sessions flap continuously when the PTX Series router is in the middle hop. [PR1291340](#)
- The rpd process crashes and generates core files multiple times when you receive an OPEN message from an existing BGP peer. [PR1299054](#)
- With BGP labeled unicast MPLS fast reroute in an inter-AS scenario, a very high fast reroute time is visible once the link is up. [PR1307258](#)

Resolved Issues: 17.3R1

Layer 2 Features

- All the XML duplications and unformatted output are addressed. For Example, histogram was just declared as a element inside pfkey container, with this change a new container is defined for histogram. [PR1271648](#)

SEE ALSO

New and Changed Features 178
Changes in Behavior and Syntax 186
Known Behavior 189
Known Issues 190
Documentation Updates 195
Migration, Upgrade, and Downgrade Instructions 196
Product Compatibility 200

Documentation Updates

There are no errata or changes in Junos OS Release 17.3R2 documentation for PTX Series.

SEE ALSO

New and Changed Features	178
Changes in Behavior and Syntax	186
Known Behavior	189
Known Issues	190
Resolved Issues	192
Migration, Upgrade, and Downgrade Instructions	196
Product Compatibility	200

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 196
- Upgrading a Router with Redundant Routing Engines | 197
- Basic Procedure for Upgrading to Release 17.3 | 197

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 14.2 to Release 16.1. 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://support.juniper.net/support/eol/junos.html>.

Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Basic Procedure for Upgrading to Release 17.3

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [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 17.3R1:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system 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 17.3R1 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-17.3  
R1.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.3  
R1.SPIN-export-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 17.3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

SEE ALSO

New and Changed Features 178
Changes in Behavior and Syntax 186
Known Behavior 189
Known Issues 190
Resolved Issues 192
Documentation Updates 195
Product Compatibility 200

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 200](#)

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 178
Changes in Behavior and Syntax 186
Known Behavior 189
Known Issues 190
Resolved Issues 192
Documentation Updates 195
Migration, Upgrade, and Downgrade Instructions 196

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- [New and Changed Features | 202](#)
- [Changes in Behavior and Syntax | 218](#)
- [Known Behavior | 221](#)
- [Known Issues | 227](#)
- [Resolved Issues | 233](#)
- [Documentation Updates | 239](#)
- [Migration, Upgrade, and Downgrade Instructions | 240](#)
- [Product Compatibility | 253](#)

These release notes accompany Junos OS Release 17.3R2 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Caveat: Juniper Networks does not recommend configuring and deploying EVPN-VXLAN on QFX Series platforms running Junos OS 17.3R2.

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

New and Changed Features

IN THIS SECTION

- [Release 17.3R2 New and Changed Features | 203](#)
- [Release 17.3R1 New and Changed Features | 205](#)

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 17.3R2: QFX5100, QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016.

Release 17.3R2 New and Changed Features

EVPNs

- **EVPN-VXLAN with MPLS as transport layer (QFX10000 line switches)**—Starting with Junos OS Release 17.3R2, Ethernet VPN-Virtual Extensible LANs (EVPN-VXLANs) are supported with MPLS as the transport layer.

At present, QFX 10000 switches provide Layer 2 and Layer 3 VXLAN gateway functions for bare-metal server (BMS) or Virtual Machines (VMs) connected to it by means of a switch network or top-of-rack through an IRB interface. It also supports inter-DC connectivity via Type-5. The current transport layer support is IP. The feature adds MPLS as a transport for Layer 2 VXLANs with EVPN type-5 gateway functionality only. Layer 3 IRB VXLAN gateways will continue to use IP as the transport layer, even if MPLS is configured.

Interfaces and Chassis

- **Support for Static link protection on Aggregated interfaces (QFX5100 switches)**—Starting in Junos OS release 17.3R1, you can enable link protection on a specified static Label-Switched Paths (LSP). You can designate a primary and backup physical link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router. When the primary link fails, traffic is routed through the backup link.

IP Tunneling

- **IPv6 GRE tunneling support (QFX10002, QFX10008, and QFX10016)**—Starting with Junos OS Release 17.3R2, Junos OS support IPv6 Generic routing encapsulation (GRE) tunnels in QFX10000 line switches..

Management

- **Support for the Junos Telemetry Interface (QFX5110 switches)**—Starting with Junos OS Release 17.3R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various network elements without involving polling. On QFX5110 switches, only gRPC streaming of statistics is supported. UDP streaming is not supported.

The following sensors are supported:

- Chassis components
- Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol
- Network Discovery Protocol table state

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.](#)]

Multicast

- **Support for next-generation multicast VPN (QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 17.3R2, QFX10002, QFX10008, and QFX10016 switches support Multiprotocol BGP (MBGP) next-generation multicast VPNs with the following provider tunnel types:
 - Ingress replication provider tunnels
 - RSVP-Traffic Engineering (RSVP-TE) point-to-multipoint (P2MP) provider tunnels
 - Multipoint LDP P2MP provider tunnels

Virtual Chassis

- **Virtual Chassis support (QFX5200 switches)**—Starting in Junos OS Release 17.3R2, QFX5200-32C switches can be interconnected into a Virtual Chassis as one logical device managed as a single chassis. A QFX5200 Virtual Chassis can contain up to 3 members that must be QFX5200-32C switches, with no mixed mode support. Any non-channelized 40-Gbps QSFP+ ports can be configured as Virtual Chassis ports (VCPs) to interconnect member switches. As of Junos OS Release 17.3R2-S4, 100-Gbps QSFP28 ports can also be configured as Virtual Chassis ports (VCPs).

Configuration and operation are the same as for other QFX Series Virtual Chassis.

[See [Understanding QFX Series Virtual Chassis.](#)]

- **Virtual Chassis and Virtual Chassis Fabric (VCF) support (QFX5100-48T)**—Starting with Junos OS Release 17.3R2, QFX5100-48T switches are supported as part of a QFX5110 Virtual Chassis or VCF. As a result, a QFX5110 Virtual Chassis or VCF can contain a combination of QFX5110 switches with any of the following QFX5100 switches: QFX5100-24Q, QFX5100-48S, QFX5100-48T, and QFX5100-98S. A QFX5110 Virtual Chassis can contain up to a total of 10 devices, with any QFX5100 switches recommended to be configured in the linecard role only. A QFX5110 VCF can contain up to a total of 20 member devices, where spine members must be QFX5110-32Q switches, and any QFX5100 switches or QFX5110 switches can be leaf members.



CAUTION: Any QFX5100 switches running a “-qfx-5-” Junos OS software image *must* first be upgraded to a “-qfx-5e-” image (using the USB method) to successfully join a mixed QFX5110 Virtual Chassis or VCF.

[See [Understanding QFX Series Virtual Chassis](#) and [Understanding QFX Virtual Chassis Fabric Components.](#)]

Release 17.3R1 New and Changed Features

Class of Service (CoS)

- **Enhanced Transmission Selection (ETS) support (QFX10000 line switches)**—Beginning with Junos OS Release 17.3R1, ETS is supported on QFX10000 Series devices, compliant with IEEE 802.1Qaz/D0.1. ETS support enables the definition of multiple priority groups at each egress port of the device. Priority queues are combined into priority groups, enabling the application of similar congestion control capabilities to all queues within a group.

[See [Understanding CoS Hierarchical Port Scheduling \(ETS\)](#).]

EVPNs

- **Support of Layer 3 connectivity in an EVPN-VXLAN topology (QFX5110)**—Starting with Junos OS Release 17.3R1, you can deploy a QFX5110 switch as a Layer 3 Virtual Extensible LAN (VXLAN) gateway in an EVPN-VXLAN topology with a two-layer IP fabric or an IP fabric that is collapsed to one layer. In this role, the QFX5110 switch provides Layer 3 connectivity between physical (bare-metal) servers and virtual machines (VMs) within a data center. On QFX5110 switches, you can configure integrated routing and bridging (IRB) interfaces that route packets between VLANs. While creating an IRB interface, you can configure the interface as a default Layer 3 gateway, which physical servers in one VLAN use to communicate with physical servers or VMs in another VLAN.

[See [Example: Configuring a QFX5110 Switch as a Layer 3 VXLAN Gateway in an EVPN-VXLAN Topology with a Two-Layer IP Fabric](#) and [Example: Configuring a QFX5110 Switch as Layer 2 and 3 VXLAN Gateways in an EVPN-VXLAN Topology with a Collapsed IP Fabric](#).]

- **Support for multiple routing instances of type Virtual Switch and EVPN, VLAN-based service on the EVPN routing instance, and VLAN-aware service on the Virtual Switch routing instance (QFX10000 line switches)**—Starting with Junos OS Release 17.3R1, you can configure both EVPN and Virtual Switch routing instances. EVPN routing instance supports VLAN-based service. It includes only a single broadcast domain and there is a one-to-one mapping between a VNI and MAC-VRF. Up to 100 EVPN routing instances are supported. The Virtual Switch instance supports VLAN-aware service, and up to 10 Virtual Switch routing instances are supported. Each Virtual Switch routing instance can have up to 4094 VLANs, but the total number of VLANs across the Virtual Switch routing instances cannot exceed the system limitation.

NOTE: If you create VLANs that are not part of a routing instance, they become part of the Default Switch routing instance.

- **EVPN Proxy ARP and ARP Suppression (QFX10000 line switches)**—Starting with Junos OS Release 17.3R1, QFX10000 switches that function as provider edge (PE) devices in an 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](#).]

- **Support for external multicast router for EVPN with IGMP snooping (QFX10000)**—Starting with Junos OS Release 17.3R1, you can configure a provider edge (PE) switch running Ethernet VPN (EVPN) to send and receive multicast traffic to an external multicast router. This implementation supports the forwarding of inter-VLAN multicast traffic without having to configure IRB interfaces. Traffic is forwarded through a Layer 3 multicast protocol such as Protocol Independent Multicast (PIM). To enable the PE switch to receive multicast traffic from the multicast router, include the **multicast-router-interface** statement at the **[edit protocols igmp-snooping vlan *vlan-name* interface *interface-name*]** hierarchy level.

Support for forwarding inter-VLAN and intra-VLAN multicast traffic in an EVPN-VXLAN environment with IRB interfaces was introduced on QFX10000 switches in Junos OS Release 17.2R1.

[See [multicast-router-interface \(IGMP Snooping\)](#).]

- **Support for external Layer 3 multicast device for EVPN with IGMP snooping (QFX10000)**—Starting with Junos OS Release 17.3R1, you can connect an Ethernet VPN (EVPN) provider edge switch to an external Layer 3 device running a multicast protocol such as Protocol Independent Multicast (PIM). In this implementation, one or more provider edge switches configured with EVPN are connected to an external, that is, gateway, multicast device through a Layer 2 VLAN. To enable the PEs to forward traffic to the external domain, configure PIM-to-IGMP translation by including the **pim-to-igmp-proxy upstream-interface *irb-interface-name*** statements at the **[edit routing-options multicast]** hierarchy level. Additionally, this implementation supports configuring PIM on the IRB interfaces on the PE so that it functions only to forward inter-VLAN traffic within the data center. This means that you do not need to configure a PIM rendezvous point because forming PIM adjacencies is not required. The gateway device only needs to view the data center as a Layer 2 multicast domain. Include the new **passive** statement at the **[edit protocols pim]** hierarchy level to configure PIM to perform only inter-VLAN forwarding of multicast traffic.

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

General Routing

- **Commit process split into two steps (QFX Series)**—Starting in Junos OS Release 17.3R1, new configuration statements are introduced for **commit** to split the commit process into two steps. These configuration statements are **prepare** and **activate**.

In the first step, known as preparation stage, **commit prepare** validates the configurations and then creates the necessary files and database entries so that the validated configurations can be activated at a later stage.

In the second step, referred to as the activation stage, **commit activate** activates the previously prepared commit. A new configuration statement, **prepared**, is added to **clear system commit**, which clears the prepared commit cache

This feature enables you to configure a number of Junos OS devices and simultaneously activate the configurations. This approach is helpful in time-critical scenarios.

[See [Commit Preparation and Activation Overview](#).]

High Availability (HA) and Resiliency

- **Support for VRRP over IRB interfaces (QFX5100 Virtual Chassis and Virtual Chassis Fabric)**—Starting in Junos OS Release 17.3R1, you can configure Virtual Router Redundancy Protocol Version 3 (VRRPv3) for an IPv4 or IPv6 IRB interface on a QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF). The Virtual Chassis or VCF can act as the master or backup switch in a VRRP group, and the IRB interface forwards traffic sent to the configured VRRP virtual address that corresponds to the default gateway for the VLAN. Use the **vrrp-group** or **vrrp-inet6-group** configuration statement in the [edit interfaces irb unit *logical-unit-number* family (inet | inet6) address *address*] statement hierarchy on the Virtual Chassis or VCF as part of the IRB interface configuration.

[See [Configuring Basic VRRP Support for QFX](#) and [Configuring IRB Interfaces](#).]

