

Release Notes: Junos[®] OS Release 17.2R1 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion Junos OS 17.2R1 Release Notes

3 September 2020

Contents	Introduction 12
	Junos OS Release Notes for ACX Series 12
	New and Changed Features 13
	Hardware 14
	Interfaces and Chassis 14
	Management 14
	Network Management and Monitoring 14
	Operation, Administration, and Maintenance (OAM) 15
	Routing Protocols 15
	Software Installation and Upgrade 15
	Timing and Synchronization 15
	Changes in Behavior and Syntax 16
	Interfaces and Chassis 17
	Management 17
	Known Behavior 18
	High Availability (HA) and Resiliency 18
	Known Issues 19
	Layer 2 Features 19
	Resolved Issues 19

Documentation Updates | 20

Migration, Upgrade, and Downgrade Instructions | 20

 Upgrade and Downgrade Support Policy for Junos OS Releases | 21

Product Compatibility | 22

 Hardware Compatibility | 22

Junos OS Release Notes for EX Series Switches | 23

New and Changed Features | 23

 Authentication, Authorization and Accounting (AAA) (RADIUS) | 24

 EVPNs | 25

 Layer 3 Features | 25

 Management | 27

 Multicast | 27

 Network Management and Monitoring | 27

 System Management | 27

Changes in Behavior and Syntax | 28

 IP Tunneling | 28

 Management | 28

Known Behavior | 29

 High Availability (HA) and Resiliency | 30

Known Issues | 30

 General Routing | 31

 High Availability (HA) and Resiliency | 31

 Infrastructure | 31

 Junos Fusion Enterprise | 31

 Platform and Infrastructure | 31

 Virtual Chassis | 32

Resolved Issues | 32

 Resolved Issues:17.2R1 | 33

Documentation Updates | 34

Migration, Upgrade, and Downgrade Instructions | 34

 Upgrade and Downgrade Support Policy for Junos OS Releases | 35

Product Compatibility | 35

 Hardware Compatibility | 36

Junos OS Release Notes for Junos Fusion Enterprise | 36

New and Changed Features | 37

Interfaces and Chassis | 38

Layer 2 Features | 38

Changes in Behavior and Syntax | 39

Known Behavior | 40

Known Issues | 40

Resolved Issues | 41

Documentation Updates | 41

Migration, Upgrade, and Downgrade Instructions | 42

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 42

Upgrading an Aggregation Device with Redundant Routing Engines | 44

Preparing the Switch for Satellite Device Conversion | 44

Converting a Satellite Device to a Standalone Switch | 46

Upgrade and Downgrade Support Policy for Junos OS Releases | 48

Downgrading from Release 17.2 | 48

Product Compatibility | 50

Hardware and Software Compatibility | 50

Hardware Compatibility Tool | 50

Junos OS Release Notes for Junos Fusion Data Center | 51

New and Changed Features | 51

Junos Fusion Data Center | 52

Hardware | 53

Class of Service (CoS) | 53

High Availability (HA) and Resiliency | 53

Interfaces | 53

Layer 2 Protocols | 58

Layer 3 Protocols | 59

Multicast Protocols | 60

Network Management and Monitoring | 62

Routing Policy and Firewall Filters | 63

Storage | 63

Software Installation and Upgrade | 64

	Software Licensing	65
Changes in Behavior and Syntax	66	
	Junos Fusion	66
Known Behavior	67	
	Junos Fusion Data Center	67
Known Issues	68	
	Junos Fusion	68
Resolved Issues	69	
	Resolved Issues: 17.2R1	69
Documentation Updates	69	
Migration, Upgrade, and Downgrade Instructions	70	
	Basic Procedure for Upgrading an Aggregation Device	71
	Preparing the Switch for Satellite Device Conversion	72
	Autoconverting a Switch into a Satellite Device	75
	Manually Converting a Switch into a Satellite Device	78
	Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology	80
	Configuring Satellite Device Upgrade Groups	81
	Converting a Satellite Device to a Standalone Device	82
	Upgrade and Downgrade Support Policy for Junos OS Releases	85
	Downgrading from Release 17.2R1	85
Product Compatibility	86	
	Hardware Compatibility	86
Junos OS Release Notes for Junos Fusion Provider Edge	87	
	New and Changed Features	88
	Class of Service (CoS)	88
	Junos Fusion	88
Changes in Behavior and Syntax	89	
Known Behavior	90	
Known Issues	90	
Resolved Issues	91	
Documentation Updates	91	

Migration, Upgrade, and Downgrade Instructions | 92

Basic Procedure for Upgrading an Aggregation Device | 92

Upgrading an Aggregation Device with Redundant Routing Engines | 94

Preparing the Switch for Satellite Device Conversion | 95

Converting a Satellite Device to a Standalone Device | 96

Upgrading an Aggregation Device | 98

Upgrade and Downgrade Support Policy for Junos OS Releases | 99

Downgrading from Release 17.2 | 99

Product Compatibility | 100

Hardware Compatibility | 100

Junos OS Release Notes for MX Series 5G Universal Routing Platforms | 101

New and Changed Features | 101

Hardware | 102

Class of Service (CoS) | 103

EVPNs | 104

Forwarding and Sampling | 106

General Routing | 107

High Availability (HA) and Resiliency | 108

Interfaces and Chassis | 109

Layer 2 VPN | 112

Management | 112

MPLS | 114

Multicast | 115

Network Management and Monitoring | 116

Operation, Administration, and Maintenance (OAM) | 118

Routing Policy and Firewall Filters | 119

Routing Protocols | 119

Services Applications | 121

Software Defined Networking (SDN) | 123

Subscriber Management and Services | 124

Virtual Chassis | 131

Changes in Behavior and Syntax | 132

EVPNs | 132

Forwarding and Sampling | 133

High Availability (HA) and Resiliency	135
Interfaces and Chassis	135
IP Tunneling	138
Management	138
MPLS	139
Network Management and Monitoring	139
Routing Protocols	140
Services Applications	141
Subscriber Management and Services	141
User Interface and Configuration	142
Known Behavior	143
General Routing	145
High Availability (HA) and Resiliency	145
Network Management and Monitoring	145
Interfaces and Chassis	145
User Interface and Configuration	145
Software Defined Networking (SDN)	146
Subscriber Management and Services	146
Known Issues	147
Class of Service (CoS)	148
EVPNs	148
Forwarding and Sampling	148
General Routing	148
Infrastructure	151
Interfaces and Chassis	151
Layer 2 Features	151
Layer 2 Ethernet Services	151
MPLS	152
Multicast	152
Network Management and Monitoring	152
Platform and Infrastructure	152
Routing Protocols	153
Services Applications	154
Subscriber Access Management	154

VPNs	154
Resolved Issues	154
Resolved Issues: 17.2R1	155
Documentation Updates	164
Subscriber Management Access Network Guide	165
Subscriber Management Provisioning Guide	165
Migration, Upgrade, and Downgrade Instructions	166
Basic Procedure for Upgrading to Release 17.2	166
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)	167
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)	169
Upgrade and Downgrade Support Policy for Junos OS Releases	171
Upgrading a Router with Redundant Routing Engines	171
Downgrading from Release 17.2	172
Product Compatibility	173
Hardware Compatibility	173
Junos OS Release Notes for NFX Series	174
New and Changed Features	174
Hardware	175
Juniper Device Manager	176
Junos Control Plane	177
vSRX	179
Changes in Behavior and Syntax	179
Known Behavior	180
Juniper Device Manager	180
Known Issues	181
Infrastructure	181
IPsec	181
Juniper Device Manager	181
Junos Control Plane	183
vSRX	184
Resolved Issues	185
Juniper Device Manager	185
Junos Control Plane	185

Documentation Updates | **186**

Release Notes | **186**

Migration, Upgrade, and Downgrade Instructions | **187**

Upgrade and Downgrade Support Policy for Junos OS Releases | **187**

Upgrading Using Unified ISSU | **187**

Upgrading a Router with Redundant Routing Engines | **188**

Basic Procedure for Upgrading to Release 17.2 | **188**

Product Compatibility | **191**

Hardware Compatibility | **192**

Junos OS Release Notes for PTX Series Packet Transport Routers | **192**

New and Changed Features | **193**

Hardware | **194**

Class of Service (CoS) | **195**

Forwarding and Sampling | **196**

General Routing | **196**

High Availability (HA) and Resiliency | **196**

Interfaces and Chassis | **197**

IPv6 | **201**

Layer 2 Features | **201**

Layer 3 Features | **201**

Management | **202**

MPLS | **205**

Multicast | **207**

Network Management and Monitoring | **207**

Routing Policy and Firewall Filters | **208**

Routing Protocols | **209**

Security | **211**

Services Applications | **211**

Software Installation and Upgrade | **211**

User Interface and Configuration | **211**

VPNs | **212**

Changes in Behavior and Syntax | **213**

Interfaces and Chassis | **213**

Management | **214**

Known Behavior | 215**Hardware | 216****High Availability (HA) and Resiliency | 216****Known Issues | 217****General Routing | 217****Platform and Infrastructure | 218****Resolved Issues | 218****Resolved Issues: 17.2R1 | 219****Documentation Updates | 220****Migration, Upgrade, and Downgrade Instructions | 220****Basic Procedure for Upgrading to Release 17.2 | 220****Upgrade and Downgrade Support Policy for Junos OS Releases | 223****Upgrading Using Unified ISSU | 224****Upgrading a Router with Redundant Routing Engines | 224****Product Compatibility | 225****Hardware Compatibility | 225****Junos OS Release Notes for the QFX Series | 226****New and Changed Features | 226****Hardware | 227****Authentication, Authorization, and Accounting (AAA) (RADIUS) | 228****Class of Service (CoS) | 228****Dynamic Host Configuration Protocol (DHCP) | 229****EVPNs | 229****Infrastructure | 231****Interfaces and Chassis | 231****IPv4 | 233****IPv6 | 233****Layer 2 Features | 233****Layer 3 Features | 234****Management | 234****MPLS | 237****Multicast | 238****Network Management and Monitoring | 239****Port Security | 240**

Routing Protocols	241
Security	243
Services Applications	244
Software Defined Networking (SDN)	244
Software Installation and Upgrade	245
Software Licensing	245
System Management	245
VLAN Infrastructure	246
Changes in Behavior and Syntax	247
Class of Service (CoS)	248
Management	248
Known Behavior	249
High Availability (HA) and Resiliency	250
Virtual Chassis	250
Known Issues	250
General Routing	251
High Availability (HA) and Resiliency	252
Infrastructure	252
Interfaces and Chassis	252
IPsec	252
Layer 2 Features	252
MPLS	252
Open vSwitch Database Management Protocol (OVSDB)	253
OpenFlow	253
Platform and Infrastructure	253
Routing Protocols	254
Resolved Issues	254
General Routing	255
Layer 2 Features	255
Routing Protocols	256
Documentation Updates	256
Migration, Upgrade, and Downgrade Instructions	257
Upgrading Software on QFX Series Switches	257
Installing the Software on QFX10002 Switches	259

Performing an In-Service Software Upgrade (ISSU)	259
Preparing the Switch for Software Installation	260
Upgrading the Software Using ISSU	260
Product Compatibility	263
Hardware Compatibility	263
Third-Party Components	264
Upgrading Using Unified ISSU	264
Compliance Advisor	264
Finding More Information	264
Requesting Technical Support	265
Self-Help Online Tools and Resources	265
Opening a Case with JTAC	266
Revision History	266

Introduction

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

These release notes accompany Junos OS Release 17.2R1 for the ACX Series, EX Series, Junos Fusion Enterprise, Junos Fusion Data Center, Junos Fusion Provider Edge, MX Series, PTX Series, and QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- New and Changed Features | 13
- Changes in Behavior and Syntax | 16
- Known Behavior | 18
- Known Issues | 19
- Resolved Issues | 19
- Documentation Updates | 20
- Migration, Upgrade, and Downgrade Instructions | 20
- Product Compatibility | 22

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Hardware | 14
- Interfaces and Chassis | 14
- Management | 14
- Network Management and Monitoring | 14
- Operation, Administration, and Maintenance (OAM) | 15
- Routing Protocols | 15
- Software Installation and Upgrade | 15
- Timing and Synchronization | 15

This section describes the new features or enhancements to existing features in Junos OS Release 17.2R1 for ACX Series Universal Metro Routers.

Hardware

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

Interfaces and Chassis

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

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

Management

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

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

Network Management and Monitoring

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

sFlow technology uses the following two sampling mechanisms:

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

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

Operation, Administration, and Maintenance (OAM)

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

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

Routing Protocols

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

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

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

Software Installation and Upgrade

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

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

Timing and Synchronization

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

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

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

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

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

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

SEE ALSO

Changes in Behavior and Syntax	16
Known Behavior	18
Known Issues	19
Resolved Issues	19
Documentation Updates	20
Migration, Upgrade, and Downgrade Instructions	20
Product Compatibility	22

Changes in Behavior and Syntax

IN THIS SECTION

- [Interfaces and Chassis](#) | [17](#)
- [Management](#) | [17](#)

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 17.2R1 for the ACX Series Universal Metro Routers.

Interfaces and Chassis

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

Management

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

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

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

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

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

SEE ALSO

[New and Changed Features | 13](#)

[Known Behavior | 18](#)

[Known Issues | 19](#)

[Resolved Issues | 19](#)

[Documentation Updates | 20](#)

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

[Product Compatibility | 22](#)

Known Behavior

IN THIS SECTION

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

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 for the ACX Series.

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

High Availability (HA) and Resiliency

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

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

SEE ALSO

[New and Changed Features | 13](#)

[Changes in Behavior and Syntax | 16](#)

[Known Issues | 19](#)

[Resolved Issues | 19](#)

[Documentation Updates | 20](#)

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

[Product Compatibility | 22](#)

Known Issues

IN THIS SECTION

- [Layer 2 Features | 19](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R1 for the ACX Series Universal Metro Routers.

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

Layer 2 Features

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

SEE ALSO

[New and Changed Features | 13](#)

[Changes in Behavior and Syntax | 16](#)

[Known Behavior | 18](#)

[Resolved Issues | 19](#)

[Documentation Updates | 20](#)

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

[Product Compatibility | 22](#)

Resolved Issues

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

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

SEE ALSO

New and Changed Features 13
Changes in Behavior and Syntax 16
Known Behavior 18
Known Issues 19
Documentation Updates 20
Migration, Upgrade, and Downgrade Instructions 20
Product Compatibility 22

Documentation Updates

There are no errata or changes in Junos OS Release 17.2R1 for the ACX Series documentation.

SEE ALSO

New and Changed Features 13
Changes in Behavior and Syntax 16
Known Behavior 18
Known Issues 19
Resolved Issues 19
Migration, Upgrade, and Downgrade Instructions 20
Product Compatibility 22

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

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

SEE ALSO

New and Changed Features 13
Changes in Behavior and Syntax 16
Known Behavior 18
Known Issues 19
Resolved Issues 19
Documentation Updates 20
Product Compatibility 22

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 22](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

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

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- New and Changed Features | 23
- Changes in Behavior and Syntax | 28
- Known Behavior | 29
- Known Issues | 30
- Resolved Issues | 32
- Documentation Updates | 34
- Migration, Upgrade, and Downgrade Instructions | 34
- Product Compatibility | 35

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Authentication, Authorization and Accounting (AAA) (RADIUS) | 24
- EVPNs | 25
- Layer 3 Features | 25
- Management | 27
- Multicast | 27
- Network Management and Monitoring | 27
- System Management | 27

This section describes the new features and enhancements to existing features in Junos OS Release 17.2R1 for the EX Series.

NOTE: The following EX Series switches are supported in Release 17.2R1: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.2R1, J-Web is supported on the EX4300 and EX4600 switches in both standalone and Virtual Chassis setup.

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

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

For details about the J-Web distribution model, see [Release Notes: J-Web Application Package Release 17.2A1 for EX4300 and EX4600 Switches](#).

Authentication, Authorization and Accounting (AAA) (RADIUS)

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

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

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

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

EVPNs

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

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

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

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

Layer 3 Features

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

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

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

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

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

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

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

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

Management

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

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

Multicast

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

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

Network Management and Monitoring

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

System Management

- **Dynamic Power Management in EX9200 Switches**—Starting with Junos OS Release 17.2R1, EX9200 switches support dynamic power management. See [System Services on EX9200 Switches](#).

SEE ALSO

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

[Known Behavior | 29](#)

[Known Issues | 30](#)

[Resolved Issues | 32](#)

[Documentation Updates | 34](#)

[Migration, Upgrade, and Downgrade Instructions | 34](#)

[Product Compatibility | 35](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [IP Tunneling | 28](#)
- [Management | 28](#)

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 17.2R1 for the EX Series.

IP Tunneling

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

Management

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

namespace for the <configuration> element in the RPC reply to the same namespace as in the corresponding YANG model.

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

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

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

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

SEE ALSO

New and Changed Features	 23
Known Behavior	 29
Known Issues	 30
Resolved Issues	 32
Documentation Updates	 34
Migration, Upgrade, and Downgrade Instructions	 34
Product Compatibility	 35

Known Behavior

IN THIS SECTION

- [High Availability \(HA\) and Resiliency](#) | 30

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 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.

High Availability (HA) and Resiliency

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

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

SEE ALSO

New and Changed Features	 23
Changes in Behavior and Syntax	 28
Known Issues	 30
Resolved Issues	 32
Documentation Updates	 34
Migration, Upgrade, and Downgrade Instructions	 34
Product Compatibility	 35

Known Issues

IN THIS SECTION

- [General Routing](#) | [31](#)
- [High Availability \(HA\) and Resiliency](#) | [31](#)
- [Infrastructure](#) | [31](#)
- [Junos Fusion Enterprise](#) | [31](#)
- [Platform and Infrastructure](#) | [31](#)
- [Virtual Chassis](#) | [32](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R1 for the EX Series.

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

General Routing

- On EX4300 switches, when 802.1X single-suplicant authentication is initiated, multiple "EAP Request Id Frame Sent" packets might be sent. [PR1163966](#)
- 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](#)
- On a dot1x-enabled interface, sometimes when you log in, log off, and then log in within a short interval (within subseconds), the logical interface plus the bridge domain or VLAN remain in a pending state, and you will not be able to access the network. As a workaround, restart the I2-learning process to recover the port/interface from the problematic state. [PR1230073](#)

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)
- On an EX4300 or a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)

Infrastructure

- On an EX4300 egress VLAN-based firewall filter on a Q-in-Q interface, after a switch reboot, firewall counters might not increment as expected. [PR1165450](#)

Junos Fusion Enterprise

- On a Junos Fusion Enterprise, Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) fast start does not work. [PR1171899](#)

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](#)
- 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](#)

Virtual Chassis

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

SEE ALSO

New and Changed Features 23
Changes in Behavior and Syntax 28
Known Behavior 29
Resolved Issues 32
Documentation Updates 34
Migration, Upgrade, and Downgrade Instructions 34
Product Compatibility 35

Resolved Issues

IN THIS SECTION

- [Resolved Issues:17.2R1 | 33](#)

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:17.2R1

Interfaces and Chassis

- ISSU will not work from 15.1R to later images (like 15.1F, 16.1R2, ...), when the router is equipped with QSFP/CXP/CFP2 optics on MPC3E/MPC4E/MPC5E/MPC6E/MPC 3D 16x10GE cards. This is because a dark window issue is fixed for SFPP/QSFP/CXP/CFP2 optics in 16.1 and 15.1F, which makes the 15.1R image ISSU incompatible with later images. Doing ISSU on the incompatible image from 15.1R to later versions might result in the line card crash as a side effect. [PR1216924](#)

Network Management and Monitoring

- On EX4600VC with default SNMPv3 engine ID configuration, it may happen that after the reboot of the VC, SNMPv3 engine ID changes the value, even though the routing engine master switchover didn't happen, causing the authentication of the SNMP user to fail. Re-configuration of SNMP user details will solve this issue. With Junos version which contains the fix for this issue, SNMPv3 engine ID will not change after the reboot of the VC. However, the SNMPv3 engine ID will still change if the following sequence of the events occurs: * routing engine master switchover * reboot of the VC. To prevent this to happen it is recommended to configure SNMPv3 to use a predefined suffix by issuing the following line: `set snmp engine-id local <engine-id-suffix>` This will force the system to generate the same engine ID for all routing engines on all members. When doing a global switchover, the engine ID will not change and SNMPv3 authentication will work as before. [PR1256166](#)

Platform and Infrastructure

- In a L2PT scenario using an EX4300 Series switch, the egress PE (provider edge) device(EX4300) sends out frames towards the CE (customer edge) device with a duplicated frame of destination MAC address of 01:00:0c:cd:cd:d0 which is rewritten by the ingress PE device period.[PR1251391](#)

Port Security

- On a EX4600 connected either back to back or in a Virtual Chassis that has MACsec configured, a MACsec tunnel comes up, but after sometime or when removing optics (for failure testing) the MACsec tunnel will not come up again. [PR1234447](#)

Routing Protocols

- On an EX9200 switch, when you perform an in-service software upgrade (ISSU) while protocol sessions are active, the protocols might go down and come back up again, which can cause traffic loss for up to 40 seconds. [PR1247937](#)

SEE ALSO

[New and Changed Features | 23](#)

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

Known Behavior	 29
Known Issues	 30
Documentation Updates	 34
Migration, Upgrade, and Downgrade Instructions	 34
Product Compatibility	 35

Documentation Updates

There are no errata or changes in Junos OS Release 17.2R1 for the EX Series switches documentation.

SEE ALSO

New and Changed Features	 23
Changes in Behavior and Syntax	 28
Known Behavior	 29
Known Issues	 30
Resolved Issues	 32
Migration, Upgrade, and Downgrade Instructions	 34
Product Compatibility	 35

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 35

This section contains the upgrade and downgrade support policy for Junos OS for the 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. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

SEE ALSO

New and Changed Features 23
Changes in Behavior and Syntax 28
Known Behavior 29
Known Issues 30
Resolved Issues 32
Documentation Updates 34
Product Compatibility 35

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 36](#)

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	 23
Changes in Behavior and Syntax	 28
Known Behavior	 29
Known Issues	 30
Resolved Issues	 32
Documentation Updates	 34
Migration, Upgrade, and Downgrade Instructions	 34

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

●	New and Changed Features	 37
●	Changes in Behavior and Syntax	 39
●	Known Behavior	 40
●	Known Issues	 40
●	Resolved Issues	 41
●	Documentation Updates	 41

- Migration, Upgrade, and Downgrade Instructions | 42
- Product Compatibility | 50

These release notes accompany Junos OS Release 17.2R1 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at <https://www.juniper.net/documentation/software/junos/>.

New and Changed Features

IN THIS SECTION

- Interfaces and Chassis | 38
- Layer 2 Features | 38

This section describes the new features and enhancements to existing features in Junos OS Release 17.2R1 for Junos Fusion Enterprise.

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

Interfaces and Chassis

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

To explicitly configure full duplex:

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

To verify a half-duplex setting:

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

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

Layer 2 Features

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

In a JFE PVLAN topology:

- Multiple satellite devices can be clustered into a group and cabled into the JFE as a group instead of as individual satellite devices.
- Aggregation device native ports or satellite device extended ports can act as promiscuous ports, isolated ports, or community VLAN ports.

- The promiscuous port can be attached to a core switch or router through physical interfaces or aggregated links.
- PVLANs are supported in dual aggregation device JFEs.

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

SEE ALSO

Changes in Behavior and Syntax	 39
Known Behavior	 40
Known Issues	 40
Resolved Issues	 41
Documentation Updates	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 50

Changes in Behavior and Syntax

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 17.2R1 for Junos Fusion Enterprise.

SEE ALSO

New and Changed Features	 37
Known Behavior	 40
Known Issues	 40
Resolved Issues	 41
Documentation Updates	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 50

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 for Junos Fusion Enterprise.

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

SEE ALSO

New and Changed Features	 37
Changes in Behavior and Syntax	 39
Known Issues	 40
Resolved Issues	 41
Documentation Updates	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 50

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.2R1 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.

SEE ALSO

New and Changed Features	 37
Changes in Behavior and Syntax	 39
Known Behavior	 40
Resolved Issues	 41
Documentation Updates	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 50

Resolved Issues

There are no fixed issues in Junos OS Release 17.2R1 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.

SEE ALSO

New and Changed Features	 37
Changes in Behavior and Syntax	 39
Known Behavior	 40
Known Issues	 40
Documentation Updates	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 50

Documentation Updates

There are no errata or changes in Junos OS Release 17.2R1 for Junos Fusion Enterprise documentation.

SEE ALSO

New and Changed Features	 37
Changes in Behavior and Syntax	 39
Known Behavior	 40
Known Issues	 40
Resolved Issues	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 50

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading Junos OS on an Aggregation Device | 42
- Upgrading an Aggregation Device with Redundant Routing Engines | 44
- Preparing the Switch for Satellite Device Conversion | 44
- Converting a Satellite Device to a Standalone Switch | 46
- Upgrade and Downgrade Support Policy for Junos OS Releases | 48
- Downgrading from Release 17.2 | 48

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 17.2R1:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.
4. Click on the **Junos Fusion EX9200 (Enterprise)** title to expand the list of supported releases.
5. Click on the release number (the software version that you want to download) from the list.
6. Select the aggregation device software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **junos-install** package on the aggregation device.