Interfaces and Chassis

- **Increased number of link aggregation groups (LAGs) (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.3R1, you can configure up to 1000 LAGs on QFX10008 and QFX10016 switches. 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.
- **Short-reach mode (QFX5100-48T switch)**—Allows you to use short cable lengths (less than 10 meters) for copper-based 10-Gigabit Ethernet interfaces. Enabling short-reach mode reduces power consumption on these interfaces. You can configure short-reach mode for individual interfaces and for a range of interfaces. Enable short-reach mode for individual interfaces by including the enable statement at the [edit chassis fpc <slot-number> pic <slot-number>] hierarchy. Enable short-reach mode for a range of interfaces by including the enable statement at the [edit chassis fpc <slot-number> pic port-range <port low> <port high>] hierarchy.
- **IEEE 1588v2 Precision Time Protocol (PTP) Boundary Clock (QFX10002 switches)**—Starting with Junos OS Release 17.3R1, a boundary clock, which has multiple network connections, can act as a source

(master) or destination (slave) for synchronization messages. The boundary clock intercepts and processes all Precision Time Protocol (PTP) messages and passes all other traffic. The best master clock algorithm (BMCA) is used by the boundary clock to select the best clock from configured acceptable masters. You can configure a port as a boundary slave or as a boundary master. To configure a boundary clock, include the **boundary** statement at the **[edit protocols ptp clock-mode]** hierarchy level.

[See [IEEE 1588v2 PTP Boundary Clock Overview](#).]

- **Auto-channelization of interfaces (QFX5200 switch)**—Starting in Junos OS Release 17.3, you can use the auto-channelization feature to divide and channelize data automatically by detecting the cable type. The mode and number of channels are decided based on the channel link status. On QFX5200, auto-channelization supports three modes of operation with unique port settings:
 - When 4x10G split cables are connected, the 40G port auto-channelizes to four 10G channels.
 - When 2x50G split cables are connected, the 100G port auto-channelizes to two 50G channels.
 - When 4x25G split cables are connected, the 100G port auto-channelizes to four 25G channels.
- **Support for consistent load balancing for ECMP groups (QFX10000 line switches)**—Starting with Junos OS Release 17.3R1 on QFX10000 switches, you can prevent the reordering of flows to active paths in an ECMP group when one or more paths fail. Only flows that are on inactive paths are redirected. This feature applies only to Layer 3 adjacencies learned through external BGP connections. It overrides the default behavior of disrupting all existing, including active, TCP connections when an active path fails. Include the **consistent-hash** statement at the **[edit policy-options policy-statement policy-statement-name then load-balance]** hierarchy level. You must also configure a global per-packet load-balancing policy.

[See [Understanding Consistent Load Balancing Through Resilient Hashing on ECMP Groups](#).]

- **CL74 FEC support for 25-gigabit and 50-gigabit channel speeds (QFX5200 switches)**—Starting with Junos OS Release 17.3, you can disable or reenab le clause 74 (CL74)—as well as CL91—forwarding error correction (FEC) support on QFX5200 switches. FEC CL91 is supported for the 100-gigabit port speed and FEC CL74 is supported for both 25-gigabit and 50-gigabit port speeds. FEC CL91 is enabled by default for the 100-gigabit port speed; when the ports are channelized either in 4x25-gigabit or 2x50-gigabit, FEC CL74 is enabled.

- To disable the FEC mode:

```
[edit]
set interfaces interface-name together-options fec none
```

- To reenab le the FEC mode:

```
[edit]
delete interfaces interface-name together-options fec none
```

or

```
[edit]
  set interfaces interface-name gigether-options fec (fec74|fec91)
```

- To check FEC status:

```
show interfaces interface-name
```

The output for the show command will list FEC statistics for a particular *interface-name*, including the FEC corrected errors count, the FEC uncorrected errors count, and the type of FEC that was disabled or enabled.

Layer 2 Features

- **Support to exclude IRB Interfaces from state calculations (QFX5100)**—Starting with Junos OS Release 17.3R1, you can exclude a trunk or access interface from the state calculations for an IRB interface for member VLANs. An IRB interface typically has multiple ports in a single VLAN. Excluding trunk and access interfaces from state calculations means that as soon as the port specifically assigned to the VLAN goes down, the IRB interface for the VLAN is marked as down. Include the **autostate-exclude** statement at the **[edit interfaces ether-options]** hierarchy level. This feature was previously introduced in Junos OS Release 14.1X53-D40.

[See [Excluding an IRB Interface from State Calculations.](#)]

- **Increases number of vmembers to 256k for integrated routing and bridging interfaces and aggregated Ethernet interfaces (QFX10000 line switches)**—To calculate how many interfaces are needed to support 4,000 VLANs, for example, divide the number of vmembers (256,000 is the upper limit) by the number of configured VLANs (4,000). In this case, 64 interfaces are required.

Management

- **Enhancements to BGP peer sensors for Junos Telemetry Interface (QFX5110, QFX5200, and QFX10000)**—Starting with Junos OS Release 17.3R1, telemetry data streamed through gRPC for BGP peers is reported separately for each routing instance. To export data for BGP peers, you must now include the following path in front of all supported paths:

```
/network-instances/network-instance/[name_'instance-name']/protocols/protocol/
```

Additionally, the following paths are also now supported:

- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/accepted`
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/snmp-peer-index`
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/queues/output`
- `/network-instances/network-instance/protocols/protocol/`

bgp/neighbors/neighbor/afi-safis/afi-safi/state/queues/input

- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/state/ImportEval**
- **/network-instances/network-instance/protocols/protocol/
bgp/neighbors/neighbor/state/ImportEvalPending**

Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Support for LSP events and properties sensor for Junos Telemetry Interface (QFX5110 and QFX5200)**—Starting with Junos OS Release 17.3R1, 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 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. This sensor was previously supported only on QFX10000 switches, MX Series routers, and PTX Series routers.

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

- **Support for the Junos Telemetry Interface (QFX5110 switches)**—Starting with Junos OS Release 17.3R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various network elements without involving polling. On QFX5110 switches, only gRPC streaming of statistics is supported. UDP streaming is not supported.

The following sensors are supported:

- Chassis components
- Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol
- Network Discovery Protocol table state

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.](#)]

- **Support for the Junos Telemetry Interface (QFX5110)**—Starting with Junos OS Release 17.3R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various network elements without involving polling on QFX5110 switches. Only gRPC streaming of statistics is supported on QFX5110 switches. UDP streaming is not supported.

The following sensors are supported:

- BGP peers
- RSVP interface events
- Memory utilization for routing protocol tasks
- Label-switched-path events and properties
- Ethernet interfaces enabled with the Link Layer Discovery Protocol

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.

Support for the Junos Telemetry Interface was introduced on QFX10000 and QFX5200 switches in Junos OS Release 17.2R1.

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

Multicast

- **Support for static multicast route leaking for VRF and virtual-router instances (QFX5100 switches)**—Starting with Junos OS Release 17.3R1, you can configure your switch to share IPv4 multicast routes among different virtual routing and forwarding (VRF) instances or different virtual-router instances. 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 initially introduced in Junos OS Release 14.X53-D40.

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

MPLS

- **Support for Layer 2 circuit on aggregate interfaces (QFX10000 switches)**—Starting in Junos OS release 17.3R1, you can configure a Layer 2 circuit on aggregate interfaces. You can apply input and output VLAN tags for pop, swap, and push label operations on the VLAN-CCC interface. VLAN tags are applied when traffic is sent to and from the Layer 2 circuit interface. These operations are performed only on the outer TAG. The pop VLAN tag removes the VLAN tag from the top of the VLAN tag stack. The push

VLAN tag adds a new outer VLAN tag, and the swap VLAN tag replaces the existing outer VLAN tag with the new VLAN tag. This feature provides interoperability between Layer 2 services with a distinct VLAN at the local or remote end, or for instances where the Layer 2 service comes with a certain VLAN, but the remote peer has a different VLAN or no VLAN.

[See [CCC Overview](#) .]

- **VRF support in IRB interfaces in a Layer 3 VPN (QFX5100 and QFX5100 Virtual Chassis)**—Starting in Junos Release 17.3R1, you can configure IRB interfaces under virtual routing and forwarding (VRF) in a VPN Layer 3 network. IRB interfaces enable a switch to recognize which packets are being sent to local addresses so that they are bridged whenever possible and are routed only when needed. This same functionality applies, when IRB interfaces are part of routing instances or VRF. Virtual routing instances allows you to divide the switch into multiple independent virtual routers, each with its own routing table. This increases functionality by allowing network paths to be segmented without using multiple devices. Because traffic is automatically segregated, VRF also increases network security and can eliminate the need for encryption and authentication. Internet service providers often take advantage of VRF to create separate VPNs for their customers.

[See [Understanding Virtual Routing and Forwarding Tables](#) .]

- **Support for BGP MPLS-based Ethernet VPN (QFX10000 switches)**—Starting with Junos OS Release 17.3R1, you can use MPLS-based Ethernet VPN (EVPN) to route MAC addresses using BGP over an MPLS core network. An EVPN enables you to connect dispersed customer sites using a Layer 2 virtual bridge. As with other types of VPNs, an EVPN consists of customer edge (CE) devices (host, router, or switch) connected to a provider edge (PE) router or switch. The QFX10000 acts as a PE switch at the edge of the MPLS infrastructure. The switch can be connected by an MPLS Label Switched Path (LSP) which provides the benefits of MPLS technology, such as fast reroute and resiliency. You can deploy multiple EVPNs within a service provider network, each providing network connectivity to a customer while ensuring that the traffic sharing on that network remains private.

[See [EVPN Overview](#) .]

Operation, Administration, and Maintenance

- **Junos daemons to natively emit JSON output (QFX Series)**—Starting with Junos OS Release 17.3R1, the operational state emitted by the daemons is supported in JSON format as well as XML format. To configure JSON format, specify the following CLI command: **set system export-format state-data json compact**. To specify JSON format for specific command output, include **display json** in specific CLI commands.
- **Junos OpenConfig to support operational models for VLANs (QFX Series)**—Starting with Junos OS Release 17.3R1, support has been added for an OpenConfig YANG model for VLANs via the addition of **openconfig-vlan.yang**, revision 1.0.2. This provides a unified view for the network agent to retrieve operational state from JUNOS daemons for VLANs.

Port Security

- **MAC-limiting support (QFX10000 switches)**—Starting in Junos OS Release 17.3R1, you can configure MAC limiting on QFX10000 line switches. MAC limiting enhances port security by limiting the number of MAC addresses that can be learned within a VLAN. Limiting the number of MAC addresses protects the switch from flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). Flooding occurs when the number of new MAC addresses that are learned causes the Ethernet switching table to overflow, and previously learned MAC addresses are flushed from the table. The switch then reverts to flooding the previously-learned MAC addresses, which can impact performance and introduce security vulnerabilities.

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

[See [Understanding MAC Limiting and MAC Move Limiting for Port Security](#).]

- **IP source guard (QFX5100, QFX5110, QFX5200)**—Starting with Junos OS Release 17.3R1, you can configure the IP source guard access port security feature to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, it discards the packet.

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

Routing Protocols Policy and Firewall Filters

- **Flexible Ethernet Support (QFX10000 switches)**—Starting in Junos OS release 17.3R1, you can configure inet, inet6, or vlan-circuit cross-connect (CCC) connections on a physical or aggregate ethernet interface. This allows you to set different forwarding rules for tagged and untagged traffic on the same interface. For example, you can forward tagged packets over the l2circuit and route untagged traffic normally in the native vlan mode.

All logical devices that are under the flexible vlan tagging are identified by their vlan-id configuration. For untagged traffic, the association to the corresponding logical device is derived using the native vlan id configuration on the physical device. For traffic without a vlan tag, the default vlan id (native vlan id) is used to derive the layer2 domain.

Routing Protocols

- **Support for BGP Large Communities (QFX Series)**—Starting with Junos OS 17.3R1, BGP community is enhanced to support BGP large community that uses 12-byte encoding where the most significant 4-bytes encode autonomous system number or global administrator and the remaining two 4-bytes encode operator defined local values. Currently, BGP normal community (4-byte) and BGP extended community (6-byte) provide limited support for BGP community attributes after the introduction of 4-byte autonomous system number. Configure the large BGP community attributes under **[edit policy-options community community-name members]** hierarchy level and under **[edit routing-options static route route community]** hierarchy level with keyword **large** followed by three 4-byte unsigned integers separated by colons. The attributes are represented as large:autonomous system number:local value 1:local value2.
- **Support for segment routing for IS-IS (QFX5110 and QFX5200)**—Starting with Junos OS Release 17.3R1, 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 index-range** 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 **prefix-segment index number** statement at the **[edit policy options policy-statement policy-name then]** 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.

This feature was introduced on QFX5100 and QFX10000 switches in Junos OS Release 17.2R1.

[See [Understanding Source Packet Routing](#).]

- **BGP precision-timer support for reducing BGP hold-time (QFX5100, QFX5100 Virtual Chassis, QFX5110, QFX5200, QFX10000)**—Starting in Junos OS Release 17.3R1, you can use BGP precision timers to enable BGP sessions to send frequent keepalive messages with hold times as short as 10 seconds. The hold time is the maximum time allowed to elapse between successive keepalive messages that BGP receives from a peer. The default hold time is 90 seconds; the default frequency for keepalive messages is 30 seconds. More frequent keepalive messages and shorter hold times might be desirable in large-scale deployments with many active sessions. When you set a **hold-time** value to less than 20 seconds, we recommend that you also configure the **BGP precision-timers** statement, so that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the **precision-timers** statement is included, keepalive messages are generated in a dedicated kernel thread, thus helping to prevent BGP session flaps.

[See [precision-timers](#).]

- **Support for 128 equal-cost paths for BGP multipath (QFX10000)**—Starting with Junos OS Release 17.3R1, you can configure a maximum of 128 equal-cost paths for external BGP peers. Previously, the maximum number supported was 64. For MPLS routes, the maximum number of equal-cost paths you can configure remains unchanged at 64. To specify 128 equal-cost paths for external BGP peers, include the **maximum-ecmp 128** statement at the **[edit chassis]** hierarchy level. You must also configure a routing policy that exports routes from the routing table into BGP. Define a routing policy by including the **policy-statement policy-name** set of statements at the **[edit policy-options]** hierarchy level. Apply the policy to routes exported to the forwarding table by including the **export policy-name** statement at the **[edit routing-options forwarding-table]** hierarchy level.

[See [maximum-ecmp](#).]

NOTE: This feature is released but not supported in Junos OS Release 17.3R1.

- **Support for segment routing for OSPF (QFX5110 and QFX5200)**—Starting with Junos OS Release 17.3R1, 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.

This feature was introduced on QFX5100 and QFX10000 switches in Junos OS Release 17.2R1.

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

- **Support for alternate loop-free routes for IS-IS and OSPF (QFX10000)**—Starting in Junos OS Release 17.3R1, this feature adds fast reroute capability for IS-IS and OSPF. Junos OS precomputes loop-free backup routes for all IS-IS or OSPF routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. A loop-free path is one that does not traverse the router to reach a given destination. That is, a neighbor that already forwards traffic to the router is not used as a backup route to that destination.

You can enable support for alternate loop-free routes on any IS-IS or OSPF interface. To provide this support automatically for LDP label-switched paths (LSPs), you must also enable LDP on any interface for which you enabled support for loop-free alternate routes. In addition, you can extend backup coverage to include RSVP LSP paths.

Junos OS provides two mechanisms to enable fast reroute for IS-IS or OSPF using alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS or OSPF interface, the software creates an alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection. It supports fast rerouting of user traffic over one mission-critical link. Node-link protection establishes an alternate path through a different router altogether.

[See [Loop-Free Alternate Routes for OSPF Overview](#), [Example: Configuring Link and Node Protection for IS-IS Routes](#).]

- **Support for alternate loop-free routes for IS-IS and OSPF (QFX5110 and QFX5200)**—Starting in Junos OS Release 17.3R1, this feature adds fast reroute capability for IS-IS and OSPF. Junos OS precomputes loop-free backup routes for all IS-IS or OSPF routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. A loop-free path is one that does not traverse the router to reach a given destination. That is, a neighbor that already forwards traffic to the router is not used as a backup route to that destination.

You can enable support for alternate loop-free routes on any IS-IS or OSPF interface. To provide this support automatically for LDP label-switched paths (LSPs), you must also enable LDP on any interface for which you enabled support for loop-free alternate routes. In addition, you can extend backup coverage to include RSVP LSP paths.

Junos OS provides two mechanisms to enable fast reroute for IS-IS or OSPF using alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS or OSPF interface, the software creates an alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection.

It supports fast rerouting of user traffic over one mission-critical link. Node-link protection establishes an alternate path through a different router altogether.

[See [Loop-Free Alternate Routes for OSPF Overview](#), [Example: Configuring Link and Node Protection for IS-IS Routes](#).]

- **Support for BGP link-state distribution extensions for segment routing (QFX5110 and QFX5200)**—Starting in Junos OS Release 17.3R1, BGP link-state distribution extensions export segment-routing topology information to software-defined networking controllers. Although controllers can obtain the topology information by either being a part of an interior gateway protocol (IGP) domain or through BGP link-state distribution, the latter provides a more scalable mechanism for exporting this information. BGP link-state distribution is supported on inter-domain networks. This feature is useful in networks that are moving to segment routing at the transport layer but also have RSVP deployed. Include the **ipv4-prefix** statement at the **[edit policy-options policy-statement policy-name term term-name from traffic-engineering]** hierarchy level. This feature was introduced in Junos OS Release 17.2R1 on MX Series and PTX Series routers and on QFX5100 and QFX10000 switches.

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

- **Routing protocol process (rpd) recursive resolution over multipath (QFX Series)**—Starting in Junos OS Release 17.3R1, when a BGP prefix that has a single protocol next hop is resolved over another BGP prefix that has multiple resolved paths (unilist), all the paths are selected for protocol next-hop resolution. In prior Junos OS releases, only one of the paths is picked for protocol next-hop resolution. This new feature benefits densely connected networks where BGP is used to establish infrastructure connectivity such as WAN networks with high equal-cost multipath and seamless MPLS topology.

To configure recursive resolution over multipath, define a policy that includes the **multipath-resolve** action at the **[edit policy-options policy-statement policy-name then]** hierarchy level and import the policy at the **[edit routing-options-resolution rib rib-name]** hierarchy level.

[See [Configuring Recursive Resolution over BGP Multipath](#).]

Virtual Chassis

- **Virtual Chassis and Virtual Chassis Fabric (VCF) support (QFX5110)**—Starting with Junos OS Release 17.3R1, QFX5110 switches can be interconnected into a Virtual Chassis or VCF and operate as one logical device managed as a single chassis, as follows:
 - QFX5110 Virtual Chassis: Up to 10 members, all QFX5110 switches or in combination with QFX5100 switches. We recommend using QFX5110 switches in the master and backup Routing Engine roles, and QFX5100 switches only in the linecard role.
 - QFX5110 VCF: Up to 20 members, all QFX5110 switches or in combination with QFX5100 switches. Spine members must be QFX5110-32Q switches.
 - A QFX5110 Virtual Chassis or VCF can contain QFX5110-32Q, QFX5110-48S, QFX5100-24Q, QFX5100-48S, and QFX5100-98S switches. The same software image runs on QFX5110 or QFX5100 switches in a Virtual Chassis or VCF, and you do not need to configure the switches into mixed mode.



CAUTION: Any QFX5100 switches running a “-qfx-5-” Junos OS software image *must* first be upgraded to a “-qfx-5e-” image (using the USB method) to successfully join a mixed QFX5110 Virtual Chassis or VCF.

- Any (non-channelized) 100-Gbps or 40-Gbps QSFP28 ports, 40-Gbps QSFP+ ports, or 10-Gbps SFP+ ports can be Virtual Chassis ports (VCPs).

[See [Understanding QFX Series Virtual Chassis](#) and [Understanding QFX Virtual Chassis Fabric Components](#).]

SEE ALSO

Changes in Behavior and Syntax	 218
Known Behavior	 221
Known Issues	 227
Resolved Issues	 233
Documentation Updates	 239
Migration, Upgrade, and Downgrade Instructions	 240
Product Compatibility	 253

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\)](#) | [219](#)
- [EVPNs](#) | [219](#)
- [Management](#) | [219](#)
- [Network Management and Monitoring](#) | [220](#)
- [Virtual Chassis](#) | [221](#)
- [VLAN Infrastructure](#) | [221](#)

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 17.3R2 for the QFX Series.

Class of Service (CoS)

- When you configure a **transmit-rate**, you must also configure a **guaranteed-rate** at **traffic-control-profiles**. If you commit a configuration of a **transmit-rate** without a **guaranteed-rate**, a warning message is displayed and the default scheduler map is applied.

EVPNs

- On QFX10000 switches running Junos OS Release 17.3R3 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 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.

Management

- **Changes to custom YANG RPC syntax (QFX Series)**—Starting in Junos OS Release 17.3, custom YANG RPCs have the following changes in syntax:
 - The **junos:action-execute** statement is a substatement to **junos:command**. In earlier releases, the **action-execute** and **command** statements are placed at the same level, and the **command** statement is optional.
 - The CLI formatting for a custom RPC is defined within the **junos-odl:format** statement, which takes an identifier as an argument. In earlier releases, the CLI formatting is defined using a container that includes the **junos-odl:cli-format** statement with no identifier.
 - The **junos-odl:style** statement defines the formatting for different styles within the statement. In earlier releases, the CLI formatting for different styles is defined using a container that includes the **junos-odl:cli-format** and **junos-odl:style** statements.
- **Enhancement to show agent sensors command (QFX Series)**—Starting with Junos OS Release 17.3R1, the **show agent sensors** command, which displays information about Junos Telemetry Interface sensors, displays the default value of **0** for the **DSCP** and **Forwarding-class** values. Previously, the displayed default value for these fields was **255**. The default value is displayed when you do not configure a DSCP or forwarding-class value for a sensor at the **[edit services analytics export-profile *profile-name*]** hierarchy level.

[See [export-profile](#) and [show agent sensors](#).]

Network Management and Monitoring

- **Enhancement to about-to-expire logic for license expiry syslog messages (QFX Series)**—As of Junos OS Release 17.3R1, the logic for multiple capacity type licenses and when their expiry raises alarms was changed. Before, the behavior had alarms and syslog messages for expiring licenses raised based on the highest validity, which would mislead users in the case of a license expiring earlier than the highest validity license. The new behavior has the about-to-expire logic based on the first expiring license.
- **Change in default log level setting (QFX Series)**—In Junos OS Release, 17.3R2, 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](#).

- **SNMP syslog messages changed (QFX Series)**—In Junos OS Release 17.3R1, 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](#).]

Virtual Chassis

- **Adaptive load balancing (ALB) feature (Virtual Chassis Fabric)**—Starting in Junos OS Release 17.3R2, 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](#).]

VLAN Infrastructure

- **LAG interface flaps while adding/removing a VLAN**—From Junos OS Release 17.3 or later, the LAG interface flaps while adding or removing a VLAN. The flapping happens when a low speed SFP is plugged into a relatively high speed port. To avoid flapping, configure the port speed to match the speed of the SFP.

SEE ALSO

New and Changed Features 202
Known Behavior 221
Known Issues 227
Resolved Issues 233
Documentation Updates 239
Migration, Upgrade, and Downgrade Instructions 240
Product Compatibility 253

Known Behavior

IN THIS SECTION

- [Class of Service \(CoS\) | 222](#)
- [EVPNs | 222](#)
- [High Availability \(HA\) and Resiliency | 223](#)
- [Interfaces and Chassis | 223](#)

- [Layer 2 Features | 223](#)
- [Layer 3 Features | 224](#)
- [Platform and Infrastructure | 224](#)
- [Routing Protocols | 226](#)
- [Virtual Chassis | 226](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 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.

Class of Service (CoS)

- On QFX10000 line switches, oversubscribing all 8 queues configured with the **transmit rate exact** statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level might result in less than 100 percent utilization of port bandwidth.

[See [transmit-rate](#).]

EVPNs

- A provider edge (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 device. The IGP instance running in the VRF on the PE might 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 device. [PR977945](#)
- EVPN/VXLAN implementations support up to 100 EVPN VLAN-based routing instances. Above 100 instances, MAC learning may behave incorrectly. [PR1287644](#)
- A QFX10000 switch running Junos OS Release 17.3Rx software might experience a small and continuous traffic loss under the following conditions:
 - The switch is configured as a Layer 2 and/or Layer 3 VXLAN gateway in an EVPN-VXLAN topology with either a two-layer or collapsed IP fabric.
 - The switch has default ARP and MAC aging timer values.

Under these conditions, the following types of traffic flows might be impacted:

- Bidirectional Layer 3 traffic in a multihomed topology.

- Unidirectional Layer 3 traffic in a single-homed topology.

Note that this issue does not impact bidirectional Layer 3 traffic in a single-homed topology.