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

Customers in the United States and Canada, use the following commands.

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-17.2R1.13.tgz
```

All other customers, use the following commands.

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-17.2R1.13-limited.tgz
```

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos Fusion Enterprise. See [Configuring or Expanding a Junos Fusion Enterprise](#).

For satellite device hardware and software requirements, see [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#).

Use the following command to install Junos OS on a switch before converting it into a satellite device:

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

NOTE: The following conditions must be met before a Junos switch that is running Junos OS Release 17.2R1 can be converted to a satellite device when the action is initiated from the aggregation device:

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

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

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

```
[edit]
user@satellite-device# request system zeroize
```

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

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

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

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

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

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

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

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

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, or preconfiguration. See [Configuring or Expanding a Junos Fusion Enterprise](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Switch

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

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://www.juniper.net/support/downloads>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** from the menu and select the switch platform series and model for your satellite device.
4. Select the software image for your platform. For satellite device software requirements, see [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#).
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 an EX4300 switch acting as the satellite device using FPC slot 102:

```
[edit]  
user@aggregation-device> request chassis satellite install  
/var/tmp/jinstall-ex-4300-14.1X53-D35.3-domestic-signed.tgz 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. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

Downgrading from Release 17.2

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

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

To downgrade a Junos Fusion Enterprise from Junos OS Release 17.2R1, you must first downgrade the satellite software version on the satellite devices.

1. Downgrade the satellite software on the satellite devices from 3.0R1 to 2.0R1:


```
user@aggregation-device> request system software add package-name no-validate upgrade-group
cluster1
```

For example, to downgrade the satellite software to 2.0R1:

```
user@aggregation-device> request system software add satellite-2.0R1-signed.tgz no-validate
upgrade-group cluster1
```

After the satellite devices are downgraded to satellite software, they will not show as being online until the aggregation device is downgraded to a compatible software version. To check software compatibility, see <https://www.juniper.net/support/downloads/solutions/fusion/>.

2. Downgrade the aggregation device. Follow the procedure for upgrading, but replace the 17.2 **junos-install** package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

[New and Changed Features | 37](#)

[Changes in Behavior and Syntax | 39](#)

[Known Behavior | 40](#)

[Known Issues | 40](#)

[Resolved Issues | 41](#)

[Documentation Updates | 41](#)

[Product Compatibility | 50](#)

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 50](#)
- [Hardware Compatibility Tool | 50](#)

Hardware and Software Compatibility

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 37
Changes in Behavior and Syntax 39
Known Behavior 40
Known Issues 40
Resolved Issues 41
Documentation Updates 41
Migration, Upgrade, and Downgrade Instructions 42

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

- New and Changed Features | 51
- Changes in Behavior and Syntax | 66
- Known Behavior | 67
- Known Issues | 68
- Resolved Issues | 69
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 70
- Product Compatibility | 86

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Junos Fusion Data Center | 52
- Hardware | 53
- Class of Service (CoS) | 53
- High Availability (HA) and Resiliency | 53
- Interfaces | 53
- Layer 2 Protocols | 58
- Layer 3 Protocols | 59
- Multicast Protocols | 60
- Network Management and Monitoring | 62

- [Routing Policy and Firewall Filters | 63](#)
- [Storage | 63](#)
- [Software Installation and Upgrade | 64](#)
- [Software Licensing | 65](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.2R1 for Junos Fusion Data Center.

Junos Fusion Data Center

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

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

- QFX10002 switches

You can configure the following switches as satellite devices:

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

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

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

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

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

Hardware

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

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

Class of Service (CoS)

- **Class of Service Support (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports the standard Junos class of service (CoS) features and operational commands in either a single or dual aggregation device configuration. Each extended port on a satellite device is a logical extension to the aggregation device. Therefore, the default CoS policy on the aggregation device applies to each extended port. You can also create standard CoS policies for extended ports.

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

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

High Availability (HA) and Resiliency

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

[See [Understanding VRRP](#).]

Interfaces

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

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

LACP is supported in single and dual aggregation device topologies.

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

- **Increased number of aggregated Ethernet interfaces (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can configure up to 1,000 aggregated Ethernet interfaces for a Junos Fusion Data Center system. To configure, include the **device-count** statement with a value of 1000 at the **[edit chassis aggregated-devices ethernet]** hierarchy level and add member links in each bundle.
- **Automatic ICCP Provisioning of an Interchassis Link in a Junos Fusion (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, automatic ICCP provisioning of an interchassis link (ICL) simplifies configuration of a Junos Fusion with dual aggregation devices by automatically provisioning the ICCP configuration within the Junos Fusion, instead of requiring the user to manually configure all ICCP parameters.

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

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

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

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

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

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

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

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

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

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

Uplink port pinning is configured in the following steps:

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

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

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

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

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

EX4300 and QFX5100 switches configured as satellite devices have a default set of uplink interfaces. [Table 1 on page 56](#) shows the default set of uplink interfaces that UFD selects for failure detection:

Table 1: UFD Default Uplink Interfaces for Satellite Devices

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

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

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



```

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

```

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

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

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

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

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

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

Layer 2 Protocols

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

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

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

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

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

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

Configuration notes for VLAN Auto Sensing:

- VLAN auto sensing is supported on extended ports only.
- Only single VLAN tagged packets are auto sensed.

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

- **Loop detection on extended ports (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can configure a Junos Fusion Data Center system to detect and break loops of unicast traffic on downstream extended ports without configuring spanning tree protocols. Typically, the loops are caused by either miswiring or by misconfiguration. Loop detection transmits special protocol data units (PDUs) periodically, and if a PDU is received on an extended port, the loop is detected and broken. Loop detection blocks the ingress port and issues a loop detection PDU error. When a port is blocked, you need to manually bring up the interface. Loop detection only responds to detect PDUs, not BPDUs.

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

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

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

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

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

- **VSTP Enhancements (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, VSTP is supported on a QFX10000 switch acting as a single-homed aggregation device. The VSTP configuration can include native ports or extended ports in a Junos Fusion Data Center.

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

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

Layer 3 Protocols

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

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

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

Multicast Protocols

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

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

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

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

Use the following commands to display local replication information:

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

- **show multicast snooping route satellite**
- **show multicast statistics satellite**
- **show multicast summary satellite**

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

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

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

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

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

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

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

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

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

Network Management and Monitoring

- **Local port mirroring (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, Junos Fusion Data Center supports local port mirroring. Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. You can use local port mirroring to troubleshoot and monitor applications. You can mirror packets per port, and you can configure the source and mirror ports on the same satellite device.

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

- **Analyzers on extended ports (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, you can use port mirroring (analyzers) on extended ports on satellite devices in a Junos Fusion Data Center. Extended-port port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. When a port is ingress-mirrored, any packet received on that port is mirrored to the user-configured destination. When a port is egress-mirrored, any packet transmitted from that port is mirrored to your configured port-mirroring destination.

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

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

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

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

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

- **Chassis MIB support (Junos Fusion Data Center)**—Starting in Junos OS Release 17.2R1, satellite devices in a Junos Fusion topology are represented in the chassis MIB. Satellite devices are represented as FPC slots (100, 101,102,..) in the aggregation device. The support is enabled using a range of container

indexes, which enable the SNMP process to redirect SNMP requests to the chassis process or SPMD based on the first index entry.

The following tables have been implemented for satellite devices:

- jnxContainersTable
- jnxContentsTable
- jnxFilledTable
- jnxOperatingTable
- jnxFRUTable

These tables have an index consisting of four entries—CIDX, L1, L2, and L3—where CIDX is the container index, L1 is the level-1 index (slot), L2 is the level-2 index (component within the L1 slot), and L3 is the level-3 index (subcomponent of the L2 slot). For example, an entry of 104.101.2.0 would indicate CIDX 104 is the index for fans and refers to the second fan of satellite device 100 (all indexes are 1-based). The CIDX indexes for Junos Fusion are offset by 100 from regular indexes; for example, a regular CIDX 2 (power supply) is 102 for the power supply of the satellite device. Using these indexes, you can distinguish the satellite device hardware from the aggregation device hardware.

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

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

Routing Policy and Firewall Filters

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

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

Storage

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

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

[See [Understanding DCBX](#).]

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

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

Software Installation and Upgrade

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

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

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

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

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

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

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

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

Software Licensing

- **Licensing model (Junos Fusion Data Center)**—Starting with Junos OS Release 17.2R1, you need to install a Junos Fusion license in addition to any other feature licenses that you install to track and activate the following models that are shipped with satellite software. These models can only be used as satellite devices:
 - QFX5100-48SH-AFO
 - QFX5100-48SH-AFI
 - QFX5100-48TH-AFO
 - QFX5100-48TH-AFI

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

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

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

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

- QFX10K-C1-JFS-32
- QFX10K-C1-JFS-64

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

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

SEE ALSO

Changes in Behavior and Syntax 66
Known Behavior 67
Known Issues 68
Resolved Issues 69
Documentation Updates 69
Migration, Upgrade, and Downgrade Instructions 70
Product Compatibility 86

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 17.2R1 or later for Junos Fusion Data Center.

Junos Fusion

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

SEE ALSO

New and Changed Features 51
Known Behavior 67
Known Issues 68

[Resolved Issues | 69](#)

[Documentation Updates | 69](#)

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

[Product Compatibility | 86](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 for Junos Fusion Data Center.

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

Junos Fusion Data Center

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

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

SEE ALSO

New and Changed Features	 51
Changes in Behavior and Syntax	 66
Known Issues	 68
Resolved Issues	 69
Documentation Updates	 69
Migration, Upgrade, and Downgrade Instructions	 70
Product Compatibility	 86

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 17.2R1 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

There are no known issues in this release.

SEE ALSO

New and Changed Features	 51
Changes in Behavior and Syntax	 66
Known Behavior	 67
Resolved Issues	 69
Documentation Updates	 69
Migration, Upgrade, and Downgrade Instructions	 70
Product Compatibility	 86

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R1 | 69](#)

This section lists the issues fixed in the Junos OS Release 17.2R1 for Junos Fusion Data Center.

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

Resolved Issues: 17.2R1

There are no issues resolved in this release.

SEE ALSO

New and Changed Features 51
Changes in Behavior and Syntax 66
Known Behavior 67
Known Issues 68
Documentation Updates 69
Migration, Upgrade, and Downgrade Instructions 70
Product Compatibility 86

Documentation Updates

This section lists the errata or changes in Junos OS Release 17.2R1 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	51
Changes in Behavior and Syntax	66
Known Behavior	67
Known Issues	68
Resolved Issues	69
Migration, Upgrade, and Downgrade Instructions	70
Product Compatibility	86

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 71
- Preparing the Switch for Satellite Device Conversion | 72
- Autoconverting a Switch into a Satellite Device | 75
- Manually Converting a Switch into a Satellite Device | 78
- Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 80
- Configuring Satellite Device Upgrade Groups | 81
- Converting a Satellite Device to a Standalone Device | 82
- Upgrade and Downgrade Support Policy for Junos OS Releases | 85
- Downgrading from Release 17.2R1 | 85

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

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

Basic Procedure for Upgrading an Aggregation Device

When upgrading or downgrading Junos OS, always use the **jinstall** package. Use other packages (such as the **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 17.2R1 is different from previous Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to access the Junos Fusion Hardware and Software Compatibility page.
4. Click on the **Junos Fusion Data Center (QFX10000)** title to expand the list of supported releases.
5. Click on the release number (the software version that you want to download) from the list.
6. Select the aggregation device software package for the release.
7. Review and accept the End User License Agreement.

8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the aggregation device.

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

For customers in the United States, Canada, and worldwide, use the following command.

```
user@host> request system software add reboot
source/jinstall-host-qfx-10-f-x86-64-17.2R1.13-secure-signed.tgz
```

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.2R1 **jinstall** package, you cannot issue the **request system software rollback** command to return to the previously installed software. Instead you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos Fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target switch to an interim Junos OS software version that can be

converted to satellite software. [Table 2 on page 73](#) shows a support matrix that maps Junos OS software used in aggregation devices to the compatible preconverted switch software and satellite device software.

Table 2: Aggregation Device Junos OS Software Compatibility with Satellite Software

Aggregation Device Version	Switch Version (preconversion)	Satellite Device Software Version
Junos OS Release 17.2R1	Junos OS Release 14.1X53-D43 or later	3.1R1

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add reboot
source/jinstall-qfx-5-14.1X53-D43.7-domestic-signed.tgz
```

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

1. Log in to the device using the console port.
2. Clear the device and set it to a factory-default state:

```
[edit]
user@satellite-device# request system zeroize
```

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

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

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

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

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

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

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

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

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

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

Autoconverting a Switch into a Satellite Device

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

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

Before you begin:

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

To autoconvert a switch into a satellite device:

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

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

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

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

2. Log in to the aggregation device.

3. Configure the cascade ports.

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

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

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

Examples:

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

```
[edit]
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7. Enable automatic satellite conversion:

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

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

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

8. Commit the configuration:

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

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

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

Manually Converting a Switch into a Satellite Device

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

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

Before you begin:

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

To manually convert a switch into a satellite device:

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

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

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

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

Examples:

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

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

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

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

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

```
[edit]
```

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

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

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

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

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

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

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

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

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

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

6. Manually configure the switch into a satellite device:

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

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

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

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

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

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

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

Before you begin:

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

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

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

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-3.1R1.3-signed.tgz** stored in the **/var/tmp/** directory on the switch:

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

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

- To install satellite software onto a EX4300 switch, use the **satellite-ppc-3.1R1.3-signed.tgz** satellite software package.
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** is the member numbers of the satellite devices that are being added to the upgrade group. If you enter an existing upgrade group name as the **upgrade-group-name**, you add new satellite devices to the existing software upgrade group.

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

```
[edit]
```

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

NOTE:

- If the satellite device is an EX4300 switch, you install a standard, signed jinstall-ex-4300 version of Junos OS.
- The QFX5100-48SH and QFX5100-48TH models are shipped from the factory with satellite device software. You cannot convert these models to become a standalone device.
- If the satellite device is a QFX5100 switch that can be converted to a standalone device, you need to install a signed PXE version of Junos OS. The PXE version of Junos OS is the software that includes **pxe** in the Junos OS package name when you download it from the Software Center. For example, the PXE software package for Junos OS Release 14.1X53-D43 is named install-media-pxe-qfx-5-14.1X53D43.7-signed.tgz.
- Before you install a signed PXE software package on a QFX5100 switch acting as a satellite device, you must first prepare the switch in one of two ways:
 - Issue the **request system zeroize** command to revert the device to a factory-default state.
 - Configure the device by including the **auto-satellite-conversion** statement at the **[edit chassis]** hierarchy level.

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

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to access the Junos Fusion Hardware and Software Compatibility page.
4. Click on the **Junos Fusion Data Center (QFX10000)** title to expand the list of supported releases.
5. Click on the Junos OS release number associated with the aggregation device from the list.
6. Select the PXE device software package for your satellite device platform.
7. Review and accept the End User License Agreement.

8. Download the software to a local host.
9. Copy the software to the device or to your internal software distribution site.
10. 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
```

11. Commit the configuration.

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

12. 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-5-14.1X53D43.7-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-4300-14.1X53D43.7-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.

13. Wait for the reboot that accompanies the software installation to complete.
14. 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://www.juniper.net/support/eol/junos.html>.

Downgrading from Release 17.2R1

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

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

New and Changed Features 51
Changes in Behavior and Syntax 66
Known Behavior 67
Known Issues 68
Resolved Issues 69
Documentation Updates 69
Product Compatibility 86

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 86](#)

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	51
Changes in Behavior and Syntax	66
Known Behavior	67
Known Issues	68
Resolved Issues	69
Documentation Updates	69
Migration, Upgrade, and Downgrade Instructions	70

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

●	New and Changed Features	88
●	Changes in Behavior and Syntax	89
●	Known Behavior	90
●	Known Issues	90
●	Resolved Issues	91
●	Documentation Updates	91
●	Migration, Upgrade, and Downgrade Instructions	92
●	Product Compatibility	100

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- [Class of Service \(CoS\) | 88](#)
- [Junos Fusion | 88](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.2R1 for Junos Fusion Provider Edge.

Class of Service (CoS)

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

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

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

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

Junos Fusion

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

- **Support for an ingress policer**—Starting in Junos OS Release 17.2R1, Junos Fusion Provider Edge supports the use of an ingress policer to filter incoming traffic at the extended port level. This feature supports a two-color policer that allows you to limit the traffic that is received on an interface. You can configure the Layer 2 ingress policer by using the **input-policer** statement at the **[edit interfaces *interface-name* layer2-policer]** hierarchy level.

SEE ALSO

Changes in Behavior and Syntax 89
Known Behavior 90
Known Issues 90
Resolved Issues 91
Documentation Updates 91
Migration, Upgrade, and Downgrade Instructions 92
Product Compatibility 100

Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 17.2R1.

SEE ALSO

New and Changed Features 88
Known Behavior 90
Known Issues 90
Resolved Issues 91
Documentation Updates 91
Migration, Upgrade, and Downgrade Instructions 92
Product Compatibility 100

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 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	 88
Changes in Behavior and Syntax	 89
Known Issues	 90
Resolved Issues	 91
Documentation Updates	 91
Migration, Upgrade, and Downgrade Instructions	 92
Product Compatibility	 100

Known Issues

There are no known issues in the Junos OS Release 17.2R1 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	 88
Changes in Behavior and Syntax	 89
Known Behavior	 90
Resolved Issues	 91
Documentation Updates	 91
Migration, Upgrade, and Downgrade Instructions	 92
Product Compatibility	 100

Resolved Issues

There are no fixed issues in the Junos OS Release 17.2R1 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	 88
Changes in Behavior and Syntax	 89
Known Behavior	 90
Known Issues	 90
Documentation Updates	 91
Migration, Upgrade, and Downgrade Instructions	 92
Product Compatibility	 100

Documentation Updates

There are no errata or changes in Junos OS Release 17.2R1 for Junos Fusion Provider Edge documentation.

SEE ALSO

New and Changed Features	 88
Changes in Behavior and Syntax	 89
Known Behavior	 90
Known Issues	 90
Resolved Issues	 91
Migration, Upgrade, and Downgrade Instructions	 92
Product Compatibility	 100

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 92
- Upgrading an Aggregation Device with Redundant Routing Engines | 94
- Preparing the Switch for Satellite Device Conversion | 95
- Converting a Satellite Device to a Standalone Device | 96
- Upgrading an Aggregation Device | 98
- Upgrade and Downgrade Support Policy for Junos OS Releases | 99
- Downgrading from Release 17.2 | 99

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

The download and installation process for Junos OS Release 16.1R1 and later is different that for earlier Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to access the Junos Fusion Hardware and Software Compatibility page.
4. Click on the **Junos Fusion MX Series (Provider Edge)** title to expand the list of supported releases.
5. Click on the release number (the software version that you want to download) from the list.
6. Select the aggregation device software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the aggregation device.

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

- For 64-bit software:

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

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

For example:

```
user@host> request system software add validate reboot
source/junos-install-mx-x86-64-17.2R1.9-signed.tgz
```

- For 32-bit software:

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

For example:

```
user@host> request system software add validate reboot
source/junos-install-mx-x86-32-17.2R1.9-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.2R1 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-qfx-5-14.1X53-D43.7-domestic-signed.tgz
```

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

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

```
[edit]
user@satellite-device# request system zeroize
```

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

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

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

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

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

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

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

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

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See [Configuring Junos Fusion Provider Edge](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Device

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

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

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

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

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

[edit]

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

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

[edit]

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

[edit]

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

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

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

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

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

11. Wait for the reboot that accompanies the software installation to complete.
12. When you are prompted to log back 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 17.2R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R3.

Upgrade and Downgrade Support Policy for Junos OS Releases

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


You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

Downgrading from Release 17.2

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

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

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

SEE ALSO

New and Changed Features	88
Changes in Behavior and Syntax	89
Known Behavior	90
Known Issues	90
Resolved Issues	91
Documentation Updates	91

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 100](#)

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 88
Changes in Behavior and Syntax 89
Known Behavior 90
Known Issues 90
Resolved Issues 91
Documentation Updates 91
Migration, Upgrade, and Downgrade Instructions 92

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

IN THIS SECTION

- New and Changed Features | 101
- Changes in Behavior and Syntax | 132
- Known Behavior | 143
- Known Issues | 147
- Resolved Issues | 154
- Documentation Updates | 164
- Migration, Upgrade, and Downgrade Instructions | 166
- Product Compatibility | 173

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Hardware | 102
- Class of Service (CoS) | 103
- EVPNs | 104
- Forwarding and Sampling | 106
- General Routing | 107
- High Availability (HA) and Resiliency | 108
- Interfaces and Chassis | 109
- Layer 2 VPN | 112

- Management | 112
- MPLS | 114
- Multicast | 115
- Network Management and Monitoring | 116
- Operation, Administration, and Maintenance (OAM) | 118
- Routing Policy and Firewall Filters | 119
- Routing Protocols | 119
- Services Applications | 121
- Software Defined Networking (SDN) | 123
- Subscriber Management and Services | 124
- Virtual Chassis | 131

This section describes the new features and enhancements to existing features in Junos OS Release 17.2R1 for the MX Series.

Hardware

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

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

NOTE: The Routing Engines are equipped with limited encryption support only. The Junos Limited image does not have data-plane encryption and is intended only for countries in the Eurasian Customs Union because these countries have import restrictions on software containing data plane encryption. See [Junos OS Editions](#).

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

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

Class of Service (CoS)

- **Support for user-configurable traffic class map (MX routers with MPCs)** — Beginning with Junos OS Release 17.2, MX routers with MPCs support a user-configurable input priority map, known as a **traffic-class-map**, that enables you to prioritize and classify input traffic entering a Packet Forwarding Engine during ingress oversubscription. You can define traffic class maps for a packet based on DSCP, IP precedence, MPLS EXP, IEEE 802.1p, and IEEE 802.1ad CoS values and associate these CoS values with **real-time**, **network-control**, and **best-effort** traffic classes.

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

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

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

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

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

EVPNs

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

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

Unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU upgrade is only supported by dual Routing Engine platforms. Unified ISSU requires both GRES and NSR to be enabled.

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

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

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

- **Unified ISSU support for EVPN and VXLAN**—Starting in Junos OS Release 17.2R1, Junos OS supports Unified ISSU on EVPN and VXLAN. Unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU upgrade is only supported by dual Routing Engine platforms. Unified ISSU requires both GRES and NSR to be enabled.

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

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

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

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

The service adheres to the following forwarding rules:

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

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

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

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

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

- **Interconnecting data center networks over WAN (MX Series)**—Starting in Junos OS Release 17.2R1, you can interconnect data center networks running Ethernet VPN (EVPN) with Virtual Extensible LAN (VXLAN) encapsulation through a WAN running MPLS-based EVPN. This feature enables you to:
 - Connect data center edge routers over MPLS-based EVPN WAN for data center interconnections.
 - Interconnect EVPN-VXLAN and EVPN-MPLS using logical tunnel (lt-) interface on data center edge routers.

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

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

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

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

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

Unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU upgrade is only supported by dual Routing Engine platforms. Unified ISSU requires both GRES and NSR to be enabled.

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

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

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

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

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

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

Forwarding and Sampling

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

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

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

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

General Routing

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

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

- **OpenConfig: BGP Routing Table - Supporting Operational State Model (MX Series)**—Starting in Junos OS 17.2R1, the OpenConfig BGP RIB routing table supports local-rib for IPV4 and IPV6. The Openconfig-rib-bgp.yang model supports five logical RIBs per address family. There are five tables for IPV4 routes and 5 tables for IPV6 routes.
- **Support for PTP over Ethernet, hybrid mode, and G.8275.1 profile on MPC6E, MPC2E NG, MPC3E NG MPCs**—Starting in Junos OS Release 17.2R1, MPC6E, MPC2E NG, and MPC3E NG MPCs support the following features:
 - PTP over Ethernet— PTP over Ethernet enables effective implementation of packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks. PTP over Ethernet uses multicast addresses for communication of PTP messages between the slave clock and the master clock.
 - Hybrid Mode— In hybrid mode, the synchronous Ethernet equipment clock (EEC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP.
 - G.8275.1 profile— G.8275.1 is a PTP profile for applications that require accurate phase and time synchronization. It supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support and is based on the second version of PTP defined in IEEE 1588. You can configure the G.8275.1 profile by including the **profile-type g.8275.1** statement at the **[edit protocols ptp]** hierarchy level.

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

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

- **Enhancements to Precision Time Protocol feature (MX104)**—Starting in Junos OS Release 17.2R1, Precision Timing Protocol (PTP) feature in MX104 routers has been enhanced with the following changes:
 - After PTP is phase-aligned, if the system up time is less than 30 minutes and the PTP source is lost before 30 minutes, the PTP state will be moved to **freerun**. On the other hand, if the system up time is more than 30 minutes and the PTP source is lost, the PTP state will move to **holdover**.
 - If PTP is never phase-aligned and PTP source is lost, the PTP state shall move to **freerun**.