To prevent loss in these traffic flows, you must set the **aging-timer** configuration statement in the **[edit system arp]** hierarchy level so that the value is less than the value of the **global-mac-table-aging-time** configuration statement in the **[edit protocols l2-learning]** hierarchy level. [PR1309444](#)

- With VXLAN configured for 30 VXLAN VNIs, L3 Unicast traffic loss might be observed on deleting and adding back all the VXLAN VNIs. [PR1318045](#)

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an QFX5100 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](#)

Interfaces and Chassis

- On 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 required data. [PR980530](#)
- When commit a configuration change for irb from vrrp to non-vrrp and the irb address also changed to vrrp vip. Junos will lost direct route from the irb. This is a limitation, this issue was logged in other PR1191371. [PR1319124](#)

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging has been enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. [PR1075230](#)
- In EVPN-VXLAN deployment with QFX10k switches, when vxlan enabled IRB interface is configured in the same routing instance as that of the the underlay vtep tunnel and if the remote VTEP interface IP is resolved over the IRB interface using routing protocols or static route, dc-pfe cores would be generated and all the interfaces would go down. dc-pfe cores would be continuously generated until configuration is corrected. [PR1261824](#)
- When the replication tree used for flooding is reconverging, because some of the leaves have been deleted or added, there is expected to be some transient traffic loss even in leaves that have not changed. This affects only flooding and BUM traffic, not known unicast traffic. [PR1274950](#)
- When NG-MVPNis configured with RSVP provider tunnels and NSR is used, then theegress router for the tunnel might not correctly replicate some ofthe tunnel state to the backup routing engine, leading to temporarytraffic loss during NSR failover for the affected tunnels. [PR1293014](#)

Layer 3 Features

- Uneven load balancing of traffic might occur if the traffic stream changes only in the bits 0-15 of the Layer3 destination IPv6 address. This limitation may not be visible if the other parameters effecting the load balance change along with L3_DST, such as L3 source IP address, L4 source/destination ports and so on. [PR1065515](#)

Platform and Infrastructure

- Chef for Junos OS supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev resources. The netdev resource developed for interface configuration has a limitation to configuring the XE interface. The netdev interface resource determines that speed is a configurable parameter that is supported on a GE interface but not on an XE interface. Hence, the netdev interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages chef-11.10.4_1.1.*.tgz chef-11.10.4_2.0.*.tgz in all platforms {i386/x86-32/powerpc}. [PR1181475](#)
- 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](#)
- In certain interface scaling scenarios, during configuration commit/rollback, you might see an fpcx error message. You can safely ignore this message because of the FPGA monitor mechanism on DPC cards for logical interface mapping (ifl_map). Between the deletion of a physical interface and the monitoring event, this 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 the ifl_map, harmless messages might populate the log file. [PR1210877](#)
- The ptp master streams on IP and Ethernet are not supported simultaneously. [PR1217427](#)
- There is no unified ISSU from Junos OS Release 15.1 and earlier releases to Junos OS Release 16.2R1. [PR1222540](#)
- On QFX10008 switches, if you reboot a QFX10000-36Q line card or a QFX10000-30C line card with traffic running, sometimes framing errors are displayed in the CLI output. This is only a display issue. No actual framing errors have occurred, and traffic is unaffected. [PR1223330](#)
- For a LAG interface, PFE populates only the bundle statistics and not the child's IFL statistics. It always returns zero for IFL statistics. There is a limitation in the hardware which restricts the per IFL stats [PR1250870](#)
- If port speed is changed in from 25G to 100G or there are repeated changes in port speed settings, then the link may remain down. [PR1250891](#)
- To release note. User needs to restart PFE for changes to take effect. [PR1256465](#)

- 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](#)
- There is no "convenient" bit/register to read to tell if there is a false link UP going on. Again, insist to Juniper that this is not a supported scenario, then there are no provisions to address it. This is how the protocol works out, unfortunately, when you have a 40G port mistakenly connected to a 10G port. Pls. monitor bit 3.1.2 and *also* reg 1.cd09. Register 1.cd09 provides protocol code errors. When real packets come in, bits [6:4] can get set because the traffic coming into the 10G port is actually 40G traffic. These bits can also get set even without traffic going on, but only idles. So, have them test with these two registers (3.1 and 1.cd09) under their environment, so that they can declare link UP only when both registers have good values ("good" 1.cd09 value is 0x0200; "r64_sync_acq" bit). So, if 1.cd09 is different than 0x200, then you could declare link DOWN. IMPORTANT NOTE: once again, this is not a supported scenario. The IEEE standard does not address this case and doesn't require vendors to address this case either. The above workaround is our suggestion but it could potentially create other problems. Juniper's should study, evaluate and test the above idea before integrating it in the design. [PR1264489](#)
- This issue occurs when an interface comes online and both OAM protocol and MKA protocol try to establish their respective sessions. Because of contention between these two protocols, OAM takes down the interface and MKA fails to establish connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
- Device might not power up when crossover cable is used. It is advised to use straight cables. [PR1274613](#)
- Multiple instances of the DAEMON-3-JTASK_SCHED_SLIP system message may be logged when over 50,000 MACs are configured and the device attempts to establish OSPF neighbors. This has no functional impact. [PR1274706](#)
- shared-buffer maximum default for IFL Queues is 66%, independent of the shared-buffer maximum knob under IFL scheduler config. [PR1275796](#)
- On a QFX5110-32C switch, if a splitter cable is connected to a Spirent 10G CV/MX card, ports will not come up due to varied pre-empt settings for the splitter and DAC cables. There is a hardware limitation where we have no way in EEPROM to differentiate between splitter and DAC cable to apply different settings. As a workaround, use a 40G Spirent card with internal channelization on the Spirent side and manual channelization on the QFX5110-32C side. [PR1280593](#)
- On QFX10000 switches implementing EVPN-VXLAN, if the Routing Engine is repeatedly restarted on redundant gateways, then inter-vrf traffic will be dropped without notification. [PR1289091](#)
- ERPS convergence takes time after GRES switchover and hence traffic loss is observed for a brief period. [PR1290161](#)
- Traffic drop occurs on sending traffic over "et" interfaces due to CRC errors. [PR1313977](#)
- For QFX5110, there is a hardware limitation. QFX5110 can route from VxLAN (VFI) domain to VxLAN (VFI) domain only, does not support routing from VxLAN domain to non-VxLAN domain. [PR1318178](#)

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](#)
- A QFX5110 switch running Junos OS Release 17.3R1 or later software functions as both a Layer 3 VXLAN gateway and a DHCP relay in an EVPN-VXLAN topology. After a DHCP client receives and later releases an IP address on an EVPN-VXLAN integrated routing and bridging (IRB) interface configured on the QFX5110 switch, the binding between the DHCP client and the IP address might not be deleted. As a result, the next time that the DHCP client requests an IP address, the response from the DHCP server might take a few minutes. [PR1261483](#)
- QFX5110: Traffic loss of routed packets might be seen through a non-collapsed EVPN-VxLAN L3 GW, when disjoint VxLANs with IRB are provisioned and unprovisioned in bulk on it." [PR1276423](#)
- An adjacency segment identifier will not be created for IPv6-only configured interfaces. If the adjacency uses IP alone or IP+IPv6, then an IPv4 adjacency segment identifier or IPv6 adjacency segment identifier will be created. If the adjacency only uses IPv6, then no adjacency segment identifier will be created. [PR1290515](#)
- A QFX10000 switch running Junos OS Release 17.3Rx or 17.4Rx software might experience a small and continuous traffic loss under the following conditions: 1) The switch is configured as a Layer 2 and/or Layer 3 VXLAN gateway in an EVPN-VXLAN topology with either a two-layer or collapsed IP fabric, and 2) The switch has default ARP and MAC aging timer values. Under these conditions, the following types of traffic flows might be impacted: 1) Bidirectional Layer 3 traffic in a multihomed topology, and 2) Unidirectional Layer 3 traffic in a single-homed topology. Note that this issue does not impact bidirectional Layer 3 traffic in a single-homed topology. [PR1309444](#)

Virtual Chassis

- For a large VC, topology hash might have a good impact on VC stability as it reduces programming by skipping some route for intermediate topologies. However, it could delay traffic switch as we observed. By default, topology hash is on. There is hidden cli (**set virtual-chassis no-topology-hashF**) to turn it off. [PR1296196](#)
- L2/L3 traffic drop is seen after rebooting whole VC (10 member) or changing VC member list (for example, making 6 VC member from 10 VC, back to 10 member VC). [PR1314429](#)

SEE ALSO

[New and Changed Features | 202](#)

[Changes in Behavior and Syntax | 218](#)

[Known Issues | 227](#)[Resolved Issues | 233](#)[Documentation Updates | 239](#)[Migration, Upgrade, and Downgrade Instructions | 240](#)[Product Compatibility | 253](#)

Known Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 228](#)
- [EVPN | 228](#)
- [Infrastructure | 229](#)
- [Interfaces and Chassis | 229](#)
- [Layer 2 Features | 229](#)
- [Multicast | 230](#)
- [Network Management and Monitoring | 230](#)
- [Platform and Infrastructure | 230](#)
- [Routing Protocols | 232](#)

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.3R2.

Class of Service (CoS)

- On QFX5110-32C switches, throughput as per RFC 2544 is not 100 percent for some of the frame sizes when the switch is configured with mixed 10/40/100G speed ports. It is fine when tested individually with 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet ports separately. [PR1256671](#)

EVPN

- On QFX10000 line switches, jprds_dlu_alpha_add : 222 JPRDS_DLU_ALPHA KHT addition failed. [PR1258933](#)
- On QFX10000 line switches, sub-interfaces from the same physical port do not work if they are configured under the same VLAN or routing-instance. An attempt to commit such a configuration will fail for layer 2 configurations but not for EVPN/VXLAN. For EVPN/VXLAN configurations, there may be circumstances where it would be necessary to configure a sub-interface from the same physical port to support VLAN bundling [PR1278761](#)
- A new option **exclusive-mac** is added under **protocols l2-learning global-mac-move set protocols l2-learning global-mac-move exclusive-mac <mac>**. [PR1285749](#)
- When a VLAN uses an IRB interface as the routing interface, the vlan-id parameter must be set to "none" to ensure proper traffic routing. This issue is platform independent. [PR1287557](#)
- Using EVPN-VXLAN, VLAN-ID none should be used for VLANs or routing instances using an IRB as the routing interface. [PR1287565](#)
- QFX10000 line switches might drop DHCP offer unicast packets in an EVPN-VXLAN. [PR1299143](#)
- In an EVPN VXLAN scenario, a previous learned MAC address from a remote Ethernet segment Identifier (ESI) cannot be changed to local even it is connected directly. The MAC address of the host might remain as learned from ESI instead of local interface until the MAC address is aged out. [PR1303202](#)
- With VXLAN configured for 30 VXLAN VNIs, L3 Unicast traffic loss may be observed on deleting and adding back all the VXLAN VNI's. [PR1318045](#)
- There might be a traffic loss on ingress PE if the EVPN MPLS is configured latter on remote PE or from the working condition EVPN MPLS is disabled and enabled latter. [PR1319770](#)
- In an EVPN-VXLAN environment, the control-plane stress test consists of four QFX10000s as leaf nodes and four QFX5200s as TORs. These switches connect together in a typical configuration by servicing a combination of 4000 VLANs, IRBs, and VNIs. The total number of MAC addresses are 174,000. The stress test is performed by restarting routing protocol process (rpd) on each of the leaf nodes during every 90 second intervals. Eventually, one of the leaf node's rpd crashes. The core files could be seen by executing the CLI command **show system core-dumps**. [PR1320408](#)
- In EVPN-VXLAN topology, when an aggregate Ethernet (ae) interface disabled, there might be partial traffic dropped flowing through a given AE interface till the disabled AE interface comes back up. [PR1321269](#)

- Ultimat16: EVP-VXLAN: Traffic loss is seen because ARP is not synced after ARP timer expiry. [PR1322288](#)
- After deleting a VXLAN configuration and before adding it back, you need to wait for the remote VTEP aging timer to expire. The duration to wait is specified in the following command (or default timeout is 600 seconds): **set groups qinvxlan vlans VXLAN1 vxlan unreachable-vtep-aging-timer <timeout in seconds>.** [PR1322673](#)
- On the QFX10016 EVPN-VXLAN scaled testbed; it takes up to 3 minutes for traffic to converge when configuration related to a Tenant (5 IRBs/VLAN) is added. [PR1323042](#)

Infrastructure

- When the configuration statement **set system ports console log-out-on-disconnect** is enabled, the Junos OS eventd process (daemon) blocks the console-open(). However, during this stage with the syslog console configured (always logs on console), any logging continues even if the console session is ended. When the console logging continues to be in the waiting status, the eventd syslog rotation freezes and some processes directly involved in logging in to the system would also go into the wait status, causing undesirable behavior. [PR1253544](#)
- The QFX5110 might exhibit unexpected behavior or get stuck during the boot process. This is caused by noise on the console. Solution: If the switch gets stuck during the boot process, type "autoboot". Hard reboot if the above does not work. [PR1318168](#)

Interfaces and Chassis

- On QFX5100 switches, with MAC and ARP inside an IFA block, an error message that says an IRB interface and an AE 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

- When using PTP BC applications on QFX10002, the forwarding path for a directly connected device is not automatically present and is not triggered by the PTP packets generated by the QFX10002. As a workaround, either create the forwarding entries by configuring a routing protocol such as OSPF on the interface or add a static ARP entry for the remotely connected PTP device. [PR1275327](#)
- When there are multiple logical units on a lag (ae) interface, ingress pop might not work when the configuration is changed on the interface and rolled back. [PR1331722](#)
- When we do local switching between two ports on the same VXLAN, with QinQ configured on the access ports, the push operation may not work in 17.3R2. [PR1332346](#)

Multicast

- With IGMP snooping enabled in EVPNvVXLAN configuration, when a large number of IGMP leaves are received at the same time, some of the leaves might not get processed, resulting in the IGMP group state lingering till it eventually times out on group membership interval timeout. [PR1327980](#)

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 1. To help debug physical link UP issues, SNMP_TRAP_LINK_UP events will be logged by default. [PR1287244](#)