- While operating in PTP Hybrid mode, the state of PTP will be in **holdover** for 8 days after a PTP clock source is lost but a valid high stratum SyncE source is present.
- PTP state will transition to **holdover** irrespective of the current state of **acquiring** or **phase aligned** as long as PTP was phase-aligned once and system uptime was more than 30 minutes.
- **New command to display upstream and downstream clock information (MX104)**—Starting with Junos OS Release 17.2R1, a new show command, *show ptp all-master-clock*, is introduced to display all the upstream master information and clock advertised to downstream. This command is supported only on MX104 routers.
- **OpenConfig: Supporting the BGP Model in Junos OS (MX Series)**—Starting in Junos OS 17.2R1, the configuration leafs defined in the **openconfig-bgp.yang** and **openconfig-bgp-multiprotocol.yang** files are supported.

High Availability (HA) and Resiliency

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

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

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

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

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

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

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

Unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

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

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

Interfaces and Chassis

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

The following key Junos OS features are supported:

- Basic Layer 2 features including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
- Class of service (CoS)
- Firewall filters and policers
- Integrated routing and bridging (IRB)
- Interoperability with existing MPCs (excluding the Application Services Modular Carrier Card, or AS-MCC)
- Layer 2 protocols
- Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs
- Layer 3 routing protocols and MPLS
- Layer 3 services supported on MS-MIC and MS-MPC (for example, CGNAT, IP Security, inline active flow monitoring) and inline services
- Multicast forwarding
- Port mirroring
- Spanning-tree protocols, such as STP, MSTP, RSTP, and VSTP
- Synchronous Ethernet and Precision Time Protocol (IEEE 1588)
- Tunneling
- Graceful Routing Engine Switchover (GRES) and Non Stop Routing (NSR)

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

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

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

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

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

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

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

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

Junos OS support for FRU management of MX2008 routers—In Junos OS Release 15.1F7 and 17.2R1, Junos OS supports the MX2008 router (model number: CHAS-MX2008). The Junos OS chassis management software for the MX2008 routers provides enhanced environmental monitoring and field-replaceable unit (FRU) control.

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

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

The MX2008 cooling system contains two fan trays, with six fans in each. The fan trays can be installed at or removed from the back of the chassis, which allows the space in the front to be used for cable

management. The MX2008 supports temperature thresholds for each temperature sensor, which enables the router to precisely control the cooling, raise alarms, and shut down a FRU.

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

- **Limited encryption Junos OS image and boot restriction (MX Series)**—Starting with Junos OS Release 17.2R1, the MX240, MX480, MX960, MX2010, and MX2020 routers with the Routing Engines RE-S-X6-64G-LT and RE-MX2K-X8-64G-LT support only Junos Limited image. The Junos Limited image does not have data-plane encryption and is intended only for countries in the Eurasian Customs Union because these countries have import restrictions on software containing data plane encryption. Unlike the Junos Worldwide image, the Junos Limited image supports control plane encryption through Secure Shell (SSH) and Secure Sockets Layer (SSL), thus allowing secure management of the system. The Routing Engines are restricted to boot only the Junos Limited image.
- **Enhancement to ambient-temperature statement (MX Series)**—In Junos OS Release 15.1F4 and later, the default ambient temperature is set at 40° C on MX480, MX960, MX2010, and MX2020 Universal Routing Platforms. You can override ambient temperature by setting the temperature at 55° C or 25° C.

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

When a router restarts, the system adjusts the power allocation or the provisioned power for the line cards on the basis of the configured ambient temperature. If enough power is not available, a minor chassis alarm is raised. However, the chassis continues to run with the configured ambient temperature. You can configure a new higher ambient temperature only after you make more power available by adding new power supply modules or by taking a few line cards offline. By using the provisioned power that is saved by configuring a lower ambient temperature, you can bring more hardware components online.

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

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

Layer 2 VPN

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

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

Management

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

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

- **Support for LSP events and properties sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, you can export statistics for LSP events and properties through the Junos Telemetry Interface. Only gRPC streaming for this sensor is supported. You can export statistics for ingress point-to-point LSPs, point-to-multipoint LSPs, bypass LSPs, and dynamically created LSPs. To export data through gRPC, use the **/mpls/lsp/** or **/mpls/signal-protocols/** set of OpenConfig subscription paths. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of the Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models. To enable the LSP events and properties sensor, you must also include the **lsp-telemetry** statement at the **[edit routing-options]** hierarchy level.

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

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

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

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

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

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

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

- **Support for LSP statistics for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, you can stream telemetry data for LSPs through UDP and gRPC. To provision an LSP statistics sensor for UDP streaming, include the **resource /junos/services/label-switched-path/usage/** statement at the **[edit services analytics sensor sensor-name]** hierarchy level. Use the **mpls/lsp/constrained-path/tunnels/tunnel/** path to provision a sensor for streaming LSP statistics through gRPC. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions. For both UDP and gRPC streaming, you must also configure the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level. Additionally, MX Series routers should operate in enhanced mode. Support for the LSP statistics sensor was previously introduced in Junos OS Release 15.1F6 and Junos OS Release 16.1R4.

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

- **Support for device family and release in Junos OS YANG modules (MX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command**

module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

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

MPLS

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

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

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

To configure SPRING-TE for PCEP:

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

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

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

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

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

- **IPv6 support for static egress LSPs (MX Series)**—Starting in Junos OS Release 17.2R1, static LSPs on the egress router can be configured with IPv6 as the next-hop address for forwarding IPv6 traffic.

Previously, only IPv4 static LSPs were supported. The IPv6 static LSPs share the same transit, bypass, and static LSP features of IPv4 static LSPs.

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

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

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

NOTE:

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

Multicast

- **Improved multicast performance using distributed IGMP (MX Series)**—Starting in Junos OS Release 17.2R2, you can improve multicast performance by using the distributed Internet Group Management Protocol (IGMP). Distributed IGMP moves IGMP processing from the Routing Engine to the Packet Forwarding Engine. When you configure distributed IGMP, join and leave events are processed across multiple Modular Port Concentrators (MPCs) on the Packet Forwarding Engine. Instead of being processed through a centralized routing protocol process (rpd) on the Routing Engine, this improves performance and decreases join and leave latency.

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

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

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

[See [Understanding IGMP](#).]

Network Management and Monitoring

- **MIB enhancement for jnxPPPoESubIfTable and jnxSubscriberTable tables (MX Series)**—Starting in Junos OS Release 17.2, you can correlate information between the jnxPPPoESubIfTable and jnxSubscriberTable tables. Prior to Junos OS Release 17.2, you could not correlate information between the two tables because they are indexed differently. Now, the jnxPPPoESubIfTable can provide a subscriber session ID, which corresponds to each PPPoE session. This ID can be used to correlate information in the jnxSubscriberTable. Additionally, the physical interface and underlying interface names for a subscriber session are now available in the jnxSubscriberTable.
- **New indicators for the jnxLEDState MIB (MX960, MX2020, and MX2010 routers)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E include the following indicators for the jnxLEDState MIB object in the jnxLEDEntry MIB table:
 - off—Offline, not running
 - blinkingGreen—Entering state of ok, good, normally working
- **Support for kernel features on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E support the following features:
 - Addressing the IPv6 NDP DoS issue —You can address the IPv6 Neighbor Discovery Protocol (NDP) denial-of-service (DoS) issue at the Routing Engine by using NDP inspection or protection to prioritize NDP activities on the Routing Engine.
 - Maximum period for autogeneration of keepalives by the kernel using precision timer feature—Precision timers in the kernel automatically generate keepalives on behalf of BGP for a specified maximum period of time after a switchover event from standby to master.
 - IPv6 support for traceroute with AS number lookup—IPv6 is supported for traceroute with the **as-number-lookup** option. Traceroute is an application used to display a list of routers between the device and a specified destination host.
 - Targeted aggregated Ethernet distribution—You can direct traffic through specified links of a logical interface of an aggregate Ethernet bundle that is configured without link protection. By configuring targeted aggregated Ethernet distribution, you can create distribution lists consisting of specific child member links.
 - Reduction in the number of IPCs between master agent and subagent- The SNMP GetBulk requests are converted to AgentX GetNext for the repetitions specified in the request. This might result in several inter-process communication (IPCs) between the master agent snmpd and subagent AgentX in proportion to the number of max-repetitions specified in the GetBulk request. The number of IPCs between the master agent and subagent can be reduced by translating GetBulk requests with a high max-repetitions count to a single request between the master agent snmp and the subagent AgentX.

- I3-level liveness detection mechanism for child links of ethernet LAG interface.
- Match-string functionality for efficient syslog message filtering.
- **Support for features on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E support the following features:
 - LDP in an IPv6 network only, and in an IPv6 or IPv4 dual-stack network.
 - The IS-IS protocol can restrict flooding of LSAs to control sharing of routes between multiple level-2 metro ring networks.
 - For routers operating in Enhanced IP Network Services mode, you can configure a threshold that triggers fast failover in next-generation MVPNs with hot-root standby on the basis of aggregate flow rate.
 - Control word feature for LDP VPLS and FEC 129 VPLS.
 - You can specify route prefix priority of high or low through the existing import policy in protocols. Through priority, you can control the order in which the routes get updated from LDP/OSPF to RPD, and RPD to kernel.
 - RSVP with traffic engineering (RSVP-TE) protocol extensions for fast reroute (FRR) facility protection to allow greater scalability of LSPs and faster convergence times.
 - The Junos OS implementation of MPLS RSVP-TE is scaled to enhance the usability, visibility, configuration, and troubleshooting of label-switched paths (LSPs).
 - Tables and objects defined in RFC 5132, *IP Multicast MIB*, except the ipMcastZoneTable table.
 - Agent Capabilities MIB provides information about the implementation characteristics of an Agent subsystem in a network management system.
 - You can prioritize BGP route updates by using output queues.
 - Flow-aware transport (FAT) label for BGP-signaled pseudowires such as Layer 2 VPN and VPLS.
 - The NLRI format available for BGP VPN multicast is changing from the existing format of SAFI 128 to SAFI 129 as defined in RFC 6514.
 - You can use the **import-labeled-routes** statement at the **[edit routing-instances routing-instance-name protocols vpls]** hierarchy level to specify one or more nondefault routing instances where you want MPLS pseudowire labeled routes to be leaked from the mpls.0 path routing table in the master routing instance.
 - You can configure BGP-ORR with IS-IS as the interior gateway protocol (IGP) on a route reflector to advertise the best path to the BGP-ORR client groups by using the shortest IGP metric from a client's perspective, instead of the route reflector's view.
- **RPM timestamping extension on MPC7E, MPC8E, and MPC9E line cards (MX Series)**—In Junos OS Release 17.2R1, MPC7E, MPC8E, and MPC9E support timestamping of RPM probes in the Packet Forwarding Engine host processor. You can enable this feature by including the **hardware-timestamp** statement at the **[edit services rpm probe probe-name test test-name]** hierarchy level.

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

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

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

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

Operation, Administration, and Maintenance (OAM)

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

Routing Policy and Firewall Filters

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

Routing Protocols

- **Support for BGP link-state distribution with SPRING extensions (MX Series)**—Starting in Junos OS Release 17.2R1, BGP link-state extensions export source packet routing in networking (SPRING) topology information to software-defined networking controllers. Controllers can get the topology information by either being a part of an interior gateway protocol (IGP) domain or through BGP link-state distribution. BGP link-state distribution is supported on inter-domain networks and provides a scalable mechanism to export the topology information. This feature benefits networks that are moving to SPRING but also have RSVP deployed, and continue to use both SPRING and RSVP in their networks.
[See [Link-State Distribution Using BGP Overview](#).]
- **Support for SRGB in SPRING for IS-IS (MX Series in enhanced-ip Mode)**—Starting with Junos OS Release 17.2R1, you can configure the segment routing global block (SRGB) range label used by source packet routing in networking (SPRING). Currently Junos OS allows you to configure only node segment indices. The value of the start label depends on the dynamic label available in the system. The labels from this SRGB range are used for SPRING in the IS-IS domain. The labels advertised are more predictable and deterministic across the segment routing domain.

- To configure the starting index value of the SRGB label block, use the **start-label *start-label-block-value*** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.
- To configure the index range of the SRGB label block, use the **index-range *value*** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.

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

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

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

- **FIB scaling and performance enhancements (MX Series)**—Starting in Junos OS Release 17.2R1, the Packet Forwarding Engine is enhanced to scale and support a higher number of routes in the forwarding information base (FIB), also known as forwarding table. However, during graceful Routing Engine switchover (GRES), when there are ten million IPv4 routes in the forwarding table, there is traffic loss. This traffic loss is not seen when a routing protocol process (rpd) runs in warm standby mode. We currently do not support unified ISSU and NSR at this scale.
- **Support for unique AS path count (MX Series)**—Starting with Junos OS Release 17.2R1, you can configure a routing policy to determine the number of unique autonomous systems (ASs) present in the AS path. The unique AS path count helps determine whether a given AS is present in the AS path multiple times, typically as prepended ASs. In earlier Junos releases it was not possible to implement this counting behavior using the **as-path** regular expression policy. This feature permits the user to configure a policy based on the number of AS hops between the route originator and receiver. This feature ignores ASs in the **as-path** that are confederation ASs, such as **confed_seq** and **confed_set**.

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

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

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

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

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

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

Services Applications

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

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

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

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

- **Introducing the Junos OS URL filtering feature (MX Series)**—Starting in Junos OS Release 17.2, you can use URL filtering to filter which Web content is accessible to users based on a set of criteria or *template*. Blacklisted URLs are maintained in a URL database file. These URLs are resolved by the URL filtering process (url-filterd) on the Routing Engine to a list of IP addresses that are downloaded to the URL Filter Plugin (jservices-urfl), which is added to the Multiservices PIC management process (mspmamd) running on the service PIC.

- **Support for inline 6rd and 6to4 (MX2020)**—Starting in Junos OS Release 17.2R1, you can also configure inline IPv6 rapid deployment (6rd) or IPv6 to IPv4 (6to4) on an MX2020 router on MPC7Es, MPC8Es, and MPC9Es. You can use the inline capability to avoid the cost of using services PICs for required tunneling, encapsulation, and de-encapsulation processes. Anycast is supported for 6to4 using next-hop service interfaces. Hairpinning is also supported for traffic between 6rd domains.

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

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

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

[See [Inline Sampling Overview](#)]

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

[See [ALG Descriptions](#).]

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

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

- **Support for AMS warm standby on MS-MPC and MS-MIC (MX Series routers)**—Starting in Junos OS Release 17.2R1, you can use the same services interface as the backup in multiple aggregated multiservices (AMS) interfaces, resulting in an N:1 warm standby option for MS-MPCs and MS-MICs. Each warm standby AMS interface contains two members. One member is the service interface you want to protect, called the primary interface, and the other member is the secondary (backup) interface. You can use the same secondary member interface in multiple warm standby AMS interfaces.

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

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

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

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

handshake, saving up to one full round-trip time (RTT) over standard TCP. TFO support is for MS-MPC and MS-MIC.

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

[See [Configuring FlowTapLite](#).]

Software Defined Networking (SDN)

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

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

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

- MX Series router acting as DCI and Layer 2 gateway to translate VLAN traffic coming from an EVPN (a remote data center) to VXLAN traffic
- MX Series router acting as DCI to connect different OVSDB domains through EVPN
- MX Series router acting as a layer 3 gateway to route between an VXLAN domain

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

Feature Explorer family: Software Defined Networking (SDN)

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

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

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

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

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

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

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

Subscriber Management and Services

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

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

[See [PIM Overview](#).]

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

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

- **Support for mapping VLAN session termination cause (MX Series)**—Starting in Junos OS Release 17.2R1, new internal identifiers indicate the reasons that autoconfd initiates termination of individual VLAN out-of-band subscriber sessions. In earlier releases, the termination cause for a VLAN session is always 6 (administrative reset) and cannot be modified.

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

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

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

- **Subscriber termination supported in dynamic-bridged GRE tunnels (MX Series)**—Starting in Junos OS Release 17.2R1, dynamic-bridged generic routing encapsulation (GRE) tunnels are created and terminated at the broadband network gateway (BNG) to support the MX Series deployed as a Wi-Fi access gateway

model. Dynamic Host Configuration Protocol (DHCP) subscribers are transported through GRE tunnels as either VLAN-tagged or untagged. Subscriber services such as authentication, authorization, and accounting (AAA); address assignment; and class of service (CoS) are supported for individual DHCP subscribers within the GRE tunnels.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Increased business services scale (MX Series)**—Starting in Junos Release 17.2R1, Enhanced Subscriber Services Manager (ESSM) can support up to 1000 business services per subscriber PPP session and up to 8000 business services per chassis. All combinations of subscribers and services are supported within those limits; for example, 8 subscribers with 1000 services each, 100 subscribers with 80 services each, and so on.
- **Support for bulk CoA (MX Series)**—Starting with Junos OS release 17.2R1, bulk change of authorization CoA is supported for RADIUS-based subscriber services. The two new Radius VSAs introduced are:
 - 26-194 (Bulk-CoA-Transaction-Id)
 - 26-195 (Bulk-CoA-Identifier)

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

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

- **Rapid drain mode for DHCP address pools and lease timer enhancements (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the DHCP local server to stop allocating addresses from a local pool and gracefully terminate subscribers that are using addresses from that pool. When a DHCP subscriber attempts to renew the IP address from a pool configured for active drain, the DHCP local server replies with a NAK to the subscriber's T1 renewal messages, forcing a renegotiation, at which time the server allocates a new IP address from an alternative address pool that is not configured for active drain.

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

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

- **Traffic throughput improvements for MPC5 and MPC6 cards (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the **host-prefix-only** statement on the underlying demux interface for static or dynamic VLANs to improve datapath performance for DHCPv4 access models. This statement has the following requirements:

- All the DHCPv4 subscribers using the underlying interface must be brought up using a 32-bit host prefix.
- You must configure the **demux-source inet** statement. You must not configure **demux-source inet6** or **demux-source [inet inet6]**.

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

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

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

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

[See [Subscriber Access Line Information Forwarding by the LAC Overview](#), [Transmission of Tx Connect-Speed and Rx Connect-Speeds from LAC to LNS](#), and [Configuring the LAC to Report Access Line Information to the LNS](#).]

- **Support for passing Framed-Route attributes from a RADIUS server. (MX Series)**—Starting in Junos OS Release 17.2, for routers running enhanced subscriber management, tagged subscriber host routes from a RADIUS server can be passively imported to the routing table and thus advertised by BGP. The following attributes are included: **tag**, **metric**, and **preference**. To view the attributes, use the **show system subscriber-management route prefix** command.

[See [show system subscriber-management route prefix](#).]

- **MLPPP support for LNS and PPPoE subscribers (MX Series)**—Starting in Junos OS Release 17.2, Multilink PPP (MLPPP) support is provided for static and dynamic LNS (L2TP network server) and PPPoE (Point-to-Point Protocol over Ethernet) terminated and tunneled subscribers running on MX Series with access-facing MPC2 slots. The following features are supported:

- Mixed mode for customers with both MLPPP and single link PPP subscribers
- Fragmentation-maps for both static and dynamic inline service **si** interfaces
- Coexistence support for member-link IFL and the bundle IFL on different lookup engines
- Link fragmentation and interleaving (LFI) for a single-link bundle
- Fragment reordering optimization
- **Targeted distribution of subscriber traffic over aggregated Ethernet**—Starting in Junos OS Release 17.2R1, for a demux configuration whose underlying interface is an aggregated Ethernet interface, Junos OS provides targeted distribution of subscriber traffic while also allowing subscriber traffic redundancy. This ensures equal distribution of bandwidth and CoS resources among subscribers.

Service providers can now:

- Provide DPC and port redundancy for subscriber traffic.
- Apply per-subscriber hierarchical QoS and firewall filters on subscriber traffic over LAG.

NOTE: The “targeted-distribution” feature needs to be defined on all levels of the profile that require targeted functionality. For example, if you have targeted distribution enabled on **dvlan** profile and you have dynamic client profile. If targeted distribution is required on dynamic client profile, then you have to enable it.

To set targeted distribution in the demux logical interfaces configuration, use the **targeted-distribution** statement at the **[edit interfaces demux0 unit *logical-unit-number*]** hierarchy level.

To schedule an automatic periodic rebalance on an aggregated Ethernet bundle, use the **rebalance-periodic start-time <hh:mm> interval <hours>** option at the **[edit interfaces *aenumber* aggregated-ether-options targeted-options]** hierarchy level.

To provide module redundancy for demux subscribers on aggregated Ethernet bundles configured with targeted distribution, set the **logical-interface-fpc-redundancy** option at the **[edit interfaces *aenumber* aggregated-ether-options targeted-options]** hierarchy level.

To configure rebalance subscriber granularity, use the **logical-interface-fpc-redundancy rebalance-subscriber-granularity <rebalance-subscriber-granularity>** option or **logical-interface-chassis-redundancy rebalance-subscriber-granularity <rebalance-subscriber-granularity>** option at the **[edit interfaces *ae<number>* aggregated-ether-options targeted-options]** hierarchy level.

To manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled, use the **request interface rebalance <interface-name>** command.

To display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle, use the **show interfaces targeting *aex*** command.

To view status information about the specified demux interface, use **show interfaces demux0.<logical-interface-number>** command.

To set targeted distribution in the VLAN logical interface configuration, use the **targeted-distribution** statement at the **[edit interfaces interface-set <interface-set name> demux0 unit logical-unit-number]** hierarchy level.

- **Configurable grace period for unresponsive RADIUS servers (MX Series)**—Starting in Junos OS Release 17.2R1, you can use the **timeout-grace** statement at the **[edit access radius-options]** hierarchy level to configure a grace period that determines when an unresponsive RADIUS authentication server is marked as down or unreachable. When the server fails to respond to any of the attempts made for an authentication request, it times out, the time is noted, and the grace period begins. If the server is unresponsive for subsequent authentication requests, the grace period is checked each time the server times out. When the check determines that the grace period has expired, the server is marked as down or unreachable.

In earlier releases, the grace period is 10 seconds and is not configurable.

[See [Configuring a Timeout Grace Period to Specify When RADIUS Servers Are Considered Down or Unreachable](#).]

- **ANCP agent adjustment of cell overhead for ADSL, ADSL2, and ADSL2+ subscriber lines (MX Series)**—Starting in Junos OS Release 17.2R1, you can configure the ANCP agent to adjust the value it reports to CoS for downstream subscriber traffic on cell-mode DSL types (ADSL, ADSL2, and ADSL2+). The adjusted values enable CoS to more accurately adjust the effective shaping rate for the downstream subscriber traffic.

Use the following statements to specify number of bytes that are added to or subtracted from the cell overhead for the traffic: **adsl-bytes**, **adsl2-bytes**, or **adsl2-plus-bytes**. Use the **show ancp cos** command to view the adjustment configuration and the last updated values sent to CoS. The **show class-of-service interface interface-name** command displays the adjusted overhead values CoS has received from the ANCP agent.

[See [Configuring the ANCP Agent to Report Traffic Rates to CoS](#).]

Virtual Chassis

- **VCP link hashing enhancements(MX Series)**—Starting in Junos OS Release 17.2R1, you can use Virtual Chassis port (VCP) link hashing more effectively. All links are equally utilized no matter how many VCP links are configured. This results in better load balancing and better utilization of VCP links under heavy traffic.

[See [Guidelines for Configuring Virtual Chassis Ports](#).]

- **Support for MX Virtual Chassis environment (MX Series Routers)**—Starting with Junos OS Release 17.2R1, MX240, MX480, and MX960 routers with the Routing engine RE-S-X6-64G support MX-Virtual Chassis environment.

SEE ALSO

[Changes in Behavior and Syntax | 132](#)

[Known Behavior | 143](#)

[Known Issues | 147](#)

[Resolved Issues | 154](#)

[Documentation Updates | 164](#)

[Migration, Upgrade, and Downgrade Instructions | 166](#)

[Product Compatibility | 173](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [EVPNs | 132](#)
- [Forwarding and Sampling | 133](#)
- [High Availability \(HA\) and Resiliency | 135](#)
- [Interfaces and Chassis | 135](#)
- [IP Tunneling | 138](#)
- [Management | 138](#)
- [MPLS | 139](#)
- [Network Management and Monitoring | 139](#)
- [Routing Protocols | 140](#)
- [Services Applications | 141](#)
- [Subscriber Management and Services | 141](#)
- [User Interface and Configuration | 142](#)

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 17.2R1 for MX Series.

EVPNs

- **EVPN E-tree extended community**—Starting in Junos OS Releases 16.1R5, 17.1R2, and 17.2R1 and later releases, the etree leaf indication bit and Leaf Label in EVPN E-tree extended community follows

the E-tree Extended Community as defined in the [E-TREE Support in EVPN & PBB-EVPN IET](#) IETF draft. A mixed network environment with routers running versions of Junos OS without this fix and routers with this fix would encounter unexpected forwarding behavior. Previous versions of Junos OS has incorrect label indication bit and leaf label encoding. Junos OS Release 16.1R4 has the incorrect label indication bit and leaf label encoding.