Platform and Infrastructure

- On QFX10002 switches, the **request system snapshot** command does not work. [PR1048182](#)
- While using SSH to log in to a VNF with the error message "Unrecognized command is seen." This error has no impact on the functionality. [PR1108785](#)
- On a QFX5100 Virtual Chassis, the MAC address is not learned on an AE interface configured as a VXLAN Layer 2 port and with the interface mode configured as access. The issue is observed only with AE interfaces that span multiple Virtual Chassis members and when the member node is rebooted or power cycled. [PR1112790](#)
- L3 multicast traffic does not converge to 100 percentage and a few continuous drops are observed after bringing an interface down and back up again or while an FPC comes online after FPC restart. This behavior is seen when scaling beyond 2000 VLANs or 2000 IRBs with VLAN replication configured. [PR1161485](#)
- When per-packet load balancing is removed or deleted, next hop index may change. [PR1198092](#)
- ICCP session is maintained by multihop BFD (non-distributed mode). The time interval for BFD keepalive messages is similar to GRES configuration (for example, keepAlive = 8 seconds). [PR1230576](#)
- 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](#)
- On QFX5100-48T with short-reach mode enabled on copper ports, these copper ports will flap when you commit any configuration related to routing instances. [PR1248611](#)
- On the QFX10000-12C-DWDM coherent line card, it is possible that sometimes the link flaps when MACsec is enabled on Ethernet interfaces. [PR1253703](#)
- PDT:UC1113: kernel: Next hop index allocation failed. Regular index space is exhausted. [PR1254755](#)
- 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 will not impact any

functionality. There is no action needed from the user side in response to the crash. As a workaround, install Openconfig before installing Network Agent. [PR1265815](#)

- Flexible-vlan-tagged interface allows both primary and secondary VLAN configuration on different logical units of the same interface, but might not work as expected. [PR1267160](#)
- This issue is applicable to all Virtual Chassis and Virtual Chassis Fabric combinations on the QFX5100, QFX5110, QFX3500, EX4300, and EX4600 platforms. If the **reboot** option is used with a large Virtual Chassis, some members might not be able to reboot. As a result, some members will still be running the old image and some members will be upgraded to the new image; this causes Virtual Chassis instability. As a workaround, upgrade the image using the **request system software add** command WITHOUT using the **reboot** option in Virtual Chassis and Virtual Chassis Fabric setups. [PR1273271](#)
- No CPU usage is shown in output for show chassis fpc x (x= QFX5100) in mixed Virtual Chassis Fabric. CPU utilization values show 0, because the values are being normalized. CPU utilization value increases if the idle time decreases to some extent. [PR1274665](#)
- A hostname synchronization issue occurs between the Junos OS VM instance and the Linux host on TVP platforms. [PR1283710](#)
- On QFX5100 switches, static LAG link protection switchover/revert is not working consistently. [PR1286471](#)
- On QFX10000 line platforms, with a high scale of 4000 VNIs or 200K MACs or both, if large configuration change happens with traffic flowing, then forwarding descriptor memory corruption might occur, leading to complete traffic loss on certain ports. The qualification shows that a system with 400 VNIs has been stable. However, other configurations like global MAC count and underlying MPLS LSPs can increase system load. [PR1296089](#)
- When link protection with the backup port state "down" and LACP are both configured, sometimes the primary port state becomes down without a trigger event and the backup port comes up and begins handling traffic. [PR1297596](#)
- When link protection with the backup port state "down" and LACP are configured, if the backup state "down" is removed from the configuration, what should happen is that both ports will be up and the primary port should pass all egress traffic. In some instances, however, traffic might pass through the backup port instead of the primary port. [PR1297597](#)
- Rebooting the leaf node configured with ESI might see the traffic drop at the receivers attached to the remote PE in L2L3 non-collapsed topology. [PR1319240](#)
- Traffic drop occurs on sending L3 traffic across MPLS LSP. [PR1313977](#)
- If a customer uses 100G-PSM4 optics in ULC30C line card and the firmware on the line card is listed in the following message, they need to upgrade to the new firmware. FPC 4 U-Boot Bank A: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52) CTRL FPGA 3.2 <<<<<< PORT FPGA 2.1 <<<<<< New versions below - FPC 4 U-Boot Bank A: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52) CTRL FPGA 3.3 PORT FPGA 2.2 [PR1323321](#)

- The management process (mgd) might panic after modifying AE interface members under "ethernet-switching vlan" stanza. After mgd panic, your remote session is terminated as a result. As a workaround, use commit full. [PR1325736](#)
- Traffic coming from the remote vtep PE will be dropped on a local PE, if any vlan that the remote PE is participating is deleted on the local PE. [PR1338532](#)
- Dcpfe will crash when multicast index is greater than 8191 due to array out of bound access. [PR1340053](#)

Routing Protocols

- On QFX10000 line switches, traffic drop is seen with IS-IS version 6 traffic during convergence in either of the following two scenarios:
 - 1. While doing port unshutdown (that is, bringing up the ports after bringing them down).
 - 2. While FPC comes online after doing an FPC restart. This behavior is seen while flapping one of the IS-IS version 6 sessions. [PR1190180](#)
- On QFX10000 line platforms, during route next-hop churn or earliest deadline first (EDF) job priority changes, memory corruption might occur, leading to processing issues and constant packet drop. [PR1243724](#)
- When switchover and zeroize are done in succession quickly. As "zeroize" will delete the databases, and if dfwd is going to start SIHUP processing after the zeroize, it will core as database is not present. Zeroize should be done when the system is in stable state that is, signups processing by daemons is completed. [PR1262385](#)
- On QFX5110 switches, an EVPN-VXLAN configuration using a custom-IRB MAC (same IP, same MAC profile) might not work. As a workaround, use a virtual gateway address, use a virtual-gateway address is recommended. [PR1291406](#)
- Performing GRES on the EVPN-VXLAN topology with uRPF results in total packet loss. [PR1322217](#)
- BGP as protocol strongly recommends configuration of local-address for each Multihop iBGP/eBGP peer configuration. As a recommendation local-address should be route-able lo0 address. Using loopback address reduces dependency with interfaces. Note: Multihop is by default enabled for iBGP Peers. [PR1323557](#)
- On QFX5200 Virtual Chassis, traffic loss of 0.04 percent is seen with Routing Engine switchover for the GRE tunnel scale test. [PR1323884](#)
- The dcpfe core file are generated when the master Routing Engine switchover with traffic or with the deletion of the scale configurations for the GRE test. [PR1327535](#)
- After the UFT profile change and commit the changes are not synced across all the VC members. This issue is seen intermittently one out of four try. [PR1334646](#)
- When IP move on MH host, on remote PE one or more new ARP might be missing. [PR1340051](#)

SEE ALSO

[New and Changed Features | 202](#)[Changes in Behavior and Syntax | 218](#)[Known Behavior | 221](#)[Resolved Issues | 233](#)[Documentation Updates | 239](#)[Migration, Upgrade, and Downgrade Instructions | 240](#)[Product Compatibility | 253](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2 | 233](#)

- [Resolved Issues: 17.3R1 | 237](#)

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: 17.3R2

Hardware

- The 1G copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms.
[PR1286709](#)
- ULC-60S-6Q LC on QFX10008: The port becomes unusable after inserting a third-party SFP-T optic.
[PR1294394](#)

Class of Service (CoS)

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

EVPNs

- Next-hop installation error messages are seen on QFX10000 line switches. [PR1258930](#)
- QFX10K2 VXLAN with MPLS underlay seen traffic loss at RSVP egress [PR1289666](#)
- On QFX5100 switches with EVPN-VXLAN deployed, broadcast and multicast traffic might not be sent to other switches through VTEP interfaces. [PR1293163](#)
- On QFX10000 switches with EVPN deployed, packet corruption is seen with Packet Forward Engine trap code (129) `egp.v4_chksum` when sending L3 inter-VNI traffic with the underlay vlan-tagging inet interface. [PR1295491](#)
- `df-election-type` preference statements in the `[show interfaces esi]` hierarchy level are not supported on QFX10000 running Junos OS Release 17.3R1. [PR1300093](#)
- The dynamic routing protocols might not work correctly over the IRB interface in an EVPN-VXLAN scenario with ECMP. [PR1301521](#)
- RPD crash on loading EVPN configurations in `qfx10002-72q`. [PR1305440](#)
- EVPN Proxy ARP might work properly. [PR1312672](#)

Interfaces and Chassis

- Multicast data packets are looping in MC-LAG. [PR1281646](#)
- ARP reply drop in MC-LAG scenario. [PR1282349](#)
- On QFX5100 switches, an AE interface might flap upon commit if an explicit speed is configured on an AE member interface. [PR1284495](#)
- Traffic might not be received on a 1-Gigabit Ethernet interface if autonegotiation is disabled and speed/duplex is configured on both the QFX Series switch and the peer host. [PR1292275](#)
- The 40-Gigabit Ethernet interface might not come up if a specific vendor's DAC cable is used. [PR1296011](#)
- 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 with /31 subnet in server configuration. [PR1298234](#)
- The `dcpe` process might crash and restart on MC-LAG active and standby nodes when there is ARP/NDP next-hop change. [PR1299112](#)
- Disabled 10-Gigabit Ethernet interfaces might stay up on QFX10000 line switches. [PR1300775](#)
- QFX10008/10016: Commit error is seen when configured with mixed speed. [PR1301923](#)
- On QFX5100/5110/5200 devices, IGMP snooping entries are not learnt on MCLAG peer. [PR1302620](#)

- QSFP+4x10G-IR channelized interface down between QFX5200 and PTX5000 [PR1307400](#)
- Upgrading to 16.1R5 without “redundancy-group-id-list” statement prior in ICCP leads to commit failure during bootup. [PR1311009](#)
- Core link flap might result in an inconsistent global MAC count. [PR1328956](#)

Layer 2 Features

- Feature swap-swap might not work as expected in a Q-in-Q scenario. [PR1297772](#)
- Device transmits packets that exceed the interface MTU. [PR1306724](#)
- NLB heartbeat packets might be dropped on QFX10K/PTX. [PR1322183](#)
- The DHCP Discover packets might be looped in MC-LAG and DHCP-Relay scenario. [PR1325425](#)

Layer 3 Features

- QFX5110-48S: L3 VPN traffic is dropped for some instances when EVPN-VXLAN configuration is removed and reapplied. [PR1307590](#)

Management

- QFX5110-48S: digital optical monitoring statistics cannot be received through the CLI in Junos OS Releases 15.1X53 through 17.x. [PR1305506](#)

MPLS

- QFX5100: ISSU is not supported with MPLS configuration. [PR1264786](#)
- 17.3: U8: QFX10008 is dropping the egress MPLS traffic, if the egress interface is an IRB with access L2 AE interface. [PR1279827](#)
- DHCP clients cannot get IP addresses over BGP-L3VPN. [PR1303442](#)
- LSP stop transferring/passing traffic after MPLS route is changed. [PR1309058](#)
- MPLS forwarding might not happen properly for some LSPs. [PR1319379](#)

Network Management and Monitoring

- UFT for non-local member is not shown in the CLI. [PR1243758](#)
- MACsec issue: "show security macsec statistics" command does not show expected results. [PR1283544](#)
- SNMP process is not running on QFX Series switches with incorrect source addresses. [PR1285198](#)

Platform and Infrastructure

- Traffic loss might be observed for about 10 seconds if master member FPC reboots. [PR1283702](#)
- QFX10002 and QFX10008: BFD sessions over IRB interfaces with Junos OS Releases 17.1R1, 17.1R2, 17.2R1, and 17.3R1 are centralized. [PR1284743](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)

- Storm-control flags are not set after a Routing Engine switchover. [PR1290246](#)
- OSPFv3 authentication using IPsec SA does not work if you are using IPsec to authenticate OSPFv3 neighbors on some QFX Series platforms. [PR1301428](#)
- The sflow records are missing "extendedType ROUTER" fields as well as an outbound interface for traffic that is using BGP multipath. [PR1303236](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- JDISwitchingReg : Traffic loss is observed while performing NSSU. [PR1311977](#)
- CPU utilization is around 50% without any configuration. [PR1312520](#)
- On QFX5200 Virtual Chassis, 100G port VCP not supported. [PR1314922](#)
- Transit traffic over GRE tunnel might hit the CPU and trigger a DDoS violation on the L3 next hop. [PR1315773](#)
- On an L2 next-generation switch platform (EX4300/EX4600/EX9200/QFX3500/QFX3600/QFX5100/QFX10000), l2cpd might drop core files repeatedly if an interface is connected to a VoIP product with LLDP and LLDP-MED enabled. [PR1317114](#)
- Port speed is still showing 100G instead of 50G because the physical interface (IFD) has been channelized to 50G. [PR1319884](#)
- FPCs go offline due to the error **CHASSISD_IPC_CONNECTION_DROPPED: Dropped IPC connection for FPC**. [PR1321198](#)
- EVPN Type 5: Unicast traffic is getting dropped on the backup forwarder. [PR1323907](#)
- QFX5100/EX4600/ACX5k : Major Alarm 'Fan & PSU Airflow direction mismatch' by removing management cable. [PR1327561](#)
- QFX10002: Major alarm should be cleared once the chassis has more PEM units installed than the "minimum PEM" configuration. [PR1327999](#)