- **EVPN extended community and ISID using standard IANA value**—Starting in Junos OS Release 17.2R1, the router MAC extended community and service identifier (ISID) sub-type values have been corrected to use the Internet Assigned Numbers Authority (IANA) standardized value. In Junos OS Release 17.1R1, when you configure EVPN extended community using a pure type 5 routing mode with VXLAN encapsulation, you might encounter routing issues with the router from another vendor.

Forwarding and Sampling

- If a Packet Forwarding Engine (PFE) of an FPC is affected due to fabric path wedge errors, then as part of fabric hardening actions, the affected Packet Forwarding Engine is disabled and the associated fabric also goes offline. Fabric stream wedge occurs when the ASIC of the FPC is in the stuck state, and the ingress Packet Forwarding Engine fails to send traffic to the destination Packet Forwarding Engine. When the Packet Forwarding Engine is wedged, the fabric of the Packet Forwarding Engine goes offline. The output of **show chassis fabric fpcs** and **show chassis fabric plane** commands show a new state for the Packet Forwarding Engine as **Fabric Disabled**.

```
user@router> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    ...
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    ...
FPC 1
  PFE #0
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    ...
  PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
```

```

    Plane 2: Plane enabled
    ...FPC 2
PFE #0
    : Fabric Disabled
PFE #1
    Plane 0: Plane enabled
    Plane 1: Plane enabled
    Plane 2: Plane enabled
    ...

```

You can use the **request chassis fabric pfe *pfe-number* fpc-number offline** command to offline any Packet Forwarding Engine. There is no *online* option for this statement. To bring the Packet Forwarding Engine back online, you must restart the FPC.

High Availability (HA) and Resiliency

- In Graceful Routing Engine Switchover (GRES) configuration, use only "vmhost reboot" command on MX2008 routers—In Junos OS Release 17.2R1, you must use only **vmhost reboot** command instead of **request system reboot** command on MX2008.

Interfaces and Chassis

- **Support for maximum queues configuration on MPC7E, MPC8E, and MPC9E (MX Series)**—You can configure the maximum number of queues per MPC on MPC7E, MPC8E, and MPC9E. By default, these MPCs operate in per port queuing mode.

You can use the **set chassis fpc slot-number max-queues queues-per-line-card** command to configure number of queues per MPC. The possible values for *queues-per-line-card* are **8k, 16k, 32k, 64k, 128k, 256k, 512k, or 1M**.

Per-unit scheduling and hierarchical queuing on MPC7E, MPC8E, and MPC9E are licensed features.

You cannot configure the **max-queues** and the **flexible-queuing-mode** statements at the same time. You use the **flexi-queuing-mode** statement to configure a maximum of 32,000 queues per MPC.

If the **max-queues** statement is *not* configured, which is the default mode, the MPC starts with a message similar to the following:

FPC 0 supports only port based queuing. A license is required for per-VLAN and hierarchical features.

If the **max-queues** statement is configured and the value is less than or equal to 32,000, the MPC starts with a message similar to the following:

FPC 0 supports port based queuing and is configured in 16384 queue mode. A limited per-VLAN queuing license is required for per VLAN and hierarchical queuing features.

If the **max-queues** statement is configured and the value is greater than 32,000, the MPC starts with a message similar to the following:

FPC 0 supports port based queuing and is configured in 524288 queue mode. A full scale per-VLAN queuing license is required for per VLAN and hierarchical queuing features.

[See [Understanding Hierarchical Scheduling for MIC and MPC Interfaces](#) and [Flexible Queuing Mode Overview](#).]

- **Changes to show interfaces *interface-name* extensive Output**—Starting in Junos OS Releases 15.1R7, 16.1R5, 16.2R2, 17.1R2, and 17.2R1, the **MAC Control Frames** field of the **show interface *interface-name* extensive** command for a specified 10-Gigabit Ethernet interface displays a value of zero. In previous releases, the value for this field was calculated. Because of continuous traffic and as a result of the calculations, the value displayed for this field changed continuously.
- **Displaying accurate value of estimated BER in show interfaces (10-Gigabit Ethernet) command**—During autorecovery, when **show interfaces** command for 10-Gigabit Ethernet interface is executed, the

Estimated BER field displays **Recovery Under Progress** instead of **<= 1E-16**, as the estimated BER is not known during autorecovery.

Before:

```
Physical interface: xe-5/1/0, Enabled, Physical link is Down
Interface index: 311, SNMP ifIndex: 1503
Description: XX - ENNI LAG to PE-13 xe-11/3/1
Link-level type: Flexible-Ethernet, MTU: 9130, MRU: 9138, LAN-PHY mode,
Speed: 10Gbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: None,
Source filtering: Disabled, Flow control: Disabled
Pad to minimum frame size: Disabled
Device flags      : Present Running Down
Interface flags: Hardware-Down Link-Layer-Down SNMP-Traps Internal: 0x4004000
CoS queues        : 8 supported, 8 maximum usable queues
Schedulers        : 0
Current address: 00:17:cb:d4:67:c7, Hardware address: 00:17:cb:d4:66:c4
Last flapped      : 2016-01-12 13:37:33 EST (00:06:56 ago)
Input rate         : 0 bps (0 pps)
Output rate        : 0 bps (0 pps)
Active alarms      : LINK
Active defects     : LINK
PCS statistics
  Bit errors                Seconds
  Errored blocks            7
  Errored blocks            8956
Link Degradation :
  Link Monitoring           : Enable
  Link Degradation Set Threshold : 1E-8
  Link Degradation Clear Threshold : 1E-12
  Link Degradation War Set Threshold : 1E-9
  Link Degradation War Clear Threshold : 1E-11
  Estimated BER              : <= 1E-16
  Link-degradation event      : Seconds          Count
  State                      5521                2
  Defect Active
Interface transmit statistics: Disabled
```

After:

```
Physical interface: xe-5/1/0, Enabled, Physical link is Down
Interface index: 311, SNMP ifIndex: 1503
Description: XX - ENNI LAG to PE-13 xe-11/3/1
Link-level type: Flexible-Ethernet, MTU: 9130, MRU: 9138, LAN-PHY mode,
```



```

Speed: 10Gbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: None,
Source filtering: Disabled, Flow control: Disabled
Pad to minimum frame size: Disabled
Device flags   : Present Running Down
Interface flags: Hardware-Down Link-Layer-Down SNMP-Traps Internal: 0x4004000
CoS queues     : 8 supported, 8 maximum usable queues
Schedulers    : 0
Current address: 00:17:cb:d4:67:c7, Hardware address: 00:17:cb:d4:66:c4
Last flapped   : 2016-01-12 13:37:33 EST (00:06:56 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : LINK
Active defects : LINK
PCS statistics
  Bit errors           Seconds
  Errored blocks       8956
Link Degradate :
  Link Monitoring      : Enable
  Link Degradate Set Threshold : 1E-8
  Link Degradate Clear Threshold : 1E-12
  Link Degradate War Set Threshold : 1E-9
  Link Degradate War Clear Threshold : 1E-11
  Estimated BER        : Recovery Under Progress
  Link-degrade event   : Seconds          Count
State
                          5521          2
Defect Active
Interface transmit statistics: Disabled

```

[See [show interfaces \(10-Gigabit Ethernet\)](#).]

- **Aggregate Ethernet (AE) IFL (logical interface) targeted distribution feature now provides four level of prioritization**—Starting in Junos OS Release 17.2R1 and later, the aggregate Ethernet logical interface targeted distribution feature supports four levels of prioritization. If you configure all three distribution lists---primary, backup, and standby---then Junos OS will not implicitly add member interfaces to these distribution lists. That is, if any member interface is not defined in either of the configured lists, then it will be assigned a weight higher than the standby list weight and thus will be used only when all the interfaces in all three configured lists are down. This provides four levels of prioritization.

Previously, traffic would fail over to the standby links when both primary and backup links failed.

IP Tunneling

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

Management

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

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

- **Enhancement to the Junos Telemetry Interface (MX Series)**—Starting in Junos OS Release 17.2R1, the values displayed in the **oper-status** field for data streamed through gRPC for the physical interfaces sensor have changed.

The following values are now displayed to indicate the operational status of an interface:

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

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

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

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.2R1, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is **/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/**

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

MPLS

- **Bandwidth underflow sample on LSPs (MX Series)**—Starting in Junos OS Release 14.1R9, 15.1R7, 16.1R5, 16.2R3, and 17.2R2, all zero value bandwidth samples are considered as underflow samples, except for the zero value samples that arrive after an LSP comes up for the first time, and the zero value samples that arrive first after a Routing Engine switchover.

Network Management and Monitoring

- **Hard-coded RFC 3635 MIB OIDs updated (MX Series)**—Starting in Junos OS Release 17.2, the following RFC 3635 MIB OIDs have been updated as default values:
 - dot3StatsFCSErrors and dot3HCStatsFCSErrors, framing errors
 - dot3StatsInternalMacReceiveErrors and dot3HCStatsInternalMacReceiveErrors, MAC statistics: Total errors (Receive)
 - dot3StatsSymbolErrors and dot3HCStatsSymbolErrors, code violations
 - dot3ControlFunctionsSupported, flow control
 - dot3PauseAdminMode, flow control
 - dot3PauseOperMode, auto-negotiation
- **MIB buffer overruns only be counted under ifOutDiscard (MX Series)**---The change done via PR 1140400 Introduced a CVBC where qdrops (buffer overruns) were counted under ifOutErrors along with ifOutDiscards. This is against RFC 2863 where buffer overruns should only be counted under ifOutDiscards and not under ifOutErrors. In Junos OS Release 17.2R1, this is now fixed.
- **Update to SNMP support of apply-path statement**---In Junos OS Release 17.2, SNMP implementation for the apply-path configuration statement supports only two lists:
 - **apply-path "policy-options prefix-list <list-name> <*>"**
This configuration has been supported from day 1.
 - **apply-path "access radius-server <*>"**
This configuration is supported as of this release.
- **Enhancement to SMNPv3 traps for contextName field (MX Series)**—Starting in Junos OS Release 17.2, the contextName field in SNMPv3 traps generated from a non-default routing instance, is populated with the same routing-instance information as is given in SNMPv2 traps. SNMPv2 traps provide the routing-instance information as context in the form of context@community. This information gives the network monitoring system (NMS) the origin of the trap, which is information it might need. But in SNMPv3, until now, the contextName field was empty. For traps originating from a default routing

instance, this field is still empty, which now indicates that the origin of the trap is the default routing instance.

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

See the [SNMP MIB Explorer](#).

Routing Protocols

- **IPv6 neighbor reachability stale time range modified**—Starting with Junos OS Release 17.2R1, the stale time range of IPv6 neighbor reachability confirmation has changed from [1..1200] to [1..18000]. You can configure **nd6-stale-time** of upto 5 hours at the [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level.
- **Range of flow route rate-limit modified**—Starting with Junos OS Release 16.1R4, the range of flow route rate-limit is modified from [9600..1000000000000] to [0..1000000000000]. The following rate limits triggers the following actions:

Rate-limit	Actions
0	discard
1-999	0 kbps
1000-1000000000000	corresponding value in kbps

- **New option to force a Junos router to advertise a zero-length next-hop address in BGP routes for flowspec families**—Beginning with Junos OS Release 17.2R1, you can force a Junos OS router to advertise flow route updates with zero-length next-hop address even when a valid next-hop address is present in the local routing table. This option provides backward-compatibility with earlier Junos OS releases that flap BGP sessions on receiving a non-zero length next-hop address. Junos OS assigns a **Fictitious** type next-hop to flowspec routes received with a zero-length next-hop address. To advertise zero-length next-hop addresses, configure this new option **strip-nexthop** at the [edit protocols bgp family (inet | inet-vpn | inet6 | inet6-vpn) flow] hierarchy level.

When **strip-nexthop** is not configured, Junos OS advertises a non-zero length next-hop address (if one exists) for flowspec family routes just as it does for other address families.

[See [strip-nextthop](#).]

- **Format of session up time modified in show bfd session detail output**—Starting in Junos OS 17.2R1, the output of **show bfd session detail** includes the seconds in the session up time field. In earlier Junos OS releases, the session up time was displayed as **1w1d hh:mm**; the seconds were omitted when the up time was more than 24 hours. The modified format of the **session up time** is **1w1d hh:mm:ss**.

[See [show bfd session](#).]

Services Applications

- **Change in behavior of IKE negotiation (MX Series)**—Starting in Junos OS Release 17.2R1, when you commit an IPsec configuration that includes **establish-tunnels immediately** at the **[edit services ipsec-vpn]** hierarchy level, the service set might take up to 30 seconds to initiate IKE negotiations.

Subscriber Management and Services

- **Changes to flat-file accounting statistics collection when a service deactivation fails (MX Series)**—Starting in Junos OS Release 17.2, the collection of accounting statistics when an ESSM service is deactivated has changed. When the deactivation is initiated by a Change of Authorization (CoA) message, **essmd** sends a stop request to the accounting daemon (**pfed**), which writes the stop record and marks the statistics values at that time as a new baseline value.

When the commit for the new configuration succeeds, the logical interface on which the service was deactivated is deleted.

When the commit fails, the service is restored rather than deactivated and the logical interface is not deleted. In this case, **essmd** requests the accounting daemon (**pfed**) to resume flat-file accounting for the service. The accounting daemon (**pfed**) writes an accounting start record, then resumes writing interim accounting records, where the interim statistics equal the current value minus the baseline value.

In earlier releases, if the service deactivation fails and the service is restored on the logical interface, no interim accounting statistics are collected for the interval since the stop record was written, resulting in inaccurate values.

- **DNS servers displayed by the show subscribers extensive command (MX Series)**—Starting in Junos OS 17.2, the display of DHCP domain name servers (DNS) by the **show subscribers extensive** command has changed. When DNS addresses are configured at multiple levels, the command displays only the preferred address according to this order of precedence: RADIUS > access profile > global access. The command does not display DNS addresses configured as DHCP local pool attributes.

DNS addresses from RADIUS appear in the following fields: Primary DNS Address, Secondary DNS Address, IPv6 Primary DNS Address, IPv6 Secondary DNS Address.

DNS addresses from the access profile or the global access configuration appear in the following fields: Domain name server inet, Domain name server inet6.

In earlier releases, the command displays only DHCP DNS addresses provided by RADIUS.

- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 17.2, the **show subscribers extensive** command displays the **IPv6 Interface Address** field only when the dynamic profile includes the \$junos-ipv6-address predefined variable.

In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

- **Change to DHCP option 82 suboptions support to differentiate duplicate clients (MX Series)**—Starting in Junos OS Release 17.2, only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are considered when this information is used to identify unique clients in a subnet. Other suboptions, such as Vendor-Specific (suboption 9), are ignored.

[See [DHCPv4 Duplicate Client In Subnet Overview](#).]

- **Default L2TP resynchronization method changed and statement deprecated (MX Series)**—Starting in Junos OS Release 17.2, the default resynchronization method for L2TP peers in the event of a control connection failure is changed to silent failover. In earlier releases, the default method is failover-protocol-fall-back-to-silent-failover. The silent failover method is preferred because it does not keep tunnels open without traffic flow, waiting for the failed peer to recover and resynchronize. You can use the new **failover-resync** statement at the **edit services l2tp tunnel** hierarchy level to specify either failover protocol or silent failover as the resynchronization method.

Because silent failover is now the default, the **disable-failover-protocol** statement is no longer needed and has been deprecated. If you upgrade to this release with a configuration that includes this statement, it is supported, but the CLI notifies you it is deprecated.

[See [L2TP Failover and Peer Resynchronization](#).]

User Interface and Configuration

- **Enhancements to the show chassis fpc errors command to display the PFE enable or disable status (MX Series)**—The **show chassis fpc errors** command output is enhanced to include information about the state of the packet forwarding engine (PFE).

```
user@host> show chassis fpc errors
```

```
FPC  Level  Occurred  Cleared  Threshold  Action-Taken  Action
1   Minor      0         0        10         0    LOG|
    Major      0         0         1         0  GET STATE|CM ALARM|DISABLE PFE
    Fatal      0         0         1         0    RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
    | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
2   Minor      0         0        10         0    LOG|
    Major      0         0         1         0  GET STATE|CM ALARM|DISABLE PFE
```

```
Fatal      0      0      1      0  RESET
Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
| pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
3  Minor      0      0     10      0  LOG|
Major      0      0      1      0  GET STATE|CM ALARM|DISABLE PFE
Fatal      0      0      1      0  RESET
Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
| pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
5  Minor      0      0     10      0  LOG|
Major      0      0      1      0  GET STATE|CM ALARM|DISABLE PFE
Fatal      0      0      1      0  RESET
Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
| pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
```

SEE ALSO

- [New and Changed Features | 101](#)
- [Known Behavior | 143](#)
- [Known Issues | 147](#)
- [Resolved Issues | 154](#)
- [Documentation Updates | 164](#)
- [Migration, Upgrade, and Downgrade Instructions | 166](#)
- [Product Compatibility | 173](#)

Known Behavior

IN THIS SECTION

- [General Routing | 145](#)
- [High Availability \(HA\) and Resiliency | 145](#)
- [Network Management and Monitoring | 145](#)
- [Interfaces and Chassis | 145](#)
- [User Interface and Configuration | 145](#)

- Software Defined Networking (SDN) | 146
- Subscriber Management and Services | 146

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 for MX Series..

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

General Routing

- **Multiprotocol extensions capability code in notification message**—Starting in Junos OS Release 17.2, when a BGP speaker terminates a peering session, because the peer does not support Multiprotocols Extensions for BGP-4, it sends a notification message that contains the multiprotocol extensions capability as per the standard. In earlier releases, the BGP peer sends a notification message that contains internal code for unsupported NLRIs.

High Availability (HA) and Resiliency

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

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

Network Management and Monitoring

- **SNMP traps for certain interfaces in Admin Down state (MX Series)**—SNMP traps are generated when an interface that supports the Digital Optical Monitoring (DOM) MIB is placed in an administrative down state. This behavior informs the operator of any interface fault, alarm, or threshold condition.

Interfaces and Chassis

- **An additional commit is required when reusing Virtual IP on an interface as an interface address (MX Series)**—When you reuse a virtual IP address on an interface as an interface address, you must first delete the virtual IP address configuration and commit the configuration. You must then add the interface address configuration in a subsequent commit.

User Interface and Configuration

- **Modification to configurable link degrade threshold values (MX Series)**—Starting with Junos OS Release 15.1F7 and 16.1R1, the values of the user configurable link degrade thresholds, have to be configured as per the following guidelines:
 - **set threshold value** must be greater than **warning set threshold value**
 - **set threshold value** must be greater than **clear threshold value**
 - **warning set threshold value** must be greater than **warning clear threshold value**

If the threshold values are not configured as per these guidelines, the configuration fails and a Commit Error message is displayed.

Software Defined Networking (SDN)

- When the BSYS master Routing Engine is rebooted or shut down, the JDM-to-JDM communication, including the commit sync operation, fails. To work around this issue, commit the JDM configurations on server0 and server1 separately.
- If the GNF console remains idle for a long duration (for example, more than 10 minutes), the console might stop responding.
- Pings to the peer JDM might fail even when the connection status is shown to be up. Also, the **show server connections** command might show JDM-to-JDM ping failure issues. These ping failure issues occur when connections from the Control Board to the servers are mapped incorrectly at the JDM. Correct the mapping by verifying the connections.
- The JDM operational command **show virtual-network-functions** might sometimes show the value of the **Liveness** field as **Down** even when the GNF is up and reachable.
- The GNF VM's fxp0 interface might get slower and stop forwarding packets occasionally. When this occurs, disable the fxp0 interfaces and enable it again.

Subscriber Management and Services

- If a graceful Routing Engine switchover (GRES) is triggered by an operational mode command, the state of aggregated services interfaces (ASIs) are not preserved. For example:

```
request interface <switchover | revert> asi-interface
```

However, if GRES is triggered by a CLI commit or FPC restart or crash, the backup Routing Engine updates the ASI state. For example:

```
set interface si-x/y/z disable
commit
```

Or:

```
request chassis fpc restart
```

SEE ALSO

New and Changed Features	101
Changes in Behavior and Syntax	132
Known Issues	147
Resolved Issues	154
Documentation Updates	164
Migration, Upgrade, and Downgrade Instructions	166
Product Compatibility	173

Known Issues

IN THIS SECTION

- Class of Service (CoS) | 148
- EVPNs | 148
- Forwarding and Sampling | 148
- General Routing | 148
- Infrastructure | 151
- Interfaces and Chassis | 151
- Layer 2 Features | 151
- Layer 2 Ethernet Services | 151
- MPLS | 152
- Multicast | 152
- Network Management and Monitoring | 152
- Platform and Infrastructure | 152
- Routing Protocols | 153
- Services Applications | 154
- Subscriber Access Management | 154
- VPNs | 154

This section lists the known issues in hardware and software in Junos OS Release 17.2R1 for MX Series.

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

Class of Service (CoS)

- In Junos OS Release 17.2R1, when a cascade port (CP) is configured, CoS resources are allocated to it and corresponding CoS parameters applied on extended ports are scaled. This is done irrespective of the cascade port. If a configured cascade port goes DOWN, nothing is done. [PR1262320](#)
- In Junos OS 17.2 Release, egress rate limit at extended port does not work properly if we have rate limit configuration applied at the Physical interface by traffic-control-profile-remaining and also at some of the extended port of Logical interface by explicit traffic-control-profile in hierarchical-scheduler mode. [PR1271719](#)

EVPNs

- On a EVPN MPLS - VXLAN stitching scenario, the IRB interface might not respond to ARP requests. [PR1272974](#)

Forwarding and Sampling

- It is known that the policing filter application to the LSP is catastrophic. Any active LSP carrying traffic when applied a policing filter tears down and resignals and drops traffic for about 2 seconds. In Junos OS Release 16.1R1, it takes up to 30 seconds for the LSP to come up because: 1. Creation of the policing filter and application of the same to the LSP through configuration in the same commit sequence. 2. Load override of a configuration file that has policing filter and policing filter application to the LSP followed by commit. [PR1160669](#)

General Routing

- In case of the presence of l2tp subscribers, all FPCs except the card that hosts subscribers will report a log message `jnh_if_get_input_feature_list(9723): Could not find ifl state` after every subscriber's login attempt. [PR1140527](#)
- 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 XE interface. Netdev-interface resource assumes that 'speed' is a configurable parameter which 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](#)
- Junos OS might improperly bind Packet Forwarding Engine ukernel application sockets after unified ISSU due to a bug in IP->TNP fallback logic. Because of that bug, threads running on the ukernel that relay on UDP sockets can experience connectivity issues with host, which in turn can lead to problems. For instance, Simple Network Time Protocol (SNTP) client might fail to synchronize time, which in turn

might lead to other problems such as failure in adjacency formation for HMAC authenticated protocols. [PR1188087](#)

- Few sessions are always dropped during session setup with IPsec, consistently seen above 1M sessions. [PR1204566](#)
- PTP master streaming on IP and Ethernet interfaces is not supported on both simultaneously. Master streams on IP and Ethernet not supported simultaneously. [PR1217427](#)
- There is no ISSU from Junos OS Release 15.1 and older releases to Junos OS Release 16.2R1. [PR1222540](#)
- 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 packet time error can be greater than +/-100 sec in these MICs. [PR1226080](#)
- Changing the virtual switch type from IRB type to regular bridge, interfaces under openflow protocol werel removed. Openflow daemon failed to program any flows. [PR1234141](#)
- When non-Juniper Networks SFP is used in MIC-3D-20GE-SFP-E or MIC-3D-20GE-SFP-EH MIC, the ISR 2 (MIC error interrupt) might be running off over 2.5 seconds due to an unknown reason, and then the FPC host the MIC might be restarted and crashed. The fix add interrupt throttling for MIC interrupt and restarting the MIC if interrupts are more than the threshold (> 2500 per 5min). [PR1235475](#)
- DNS server IP addresses are not present in the output of **show subscribers extensive** for DHCP subscribers if the DNS configuration is provided from the access-profile or pool. In case of such data is provided from the RADIUS, the output is correct. The issue is cosmetic: DNS addresses are provided to subscribers. [PR1237525](#)
- The measured noise-transfer gain is around -40 dB, which is higher than the standard metric.[PR 1240054](#)
- 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 upto +/-450 nsec in these MICs.[PR1243646](#)
- Following MICs in ALOHA MPCs (MPC2E/3E-NG) are non PHY-Timestamping capable. MIC-3D-4XGE-XFP MIC3-3D-10XGE-SFPP MIC-3D-2XGE-XFP MIC-3D-20GE-SFP The dynamic time error can be greater than 40 nsec for LF & 70 nsec for HF.[PR1243871](#)
- When you activate and deactivate Route Target per bridge domain in EVPN, the RPD process might crash resulting in traffic loss. We recommend not to toggle this configuration in Junos OS Release 17.1R1. [PR1244956](#)
- When having interfaces configured with encapsulation ethernet-ccc and corresponding ifl 0 with family CCC, along with the same interfaces being also part of an interface-switch connection (under [protocol connection interface-switch] configuration hierarchy), if one or both interfaces are deactivated and later activated again, one or both might remain in down state. [PR1262382](#)
- 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](#)