Port Security

- Proxy-ARP and ARP suppression are not yet supported for the QFX10000 line. [PR1293707](#)

Routing Policy and Firewall Filters

- The rpd might crash if **vrf-target auto** is configured under routing-instance. [PR1301721](#)

Routing Protocols

- 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](#)
- FBF with next-ip/next-ip6/next-interface is not working. [PR1289642](#)
- Remotely received traffic is not flooded to AC on FPC 1 when FPC 0 was offlined. [PR1290500](#)
- IPv6 multicast traffic drop occurs in PIM SSM scenario. [PR1292519](#)

- On QFX5100, the fxpc process generates a core file. [PR1294033](#)
- The dcpfe process might crash after a period of idle time on QFX10000 switches. [PR1294055](#)
- 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](#)
- Observed **mcsnood** 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](#)
- Packets drop is seen when programming for GRE traffic. [PR1308438](#)

Software Installation and Upgrade

- After upgrading the QFX5100-96s-8q to Junos OS Release 16.1R4 from Junos OS Release 15.1R4, showing commit warning `"/boot/ffp.cookie+`". [PR1283917](#)

Virtual Chassis

- QFX5100 TVP: Not able to load TVP image on top of a non-TVP 5100 image while adding a QFX5100 switch to the Virtual Chassis. [PR1248145](#)
- QFX5200: New apply group not applying to the Virtual Chassis after a reboot. [PR1305520](#)
- QFX-VC: Sometimes seeing that Multicast packets received 2x 3x times than expected. [PR1306239](#)
- QFX5110 VC/VCF: VC members reboot before all members have image installed. [PR1309103](#)
- Some log messages are seen on the QFX5110 platform when plugging in an SFP-SX. [PR1311279](#)

Resolved Issues: 17.3R1

General Routing

- On QFX10000 line 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](#)
- VLAN association is not being updated in the Ethernet switching table when the device is configured in single supplicant mode. [PR1283880](#)
- Hostname synchronization from Junos VM instance to Linux Host in TVP Platforms (QFX). [PR1283710](#)

Interfaces and Chassis

- Interfaces randomly do not come up after a line card restart. [PR1262839](#)
- On QFX5100 switches, a 40G interface may keep flapping when a 5M DAC cable is inserted. [PR1273861](#)

- On QFX10000 switches, there may be an ot- link flap whenever there is an optics TCA alarm, however there is no loss of signal and no traffic loss observed. [PR1279351](#)
- FEC disabled by default on 100G-LR optics for QFX5200 switches. [PR1286389](#)

Layer 2 Features

- All the XML duplications and unformatted output are addressed. For Example, histogram was just declared as a element inside pfkey container, with this change a new container is defined for histogram. [PR1271648](#)

Port Security

- On QFX10000 switches, MACsec sessions are not coming up on a Layer 3 sub-interface. [PR1282995](#)

Routing Protocols

- When static Link protection mode configured back up state as down, primary port is going to down state instead of secondary port while secondary is at up state. [PR1276156](#)
- UDP traffic with destination port 520 and 521 is discarded on QFX5110 switches after a Junos OS upgrade. [PR1287271](#)
- 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 may be lost during MAC moves. [PR1291118](#)

System Management

- 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](#)

VXLAN

- Two new CLI commands are added: **set forwarding-options vxlan-routing next-hop *number*** ; **set forwarding-options vxlan-routing interface-num *number***. These commands are applicable only for QFX5110 switches. [PR1259323](#)

SEE ALSO

[New and Changed Features | 202](#)

[Changes in Behavior and Syntax | 218](#)

[Known Behavior | 221](#)

[Known Issues | 227](#)

[Documentation Updates | 239](#)

[Migration, Upgrade, and Downgrade Instructions | 240](#)

[Product Compatibility | 253](#)

Documentation Updates

IN THIS SECTION

- [Traffic Management User Guide for the QFX Series | 239](#)

This section lists the errata and changes in Junos OS Release 17.3R2 for the QFX Series switches documentation.

Traffic Management User Guide for the QFX Series

- **Consolidation of the Traffic Management User Guide for QFX Series and EX4600 Switches (QFX Series)**—Starting in Junos OS Release 17.3R1, the following three traffic management guides are consolidated into one user guide:
 - Traffic Management User Guide for QFX Series
 - Traffic Management User Guide for QFX 10000 Series
 - Traffic Management User Guide for EX4600 Switches

[See [Traffic Management User Guide for QFX Series and EX4600 Switches.](#)]

SEE ALSO

New and Changed Features 202
Changes in Behavior and Syntax 218
Known Behavior 221
Known Issues 227
Resolved Issues 233
Migration, Upgrade, and Downgrade Instructions 240
Product Compatibility 253

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 240
- Installing the Software on QFX10002 Switches | 243
- 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 | 243
- Installing the Software on QFX10008 and QFX10016 Switches | 245
- Performing a Unified ISSU | 249
- Preparing the Switch for Software Installation | 250
- Upgrading the Software Using Unified ISSU | 250

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://support.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 **17.3** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 17.3 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-10-f-x86-64-17.3
-R2.n-secure-signed.tgz reboot
```

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

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

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 17.3 **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-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 17.3R2.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-f-x86-64-17.3  
-R2.n-secure-signed.tgz reboot reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-17.3  
-R2.n-secure-signed.tgz reboot 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://support.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://support.juniper.net/support/>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-x86-64-17.3
-R2.n-secure-signed.tgz reboot
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the **request system reboot** command:

```
user@switch> request system reboot
```

NOTE: You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete <package-name>** command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

```
user@switch> show version
```

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
Routing Engine status:
Slot 1:
  Current state           Master
  Election priority       Backup (default)
```

14. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-x86-64-17.3
-R2.n-secure-signed.tgz reboot
```

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 for upgrading to Junos OS Release 17.3R2 from 17.1R1 or later. Upgrading to 17.3R2 from releases prior to 17.1R1 is not supported. For example, upgrading from Junos OS Release 14.1X53 to 17.3R2 is not supported.

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 250](#)
- [Upgrading the Software Using Unified ISSU on page 250](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:

NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the **request system snapshot** command.

Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, `/jinstall-host-qfx-10-f-x86-64-17.3-R2.n-secure-signed.tgz` **reboot**.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting /jinstall-host-qfx-5-f-x86-64-17.3
-R2.n-secure-signed.tgz reboot ...
Install jinstall-host-qfx-5-f-x86-64-17.3
-R2.n-secure-signed.tgz completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
```

```

Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item              Status              Reason
  FPC 0             Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

SEE ALSO

New and Changed Features 202
Changes in Behavior and Syntax 218
Known Behavior 221
Known Issues 227
Resolved Issues 233
Documentation Updates 239
Product Compatibility 253

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 253

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 202
Changes in Behavior and Syntax 218
Known Behavior 221

[Known Issues | 227](#)

[Resolved Issues | 233](#)

[Documentation Updates | 239](#)

[Migration, Upgrade, and Downgrade Instructions | 240](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [New and Changed Features | 254](#)
- [Changes in Behavior and Syntax | 260](#)
- [Known Behavior | 260](#)
- [Known Issues | 261](#)
- [Resolved Issues | 267](#)
- [Documentation Updates | 270](#)
- [Migration, Upgrade, and Downgrade Instructions | 270](#)
- [Product Compatibility | 274](#)

These release notes accompany Junos OS Release 17.3R2 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

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

Release 17.3R2 New and Changed Features

There are no new features in Junos OS Release 17.3R2 for the SRX Series devices.

Release 17.3R1 New and Changed Features

IN THIS SECTION

- [Flow and Processing | 255](#)
- [IDP | 256](#)
- [Interfaces and Chassis | 257](#)
- [Junos OS XML API and Scripting | 257](#)
- [Layer 2 Features | 257](#)
- [Management | 258](#)
- [Network Security | 258](#)
- [Software Installation and Upgrade | 258](#)
- [User Interface and Configuration | 259](#)

Junos OS Release 17.3R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/4200, SRX5400, SRX5600, and SRX5800. Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 12.1X44 through 15.1X49-D75. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D75 are not available in 17.3R1.

New features for security platforms in Junos OS Release 17.3R1 include:

Flow and Processing

- **ECMP reverse traffic support (SRX Series)**—Starting with Junos OS Release 17.3R1, you can enable ECMP support for reverse traffic. In this case, the SRX Series device uses a hash algorithm to determine the interface to use for reverse traffic in a flow. If you do not enable this feature, the SRX Series device selects a route in the ECMP set to the incoming interface for reverse traffic, which is the default behavior. [See [Understanding ECMP Flow-Based Forwarding for Reverse Traffic on SRX Series Devices and vSRX.](#)]
- **TCP out-of-state packet drop logging (SRX Series)**—Starting in Junos OS Release 17.3R1, SRX Series devices support logging of unsynchronized TCP out-of-state packets that are dropped by the flow module.

Within any packet-switched network, when demand exceeds available capacity, the packets are queued up to hold the excess packets until the queue fills, and then the packets are dropped. When TCP operates across such a network, it takes any corrective actions to maintain error-free end-to-end communications.

This feature enables packet recovery by logging the out-of-sync packets for error-free communication, and avoids database servers going out of sync.

TCP packet drop logging occurs when:

- TCP packets that trigger session creation are not synchronized.
- TCP three-way handshake in flow fails.
- TCP sequence check in flow fails.
- TCP SYN packets are received in TCP FIN state.

The unsynchronized TCP out-of-state packet drop log is a packet-based log, not a session-based log.

NOTE: TCP packets that are dropped by TCP-proxy and IDP are not logged.

[See [TCP Out-of-State Packet Drop Logging Overview](#).]

IDP

- **IPS signature package update (SRX Series and vSRX instances)**—Starting with Junos OS Release 17.3, when you upgrade from Junos OS Release 12.3X48 or 15.1X49 to Junos OS Release 17.3 or downgrade from Junos OS Release 17.3 to Junos OS Release 12.3X48 or 15.1X49, you must update the IPS signature package to avoid any IDP configuration commit failures. Update the IPS signature package by:
 - Downloading the IPS signature package.
 - Installing the IPS signature package update when the download completes.

NOTE: When you upgrade from Junos OS Release 15.1X49 to Junos OS Release 17.3, the following warning message is displayed:

```
WARNING: A full install of the security package is required after reboot.
WARNING: Please perform a full update of the security package using
WARNING: "request security idp security-package download full-update"
WARNING: followed by
WARNING: "request security idp security-package install"
```

[See [Downloading and Installing the IPS Signature Package from an Older Junos OS Release Version to Newer Junos OS Release Version.](#)]

Interfaces and Chassis

- **Promiscuous mode support (SRX5400, SRX5600, SRX5800)**—Promiscuous mode function is supported on the SRX5000 line MPC (SRX5K-MPC) on 1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces on the MICs.

By default, an interface enables MAC filtering. You can configure promiscuous mode on the interface to disable MAC filtering. When you delete the promiscuous mode configuration, the interface will perform MAC filtering again. You can change the MAC address of the interface even when the interface is operating in promiscuous mode. When the interface is operating in normal mode again, the MAC filtering function on MPC uses the new MAC address to filter packets.

[See [Understanding Promiscuous Mode on Ethernet Interfaces.](#)]

Junos OS XML API and Scripting

- **Support for Python language for commit, event, op, and SNMP scripts (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX instances)**—Starting in Junos OS Release 17.3R1, you can author commit, event, op, and SNMP scripts in Python on devices that include the Python extensions package in the software image. Creating automation scripts in Python enables you to take advantage of Python features and libraries as well as leverage Junos PyEZ APIs supported in Junos PyEZ Release 1.3.1 and earlier releases to perform operational and configuration tasks on devices running Junos OS. To enable execution of Python automation scripts, which must be owned by either root or a user in the Junos OS **super-user** login class, configure the **language python** statement at the **[edit system scripts]** hierarchy level, and configure the filename for the Python script under the hierarchy level appropriate to that script type. Supported Python versions include Python 2.7.

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

Layer 2 Features

- **LACP support in Layer 2 transparent mode (SRX5400, SRX5600, and SRX5800)**—Starting with Junos OS Release 17.3R1, LACP is supported in Layer 2 transparent mode in addition to existing support in Layer 3 mode.

When the SRX Series device uses LACP to bundle the member links, it creates high-speed connections, also known as *fat pipe*, with peer systems. Bandwidth can be increased by adding member links. Increased bandwidth is especially important for redundant Ethernet (reth) and aggregated Ethernet (ae) interfaces. LACP also provides automatic determination, configuration, and monitoring member links.

LACP is compatible with other peers that run the 802.3ad LACP protocol. It automatically binds member links without manually configuring the LAG, thereby avoiding errors.

NOTE: Tentative sessions are created for all interfaces in a particular VLAN. If there is plenty of one-way traffic, numerous tentative sessions are created. When sessions reach the maximum limit, vector fails and packet loss might be seen.

Management

- **Support for adding non-native YANG modules to the Junos OS schema (SRX345, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX instances)**—Starting in Junos OS Release 17.3R1, you can load custom YANG models on devices running Junos OS to add data models that are not natively supported by Junos OS but can be supported by translation. Doing this enables you to extend the configuration hierarchies and operational commands with data models that are customized for your operations. The ability to add data models to a device is also beneficial when you want to create device-agnostic and vendor-neutral data models that enable the same configuration or RPC to be used on different devices from one or more vendors. You can load custom YANG modules by using the **request system yang add** operational command.

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

Network Security

- **Maximum number of security policies increased (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 17.3R1, the maximum number of security policies for SRX5400, SRX5600, and SRX5800 devices has increased from 80,000 to 100,000.

[See [Best Practices for Defining Policies on SRX Series Devices](#).]

Software Installation and Upgrade

- **Support for FreeBSD version 10 for Junos OS (SRX5800, SRX5600, SRX5400)**—Starting with Junos OS Release 17.3R1, on the SRX5000 line of devices, FreeBSD version 10 is the underlying operating system for Junos OS. Junos OS with upgraded FreeBSD is based on an upgraded FreeBSD kernel instead of older versions of FreeBSD. The newer FreeBSD kernel base provides Junos OS with sophisticated processing, efficiency, and security.

NOTE: On SRX5000 line of devices, use **no-validate** flag at the **request system software add <filename> no-validate** command to upgrade or downgrade between Junos OS Release 17.3 and the previous releases.

NOTE: Along with the upgraded FreeBSD, the System Snapshot feature has been enhanced on the SRX5000 line of devices. For more details, see *Understanding Junos OS with Upgraded FreeBSD Snapshots* topic in [Understanding Junos OS with Upgraded FreeBSD for SRX5400, SRX5600, and SRX5800 Devices](#)

[See [Understanding Junos OS with Upgraded FreeBSD.](#)]

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices and vSRX instances)**—Starting in Junos OS Release 17.3R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database, which is an alternate configuration database that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. Junos OS provides a default instance and up to eight user-defined instances of the ephemeral configuration database. The device's active configuration is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database. Ephemeral configuration data is volatile and is deleted upon rebooting the device.

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

SEE ALSO

[Changes in Behavior and Syntax | 260](#)

[Known Behavior | 260](#)

[Known Issues | 261](#)

[Resolved Issues | 267](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

[Product Compatibility | 274](#)

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 from Junos OS Release 17.3R2 for SRX Series devices.

SEE ALSO

[New and Changed Features | 254](#)

[Known Behavior | 260](#)

[Known Issues | 261](#)

[Resolved Issues | 267](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

[Product Compatibility | 274](#)

Known Behavior

IN THIS SECTION

- [CLI | 261](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.3R2 for SRX Series devices.

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

CLI

- On SRX5400, SRX5600, and SRX5800 devices, the CLI hangs while displaying CA profile group. This CA profile group contains CA certificates with 100's of certificates, and the CLI times out as PKI needs excessive time to handle such requests. Instead of displaying entire CA group, you can display the individual CA profile inside the CA group to avoid this problem. [PR1276619](#)

SEE ALSO

[New and Changed Features | 254](#)

[Known Issues | 261](#)

[Resolved Issues | 267](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

Known Issues

IN THIS SECTION

- [Authentication and Access Control | 262](#)
- [Chassis Cluster | 262](#)
- [Class of Service \(CoS\) | 263](#)
- [CLI | 263](#)
- [Flow-Based and Packet-Based Processing | 263](#)
- [Interfaces and Routing | 264](#)
- [J-Web | 264](#)
- [Network Address Translation \(NAT\) | 264](#)
- [Network Management and Monitoring | 264](#)
- [Platform and Infrastructure | 265](#)
- [Routing Policy and Firewall Filters | 265](#)
- [Routing Protocols | 265](#)
- [System Logs | 265](#)
- [VPNs | 265](#)

This section lists the known issues in hardware and software in Junos OS Release 17.3R2.

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

Authentication and Access Control

- For a security policy with HTTP pass-through firewall authentication being configured, we recommend that you configure web-redirect for HTTP pass-through firewall authentication instead of using direct HTTP pass-through firewall authentication because web browser may automatically carry credential in subsequential request to target web-server. [PR1230447](#)
- A user session is disconnected due to aging out of a fwauth entry in spite of an existing session. [PR1265571](#)
- On SRX Series devices with user firewall feature, the users sometimes fails to authenticate from LDAP server and gets the authorized group though the group mapping shows correctly for that particular user. [PR1282744](#)

Chassis Cluster

- On SRX Series devices in a chassis cluster, the synchronization monitoring configuration might fail if the following configuration is enabled: set system encrypt-configuration-files. The synchronization monitoring configuration failure might result in disabling the secondary node after reboot. [PR1235628](#)
- On SRX1500, SRX4100, SRX4200, and vSRX chassis clusters, invalid MAC address might be allocated for certain ae interfaces, which results in the interfaces not coming up. [PR1270166](#)

Class of Service (CoS)

- On all SRX Series devices, if the action of forwarding-class is configured in the output direction on a firewall filter, the host outbound traffic matching the same term of this firewall filter will be blocked. [PR1272286](#)

CLI

- SRX5400, SRX5600, and SRX5800 devices CLI hangs while displaying CA profile group. This CA profile group contains CA certificates with 100's of certificates and CLI times out as PKId needs excessive time to handle such requests. Instead of displaying entire CA group, you can display individual CA profile inside CA group to avoid this problem. [PR1276619](#)

Flow-Based and Packet-Based Processing

- SRX1500 devices may power-off unexpectedly due to incorrect device temperature readings which reported a too high temperature, leading to an immediate pro-active power-off of the device to protect the device from overheating. However in these cases the temperature was not actually too high and a power-off would not be required. When this occurs, the following log message is shown in file /var/log/hostlogs/lcmd.log: Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor [PR1241061](#)
- On SRX3000 line and SRX5000 line platforms, cold-sync will fail when SPC is stuck and traffic loss occurs. [PR1240983](#)
- On all branch SRX Series devices, the set system ports console insecure feature does not work as expected and fails to prevent non-root users from performing password recovery by using the console. This vulnerability might allow a non-root user with physical access to the console port to gain full administrative privileges. Refer to JSA10683 for more information. [PR1241006](#)
- On SRX300 line platforms in ethernet-switching mode, STP change state might not be pushed into the Packet Forwarding Engine. [PR1259286](#)
- On SRX 300 line devices, the Packet Forwarding Engine CPU utilization reaches 100% in every 10 minutes even though the session count has not been increased. [PR1284971](#)
- On SRX Series devices with chassis cluster enabled, the issue occurs when multicast traffic goes across logical systems. The ingress interface of the multicast session in the first logical system is reth2.0 which belongs to redundancy group 2. Redundancy group 2 is active on node 1. The ingress interface of multicast session in the second logical system will be PLT interface which belongs to redundancy group 1. Redundancy group 1 is active on node 0. So the multicast session in the second logical system will be active on node 0. It will cause multicast session active/backup not aligned with traffic forwarding. The workaround is to make RG-1 and RG-2 active on the same node. [PR1295893](#)

Interfaces and Routing

- IRB interface cannot be disabled or enabled in RPM. [PR1219570](#)
- On SRX5400, SRX5600, and SRX5800 devices, Stream Control Transmission Protocol (SCTP) packet has incorrect SCTP checksum after the payload is translated by the device. [PR1310141](#)
- The connection between the SRX Series device and JIMS times out. The solution is to enforce keepalives on the TCP connection. The SRX Series code is being modified to enforce keepalives. In the meantime the JIMS 1.0.1R1 release has a workaround to avoid the issue: JIMS server will enforce keepalive (likely depends on SRX series fix as well) Connection limit will be raised to 1000 from 10 [PR1311446](#)

J-Web

- On SRX Series devices, DHCP relay configuration under Configure > Services > DHCP > DHCP Relay page is removed from J-Web in Junos OS Release 15.1X49-D60. The same DHCP relay can be configured using the CLI. [PR1205911](#)
- There is an issue with J-Web when, at the monitoring tab for the policies, you select to see the configured global context policy. The result is a empty table where it should contain the global policies logs/statistics. Also, in the same monitor policy view, another problem came up when selecting 'all' policies in the context drop-down menu. It is not possible to either get to the next page or view more than 8 policies on one page. [PR1318118](#)

Network Address Translation (NAT)

- The configuration commit check might not detect a configuration error where a source NAT pool contains no address lines except a deactivated address line. This might allow the source NAT pool to be committed without addresses and lead to core files when traffic utilizes such misconfigured NAT pools. This specifically occurs when removing address statements from a NAT pool and leaving only a preexisting deactivated address statement in the NAT pool. [PR1300019](#)

Network Management and Monitoring

- On SRX Series devices, the syslog messages from the secondary node might not reach the syslog server when reth I/F is source interface for syslog. This issue does not impact traffic. [PR1252128](#)
- On SRX1400, SRX3400, and SRX3600 devices with a NP-IOC card installed, the data-plane related to the NP-IOC card might be stuck, which might cause child interfaces to be removed from the ae/reth LAG when the LACP is enabled. [PR1285011](#)

Platform and Infrastructure

- Starting with Junos OS Release 15.1X49-D30 and later, on SRX5400, SRX5600, and SRX5800 devices, some CLI commands are missed in the Request Support Information (RSI) script. [PR1236874](#)
- On SRX Series devices, the routes activated by IP-Monitoring is not getting cleared after the probe status change from Fail to Pass. The show services ip-monitoring status shows the route NOT-APPLIED but show route might show ip-monitoring route active (Static route with preference 1). [PR1263078](#)

Routing Policy and Firewall Filters

- On SRX Series devices, when there is at least one policy using the range address in a zone, the network security daemon (NSD) crashes after executing show security shadow-policies command. [PR1232736](#)
- On SRX Series devices, when you use Integrated User Firewall (IUF), the useridd might consume high CPU usage. The traceoptions of IUF might have lots of UGCALC_AD_MEMBER_UPDATE messages. [PR1280783](#)
- On SRX Series devices with User Firewall feature, under some conditions, a core file of flowd or useridd might be triggered. [PR1299494](#)

Routing Protocols

- On SRX Series devices, RIP will be supported in packet-to-packet DC mode on st0 interfaces. [PR1141817](#)
- On a chassis with BMP configured, if the rpd termination timeout is happening while the BMP main task has failed to terminate and delete itself (seen when rpd is gracefully terminated), the rpd might crash. [PR1315798](#)

System Logs

- RT_SCTP_DATA_MSG_M3UA_SI SCTP messages are not logged in sctp_syslog messages. [PR1268849](#)

VPNs

- On the hub side, autoVPN tunnel fails to come up if establish immediately is configured. Since establish immediately is not needed on the hub side, there is no impact if establish immediately is not configured on the hub side. [PR1160948](#)
- On branch SRX Series devices in HA mode, VPN-monitoring with optimized option is configured and traffic goes through the IPsec tunnel. The VPN-monitoring status will be displayed as down after RGO failover. [PR1203723](#)