- Due to transient hardware error conditions, only syslog events XMCHIP(x) FI: Cell underflow at the state stage - Stream 0, Count 65535 are reported, which is a sign of fabric stream wedge. Additional traffic flow register pointers are validated and, if stalled, a new CMERROR alarm is raised "XMCHIP(x) FI: Cell underflow errors with reorder engine pointers stalled - Stream 0, late_cell_value 65535, max_rdr_ptr 0x6a9, reorder_ptr 0x2ae". [PR1264656](#)
- Due to transient hardware events, fabric stream might report **CPQ1: Queue underrun indication - Queue <q#>** in continuous occurrence. For each such event, all fabric traffic is queued for this Packet Forwarding Engine reporting the error and causes a very high amount of fabric drops. [PR1265385](#)
- For very large tables or very slow consumers, the BgpRouteMonitorRegister() might not send the operation BgpRouteMonitorOper.END_OF_RIBS after all of the initial updates are sent. [PR1265427](#)
- When In Service Software Upgrade (ISSU) is performed under scaled scenarios where the Packet Forwarding Engine Nexthop memory uses more than 4 Million Dwords, PPE traps and traffic loss might be observed during software-sync phase until the end of hardware-sync. [PR1267680](#)
- With the Routing protocols like LDP, mcast, configured on the router and if IPLC is present in the system, the show statistics commands issued from the CLI will timeout. Eg: set protocols mpls interface all set protocols ldp interface all set protocols <routing protocol> statistics or traffic-statistics" will timeout. eg: show ldp traffic-statistics eg: show multicast route extensive etc Thanks Praabhu [PR1275234](#)
- Routing Engine based CPCD does not work in vMX / MX86. [PR1276016](#)
- Multicast traffic when using iflsets in Universal Call Admission Control policy mode, does not work as expected in certain use cases and bbe-smgd might generate a core file. [PR1278543](#)
- Following content should be added to the release notes: Establishment of IPsec SAs 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< -----8< restart ipsec-key-management -----8< -----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< -----8< clear services ipsec-vpn ipsec security association <service set> -----8< -----8< -----8< [PR1281223](#)
- 1. If GRES/NSR is enabled, it is possible that the **request vmhost reboot** command on master RCB (which becomes backup) might result in the new backup RCB taking extra 5 minutes (approx. boot up time is 10-11 minutes) to be completely up 2. After GRES switchover, there might exist a small window of 4-5 minutes over and above the regular window when backup RCB will not be available for an immediate back-to-back switchover. [PR1282500](#)

Infrastructure

- On Junos OS devices with legacy Free BSD version 6.X, the devices might crash and reboot if there is a defect in the Junos SDK based multithreaded application, which has been used. [PR1259616](#)

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 perform a commit and then add the interface address configuration in the following commit. [PR1191371](#)
- For MX2008, MX2010, and MX2020 platforms, on removing MPC, Link Error column in show chassis fabric summary extended output shows YES for all fabric planes. Whereas, on offline of MPC using CLI command, output shows correctly. [PR1214611](#)
- Junos OS upgrade involving releases 14.2R5 (and above in 14.2 maintenance releases) and 16.1 above mainline releases with CFM configuration can cause CFMD core post upgrade. This is due the old version of /var/db/cfm.db [PR1281073](#)
- On a Junos OS based platform, sometimes it might occur that FPC is stuck in offline state with the reason "Restarted by cli command" after restarting the FPC immediately after restarting chassisd. This occurs due to the fact that it takes sometime for the system to stabilize after chassisd restarts. Though chassisd would provide the FPC status and be able to accept the commands but in the backend it would be doing many initialization. So wait until all the PIC status are also available before issuing any command that makes fpc online, offline and restart. [PR1275530](#)

Layer 2 Features

- On routers running Junos OS with graceful Routing Engine Switchover GRES enabled, if VPLS is configured with a dynamic-profile association, some traffic is observed when the Routing Engine switches from master to standby. This is due to a change in the underlying database that handles the dynamic-profile sessions, which causes the VPLS connection to destroy and re-create after a Routing Engine switchover. [PR1220171](#)

Layer 2 Ethernet Services

- When MSTP is configured under routing-instance, both the primary and standby VPLS pseudowires are struck in ST state due to a bug in the software. That has been fixed and now the PW status is set correctly. [PR1206106](#)
- After changing the underlying physical interface for a static VLAN demux interface, the NAS port ID is still formed based on the previous physical interface. [PR1255377](#)

MPLS

- When a container LSP has more than 10 member LSPs only the first 10 LSPs are shown in the **show mpls container-lsp name <lsp-name> statistics** output. [PR1267774](#)
- The throughput measurement might be inaccurate when doing performance measurement on a MPLS label-switched-path. [PR1274822](#)

Multicast

- Under some extremely rare condition, OSPF neighbors might not come up if PIM is also configured in the same routing instance. PIM might install a multicast forwarding entry that prevents OSPF from receiving hello protocol messages. This happens in a rare situation when router control plane is under extreme load just after enabling OSPF and PIM. [PR1279682](#)

Network Management and Monitoring

- Symptom: "MIB2D_RTSLIB_READ_FAILURE: rtslib_iflm_snmp_pointchange" syslog during configuration restore. Cause: mib-process daemon send requests to kernel to update snmp ifIndex for the interfaces that it is learning. If this interface was already deleted from kernel, the above syslogs could be seen. This interface learning by mib-process daemon will happen later, once kernel sends the ADD notification for these interfaces. There is no impact due to this syslog during the configuration restore scenario. [PR1279488](#)

Platform and Infrastructure

- FPC reports the following errors and the FPC is not able to connect any subscriber **Pkt Xfer:** WEDGE DETECTED IN PFE 0 TOE host packet transfer: %PFE-0: reason code 0x1**. Also the MQ FI might be wedged and below log can be seen: **Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) FI Reorder cell timeout Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) FI Enqueueing error, type 1 seq 404 stream 0 Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) MALLOC Pre-Q Reference Count underflow - decrement below zero.** [PR873217](#)
- In configurations with IRB interfaces, during times of interface deletion, such as an FPC reboot, the Packet Forwarding Engine might log errors stating "nh_ucast_change:291Referenced l2ifl not found". This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- On MX Series, frequent configuration edit from an ephemeral database OR issuing command **user@router> show ephemeral-configuration | display merge** might cause one of the daemons (dcd, rpd, dfwd, pfed, cosd, sampled) to generate core files. [PR1249979](#)
- In affected release with bridge over GRE feature configured, the traffic failed to send to the GRE underlying interface once bouncing the gr- interface. [PR1255706](#)

- In scaled configurations with Distributed IGMP configured (above 1000 subscribers per FPC joining 1000 multicast groups) FPCs might crash. [PR1270928](#)
- With unified In Service Software Upgrade (ISSU), it is expected to see momentary traffic loss. In EVPN ETREE, in addition to traffic loss, the known unicast frames might be flooded for around 30 seconds during ISSU before all forwarding states are restored. This issue does not affect BUM - Broadcast, Unknown Unicast, and Multicast - traffic. As a workaround, Non-Stop Bridging (NSB) can be configured [set protocols layer2-control nonstop-bridging]. This reduces traffic flood to around 10 seconds in a moderate setup. [PR1275621](#)
- Syncing of ephemeral default database will not work between master and backup. This is a use case only for DT as they are the only ones who use Ephemeral. Syncing of instance ephemeral database is not impacted by this issue. [PR1279653](#)

Routing Protocols

- 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 the context of a large number of configured VPNs, routes changing in the midst of a BGP path-selection configuration change can sometimes lead to generate an RPD core file. This core file has been seen with the removal of the "always-compare-med" option. [PR1213131](#)
- The commit for PR 1252151 changed a behavior of the BMP monitoring. Before that change, the BMP session would send both peer down events as well as route withdrawals when a peer monitoring was disabled through a configuration event. After that commit, only the peer down events are sent. [PR1265783](#)
- On rpd crash with "switchover-on-routing-crash" enabled on box, live vmcores might be seen on both Routing Engines without an impact on the system. [PR1267796](#)

Services Applications

- Business services are activated and a Routing Engine switchover is performed. In this case, if we try to deactivate the Business services (aka ESSM subscribers) by logging out the parent PPP session, the business services are getting stuck in terminating state. Business services that have LI applied are stuck and the services not having LI are logged out successfully. [PR1280074](#)

Subscriber Access Management

- Subscribers stuck in terminated state during PPPoE login/logout test. [PR1262219](#)

VPNs

- It only occurs when doing VRF localization. The routes are deleted and re-added when this is done and there is some loss. In the current scenario, however, after the routes are deleted, it was taking slightly more time to add it back and get stabilized. When the script checks for the route after VRF localization is done, it was not available, but it comes up after some time and everything should work fine. [PR1264366](#)

SEE ALSO

New and Changed Features 101
Changes in Behavior and Syntax 132
Known Behavior 143
Resolved Issues 154
Documentation Updates 164
Migration, Upgrade, and Downgrade Instructions 166
Product Compatibility 173

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R1 | 155](#)

This section lists the issues fixed in the Junos OS 17.2R1 Release for MX Series.

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

Resolved Issues: 17.2R1

Class of Service (CoS)

- If the hidden command **show class-of-service queue-consumption** is executed many times (in this case, for 100 times), in rare cases, the cosd process might crash with a core file generated. The core files can be seen by executing the CLI command **show system core-dumps**. [PR1066009](#)

Forwarding and Sampling

- Accounting Interim interval might be reset after performing a graceful switchover and does not take into account the last interim accounting timing before the switchover. [PR1261472](#)
- The final service stats are queried via on-demand service stats handling module of PFED (RE). When the responses are returned from PFE to Routing Engine (PFED) that are mapped to the request via the request ID as well as location offset. When there are more than one filter configured for a BBE filter service session (out of IPV4,IPV6 IN,OUT filters), more than one request will be sent to the same location (PFE) with same request ID. As long as PFE was returning the response in the order of request being sent from PFED, there is no problem observed. In this PR mentioned scaled scenario, the PFE (bulkget) is returning the service stats response in an out of order for a few service sessions which resulted in mapping the response to wrong request element (data_entry structure) Fix description
===== The response to request mapping at Routing Engine (PFED) ondemand service stats handling module should contain the filter index as well as a KEY along with other KEY params. This modification is done via this PR fix. [PR1262876](#)

General Routing

- The jsscd might crash in a static-subscribers scaling environment (for example, 112K total subscribers, 77K dhcp subscribers, 3K static-subscribers, 32K dynamic vlans). When this issue occurs, the subscribers might be lost. `abc@abc_RE0> show system core-dumps -rw-rw---- 1 root field 8088852 Jan 1 11:11 /var/tmp/jsscd.core-tarball.0.tgz` [PR1133780](#)
- The logging and reporting function (LRF) supports change of template in the middle of a session. When the template changes, LRF flushes and sends IPFIX records out for the existing template, and a new template is allocated based on the newly matched rule. [PR1174938](#)
- On MX Series platforms, the MS-MIC crash might occur. The exact trigger of the issue is unknown; normally, this issue happens very rarely without any external triggers. The crash might occur with any services configuration, with core files pointing to - Program terminated with signal 4, Illegal instruction. [PR1183828](#)
- On MX Series platform with Junos OS Telemetry Interface (JTI) enabled, and with **set routing-options lsp-telemetry** knob is configured. When SDN-telemetry (the agentd process) is disabled or continuously

restarted, certain messages repeatedly logged into syslog, the rpd and eventd processes CPU might get near 100% and eventually, the agentd also gets near 100%. When this issue happens, the agentd process is not able to accept new subscriptions, dropping all existing subscriptions. It can be triggered by restarting consecutively SDN-telemetry (the agentd process), or after device reboot. To recover from the first condition, the user needs to unconfigure "routing-options lsp-telemetry", what generates an RPD core dump if the bad state is active. To recover from the second condition, the user needs to manually restart the s SDN-telemetry. [PR1192366](#)

- During unified in-service software upgrade (ISSU) it might be noticed that the following log messages are displayed: SFP: pointer Null, sfp_set_present. This might trigger a flap in the interfaces on MX Series routers while upgrading with unified ISSU. [PR1200045](#)
- Unified ISSU is not supported on MX2008. [PR1213193](#)
- In an L2TP scenario, in rare situations, the command **show subscribers summary port extensive** output might have wrong tunneled or terminated sessions count due to an issue on populating the outputs. There is no traffic impact. Use the command **show subscribers physical-interface** instead when that count is not correct. [PR1206208](#)
- MS-MPC/MS-MIC might crash when large fragmented (larger than 2048 bytes) traffic goes through an ALG. [PR1214134](#)
- Syslog message : "fpc_pic_process_pic_power_off_config:xxxx :No FPC in slot y" is displayed on empty FPC slots with no PIC power off configured by committing configuration change under chassis hierarchy. [PR1216126](#)
- PPPoE or DHCP subscribers fail to bind due to ProcessPADIFailedUiflNotActive/SML_CLIENT_DELETE_SDB_ADD_FAILED errors. It is seen during inflight tests. [PR1221690](#)
- When CLK0 and/or CLK1 port of MX2008 is configured as building-integrated timing system (BITS) clock sources, BITS LED glows GREEN when external BITS clock is not connected to the port. When external BITS clock is connected and then removed, BITS LED glows RED. [PR1222041](#)
- The following error messages might be seen when there is a traffic from subscribers with captive-portal-content-delivery service. cpcdd[29943]: %DAEMON-3: early: bad stored heap: heap-ptr=0x0 data-ptr=0x839f742 cpcdd[29943]: %DAEMON-3: opDel: bad stored heap: heap-ptr=0x1000000 data-ptr=0x839f0aa This does not have any affect on functionality. [PR1226782](#)
- In an idle MX2008 chassis, chassisd CPU utilization oscillates between 10% to 20%. There is no functional impact. [PR1231333](#)
- On MX Series or PTX Series routers with Junos Telemetry Interface deployed, streaming telemetry stops being sent for Message Queue Telemetry Transport (mqtttd) memory that reaches to a maximum of 4G. [PR1238803](#)
- Access Node Control Protocol (ANCP) neighbors go down after a commit if any ANCP-related configuration was changed. [PR1243164](#)

- Currently MS-MIC supports a maximum of 2 million routes scale. This includes all IPv4, IPv6, and MPLS routes in the system. When the scale limit is exceeded, the forwarding database (FDB) memory will exhaust and the MS-MIC will start to drop the routes and also print logs. [PR1243581](#)
- When you activate and deactivate route target per bridge domain in EVPN, the RPD process may crash resulting in traffic loss. We recommend that not to toggle this configuration in Junos OS Release 17.1R1. [PR1244956](#)
- Network slices GNF console might hang some times when console is idle for longer duration (Greater than 10 mins) [PR1250726](#)
- On all Junos OS platforms that have the routing protocol process (rpd), if some interfaces go down that results in some peers going down or BGP-RR (route-reflector) re-advertising routes, and then the rpd process might crash. [PR1250978](#)
- From Junos OS 16.1 and later releases, KRT queue stuck on Routing Engine causing RIB and FIB to go out of sync when a route or next-hop is added to the kernel, it's a timing problem between multiple rpd threads. [PR1251556](#)
- In certain scenarios, the output of **show ancp subscriber detail** command might omit many TLVs, including mandatory Actual Net Data Upstream and Actual Net Data Downstream TLVs. [PR1252747](#)
- A 1-PPS error jump during the Synchronous Ethernet transient is within the acceptable limits. [PR1253083](#)
- On MX Series routers with MPC2E-NG/MPC3E-NG, the interfaces of these line cards might not come up when connecting to a third party transport switch. [PR1254795](#)
- In the output of **show subscribers extensive** the first IP address from the Framed Prefix (returned in Framed-IPv6-Prefix) looks to be assigned to the subscriber interface although it is not. The fix removes incorrect data. [PR1255029](#)
- In vMX platform, if a lot of subscribers login/logout activity when there are large number of flows (>500K), Multiple Riot cores might be seen. [PR1255866](#)
- VPLS MAC table is not being populated properly when checked with CLI **show vpls mac-table**, although all subscribers have traffic. Thus it is considered a cosmetic issue. [PR1257605](#)
- Unable to run **show subscribers extensive** and some other CLI commands after GRES because subscriber-management database is unavailable. The other symptoms of the bug can be messages like 'sdb.db: close: Bad file descriptor' after 'commit full' [PR1258238](#)
- In a subscriber service environment, the device control process (dcd) might restart unexpectedly during commit process after changes to ATM interface configuration. [PR1258744](#)
- Subscriber management process (bbe-smgd) might crash and generate a core file during Routing Engine mastership switchover. This issue might be seen when having VLAN-OOB subscribers and running GRES and NSR. [PR1258817](#)
- When using an AMS interface and running the **show interfaces extensive** command the subinterfaces will show only 0 for the packet counters. [PR1258946](#)

- On QSFP-40G LR4 optics, link could remain down after a flap. This issue is specific to these optics. [PR1259930](#)
- Class of Service (CoS) does not correctly classify egress L3 multicast traffic from an ingress VLAN bridge interface after a configuration change. [PR1260413](#)
- Only the first multicast IP packet was saved when waiting for a route to be resolved. This fix will save up to 20 additional IPv4 multicast packets and send all saved packets after the route is resolved. [PR1260729](#)
- In an MX Series, broadband network gateway (BNG) subscriber management environment, there could be a slight deviation in the service accounting statistics when the subscriber session terminated abruptly. [PR1260898](#)
- During multicast activation of dynamic subscribers through a service profile, the bbe-smgd process in backup Routing Engine could sometimes crash. [PR1261285](#)
- It was observed that an authenticated dynamic VLAN interface with an idle-timeout is removed if there are no subscribers on top and if "remove-when-no-subscribers" is configured at the auto-configure stanza. Such dynamic VLAN interface should only be removed after its idle timeout expires if it stayed idle during this period. [PR1262157](#)
- When subscriber-management is enabled, an unified ISSU from Junos OS release 16.1R4 to Junos OS release 16.2 or 17.1 requires that any service cards be offlined before the ISSU is started. [PR1262877](#)
- There is a problem where the MX Series router used the wrong routing table to send out the ICMP network unreachable message back to the source thus might cause some problem on the end user CPE. [PR1263094](#)
- Currently when the CoS adjustment-control-profile (ACP) is configured with radius-coa using adjust-less algorithm, cosd strictly follows the configured algorithm when (1) only service-profiles and/or CoA is used to apply rates to the subscriber flow and (2) no line rate adjustment protocols such as ANCP or protocol tags (for example, PPPoE-tags) are being used to apply updates. This results in undesirable complexity in applying service profiles in the order activated based on an ACP approach that is intended to control the comparison of a configured-rate and a line-rate, where the former represents a policy and the latter the capabilities of the access-loop. When only service profiles are in use, such that more than one service profile might be applied to the subscriber through RADIUS CoA and each service profile affects the shaping rate of the subscriber. The correct behavior is for CoS to ignore the algorithm when no line rate protocol is in use. Instead, it should use a replacement semantic (logically the algorithm "adjust-always") to apply a service profile initiated through CoA in the order received. Thus, a profile chain can be easily managed that includes the client profile and one or more service profiles. Thereby allowing predictable and intuitive revert semantics during the service-deactivation or re-activation scenarios. Once a line rate protocol such as ANCP is enabled and updates are received, only then cosd should follow the algorithm as it will then compare with the configured rate and a line rate (where the intended goal is minimum (policy rate, line rate). As a follow-on, the ACP configuration syntax will be revisited as it is unnecessarily complex for the intended use case. [PR1263337](#)

- After router reboot or the JSD process crash, sometimes the listening socket for JSD is not operational
labroot> show configuration system services extension-service request-response { grpc { clear-text { address 192.168.200.1; port 50051; } skip-authentication; } } notification { allow-clients { address 0.0.0.0/0; } } traceoptions { file grpc.log size 1m files 10; flag all; } labroot> labroot> show system connections | match 50051 labroot> <----- Socket not listening. [PR1263748](#)
- After running **show arp** with subscribers connected bbe-smgd can become unresponsive or slow to other CLI commands. [PR1264038](#)
- On MX with MS-MPC, if the ethernet frames with more than 2000 bytes of payload, the mspmand process which manages the Multi-Services PIC might crash. The traffic forwarding might be affected. [PR1264712](#)
- In some situations, MX Series L2TP access concentrator (LAC) does not encapsulate packets received from CPE in l2tp tunnel. This subscriber has a static pp0 unit configured on the LAC side. This issue is causing a permanent traffic black hole for this subscriber and leads to PPP session flaps or impossibility to establish a PPP session between CPE and LNS in case of using LCP renegotiation on LNS side. [PR1265414](#)
- For very large tables or very slow consumers, the BgpRouteMonitorRegister() may not send the operation BgpRouteMonitorOper.END_OF_RIBS after all of the initial updates are sent. [PR1265427](#)
- In a subscriber management environment, starting from Junos OS release 15.1R4, new generation architecture is used (aka Tomcat). In these new releases, if dynamic VLAN profile doesn't have IFF configuration (for example, either family PPPoE or family inet), but has firewall filter configuration, firewall filter indices will not be released after the dynamic VLAN is removed, which will eventually lead to depletion of available firewall filter indices. [PR1265973](#)
- Per IETF RFCs, IGMPv3 & MLDv2 reports not sent to IANA reserved multicast addresses 224.0.0.22(IGMP V3 ROUTERS) and ff02::16(MLD V2 ROUTERS), respectively, should be discarded. But broadband network gateway (BNG) processes these reports. With this fix, the reports will be discarded and Rx error counter updated. [PR1266309](#)
- ISSU upgrades to Junos OS Release 16.1R4-S1.x are not supported with subscriber management enabled(MX Series). For systems running subscriber management, upgrade manually by copying the target release to both Routing Engines and rebooting them at the same time (use the "request system reboot both-routing-engines" command). [PR1266317](#)
- When VSTP is enabled on a double tagged AE IFL and there is another single tagged AE IFL configured with same outer VLAN tag, then the incoming traffic on that VLAN is incorrectly hitting the AE_RESERVED_IFL_UNIT (AEx.32767) and the traffic is getting dropped. [PR1267238](#)
- It is possible to see a bbe-smgd core file under certain boundary conditions on the standby Routing Engine with certain specific configurations. Since the core file is on the standby no disruption in service is expected and system recovers from this condition. [PR1267646](#)
- The CLI configuration command **set chassis effective-shaping-rate** is enabled for the MX104. [PR1267829](#)

- The router MAC extended community is not using the IANA standardized value. This affects EVPN type 5 configured in direct-nexthop (pure type 5) mode with VXLAN encapsulation only. EVPN type 5 in gateway address mode or with MPLS encapsulation is not affected because it does not use the affected extended community. [PR1269236](#)
- Database status might remain as not ready after several SO. [PR1271306](#)

Infrastructure

- In case, SSD contains a valid permanent (non-resettable) offline-uncorrectable-sectors positive value, smartd logs on the non-zero value every 30 minutes by default. This is too frequent logging considering that there has not been a change in the value. [PR1233992](#)
- On all Junos OS platform and on the router with PIM enabled which has a local receiver, stale nexthops are present where they did not get deleted by daemons due to timing issue. [PR1250880](#)
- Legacy Junos OS Kernel might generate a core file on `userland_sysctl / sysctl_root / sysctl_kern_proc_env / panic_on_watchdog_timeout` [PR1254742](#)
- On Junos OS devices with legacy FreeBSD (version 6.X), the devices might crash and reboot if there is a defect in the Junos OS SDK based multithreading application that has been used. [PR1259616](#)

Interfaces and Chassis

- Unified ISSU will not work from Junos OS Release 15.1R to later images (like 15.1F, 16.1R2, ...), when the router is equipped with QSFP/CXP/CFP2 optics on MPC3E/MPC4E/MPC5E/MPC6E/MPC 3D 16x10GE cards. This is because a dark window issue is fixed for SFPP/QSFP/CXP/CFP2 optics in 16.1 and 15.1F, which makes the 15.1R image ISSU incompatible with later images. Doing ISSU on the incompatible image from 15.1R to later versions might result in the line card crash as a side effect. [PR1216924](#)
- t3 interfaces configured with "compatibility-mode digital-link" may fail to come up due to an incorrect subrate. To verify, **show interfaces t3-0/0/0 extensive** to display the subrate. DSU configuration: Compatibility mode: Digital Link, Scrambling: Enabled, Subrate: 4195338 Kbps <<<< // expected result: DSU configuration: Compatibility mode: Digital Link, Scrambling: Enabled, Subrate: Disabled [PR1238395](#)
- In some rare situations Ethernet Connectivity Fault Management process (cfmd) might crash when committing a configuration where CFM filter refers to a firewall policy. When hitting this issue, all CFM enabled interfaces are down. [PR1246822](#)
- SNMP-SET command fails when FPC/PIC/port have a value greater than 9 When the **snmp-set** command is issued, it encounters an error due to an incomplete port number in the command pushed:: Jan 18 10:49:53.626342 snmpd_process_nvset: talking to mgd (60001) Jan 18 10:49:53.626350 >>> xml to mgd >>> Jan 18 10:49:53.626418 RPC-REPLY ERROR: missing or invalid port number in 'et-10/0/' <<<<<<<<<< commit failed [PR1259155](#)
- On MIC-3D-20GE-SFP-E or MIC-3D-20GE-SFP, when SFP diagnostic information is being read out periodically, due to misbehaving SFP or noise on the I2C BUS, SFP thread might be hogging and watchdog

check will restart the MPC to recover. Enhancements of such handling will prevent the SFP thread hogging and MPC restart. [PR1260517](#)