- On SRX High End Series devices, if traffic-selector is configured with DPD backup gateway, the IKE redundant gateway failover fails. This may cause IPsec management daemon to restart. [PR1249908](#)
- On SRX Series devices in a chassis cluster, in a rare condition, modifying the IPsec VPN configuration might cause /var/etc/vpn_tunnel.id file mismatch between both primary node and secondary node, then the RG0 failover results in the kmd process crash on the new primary node. [PR1250178](#)
- On SRX1500 devices in a chassis cluster, IP leak might occur under the following scenarios:
 - In case of IKEv1, it is possible for an IPsec VPN tunnel to be active without an active IKEv1 phase 1 SA. Since the assigned IP address associated with an IPsec VPN tunnel (for a user) is stored in the record of phase 1 SA, if HA RG0 failover occurs while there is no active IKEv1 phase SA exist for an IPsec VPN tunnel, the assigned IP address will be released to the authd daemon when the IPsec VPN tunnel is disconnected.
 - In case a remote access IPsec VPN tunnel is cleared (for both IKEv1 and IKEv2), the assigned IP address is kept for 30 seconds before it is released back to the authd within an additional 2 minutes. If HA failover occurs during this time before the IP is received at the authd, there will be an IP address leak.
 - If a new IP is assigned by authd daemon after every user is authenticated, regardless of the user already having an IP assigned from an early authentication. In case of IKEv1, authentication occurs at every IKE phase 1 SA rekey. If the KMD daemon restarts immediately (within 2 minutes) after an IKEv1 phase 1 SA rekey, there is a possibility that the newly assigned IP has not been released to authd daemon yet.

This will lead to the leak of that IP. [PR1252181](#)

- On all SRX Series devices, when manual route-based IPsec VPN is configured, enabling VPN monitoring will cause the st0.* interface down, which results in VPN traffic dropping. [PR1259422](#)
- ADVPN shortcuts can cause kmd core files on the suggester. [PR1259844](#)
- On SRX Series devices, manual Next-Hop Tunnel Binding (NHTB) does not work on Junos OS 15.1X and 12.3X releases. The following error is displayed on the IKE traces **Internal Error: Manual NHTB add failed**. [PR1266797](#)
- On SRX Series devices, if traffic-selector is configured, the IKE redundant gateway failover fails. [PR1270000](#)
- When an SRX Series device is an initiator behind the NAT, disabling NAT on the middle router causes an immediate new negotiation failure due to an attempt with port 4500. The next attempt will succeed by using port 500. Disabling NAT and bringing down all the existing tunnels and re-establishing the tunnels with port 500 is the expected behavior solution. [PR1273213](#)
- On all SRX Series devices, if a large number of IPsec VPN tunnels are established (for example, 16k traffic-selector based tunnels are established on an SRX5600 platform), changing the configuration of IPsec VPN (for example, removing some or all these IPsec VPN tunnels and then adding them again) might result in VPN tunnels being established in the the data plane. However, the VPN configuration is already removed in the Routing Engine, which results in the kmd process crash. [PR1276058](#)

- On all SRX Series devices, when ike-ha-link-encryption is enabled, the IKE and IPsec configuration might not be pushed to data-plane. [PR1277229](#)
- On SRX5000 line platforms, you cannot load PKI local-certificate and CA certificate with cmpv2. [PR1277317](#)
- On all SRX Series devices, CRL download fails when missing content-length field in http header and CRL occupies are in at least 2 packets. [PR1278631](#)
- On all SRX Series devices, enabling or disabling CRL download in CA profile does not work as expected. [PR1280530](#)
- Commit returns a warning for NULL authentication and remains a commit error for FIPS mode. [PR1285284](#)

SEE ALSO

[New and Changed Features | 254](#)

[Resolved Issues | 267](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.3R2 | 268](#)
- [Resolved Issues: 17.3R1 | 269](#)

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: 17.3R2

Application Layer Gateways (ALGs)

- The pfed process crashes and generates core files. [PR1292992](#)

Authentication and Access Control

- SRX Series device assigns IP address 0.0.0.0 to xauth clients. [PR1315999](#)

Chassis Clustering

- The SRX1500 stops forwarding traffic randomly. [PR1277435](#)
- Duplicate RFSP IE drops the GTP packet. [PR1284311](#)
- After software upgrade, the cluster goes to short split-brain when rebooting RGO secondary, and multiple errors and issues are seen. [PR1288819](#)
- ISSU can be unsuccessful if control-link-recovery is configured. [PR1303948](#)

CoS

- On SRX1500 devices, CoS scheduler and shaping do not work on IRB interface. [PR1292187](#)

Ethernet Switching

- Ping to VRRP(VIP) address failed when VRRP on vlan-tagging. It only affects IOC2 and IOC3 cards in SRX5000 line devices. [PR1293808](#)

Flow-Based and Packet-Based Processing

- ECMP does not work for traffic with ECN enabled and with IPv6. [PR1265576](#)
- On SRX1500 devices, core files are generated when J-Flow is enabled. [PR1271466](#)
- More CPU threshold warnings are seen than in the previous releases. [PR1291506](#)
- SCTP association capacity cannot reach up to 20K. [PR1299186](#)
- The name daemon (named) might crash if SRX Series device is configured for dns-proxy. [PR1307435](#)

J-Web

- J-Web removes backslash character on source identity object when committing changes. [PR1304608](#)

Network Address Translation (NAT)

- The **show security zones detail** command causes memory leak. [PR1269525](#)

Network Management and Monitoring

- The mib2d process might crash when polling the OID ifStackStatus.0 after an a logical interface (IFL) of lo0 is deleted. [PR1286351](#)

- The **show arp no-resolve interface X** command for non-existent interface X is showing all unrelated static ARP entries. [PR1299619](#)

Platform and Infrastructure

- SRX Series device does not process traffic due to an IPv6 NA packets burst. [PR1293673](#)

Resolved Issues: 17.3R1

Interfaces and Chassis

- On SRX1500, if Junos OS Release 15.1X49-D70 or later is installed and you have a single PEM in slot 0, you will see an alarm saying PEM 1 is not present. [PR1265795](#)

Layer 2 Ethernet Services

- On SRX1500 devices, when configuring the devices to switching mode, an IRB interface located in a custom routing-instance is not reachable. [PR1234000](#)

Platform and Infrastructure

- On SRX Series devices in a chassis cluster, if sampling is used, the flowd process fails and core files are seen on both the nodes, when the route is updated through dynamic protocols such as BGP. [PR1249254](#)

Routing Policy and Firewall Filters

- Starting in Junos OS Release 15.1X49-D100, a new default application, application junos-smtps, has been added for secured e-mail traffic using port 587 or 465. To view the new default policy, use the **show configuration groups junos-defaults applications** command. [PR1273725](#)

Unified Threat Management (UTM)

- Some traffic from web-cam contain non-standard HTTP boundary format will cause SRX Series UTM/SAV to hold traffic/mbuf and later causes failover. [PR1283806](#)

VPNs

- On SRX5400, SRX5600, and SRX5800 devices, the st0 interface global counter statistics do not increment and remain zero, although traffic passes through the tunnel sub-interfaces such as st0.0 and st0.1. [PR1171958](#)

SEE ALSO

[New and Changed Features | 254](#)

[Known Issues | 261](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.3R2 for the SRX Series documentation.

SEE ALSO

[New and Changed Features | 254](#)

[Changes in Behavior and Syntax | 260](#)

[Known Behavior | 260](#)

[Known Issues | 261](#)

[Resolved Issues | 267](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

[Product Compatibility | 274](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade for Layer 2 Configuration | 270](#)
- [Upgrade and Downgrade Scripts for Address Book Configuration | 271](#)

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

Upgrade for Layer 2 Configuration

Starting with Junos OS Release 15.1X49-D10 and later, only enhanced Layer 2 CLI configurations are supported. If your device was configured earlier for Layer 2 transparent mode, then you must convert the legacy configurations to Layer 2 next-generation CLI configurations.

For details on how to migrate from Junos OS Release 12.3X48-D10 and earlier releases to Junos OS Release 15.1X49-D10 and later releases, refer to the Knowledge Base article at <https://kb.juniper.net/InfoCenter/index?page=content&id=KB30445>.

Upgrade and Downgrade Scripts for Address Book Configuration

IN THIS SECTION

- [About Upgrade and Downgrade Scripts | 271](#)
- [Running Upgrade and Downgrade Scripts | 272](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 273](#)

Beginning with Junos OS Release 12.1, you can configure address books under the **[security]** hierarchy and attach security zones to them (zone-attached configuration). In Junos OS Release 11.1 and earlier, address books were defined under the **[security zones]** hierarchy (zone-defined configuration).

You can either define all address books under the **[security]** hierarchy in a zone-attached configuration format or under the **[security zones]** hierarchy in a zone-defined configuration format; the CLI displays an error and fails to commit the configuration if you configure both configuration formats on one system.

Juniper Networks provides Junos operation scripts that allow you to work in either of the address book configuration formats (see [Figure 1 on page 272](#)).

About Upgrade and Downgrade Scripts

After downloading Junos OS Release 12.1, you have the following options for configuring the address book feature:

- **Use the default address book configuration**—You can configure address books using the zone-defined configuration format, which is available by default. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.
- **Use the upgrade script**—You can run the upgrade script available on the Juniper Networks support site to configure address books using the new zone-attached configuration format. When upgrading, the system uses the zone names to create address books. For example, addresses in the trust zone are created in an address book named **trust-address-book** and are attached to the trust zone. IP prefixes used in NAT rules remain unaffected.

After upgrading to the zone-attached address book configuration:

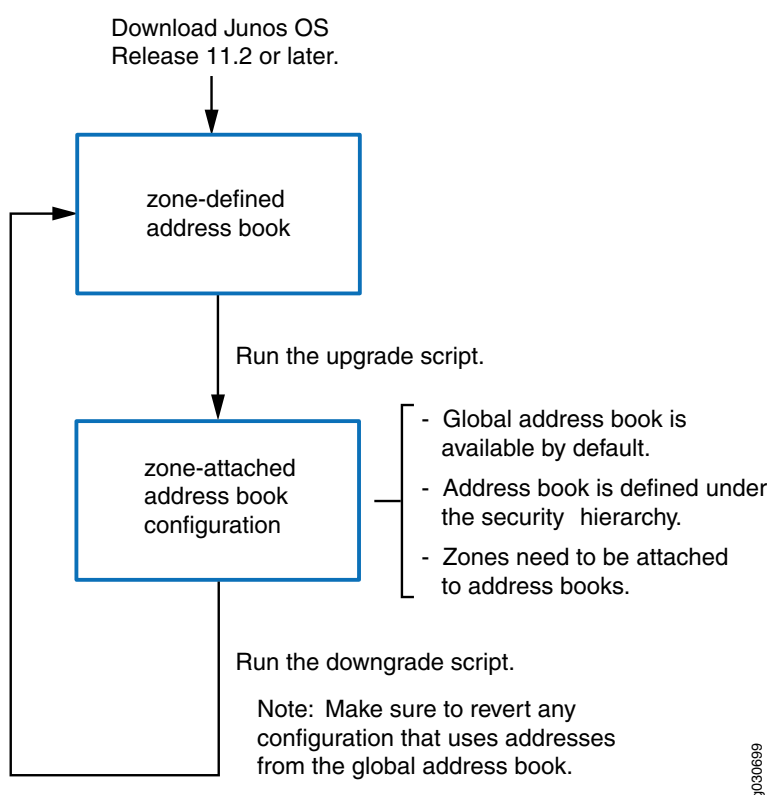
- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.

For information on how to configure zone-attached address books, see the Junos OS Release 12.1 documentation.

- **Use the downgrade script**—After upgrading to the zone-attached configuration, if you want to revert to the zone-defined configuration, use the downgrade script available on the Juniper Networks support site. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.

NOTE: Before running the downgrade script, make sure to revert any configuration that uses addresses from the global address book.

Figure 1: Upgrade and Downgrade Scripts for Address Books



Running Upgrade and Downgrade Scripts

The following restrictions apply to the address book upgrade and downgrade scripts:

- The scripts cannot run unless the configuration on your system has been committed. Thus, if the zone-defined address book and zone-attached address book configurations are present on your system at the same time, the scripts will not run.
- The scripts cannot run when the global address book exists on your system.
- If you upgrade your device to Junos OS Release 12.1 and configure logical systems, the master logical system retains any previously configured zone-defined address book configuration. The master

administrator can run the address book upgrade script to convert the existing zone-defined configuration to the zone-attached configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.

NOTE: You cannot run the downgrade script on logical systems.

For information about implementing and executing Junos operation scripts, see the *Junos OS Configuration and Operations Automation Guide*.

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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Release 12.3X48 is an EEOL release. You can upgrade from Junos OS Release 12.1X46 to Release 12.3X48 or even from Junos OS Release 12.3X48 to Release 15.1X49-D10. For upgrading from Junos OS Release 12.1X47-D15 to Junos OS Release 15.1X49-D10, ISSU is supported. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

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

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

SEE ALSO

[New and Changed Features | 254](#)

[Changes in Behavior and Syntax | 260](#)

[Known Behavior | 260](#)

[Known Issues | 261](#)

[Resolved Issues | 267](#)

[Documentation Updates | 270](#)

[Product Compatibility | 274](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 274](#)

Hardware Compatibility

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

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

SEE ALSO

[Changes in Behavior and Syntax | 260](#)

[Known Behavior | 260](#)

[Known Issues | 261](#)

[Resolved Issues | 267](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

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 [High Availability User Guide for Routing 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

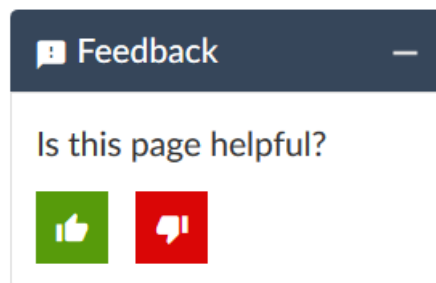
For the latest, most complete information about known and resolved issues with Junos OS, see the Juniper Networks Problem Report Search application at: <https://prsearch.juniper.net>.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales 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/assets/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://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- 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 located at <https://entitlementsearch.juniper.net/entitlementsearch/welcome.do>.

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.html>.

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/cgi-bin/docbugreport/>.

Revision History

30 September 2021—Revision 15, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 14, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 May 2020—Revision 13, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

5 December 2019—Revision 12, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 February 2019—Revision 11, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 January 2019—Revision 10, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 October 2018—Revision 9, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 July 2018—Revision 8, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 July 2018—Revision 7, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 May 2018—Revision 6, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 April 2018—Revision 5, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 February 2018—Revision 4, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 February 2018—Revision 3, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 February 2018—Revision 2, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 February 2018—Revision 1, Junos OS Release 17.3R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 November 2017—Revision 8, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 October 2017—Revision 7, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 October 2017—Revision 6, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 October 2017—Revision 5, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 September 2017—Revision 4, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 September 2017—Revision 3, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

1 September 2017—Revision 2, Junos OS Release 17.3R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 August 2017—Revision 1, Junos OS Release 17.3R1—ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

Copyright © 2018 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.