- In a subscriber scenario, when trace options are enabled with flag GRES under PPPoE, due to subscribers' username which containing formate specifier(i.e the character "%") might be unsuccessfully handle by traceoption process, Jpppd might crash. [PR1264000](#)
- Mar 2 00:14:30 CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments
fru_send_msg Global FPC 14 Mar 2 00:14:30 SCC fru_set_boolean: send: set_boolean_cmd Global FPC 14 setting hold-pic-online-for-fabric-ready on These types of messages are benign will potentially be observed with configuration changes in MXVC environment. [PR1264647](#)
- In PPPoE scenario, subscriber might get disconnected due to Keepalive failure in case CPE is adding additional Data field in PPP Echo Request. [PR1273083](#)
- dot1agCfmMepHighestPrDefect may be reported in SNMP trap with value of -1 instead of 0 on recovery after RDI. [PR1273278](#)

Layer 2 Ethernet Services

- In a large scale ISSU testing, a MPC/FPC may go offline during the FRU upgrade phase of unified ISSU. [PR1256940](#)
- During the DHCPv6 renegotiation lockout time BNG doesn't accept any DHCP Solicits with rapid commit options for further processing. This may slow down the subscriber initialization in relatively high packet drop access network segments. Fix for this PR eliminate the impact of DHCPv6 renegotiation lockout timer for DHCP Solicits with rapid commit options [PR1263156](#)

MPLS

- When multiple RSVP LSPs are in ECMP and configured with metric values, if one of the LSPs removed the metric, other LSPs in ECMP might not honor the configured metric. [PR1261961](#)
- During MBB, Next-Hop will change in PFE RSVP route doesn't request for a Next-Hop ACK before changing the route pointing to new NH When scale is high, traffic loss can be seen up to 1 second [PR1264089](#)
- When a container LSP has more than 10 member LSPs, only the first 10 LSPs will be shown in the **show mpls container-lsp name lsp-name statistics** output. [PR1267774](#)
- This core can occur for customers who have configured BGP family L2VPN in Junos OS 17.2R1 Release. This core is currently being worked upon and a fix is targeted for Junos OS 17.2R2 and later releases. This core occurs due to labels being played out of order from master Routing Engine to the backup Routing Engine. [PR1271704](#)

Network Management and Monitoring

- On all platforms, if changing the syslog configuration, the eventd process might stop sending syslog message to a configured syslog server. [PR1246712](#)

- SNMPv2 traps used to have the routing-instance information(context) in the community in the form context@community In SNMPv3, the same routing-instance information will be added to the contextName field of SNMPv3 trap. For traps originating from default routing instance, this field will be empty as it was earlier. [PR1265288](#)

Platform and Infrastructure

- NPC generated a core file with reference to [0x41490f64 in trinity_policer_free (result_ptr=0x5d671f64, nh_ptr=0x5d671f78) at ../../src/pfe/common/pfe-arch/trinity/applications/dfw/dfw_action.c:1049]. This type of NPC core file can be observed with a dynamic configuration change to policer, the processing time in attempting to update all associated policers was exceeded. [PR1071040](#)
- On MX Series with MPCs/MICs-based line cards, when the Major Error happens on MPCs except MPC7/8/9, the default CMERROR actions will be changed to raising alarm and disabling the related Packet Forwarding Engine (PFE). Without the fix, the default action was getting state and raising alarm. This enhancement will disable the PFE as quickly as possible when Major Error happens. [PR1186421](#)
- Due to a code defect related to an ephemeral database, rpd might crash if the ephemeral database is enabled. [PR1214298](#)
- A MPC/FPC might report a LUCHIP EDMEM error during unified ISSU. This may cause inconsistency or incorrect forwarding information (FIB) inside the Packet Forwarding Engine. While the MPC is in the problem state, the Packet Forwarding Engine might experience packet loss. The issue should be self corrected after the unified ISSU process is complete and the Packet Forwarding Engine learns new FIB entries. However, if the problem persist, the MPC may need to be restarted. [PR1249395](#)
- On MX Series, frequent configuration edit from ephemeral database OR issueing command 'user@router> show ephemeral-configuration | display merge' might cause one of the daemons core (dcd, rpd, dfwd, pfed, cosd, sampled) [PR1249979](#)
- Auditd might crash when there are RADIUS servers not reachable and multiple Routing Engine switchovers. When you try to send RADIUS requests to non reachable RADIUS servers, you try for maximum number of times. After the maximum number of tries is reached, you close the socket used to send radius requests. After the socket is closed, you are trying to dispatch next message resulting in crash. Auditd will get restarted automatically after it is crash, so that radius messages if any present in the queue at the time of crash will be lost. After auditd gets restarted, the next event that has to be sent to radius server, will be sent normally. [PR1250525](#)
- Running a pppoe login / logout with igmp distributed enabled but without any multicast forwarding nor igmp joins bbe-smgd daemon might crash. [PR1253036](#)
- In logical-systems environment, if there are some failure resulting that cause Routing Engine switchover(Not perform Routing Engine switchover manually), KRT queue might get stuck on the new master Routing Engine with the error "ENOENT -- Item not found". [PR1254980](#)
- During ISSU, memory from the previous image related to hash tables is not properly recycled, which leads to blocks of physical memory being left unused. The crash is triggered by an attempt to create a memory pool using one of these blocks . [PR1258795](#)

- If IX chipset-based mic(MIC-3D-20GE, for example) is used on an MPC that where has two more MIC slots, **show pfe statistics traffic detail** command could display in or out pps statistics unexpectedly. [PR1259427](#)
- After interface switch when MAC moves from one interface to another, next hop was incorrectly following MAC route ,which has been corrected via code changes . [PR1259551](#)
- When a DHCP/BOOP reply packet is received from an unnumbered interface, FUD process might fail. [PR1260623](#)
- After an unified ISSU upgrade, WRED drop profile might not be programmed correctly, resulting in incorrect WRED drop. [PR1260951](#)
- On MQ chip based MPC, some DDRIF checksum errors are observed, which might traffic blackhole.This PR also adds chassis management alarm when there is DDRIF checksum error on MPC. [PR1260983](#)
- Under certain conditions, MPC cards might crash with a core file due to interface specific filters with policer. [PR1267908](#)
- On all platforms, fast flapping of interfaces/fast changing of configurations might cause rpd crash and BGP sessions will flap very quickly. [PR1269116](#)

Port Security

- The transmit delay interval is the maximum time key server will wait before installing new TX SAK. (default value is 6sec) When MKA transmit interval set to 6 seconds, during key roll over both transmit interval and delay interval timers expires same time and new TX SAK gets installed on key server before RX SAK is installed on peer node causing traffic drop. [PR1257041](#)

Routing Protocols

- Routing protocol process (rpd) may restart unexpectedly with a reference to `ioth_session_delete_internal` () routine. [PR1261970](#)
- On an MX Series platform, if enabling IS-IS segment routing but certain interfaces are not enabled with RSVP, then it might cause corrupted TLV 22 of IS-IS (the size of the value part of the TLV exceeds 255), and it might cause rpd to crash for parsing the LSP (labeled switchover path). [PR1262612](#)
- If `vrf-table-label` is configured in CoC (Carrier of Carriers) VRF routing-instance and a direct interface route is advertised from the VRF towards CE device as BGP-LU (BGP Labeled Unicast) route, the MPLS label entry for the direct route is permanently stuck in KRT (kernel routing table) queue. [PR1263291](#)
- In MX Series, when configuring import policy of IPv6 prefix with an IPv4 next hop for a BGP neighbor, the RPD might crash continuously. The RPD crashing stops only after deletion of the policy. [PR1265224](#)
- After configuring "family inet unicast extended-nexthop", in the BGP open message sent to the peer, "Nexthop AFI=2" should be in the message instead of "Nexthop AFI=3". [PR1272807](#)

Services Applications

- With MS-MIC/MS-MPC used for NAT service, when changing the source IP address under a NAT rule term for a BASIC-NAT translation type, all future traffic hitting the NAT term will be dropped. [PR1257801](#)

- Apply-group configuration may cause KMD process crash during "commit check" process, which causes IPsec tunnel establishment failures. After this fix, apply group can be used. [PR1265404](#)

Subscriber Access Management

- Call rate performance may be impacted under heavy load if there are large numbers of small linked address pools due to a bug in the allocation traversal algorithm. [PR1264052](#)
- **show network-access aaa statistics radius detail** can display incorrect number of messages to RADIUS server in case configured RADIUS server's are continuously flapping. [PR1267307](#)

VPNs

- RPD memory leak is seen when NG-MVPN type 6 and type 7 route adds/deletes/changes. The leak is 36 byte block size on Junos OS Release 15.1 and earlier, and 44 byte block size on Junos OS Release 15.1 or later. [PR1259579](#)

SEE ALSO

[New and Changed Features | 101](#)

[Changes in Behavior and Syntax | 132](#)

[Known Behavior | 143](#)

[Known Issues | 147](#)

[Documentation Updates | 164](#)

[Migration, Upgrade, and Downgrade Instructions | 166](#)

[Product Compatibility | 173](#)

Documentation Updates

IN THIS SECTION

- [Subscriber Management Access Network Guide | 165](#)
- [Subscriber Management Provisioning Guide | 165](#)

This section lists the errata and changes in Junos OS Release 17.2R1 documentation for MX Series.

Subscriber Management Access Network Guide

- The “Configuring the L2TP Resynchronization Method” and “disable-failover-protocol (L2TP)” topics have been updated to state that you can configure the LNS to support only silent failover for peer resynchronization. This capability has been supported on both the LAC and the LNS since Junos OS Release 11.2.

Subscriber Management Provisioning Guide

- Support for the packet-triggered subscribers and policy control rule base (PTSP) feature was discontinued starting in Junos OS Release 13.1R1, but this was not reflected in the documentation. Text exclusive to PTSP has been removed from the *Broadband Subscriber Sessions User Guide*. This includes all CLI topics and the following chapters:
 - “Configuring the PTSP Feature to Support Dynamic Subscribers”
 - “Configuring the PTSP Partition to Connect to the External Policy Manager”
 - “Configuring PTSP Services and Rules”
 - “Monitoring and Managing Packet-Triggered Subscribers”

Topics for other features that refer to PTSP are updated to report the end of support.

SEE ALSO

New and Changed Features 101
Changes in Behavior and Syntax 132
Known Behavior 143
Known Issues 147
Resolved Issues 154
Migration, Upgrade, and Downgrade Instructions 166
Product Compatibility 173

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.2 | 166](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) | 167](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) | 169](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 171](#)
- [Upgrading a Router with Redundant Routing Engines | 171](#)
- [Downgrading from Release 17.2 | 172](#)

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

Starting with Junos OS Release 15.1, in some of the devices, FreeBSD 10.x is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.1 remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see [Understanding Junos OS with Upgraded FreeBSD](#).

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.1-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 17.2

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 [Upgrading Junos OS with Upgraded FreeBSD](#).

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](#).

Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x):

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.

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

10. Install the new **jinstall** package on the routing platform.

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

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.2R2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.2R2.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.2R2.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.2R2.9-limited.tgz
```

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

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

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.1) 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 6.1) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the **junos-vmhost-install-x.tgz** image and specify the name of the regular package in the **request vmhost software add** command. For more information, see VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 17.2 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.1) 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](#).

Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)

Products impacted: MX80, and MX104.

To download and install from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1):

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

<https://www.juniper.net/support/downloads/>

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.

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

- All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/jinstall-ppc-17.2R2.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot  
source/jinstall-ppc-17.2R2.x-limited-signed.tgz
```

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

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

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

Rebooting occurs only if the upgrade is successful.

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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://www.juniper.net/support/eol/junos.html>.

Upgrading a Router with Redundant Routing Engines


If the router has two Routing Engines, perform the following Junos OS installation on each Routing Engine separately to avoid disrupting network operation:

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

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

Downgrading from Release 17.2

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

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

See [Installation and Upgrade Guide](#).

SEE ALSO

New and Changed Features 101
Changes in Behavior and Syntax 132
Known Behavior 143
Known Issues 147
Resolved Issues 154
Documentation Updates 164
Product Compatibility 173

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 173](#)

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 | 101](#)

[Changes in Behavior and Syntax | 132](#)

[Known Behavior | 143](#)

[Known Issues | 147](#)

[Resolved Issues | 154](#)

[Documentation Updates | 164](#)

[Migration, Upgrade, and Downgrade Instructions | 166](#)

Junos OS Release Notes for NFX Series

IN THIS SECTION

- New and Changed Features | 174
- Changes in Behavior and Syntax | 179
- Known Behavior | 180
- Known Issues | 181
- Resolved Issues | 185
- Documentation Updates | 186
- Migration, Upgrade, and Downgrade Instructions | 187
- Product Compatibility | 191

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Hardware | 175
- Juniper Device Manager | 176
- Junos Control Plane | 177
- vSRX | 179

This section describes the new features or enhancements to existing features in Junos OS Release 17.2R1 for NFX Series devices.

Hardware

- **NFX250 platform**—The NFX250 devices constitute Juniper Network’s secure, automated, software-driven customer premises equipment (CPE) devices that deliver virtualized network and security services on demand. Leveraging Network Functions Virtualization (NFV) and built on the Juniper Cloud CPE solution, NFX250 enables service providers to deploy and service chain multiple, secure, high-performance virtualized network functions (VNFs) in a single device.

Table 3: NFX250 Models

Product Number	Specifications	Features
NFX250-S1	1.9 GHz 6-core Intel CPU 16 GB of memory and 100 GB of solid-state drive (SSD) storage Eight 1-GbE network ports, two 1-GbE RJ-45 ports that can be used as either access ports or as uplinks, two SFP ports, two SFP+ ports, one management port, and two console ports	Basic Layer 2/Layer 3
NFX250-S2	1.9 GHz 6-core Intel CPU 32 GB of memory and 400 GB of SSD storage Eight 1-GbE network ports, two 1-GbE RJ-45 ports that can be used as either access ports or as uplinks, two SFP ports, two SFP+ ports, one management port, and two console ports	Basic Layer 2/Layer 3
NFX250-LS1	1.6 GHz 4-core Intel CPU 16 GB of memory and 100 GB of solid-state drive (SSD) storage Eight 1-GbE network ports, two 1-GbE RJ-45 ports which can be used as either access ports or as uplinks, two SFP ports, two SFP+ ports, one Management port, and two Console ports	Supports up to 100 MBPS throughput Secure Router functionality for the following features: <ul style="list-style-type: none"> • IPSec VPN • NAT • Stateful Firewall • Routing services – BGP, OSPF, DHCP, IPv4 and IPv6

- **Transceivers**—NFX250 supports the following optics:
 - 10-gigabit SFP+ transceivers: EX-SFP-10GE-USR, EX-SFP-10GE-SR, EX-SFP-10GE-LR, EX-SFP-10GE-ER, EX-SFP-10GE-ZR
 - 1G-gigabit SFP transceivers: EX-SFP-1GE-SX, EX-SFP-1GE-SX-ET, EX-SFP-1GE-LX, EX-SFP-1GE-LH, EX-SFP-1GE-T, EX-SFP-1GE-LX40K, EX-SFP-GE10KT13R14, EX-SFP-GE10KT14R13, EX-SFP-GE10KT13R15, EX-SFP-GE10KT15R13, EX-SFP-GE40KT13R15, EX-SFP-GE40KT15R13,

EX-SFP-GE80KCW1470, EX-SFP-GE80KCW1490, EX-SFP-GE80KCW1510, EX-SFP-GE80KCW1530, EX-SFP-GE80KCW1550, EX-SFP-GE80KCW1570, EX-SFP-GE80KCW1590, EX-SFP-GE80KCW1610

- **Direct attach copper (DAC) cables**—NFX250 supports the following DAC cables:
 - EX-SFP-10GE-DAC-1M
 - EX-SFP-10GE-DAC-3M
 - EX-SFP-10GE-DAC-5M

Juniper Device Manager

The Juniper Device Manager (JDM) is a low-footprint Linux container that provides these functions:

NOTE: These features were previously supported in the 15.1X53-D40 and 15.1X53-D47 releases of Junos OS.

- Virtual machine (VM) life cycle management
- Device management and isolation of host OS from user installations
- NIC , single-root I/O virtualization (SR-IOV), and virtual input/output (VirtIO) interface provisioning
- Support for the Network Service Orchestrator module to connect to Network Service Activator
- Inventory and resource management
- Internal network and image management
- Service chaining—provides building blocks such as virtual interfaces and bridges for users to implement service chaining policies
- Virtual console access to VNFs including vSRX and vjunos
- Support for outbound SSH connections
- Authentication of users using TACACS+
- Support for configuring SNMP and handling SNMP queries and traps
- Enhanced CLI to support launching VNFs, service chaining VNFs, and configuring and monitoring various system parameters and statistics
- IPsec—The IPsec implementation for NFX250 platforms has been enhanced to protect the management traffic between JDM, VNFs, and the remote SDN controller and other central servers. The IPsec implementation uses AutoKey IKE with preshared keys to authenticate the participants in an IKE session; each side must configure and securely exchange the preshared key in advance. IPsec for NFX250 devices supports only traffic selector based tunnels; multiple IPsec security associations are negotiated based on multiple traffic selectors configured. Configuration of interfaces and static routes is supported.

Junos Control Plane

Junos Control Plane (JCP) is the Junos VM running on the hypervisor. By default, JCP runs as vjunos0 on NFX250. You can use JCP to configure the network ports of the NFX250 device. You can log in to JCP from JDM by using the SSH service and CLI, which is similar to the Junos OS CLI. The JCP supports the following features:

- **Link aggregation**—Link aggregation enables you to use multiple network cables and ports in parallel to increase link speed and improve redundancy.
- **Support for Layer 3 logical interfaces**—A Layer 3 logical interface is a logical division of a physical interface or an aggregated Ethernet interface that operates at the network level and that can receive and forward IEEE 802.1Q VLAN tags. You can use these interfaces to route traffic between multiple VLANs along a single trunk line that connects an NFX250 device to a Layer 2 switch. Only one physical connection is required between the NFX250 device and the switch.
- **VLAN support**—VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
- **Link Layer Discovery Protocol (LLDP) support**—LLDP enables a switch to advertise its identity and capabilities on a LAN, and to receive information about other network devices.
- **Q-in-Q tunneling support**—This feature enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.
- **Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and VLAN Spanning Tree Protocol (VSTP) support**—These protocols enable a switch to advertise its identity and capabilities on a LAN and receive information about other network devices.
- **OSPF support**—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). NFX devices support OSPFv1 and OSPFv2. You can configure OSPF at the [edit protocols ospf] hierarchy level.
- **Bidirectional Forwarding Detection (BFD) support for static routes and the OSPF and RIP protocols**—BFD uses control packets and shorter detection time limits to rapidly detect failures in a network. Hello packets are sent at a specified, regular interval by routing devices. A neighbor failure is detected when a routing device stops receiving a reply after a specified interval.
- **Virtual Router Redundancy Protocol (VRRP) support**—VRRP enables you to provide alternative gateways for end hosts that are configured with static default routes. You can implement VRRP to provide a highly available path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.
- **Internet Group Management Protocol (IGMP) support**—IGMP manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to multicast routers that are their immediate neighbors. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.

- **IGMP snooping support**—IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.
- **Protocol Independent Multicast (PIM) sparse mode support**—PIM sparse mode enables efficient routing to multicast groups with receivers that are sparsely spread over multiple networks. To configure PIM sparse mode, include the **pim** statement at the **[edit protocols]** hierarchy level.
- **SNMP support**—SNMP includes versions 1, 2, and 3 for monitoring system activity.
- **System logging (syslog) support**—Syslog enables you to log system messages into a local directory on the switch or to a system log server.
- **Port mirroring support**—Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, and correlating events.
- **Firewall filter support**—You can provide rules that define whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces.
- **Policing support**—You can use policing to apply limits to traffic flow and to set consequences for packets that exceed those limits.
- **Storm control support**—You can enable the switch to monitor traffic levels and take a specified action when a specified traffic level—called the storm control level—is exceeded, preventing packets from proliferating and degrading service. You can configure a switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs.
- **Class of service (CoS)**—When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default class-of-service(CoS) settings or the CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and any rewrite rules to re-mark packets.
- **Class-of-service (CoS) rewrite rules and classifier support**—You can use rewrite rules to set the value of the CoS bits within a packet header, so you can alter the CoS settings of incoming packets. Packet classification maps incoming packets to a particular class-of-service (CoS) servicing level. You can use classifiers to map packets to a forwarding class and a loss priority and to assign packets to output queues based on the forwarding class.
- **Secure Boot**—The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement Secure Boot.

vSRX

vSRX offers the same capabilities as Juniper Networks SRX Series Services Gateways in a virtual form factor, providing perimeter security, IPsec connectivity, and filtering for malicious traffic without sacrificing reliability, visibility, and policy control. This virtual security and routing appliance ensures reliability for each application. By default, Junos OS Release 15.1X49-D75 for vSRX version is pre-loaded on Junos OS Release 17.2R1 for NFX250 Network Services platform . Earlier versions of vSRX are not compatible with Junos OS Release version 17.2R1 release running on NFX250.

SEE ALSO

[Changes in Behavior and Syntax | 179](#)

[Known Behavior | 180](#)

[Known Issues | 181](#)

[Resolved Issues | 185](#)

[Documentation Updates | 186](#)

[Migration, Upgrade, and Downgrade Instructions | 187](#)

[Product Compatibility | 191](#)

Changes in Behavior and Syntax

There are no changes in default behavior and syntax in Junos OS Release 17.2R1 for the NFX Series.

SEE ALSO

[New and Changed Features | 174](#)

[Known Behavior | 180](#)

[Known Issues | 181](#)

[Resolved Issues | 185](#)

[Documentation Updates | 186](#)

[Migration, Upgrade, and Downgrade Instructions | 187](#)

[Product Compatibility | 191](#)

Known Behavior

IN THIS SECTION

- [Juniper Device Manager | 180](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 for the NFX Series.

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

Juniper Device Manager

- JDM shell configurations of interfaces override JDM CLI configurations. As a workaround, use the JDM CLI to configure interfaces. [PR1155749](#)
- SR-IOV interfaces do not support more than 64 VLANs on NFX250. [PR1156348](#)

SEE ALSO

[New and Changed Features | 174](#)

[Changes in Behavior and Syntax | 179](#)

[Known Issues | 181](#)

[Resolved Issues | 185](#)

[Documentation Updates | 186](#)

[Migration, Upgrade, and Downgrade Instructions | 187](#)

[Product Compatibility | 191](#)

Known Issues

IN THIS SECTION

- [Infrastructure | 181](#)
- [IPsec | 181](#)
- [Juniper Device Manager | 181](#)
- [Junos Control Plane | 183](#)
- [vSRX | 184](#)

This section lists the known issues in hardware and software in Junos OS Release 17.2R1 for the NFX Series devices.

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

Infrastructure

- You might not be able to upgrade from Junos OS releases 15.1X53-D40 and 15.1X53-D45 to Junos OS release 17.2R1. As a workaround, you can use the image file on a USB, configure the NFX device to boot from the USB, and install the upgrade. [PR1252323](#)

IPsec

- There is no CLI command to clear interface flow-statistics on ipsec-nm. [PR1216474](#)
- Initial allocation of hugepages is not guaranteed when the SRX Series packet Forwarding Engine is killed or restarted. [PR1233794](#)

Juniper Device Manager

- There might be no checks when you configure the IP address on different logical units of interfaces. The commit will go through, and will be displayed in the configuration. [PR1150512](#)
- The following commands are not supported:
 - **clear system reboot** and **clear system commit**
 - **restart gracefully**, **restart immediately**, **restart init**, and **restart soft**

- **show ethernet-switching, show version brief, show version all members, and show system services service-deployment**

[PR1154819](#)

- When you use the netconf command to display system information details such as model and OS, the system OS is displayed as QFX. [PR1160055](#)
- Ubuntu package does not successfully install on the JDM container. As a workaround, install the package passwd by using the sudo apt-get install passwd command, which enables the useradd command again. [PR1168680](#)
- When you configure a static route on JDM in enhanced-orchestration disabled mode, there might not be an explicit check to validate the IP address. [PR1173039](#)
- System host bridge uses a default MTU of 1500 and does not support jumbo frames. Currently there is no CLI to configure the MTU on the host bridge. [PR1192169](#)
- The Network Service Orchestrator module commits the configuration on JDM, Junos Control Plane, and IPsec-NM sequentially. If the commit fails on any one of these system VNFs, the Network Service Orchestrator module automatically rolls back to the older configuration on the VNF where the commit error is seen. All prior Network Service Orchestrator module configuration commits on the earlier VNFs continue to exist and not reversed. [PR1196253](#)
- There is no commit check if the PCI address is reused for different interfaces in a VNF. Stop the VNF and then add or delete interfaces. [PR1205497](#)
- Certain VNFs support hot plugging of virtio interfaces when the VNF is running. When a VLAN mapped interface is hot plugged to VNFs such as Centos, the interface is not reachable from the vjunos0 VM. As a workaround, delete the VNF configuration and recommit the complete configuration along with the new interface. [PR1213451](#)
- After enabling or disabling the ipsec-nm service on the NFX250 platform, a warning message might not be displayed asking for a consent to reboot the device. The enabling or disabling action will be effective only after the device is rebooted. Similarly, no warning is displayed when Enhanced orchestration is either enabled or disabled. [PR1213489](#)
- Pre-allocation of hugepages might not consider the available memory and proper commit check is required. Use the feature based on free system memory availability. By default, the system requires up to 6 to 7 gigabytes of memory for various operations. The system might not function properly if more memory than what is available is allocated. [PR1213944](#)
- While spawning a VNF, there might not be a commit check for the valid image type supported. [PR1221642](#)
- If a VNF requests for more memory than the available system memory, commit might go through without any errors resulting in VNF going into a shut off state. As a workaround, use the **show system visibility memory** command to check the available free memory before spawning a VNF. Alternatively, check the log files and the VNF shut off reason will be captured in the `/var/log/syslog` file. [PR1221647](#)
- The following commands are not supported:

- **show host**
- **request system software delete**
- **request system software rollback**
- **request system storage cleanup**

[PR1219972](#)

- DHCP service can be configured on custom system bridges for service chaining. There might be no commit check if the lower and higher values of the pool range are swapped. [PR1223247](#)
- If the configured TACACS+ server has an IP that can be accessed from JDM, the tacplus pam might not wait to timeout if the TACACS+ server is unreachable. [PR1224420](#)
- The Swap memory information displays incorrect values in the show system visibility jdm command output for NFX250 platforms with optimized SSD layouts. [PR1227528](#)
- With enhanced-orchestration mode enabled and routing over management configured on vSRX for WAN redundancy for critical traffic, the system CPU utilization will reach 100% if WAN link goes down and traffic is routed through out-of-band management. vSRX may not respond to ping or management requests. Egress traffic through management might be throttled. [PR1233478](#)
- Removing the IRB configuration along with the DHCP configuration on JDM and rolling back the configuration might result in the DHCP service not functioning for service chaining of VNFs. [PR1234055](#)
- Hugepages that are preconfigured through CLI are not used if a custom init-descriptor is used. [PR1245330](#)
- When a VLAN tag is configured through a JDM CLI on a VNF that is provisioned to a DPDK-enabled VM and the VM is spawned, the VLAN filtering or striping configuration on the VNF stops taking effect. Removing and recommitting the JDM VLAN ID configuration on the VNF can resolve the issue unless the system or the VNF is rebooted. [PR1251596](#)
- The **show system visibility cpu** command on JDM has the field values for IOWait and Intr always set to zero. [PR1258361](#)
- Configuring more than the available number of SR-IOV interfaces in enhanced mode might result in a state where the used MAC addresses for such interfaces are not released back to the system MAC pool on deletion of the VNF. [PR1259975](#)

Junos Control Plane

- The alarm LED will be amber for a major alarm instead of red. In the NFX250-S1E model, the alarm LED does not blink for any alarms. [PR1146307](#)
- Configuring DSCP and DSCPv6 classifiers together on a Layer 2 interface is not supported. [PR1169529](#)
- When the option **accept-source-mac mac-address** is configured on an interface and then deleted, no additional MAC's will be learnt on the interface. Only the MAC's that were earlier configured will be available. [PR1168197](#)

- When LLDP is configured on vjunos0 on an NFX250 Network Services platform, the system name TLV(5) might not be advertised. [PR1169479](#)
- There might a traffic drop in IPv4 multicast traffic on JCP when flow-control is configured on interfaces and multicast traffic is more than 400pps. [PR1191794](#)
- On an interface with family inet configured, you might not be able to configure a classifier or rewrite rules. [PR1262840](#)
- If the traffic in the out-of-band interface is heavy, the control plane connectivity might get blocked for sometime while the packets are processed. If this interruption persists, the connection between the packet forwarding engine (PFE) and control plane is cleared, which results in a PFE restart or shutdown. You must ensure that there is no heavy traffic flow in the management VLAN. [PR1270689](#)

vSRX

- On an NFX250-S1E platform running vSRX VNF, the performance of SR-IOV with UTM and IDP is lower than VirtIO with UTM and IDP. [PR1214118](#)
- If per-unit-scheduler is not configured, the IFD shaping fails and no packet is queued. [PR1264556](#)
- After configuring the IFD shaping, the ingress interface cannot receive packets. [PR1264850](#)
- The current maximum number of concurrent SIP calls is below the specified maximum limit. [PR1273356](#)

SEE ALSO

[New and Changed Features | 174](#)

[Changes in Behavior and Syntax | 179](#)

[Known Behavior | 180](#)

[Resolved Issues | 185](#)

[Documentation Updates | 186](#)

[Migration, Upgrade, and Downgrade Instructions | 187](#)

[Product Compatibility | 191](#)

Resolved Issues

IN THIS SECTION

- [Juniper Device Manager | 185](#)
- [Junos Control Plane | 185](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

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

Juniper Device Manager

- User-defined login class is not supported. [PR1155965](#)
- ping command when initiated with **record route** option on virtio interfaces, returns no results. [PR1162659](#)
- Clients that are connected via the front panel ports are assigned the default gateway 10.10.10.254 by Network Service Activator. [PR1168284](#)
- The CLI command to configure the time zone is not functional. [PR1169675](#)
- JDM SNMP traps have limited support on interfaces only. [PR1173216](#)

Junos Control Plane

- You cannot configure a transmit rate of 0 for class-of-service schedulers. [PR1158085](#)
- If a cable is not connected to the front panel RJ-45 ports, the status led will blink. [PR1168054](#)
- SFP-T transceivers are not supported. [PR1151575](#), [PR1166808](#), [PR1168203](#)

SEE ALSO

[New and Changed Features | 174](#)

[Changes in Behavior and Syntax | 179](#)

[Known Behavior | 180](#)

[Documentation Updates | 186](#)[Known Issues | 181](#)[Migration, Upgrade, and Downgrade Instructions | 187](#)[Product Compatibility | 191](#)

Documentation Updates

IN THIS SECTION

- [Release Notes | 186](#)

There are no errata or changes in Junos OS Release 17.2R1 documentation for NFX Series.

Release Notes

- The *Junos OS Release Notes* for Release 17.2R1 erroneously listed support for the NFX250-S2-R device and its supported features. This device is not currently supported.

SEE ALSO

[New and Changed Features | 174](#)[Changes in Behavior and Syntax | 179](#)[Known Behavior | 180](#)[Known Issues | 181](#)[Resolved Issues | 185](#)[Migration, Upgrade, and Downgrade Instructions | 187](#)[Product Compatibility | 191](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 187
- Upgrading Using Unified ISSU | 187
- Upgrading a Router with Redundant Routing Engines | 188
- Basic Procedure for Upgrading to Release 17.2 | 188

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. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

Upgrading Using Unified ISSU

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

upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [Understanding High Availability Features on Juniper Networks Routers](#).

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

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

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

Basic Procedure for Upgrading to Release 17.2

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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

To download and install Junos OS Release 17.2R1:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the router.

NOTE: After you install a Junos OS Release 17.2R1 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-17.2R1.13-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.2R1.13-export-signed.tgz
```

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

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

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

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

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the *Software Installation and Upgrade Guide*.

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

NOTE: Most of the existing **request system** commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the *Installation and Upgrade Guide*.

SEE ALSO

[New and Changed Features | 174](#)

[Changes in Behavior and Syntax | 179](#)

[Known Behavior | 180](#)

[Documentation Updates | 186](#)

[Known Issues | 181](#)

[Resolved Issues | 185](#)

[Product Compatibility | 191](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 192](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

[New and Changed Features | 174](#)

[Changes in Behavior and Syntax | 179](#)

[Known Behavior | 180](#)

[Documentation Updates | 186](#)

[Known Issues | 181](#)

[Resolved Issues | 185](#)

[Migration, Upgrade, and Downgrade Instructions | 187](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- [New and Changed Features | 193](#)
- [Changes in Behavior and Syntax | 213](#)
- [Known Behavior | 215](#)
- [Known Issues | 217](#)
- [Resolved Issues | 218](#)
- [Documentation Updates | 220](#)

- Migration, Upgrade, and Downgrade Instructions | 220
- Product Compatibility | 225

These release notes accompany Junos OS Release 17.2R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Hardware | 194
- Class of Service (CoS) | 195
- Forwarding and Sampling | 196
- General Routing | 196
- High Availability (HA) and Resiliency | 196
- Interfaces and Chassis | 197
- IPv6 | 201
- Layer 2 Features | 201
- Layer 3 Features | 201
- Management | 202
- MPLS | 205
- Multicast | 207
- Network Management and Monitoring | 207
- Routing Policy and Firewall Filters | 208
- Routing Protocols | 209
- Security | 211
- Services Applications | 211
- Software Installation and Upgrade | 211

This section describes the new features and enhancements to existing features in Junos OS Release 17.2R1 for the PTX Series.

Hardware

- **PTX10008 Packet Transport Router**—PTX10008 Packet Transport Router provides 3.0 Tbps per slot forwarding capacity for the service providers and cloud operators. The router provides an opportunity for the cloud and data center operators for a smooth transition from 10-Gigabit Ethernet and 40-Gigabit networks to 100-Gigabit Ethernet high-performance networks. This high-performance, 13 rack unit (13RU) modular chassis provides 24 Tbps of throughput and 16 Bpps of forwarding capacity. The PTX10008 router has eight slots for the line cards that can support a maximum of 1152 10-Gigabit Ethernet ports, 288 40-Gigabit Ethernet ports, or 240 100-Gigabit Ethernet ports.

You can deploy the PTX10008 router in the core of the network for the following functions:

- Label Switched Routing (LSR)
- IP Core Routing
- Internet Peering

PTX10008 Packet Transport Router supports two new line cards, LC1101 and LC1102. The LC1101 line card consists of thirty QSFP+ Pluggable Solution (QSFP28) cages that support 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The line card supports speed of either 40-Gbps or 100-Gbps. It also supports 10-Gigabit Ethernet by channelizing the 40-Gigabit Ethernet ports. The default port speed is 100-Gbps. If the user plugs in 40Gigabit or 4x10Gigabit optic, the appropriate port speed has to be configured manually.

The LC1102 line card consists of 36 quad small form-factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. The QSFP+ ports support , 40-Gigabit or 100-Gigabit Ethernet optical transceivers in selected ports. The default port speed on the LC1102 line card is channelized 10-Gbps. Out of these 36 ports, 12 ports are QSFP28 capable for supporting 100-Gigabit Ethernet. The

line card supports 10-Gigabit Ethernet by channelizing the 40-Gigabit ports. Channelization is supported on fiber break-out cable using standard structured cabling techniques

Class of Service (CoS)

- **Support for CoS-Based Forwarding (PTX10008)**—CoS-Based Forwarding (CBF) enables the control of next-hop selection based on a packet's class of service field. Starting with Junos OS Release 17.2R1, PTX10008 routers support CBF. CBF can only be configured on a device with eight or fewer forwarding classes plus a default forwarding class. You can implement CBF by specifying **next-hop-map** at the **[edit class-of-service forwarding-policy]** hierarchy level and then applying the **next-hop-map** at the **[edit policy-options]** hierarchy level.
- **CoS-based forwarding support for up to 16 forwarding classes (MX and PTX routers)** — Beginning with Junos OS Release 17.2, MX routers with MPCs or MS-DPCs, vMX, PTX3000 routers, PTX5000 routers, and VPTX support configuring CoS-based forwarding (CBF) for up to 16 forwarding classes. All other platforms support CBF for up to 8 forwarding classes. To support up to 16 forwarding classes for CBF on MX routers, enable **enhanced-ip** at the **[edit chassis network-services]** hierarchy level.

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

- **Support for class of service (CoS) (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, support is extended for class of service (CoS). CoS is the assignment of traffic flows to different service levels. Service providers can use router-based CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows.

On a PTX1000 router, you can divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs.

- **Support for Shaping of Traffic Exiting Third-Generation FPCs on PTX1000 Routers (PTX Series)**—Beginning with Junos OS Release 17.2R1, you can shape the output traffic of an FPC3 physical interface on a PTX1000 packet transport router so that the interface transmits less traffic than it is physically capable of carrying. Shaping on all PTX Series packet transport router interfaces has a minimum rate of 1 Gbps and an incremental granularity of 0.1 percent of the physical interface speed after that (for example, 10 Mbps increments on a 10 Gbps interface). You can shape the output traffic of a physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces interface-name]** or **[edit class-of-service traffic-control-profiles profile-name]** hierarchy level and applying the traffic control profile to an interface.

[See [shaping-rate \(Applying to an Interface\)](#)]

Forwarding and Sampling

- **Support for Bidirectional Forwarding Detection (BFD) (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, support is extended for Bidirectional Forwarding Detection (BFD). The BFD protocol uses control packets and shorter detection time limits to rapidly detect failures in a network. Hello packets are sent at a specified, regular interval by routing devices. A neighbor failure is detected when a routing device stops receiving a reply after a specified interval.

On a PTX1000 router, you can configure BFD for static routes and for the BGP, IS-IS, OSPF, PIM, and RIP protocols.

General Routing

- **OpenConfig: Supporting the BGP Model in JUNOS (PTX Series)**—Starting in Junos OS Release 17.2R1, the configuration leafs defined in the `openconfig-bgp.yang` and `openconfig-bgp-multiprotocol.yang` files are supported.
- **OpenConfig: BGP Routing Table - Supporting Operational State Model (PTX Series)**—Starting in Junos OS Release 17.2R1, the OpenConfig BGP RIB routing table supports local-rib for IPv4 and IPv6. The `openconfig-rib-bgp.yang` model supports five logical RIBs per address family. There are five tables for IPv4 routes and five tables for IPv6 routes.

High Availability (HA) and Resiliency

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

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

- **Resiliency Support for LC1101 and LC1102 (PTX10008)**—Starting with Junos OS Release 17.2R1, resiliency support is enabled for the following devices:
 - LC1101 and LC1102
 - Switch Interface Boards
- **Chassis management**—In Junos OS Release 16.1X65 and 17.2R1, the following CLI operational mode commands are supported on a PTX1000 router:
 - `show chassis hardware`
 - `show chassis temperature-thresholds`

- **show chassis environment**
- **show chassis firmware**

Interfaces and Chassis

- **Support for packet-forwarding features on LC1101 and LC1102 (PTX10008)**—Starting in Junos OS Release 17.2R1, the following key packet-forwarding features are enabled on LC1101 and LC1102 for PTX10008 routers:
 - Basic Layer 2 features and protocols
 - Class-of-service (CoS)
 - Firewall filters and policers
 - Hash enhancement
 - Large scale IPv4 and IPv6 Forwarding information base(FIB)
 - Layer 3 VPNs
 - MPLS
 - Sampling and port mirroring
- **Fabric Management Support (PTX10008)**—Starting in Junos OS Release 17.2R1, you can set up and manage the fabric connections between the Packet Forwarding Engines of LC1101 and LC1102 in the PTX10008 routers. Fabric management includes collecting fabric status and statistics, monitoring health of the hardware, and responding to CLI queries. It also tracks addition and removal of FRUs from the router and monitors faults in the data plane. It is enabled by default and can be monitored by using the following commands:
 - **show chassis fabric summary**
 - **show chassis fabric fpcs fpc fpc-slot**
 - **show chassis fabric sibs**
 - **show chassis fabric errors**
 - **show chassis fabric reachability**
- **Support for LC1101 and LC1102 with Routing and Control Board (RCB) (PTX10008)**—Starting with Junos OS Release 17.2R1, the PTX10000 Routing and Control Board (RCB) is supported on PTX10008 routers. The PTX10008 chassis can run with one or two RCBs. A fully-redundant system requires a second RCB. When two RCBs are installed, one RCB functions as the master and the second as the backup. If the master RCB is removed, the backup starts and becomes the master. The RCB integrates the control plane and Routing Engine functions into a single management unit. The RCB handles system control functions such as environmental monitoring, routing Layer 2 and Layer 3 protocols, alarm and logging functions and other functions required to manage the operation of a chassis.

- **Support for 10-Gigabit Ethernet on LC1101 - 30C line card (PTX10008)**—Starting in Junos OS Release 17.2R1, PTX10008 routers support 10-Gigabit Ethernet interfaces in addition to 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces on the LC1101 - 30C line card.

When a particular PE chip or PFE (Packet Forwarding Engine) is working in mode A to support 10-Gigabit Ethernet, ports 6, 7, 16, 17, 26, and 27 at the PE0 to PE5 level are non-operational. However, once the PE goes into mode D, these ports become operational and can operate at 40-Gigabit Ethernet, or 100-Gigabit Ethernet speed.

For 10-Gigabit Ethernet, you must configure the port using the channelization command. Because there is no port-groups option for the 100-Gigabit Ethernet line card, you must use individual port channelization commands.

In 30C line card, by default FPC comes up in mode D, when you channelize the first port in any PE, whole FPC restarts and corresponding PE comes up in mode A. Further channelization in that PE does not restart the FPC. But if you channelize some another ports in other PE, then the whole FPC restarts again. If you undo the channelization of all ports in any PE, then FPC gets restarted and corresponding PE comes up in mode D, which is the default mode.

NOTE: If any mode changes (A to D or D to A) occur at the PE, the line card automatically performs a cold reboot.

- **Support for channelizing the 40-Gigabit Ethernet ports on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, PTX1000 Packet Transport Router supports 10-Gbps, 40-Gbps, and 100-Gbps port speeds, enabling service providers to organically distribute peering points throughout the network. You can channelize four 10-Gigabit Ethernet interfaces from the 40-Gigabit Ethernet interfaces. By default, the 40-Gigabit Ethernet interfaces are named **et-fpc/pic/port**. The names of the channelized 10-Gigabit Ethernet interfaces appear in the format: **et-fpc/pic/port:channel**, where **channel** is a value from 0 through 3.

To configure a port speed of 40-Gbps or 10-Gbps, use the **set chassis fpc slot pic pic-slot port 0..71 channel-speed (10g|40g)** command.

You can also configure 24 out of the 72 ports to operate at 100-Gbps speed.

- **Support for Packet forwarding features (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, the PTX1000 Packet Transport Router supports the following key packet forwarding features:
 - Basic Layer 2 features and protocols
 - Class-of-service (CoS)
 - Firewall filters and policers
 - Hash enhancement
 - Large scaling IPv4/v6 forwarding information base (FIB)

- Layer 3 VPNs
- MPLS
- Sampling and port mirroring
- **Support for configuring multiple port speeds on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, PTX1000 Packet Transport Router supports 10-Gbps, 40-Gbps, and 100-Gbps port speeds, enabling service providers to organically distribute peering points throughout the network. To configure the port speed, use the **speed [10G | 40G | 100G]** statement at the **[edit chassis fpc slot-number pic pic-number port port-number]** hierarchy level. The default port speed is **10G**.

Support for configuring local loopback on PTX1000—In Junos OS Release 16.1X65 and 17.2R1, to enable the local loopback, use the **loopback local** configuration statement on PTX1000 interfaces. The PTX1000 supports only local loopback, not remote loopback. Configure the statement at the **[edit interfaces interface-name gigether-options]** hierarchy level.

- **Support for Aggregated Ethernet (AE) (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, support is extended for aggregated Ethernet (AE). The Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-flow load balancing.

On a PTX1000 router, you can configure the member links of an aggregated Ethernet bundle with any combination of rates—also known as mixed rates. The bandwidth that is provided by an aggregated Ethernet bundle can be utilized completely and efficiently when the links are configured with different rates.

- **Support for DCU accounting and SCU accounting on PTX1000 routers**—Starting in Junos OS Release 16.1X65 and 17.2R1, the destination class usage (DCU) accounting and source class usage (SCU) accounting are supported on PTX1000 routers.

You can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets, which are defined as source classes and destination classes. SCU counts packets sent to customers by performing lookups on the source IP address and the destination IP address. SCU accounting enables you to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. DCU counts packets from customers by performing lookups of the IP destination address. DCU accounting enables you to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.

NOTE: DCU accounting and SCU accounting are supported only if the enhanced-mode statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Support for unicast RPF (PTX1000)**—Starting in Junos OS Release 17.2R1, you can configure unicast reverse path forwarding (RPF) to reduce the impact of denial of service (DoS) attacks on PTX Series routers that have third-generation FPCs installed.

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

IPv6

- **IPv6 statistics on PTX1000, PTX3000, PTX5000, and PTX10008 with third-generation FPCs**—Starting in Junos OS Release 17.2R1, you can obtain the transit IPv6 statistics at both the physical interface and logical interface levels on third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX 5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008 by using both a CLI command and SNMP MIB counters. Use the **show interfaces statistics** command to display both physical interface and logical interface statistics. You can view only logical interface statistics if you use SNMP MIB counters. However, for aggregated Ethernet interfaces, the accounting is not done at the level of the child links and, thus, IPv6 statistics for child links are not displayed.

To start getting IPv6 statistics on third-generation FPCs, use the **route-accounting** statement at the **[edit forwarding-options family inet6]** hierarchy level. PTX Series routers with first- and second-generation FPCs do not display IPv6 statistics for physical interfaces or logical interfaces, and transit statistics on child links in aggregated Ethernet interfaces are also not taken into account.

NOTE: Egress IPv6 statistics are not taken into account in case of MPLS POP where IPv6 traffic is encapsulated within MPLS and MPLS is stripped off before the plain IPv6 traffic is forwarded.

[See [route-accounting](#) and [show interfaces extensive](#) .]

Layer 2 Features

- **Support for Layer 2 protocols**—Starting in Junos OS Release 17.2R1, Layer 2 protocols are supported on PTX10008 routers. Layer 2 protocols include Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), VLAN Spanning Tree Protocol (VSTP), Link Layer Discovery Protocol (LLDP), and so on.

Layer 3 Features

- **BGP (PTX1000)**—In Junos OS Release 16.1X65 and 17.2R1, BGP is an exterior gateway protocol (EGP) for routing traffic between autonomous systems (AS). You can configure BGP at the **[edit protocols bgp]** hierarchy level.

OSPF (PTX1000)—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). PTX1000 routers support OSPFv1, OSPFv2, and OSPFv3. You can configure OSPF at the **[edit protocols ospf]** hierarchy level.

Synchronization between OSPF and LDP (PTX1000)—LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by OSPF. If the synchronization

between LDP and OSPF is lost, the label-switched path (LSP) goes down. Therefore, LDP and IS-IS synchronization is beneficial.

To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the **ldp-synchronization** statement at the **[edit protocols ospf interface *interface-name*]** hierarchy.

IS-IS (PTX1000)—The IS-IS protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system.

Synchronization between IS-IS and LDP (PTX1000)—LDP distributes labels in non-traffic-engineered applications. Labels are distributed along the best path determined by IS-IS. If the synchronization between LDP and IS-IS is lost, the label-switched path (LSP) goes down. Therefore, LDP and IS-IS synchronization is beneficial.

To advertise the maximum cost metric until LDP is operational for LDP synchronization, include the **ldp-synchronization** statement at the **[edit protocols isis interface *interface-name*]** hierarchy.

- **Support for Layer 3 protocols**—Starting in Junos OS Release 17.2R1, Layer 3 protocols are supported on PTX10008 routers. Layer 3 protocols include the Multiprotocol Label Switching (MPLS), Layer 3 Virtual Private Network (L3VPN), Bidirectional Forwarding Detection (BFD), Layer 2 Virtual Private Network (L2VPN), Point-to-multipoint (P2MP), Fast ReRoute (FRR), Operations, Administration and Maintenance (OAM), Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP), Adaptive Load Balancing (ALB), and so on.

Management

- **Support for LSP events and properties sensor for Junos Telemetry Interface (PTX3000 and PTX5000 routers)**—Starting with Junos OS Release 17.2R1, you can export statistics for LSP events and properties through the Junos Telemetry Interface. Only gRPC streaming for this sensor is supported. You can export statistics for ingress point-to-point LSPs, point-to-multipoint LSPs, bypass LSPs, and dynamically created LSPs. To export data through gRPC, use the **/mpls/lsp/** or **/mpls/signal-protocols/** set of OpenConfig subscription paths. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of the Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models. To enable the LSP events and properties sensor, you must also include the **lsp-telemetry** statement at the **[edit routing-options]** hierarchy level.

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

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

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

- **Support for LSP statistics for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.2R1, you can stream telemetry data for LSPs through UDP and gRPC. To provision an LSP statistics sensor for UDP streaming, include the **resource /junos/services/label-switched-path/usage/** statement at the **[edit services analytics sensor sensor-name]** hierarchy level. Use the **mpls/lsp/constrained-path/tunnels/tunnel/** path to provision a sensor for streaming LSP statistics through gRPC. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions. For both UDP and gRPC streaming, you must also configure the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level. Support for the LSP statistics sensor was previously introduced in Junos OS Release 15.1F6 and Junos OS Release 16.1R4.

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

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

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

- **Support for gRPC streaming for Junos Telemetry Interface firewall filter statistics (PTX3000 and PTX5000 routers)**—Starting with Junos OS Release 17.2R1, you can use gRPC interfaces to provision sensors to subscribe to and receive firewall filter telemetry data. Traffic-class counter statistics are also collected. Use the **/junos/firewall/firewall-stats/** path to provision a sensor for firewall filter statistics. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models. OpenConfig paths are used to define telemetry parameters for data streamed through gRPC. This functionality was previously introduced in Junos OS Release 16.1R4.

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

- **Support for the Junos Telemetry Interface (PTX1000 routers)**—Starting with Junos OS Release 17.2R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for several network elements without involving polling. You can stream data through UDP or gRPC.

Only the following sensors are supported on PTX1000 routers:

- Physical interfaces statistics
- Label-switched-path (LSP) statistics
- Network processing unit (NPU) memory
- NPU memory utilization
- CPU memory

To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. To provision sensors to stream data through gRPC, use the **telemetrySubscribe RPC** to specify telemetry parameters for a specified list of OpenConfig command paths. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models.

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

- **Support for queue statistics for logical interface sensors for Junos Telemetry Interface (PTX3000 and PTX5000 routers)**—Starting with Junos OS Release 17.2R1, logical interface sensors also collect egress and ingress queue statistics. Both UDP and gRPC streaming are supported. Queue statistics, including for per-unit queuing and hierarchical queuing, are exported when a queuing structure is configured on a logical interface. To provision a logical interfaces statistics sensor for UDP streaming, include the **resource /junos/system/linecard/interface/logical/usage/** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level. To provision a sensor for gRPC streaming, include **/interfaces/interface[name='interface-name']/subinterfaces/** in the subscription path. Use the **telemetrySubscribe RPC** to define telemetry parameters for gRPC streaming. If your Juniper Networks device is running a version of Junos OS with the upgraded FreeBSD kernel, you must download the Junos Network Agent package, which provides the interfaces to manage gRPC subscriptions.

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

MPLS

- **MPLS inter-AS link protection (PTX Series)**—Starting in Junos OS Release 17.2R1, MPLS inter-AS link protection is supported. Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router will then choose an alternate link through another interface to send traffic to its destination.

For an MPLS inter-AS environment, link protection can be enabled when labeled-unicast is used to send traffic between autonomous systems (ASs). To configure link protection on an interface, the **protection** statement is introduced at the `[edit protocols bgp group group-name family inet labeled-unicast]` hierarchy level.

[See [protection](#).]

- **Support for filter-based GRE for IPv4 and IPv6 tunneling on PTX Series routers**—In Junos OS Release 16.1X65 and 17.2R1, the filter-based generic routing encapsulation (GRE) for IPv4 and IPv6 tunneling uses firewall filters to provide de-encapsulation of GRE traffic. The configuration of filter-based GRE de-encapsulation supports the **routing-instance** statement as one of the attributes.

NOTE: Configuring filter-based GRE for IPv4 and IPv6 tunneling is supported only when the enhanced-mode statement is configured at the `[edit chassis network-services]` hierarchy level.

- **MPLS inter-AS link protection (PTX1000)**—Starting in Junos OS Release 17.2R1, MPLS inter-AS link protection is supported. Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router will then choose an alternate link through another interface to send traffic to its destination.

For an MPLS inter-AS environment, link protection can be enabled when **labeled-unicast** is used to send traffic between autonomous systems (ASs). To configure link protection on an interface, the **protection** statement is introduced at the `[edit protocols bgp group group-name family inet labeled-unicast]` hierarchy level.

- **LDP (PTX1000)**—The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP enables routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths. For more information, see the *MPLS Applications User Guide for Routing Devices*.

RSVP (PTX1000)—RSVP is a resource reservation setup protocol that is used by both network hosts and routers. Hosts use RSVP to request a specific class of service (CoS) from the network for particular application flows. Routers use RSVP to deliver CoS requests to all routers along the data path. RSVP also can maintain and refresh states for a requested CoS application flow. For more information, see the *MPLS Applications User Guide for Routing Devices*.

ECMP (64-way) with configurable Layer 3 hash options (PTX1000)—Allows you to configure 64 equal-cost multipath (ECMP) next hops for RSVP and LDP LSPs. To configure the maximum limit for ECMP next hops, include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level.

To view the details of the ECMP next hops, issue the **show route** command. The **show route summary** command also shows the current configuration for the maximum ECMP limit.

MPLS capabilities (PTX1000)—MPLS provides both label edge router (LER) and label switch router (LSR) and provides the following capabilities:

- Object access method, including ping, traceroute, and Bidirectional Forwarding Detection (BFD)
- Fast reroute (FRR), a component of MPLS local protection

Both one-to-one local protection and many-to-one local protection are supported.

- Loop free alternate FRR
- 6PE and 6VPE devices
- Layer 3 VPNs for both IPv4 and IPv6
- **IPv6 tunneling over an MPLS-based IPv4 network (PTX1000)**—In Junos OS Release 16.1X65 and 17.2R1, tunneling enables you to connect IPv6 sites over an IPv4 MPLS-enabled backbone. IPv6 packets are carried over an IPv4 MPLS tunnel. To enable this service, you need to deploy provider edge (PE) routers that can run IPv4, MPLS, and BGP toward the core and IPv6 toward the edge.

[See [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks](#)]

- **Egress peer engineering of service labels (such as BGP and MPLS) and egress peer protection for BGP-LU (PTX1000)**—In Junos OS Release 16.1X65 and 17.2R1, you can enable traffic engineering of service traffic, such as MPLS LSP traffic between autonomous systems (ASs), by using BGP-labeled unicast for optimum utilization of the advertised egress routes. You can specify one or more backup devices for the primary egress AS boundary router. Junos OS installs the backup path in addition to the primary path in the MPLS forwarding table, which enables MPLS fast reroute (FRR) when the primary link fails. It provides support for the FRR protection backup scheme to perform an IP lookup to determine a new egress interface.

[See [Configuring Egress Peer Traffic Engineering by Using BGP Labeled Unicast and Enabling MPLS Fast Reroute](#)]

Multicast

- **Support for Internet multicast (PTX Series)**—Starting in Junos OS Release 17.2R1, the `mpls-internet-multicast` routing instance type uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP. Previously this feature was supported only on PTX Series routers with third-generation FPCs installed. Now this feature is supported when first-generation FPCs or second-generation FPCs are installed with third-generation FPCs on a PTX Series router.

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

NOTE: For the third-generation FPCs to interoperate with the previous FPCs, the enhanced-mode statement cannot be configured on the chassis. To support Internet Multicast, the MPLS core-facing interfaces must be third-generation FPCs.

Network Management and Monitoring

- **Support for inline active flow monitoring on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, you can use the export capabilities of inline active flow monitoring with IP Flow Information Export (IPFIX) to define a flow record template suitable for IPv4 or IPv6 traffic. The flow record template provides the flexibility for future enhancements and the ability to add new attributes to inline active flow monitoring without changing to a newer version.
- **Support for SNMP on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, SNMP versions 1, 2, and 3 are supported on the PTX1000. SNMP enables you to monitor network devices from a central location. Junos OS includes an SNMP agent that provides remote management applications with access to detailed information about the devices on the network.
- **Junos Space Service Now (PTX1000)**—In Junos OS Release 16.1X65 and 17.2R1, PTX1000 routers support Junos Space Service Now. The Junos Space Service Now is an application that runs on the Junos Space Network Management Platform to automate fault management and accelerate issue resolution.
[See [Junos Space Service Now](#)]
- **Support for accounting profiles (PTX1000)**—Starting in Junos OS Release 17.2R1, you can configure accounting profiles to collect data on PTX Series routers that have third-generation FPCs installed.

NOTE: Configuring accounting profiles is supported only when the enhanced-mode statement is configured at the `[edit chassis network-services]` hierarchy level.

- **SNMP support for monitoring tunnel statistics (PTX Series)**—Starting in Junos OS Release 17.2R1, SNMP MIB `jnxTunnelStat` supports monitoring of tunnel statistics for IPv4 over IPv6 tunnels. This is a new enterprise-specific MIB, Tunnel Stats MIB, that currently displays three counters: tunnel count in `rp`, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine. This MIB can be extended to support other tunnel statistics. The MIB is defined in `jnx-tunnel-stats.txt`. This MIB is attached to `jnxMibs`.

Routing Policy and Firewall Filters

- **Hop-limit firewall filter match condition supported (PTX Series)**—Starting in Junos OS Release 17.2R1, you can configure a firewall filter using the `hop-limit` and `hop-limit except` match conditions for IP version 6 (IPv6) traffic (family `inet6`).

NOTE: The `hop-limit` and `hop-limit except` match conditions are supported on PTX series routers when `enhanced-mode` is configured on the router.

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

- **Support for firewall filter with match conditions (PTX1000)**—Starting in Junos OS Release 16.1X65 and 17.2R1, you can configure a firewall filter with match conditions for IP version 4 (IPv4) traffic.
- **Support for the no-decrement-ttl tunneling attribute (PTX1000)**—Starting in Junos OS Release 17.2R1, you can configure the `no-decrement-ttl` tunneling attribute for filter-based generic routing encapsulation (GRE) for IPv4 and IPv6 tunneling.

NOTE: The `no-decrement-ttl` tunneling attribute is supported only when the enhanced-mode statement is configured at the `[edit chassis network-services]` hierarchy level.

Routing Protocols

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

[See [Understanding BGP Link-State with Source Packet Routing in Networking \(SPRING\) Extensions](#).]

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

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

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

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

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

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

- Support for IS-IS segment routing on PTX1000**—Starting in Junos OS Release 16.1X65 and 17.2R1, IS-IS segment routing support is enabled through MPLS. Currently, label advertisements are supported for IS-IS only. IS-IS creates an adjacency segment per adjacency, per level, and per address family (one each for IPv4 and IPv6). Junos OS IS-IS implementation allocates node segment label blocks in accordance with the IS-IS protocol extensions for supporting segment routing node segments and provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:
 - source-packet-routing**—Enable the source packet routing feature.
 - node-segment**—Enable source packet routing at all levels.
 - use-source-packet-routing**—Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and primary IS-IS source packet routing node segments.
 - no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.
- BGP advertises multiple add-paths based on community value (PTX1000)**—Beginning with Junos OS 17.2R1, you can define a policy to identify eligible multiple path prefixes based on community values. BGP advertises these community-tagged routes in addition to the active path to a given destination. If the community value of a route does not match the community value defined in the policy, then BGP does not advertise that route. This feature allows BGP to limit the number of multiple paths that are processed and not advertise more than 20 paths to a given destination. You can limit and configure the number of prefixes that BGP considers for multiple paths without actually knowing the prefixes in advance. Instead, a known BGP community value determines whether or not a prefix is advertised.
- Selective advertising of BGP multiple paths (PTX1000)**—Beginning with Junos OS Release 17.2R1, you can restrict BGP **add-path** to advertise contributor multiple paths only. Advertising all available multiple paths might result in a large overhead of processing on device memory and is a scaling consideration, too. You can limit and configure upto six prefixes that the BGP **multipath** algorithm selects. Selective advertising of multiple paths facilitates Internet service providers and data centers that use route reflector to build in-path diversity in IBGP.
- Support for BGP to carry flow-specification routes (PTX1000)**--Starting in Junos OS Release 17.2R1, BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX1000 routers that have third-generation FPCs installed. Propagating flow routes as BGP NLRI messages in essence enables the propagation of firewall filters which protects the system against denial-of-service (DOS) attacks.

Security

- **Firewall filter support (PTX10008)**—Starting in Junos OS Release 17.2R1, you can define firewall filters on the PTX10008 routers that defines whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces.

Services Applications

- **Support for inline JFlow version 9 flow templates (PTX5000 and PTX3000)**—Starting in Junos OS Release 17.2R1, 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 [Configuring Flow Aggregation to Use Version 9 Flow Templates on PTX Series Routers](#).]

Software Installation and Upgrade

- **Zero Touch Provisioning on PTX1000 routers**—Starting in Junos OS Release 17.2R1, ZTP is supported to automate the provisioning of the device configuration and software image with minimal manual intervention.

When you physically connect a router to the network and boot it with a default configuration, the router attempts to upgrade the Junos OS software image automatically and autoinstall a configuration file from the network. The router uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the router boots with the pre-installed software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version.

[See [Understanding Zero Touch Provisioning](#) and [Configuring Zero Touch Provisioning](#)]

User Interface and Configuration

- **Monitoring, detecting, and taking action on degraded physical 100-Gigabit Ethernet links to minimize packet loss (PTX1000)**—Starting with Junos OS Release 17.2R1, you can monitor physical link degradation (indicated by bit error rate (BER) threshold levels) on Ethernet interfaces, and take corrective actions if the BER threshold value drops to a value in the range of 10^{-13} to 10^{-5} .

The following new configurations have been introduced at the `[edit interfaces interface-name]` hierarchy level to support the physical link degrade monitoring and recovery feature on Junos OS:

- To monitor physical link degrade on Ethernet interfaces, configure the **link-degrade-monitor** statement.
- To configure the BER threshold value at which the corrective action must be triggered on or cleared from an interface, use the **link-degrade-monitor thresholds (set value | clear value)** statement.
- To configure the link degrade interval value, use the **link-degrade-monitor thresholds interval value** statement. The configured interval value determines the number of consecutive link degrade events that are considered before any corrective action is taken.
- To configure link degrade warning thresholds, use the **link-degrade-monitor thresholds (warning-set value | warning-clear value)** statement.
- To configure the link degrade action that is taken when the configured BER threshold level is reached, use the **link-degrade action media-based** statement.
- To configure the link degrade recovery options, use the **link-degrade recovery (auto interval value | manual)** statement.

You can view the link recovery status and the BER threshold values by using the **show interfaces interface-name** command.

VPNs

- **Layer 3 VPN (PTX1000)**—A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing Internet backbone.

In Junos OS, Layer 3 VPNs are based on RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*. This RFC defines a mechanism by which service providers can use their IP backbones to provide Layer 3 VPN services to their customers.

[See [Understanding Layer 3 VPNs](#)]

SEE ALSO

[Changes in Behavior and Syntax | 213](#)

[Known Behavior | 215](#)

[Known Issues | 217](#)

[Resolved Issues | 218](#)

[Documentation Updates | 220](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

[Product Compatibility | 225](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Interfaces and Chassis | 213](#)
- [Management | 214](#)

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 17.2R1 for the PTX Series.

Interfaces and Chassis

- **Value of `sysObjectID` now displays `jnxProductNamePTX1000 (PTX1000)`**—Starting in Junos OS Release 17.2R1 and later, the value of `sysObjectID` now displays as `jnxProductNamePTX1000` instead of `jnxProductPTX1000` (which is an incorrect value), as shown in the following example:

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNamePTX1000
```

The `sysObjectID` value is updated to `jnxProductNamePTX1000` to maintain synchronization across devices (or routers) belonging to the PTX Series.

- **Change in command outputs after a health check failure (PTX5000)**—Starting in Junos OS Release 17.2R1, when a health check fail for a PSM is detected on a PTX5000 router, until a system reboot or restart chassisd occurs, the following changes are displayed in the command outputs:
 - The output of the **show chassis environment pdu** displays the reason for the health check fail and the following information:

```
Health Check FAILED for PSM PSM_Number
```

```
PSM_Number is Present|Not OK
```

- The status of the PSM which failed the health check is set to offline and the output of **show chassis alarms** command displays the following existing alarm:

```
PDU slot PSM PSM_Number Not OK
```

After a system reboot or restart chassisd, the router checks the PSM register 0x1D bit-0:

- The output of the **show chassis environment pdu** displays the reason for the health check fail and the following information for the PSM:

```
PSM_Number is Present|Not OK
```

Management

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

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

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

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

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

- **Enhancement to the Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.2R1, the values displayed in the **oper-status** key-value field of data streamed through gRPC for the physical interfaces sensor have changed.

The following values are now displayed to indicate the operational status of an interface:

- operational status up—**UP**
 - operational status down—**DOWN**
 - operational status unknown—**UNKNOWN**
- **Enhancement to NPU memory sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.2R1, the path used to subscribe to telemetry data for network processing unit (NPU)

memory and NPU memory utilization through gRPC has changed. The new path is `/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/`

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

SEE ALSO

New and Changed Features 193
Known Behavior 215
Known Issues 217
Resolved Issues 218
Documentation Updates 220
Migration, Upgrade, and Downgrade Instructions 220
Product Compatibility 225

Known Behavior

IN THIS SECTION

- [Hardware | 216](#)
- [High Availability \(HA\) and Resiliency | 216](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 for PTX Series.

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

Hardware

- **Enhanced resiliency and system snapshot (PTX1000)**—The 2 x 64-GB SSDs installed in the PTX1000 support the **request vmhost snapshot** command, which creates a recovery snapshot of the currently running and active file system partitions, and the **request vmhost snapshot recovery** command, which recovers the primary disk from the snapshot content stored in the backup disk. In addition, the 64-GB SSDs support enhanced hardware resiliency through storage partitioning and redundancy.

Earlier versions of the PTX1000 have 2 x 32-GB M.2 SATA SSDs. PTX1000 routers with 32-GB SSDs do not support the **request vmhost snapshot** and **request vmhost snapshot recovery** commands, and do not support enhanced hardware resiliency. To determine the size of the SSDs installed in your device, issue the **show vmhost hardware** CLI command. The capacity of **Disk1** and **Disk2** is displayed in the output as **32.0 GB** if 32-GB SSDs are installed, and the capacity is displayed as **50.0 GB** if 64-GB SSDs are installed.

[See the [Junos OS Software Installation and Upgrade Guide](#).]

High Availability (HA) and Resiliency

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

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

SEE ALSO

[New and Changed Features | 193](#)

[Changes in Behavior and Syntax | 213](#)

[Known Issues | 217](#)

[Resolved Issues | 218](#)

[Documentation Updates | 220](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

[Product Compatibility | 225](#)

Known Issues

IN THIS SECTION

- General Routing | 217
- Platform and Infrastructure | 218

This section lists the known issues in hardware and software in Junos OS Release 17.2R1 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

- Sometimes qsf28+ can run into clock stretch and will be disabled with **i2c-accel sync access failed** error message. This enhancement PR will support qsf28+ clock stretch and reset mechanism for disabled qsf28+. [PR1181493](#)
- When configuring logical units to match the S-TAG (outer) of inbound frames for forwarding through L2VPN, PTX Series routers should match on S-TAG and ignore any C-Tag (inner) that might be present. Instead, it requires an exact match of S-TAG and C-TAG, therefore no traffic matches and no frames are forwarded through the L2VPN. [PR1252443](#)
- On rare occasions, upon reboot, the kernel cannot create sysfs entries for the SSD's in the system. This might result into the sytem going into Panic mode and hang. [PR1261068](#)
- If pio errors are encountered at a certain stage during ASIC initialization some data structures are not correctly updated. Later an attempt to recover from the error automatically might detect data structure inconsistencies and abort, resulting in generating a ukern core file. [PR1268253](#)
- Condition: Offline/Restart an FPC 'x' which is sending traffic to FPC 'y' * Following error messages are seen on destination FPC * It also results in corresponding alarm being set on destination FPC * Specific to Vale is the transient alarm that gets set when this condition occurs by CMERROR infra * The alarm clears off later on since the source FPC is being offlined 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](#)

- With the Routing protocols like LDP, mcast, configured on the router and if IPLC is present in the system, the show statistics commands issued from the CLI will timeout. Eg: set protocols mpls interface all set protocols ldp interface all set protocols <routing protocol> statistics or traffic-statistics" will timeout. eg: show ldp traffic-statistics eg: show multicast route extensive etc [PR1275234](#)

Platform and Infrastructure

- In scaled FIB setups, IS-IS Graceful-restart may Abort on the restarting node with T3-timer expiry log, because of hold-time expiry on IS-IS GR-helper peers. **May 20 01:22:55.992972 T3 Restart timer expired (graceful restart aborted)** This occurs because of the time taken by the routing protocol process (rpd) to learn routes from FIB on RPD-Startup, that were installed by previous-incarnation of the rpd. IS-IS does not get to initialize and send hellos because of this delay, causing holdtimer-expiry at helpers. {master}[edit] user@router# run show route forwarding-table summary | match user May 20 03:07:26 user: 801650 routes user: 403638 routes user: 6 routes. [PR1277933](#)

SEE ALSO

New and Changed Features 193
Changes in Behavior and Syntax 213
Known Behavior 215
Resolved Issues 218
Documentation Updates 220
Migration, Upgrade, and Downgrade Instructions 220
Product Compatibility 225

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R1 | 219](#)

This section lists the issues fixed in the Junos OS main release and the maintenance releases. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

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

Resolved Issues: 17.2R1

General Routing

- On PTX Series platform with Junos Telemetry Interface deployed, streaming telemetry stops being sent for message queue telemetry transport (mqtt) if memory reaches its maximum of 4G. [PR1238803](#)
- On PTX Series platform, add 'set' parameter (optional) to CLI command **request system software add**. It provides a way to install multiple software packages and software add-on packages at the same time. [PR1246675](#)
- On PTX Series platform, if running interfaces on QSFP28 PIC in 10G mode, some of the interfaces on the QSFP28 PIC might not come up after a system reboot or PIC restart. [PR1263413](#)
- This only affects PTX5000 and PTX3000 platforms with Third-generation FPCs. Software periodically monitors voltages on the FPCs to check if they are within the proper range. This change adjusts the expected values for voltages on certain power rails of the FPCs. In rare cases it is possible that a marginal FPC was operating inside the older limits but outside the new limits, in which case a new chassis alarm will be raised for that FPC. [PR1263675](#)

MPLS

- When the RPD daemon is terminating, the process of signaling the deletion of all RSVP LSPs might take so long that a watchdog timer is firing, resulting in generating an RPD core file. [PR1257367](#)

SEE ALSO

[New and Changed Features | 193](#)

[Changes in Behavior and Syntax | 213](#)

[Known Behavior | 215](#)

[Known Issues | 217](#)

[Documentation Updates | 220](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

[Product Compatibility | 225](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.2R1 documentation for PTX Series.

SEE ALSO

[New and Changed Features | 193](#)

[Changes in Behavior and Syntax | 213](#)

[Known Behavior | 215](#)

[Known Issues | 217](#)

[Resolved Issues | 218](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

[Product Compatibility | 225](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.2 | 220](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 223](#)
- [Upgrading Using Unified ISSU | 224](#)
- [Upgrading a Router with Redundant Routing Engines | 224](#)

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

Basic Procedure for Upgrading to Release 17.2

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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

To download and install Junos OS Release 17.2R2:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.

8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the router.

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

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-17.2R2.9.tgz
```

Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-17.2R2.9-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

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

NOTE: Most of the existing **request system** commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

Upgrading Using Unified ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [Understanding High Availability Features on Juniper Networks Routers](#).

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

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

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

SEE ALSO

New and Changed Features 193
Changes in Behavior and Syntax 213
Known Behavior 215
Known Issues 217
Resolved Issues 218
Documentation Updates 220
Product Compatibility 225

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 225](#)

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 193
Changes in Behavior and Syntax 213
Known Behavior 215
Known Issues 217
Resolved Issues 218
Documentation Updates 220
Migration, Upgrade, and Downgrade Instructions 220

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- New and Changed Features | 226
- Changes in Behavior and Syntax | 247
- Known Behavior | 249
- Known Issues | 250
- Resolved Issues | 254
- Documentation Updates | 256
- Migration, Upgrade, and Downgrade Instructions | 257
- Product Compatibility | 263

These release notes accompany Junos OS Release 17.2R1 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at <https://www.juniper.net/documentation/software/junos/>.

New and Changed Features

IN THIS SECTION

- Hardware | 227
- Authentication, Authorization, and Accounting (AAA) (RADIUS) | 228
- Class of Service (CoS) | 228
- Dynamic Host Configuration Protocol (DHCP) | 229
- EVPNs | 229
- Infrastructure | 231
- Interfaces and Chassis | 231
- IPv4 | 233
- IPv6 | 233

- Layer 2 Features | 233
- Layer 3 Features | 234
- Management | 234
- MPLS | 237
- Multicast | 238
- Network Management and Monitoring | 239
- Port Security | 240
- Routing Protocols | 241
- Security | 243
- Services Applications | 244
- Software Defined Networking (SDN) | 244
- Software Installation and Upgrade | 245
- Software Licensing | 245
- System Management | 245
- VLAN Infrastructure | 246

This section describes the new features for the QFX Series switches in Junos OS Release 17.2R1.

NOTE: The following QFX Series platforms are supported in Release 17.2R1: QFX5100, QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016.

Hardware

- **QFX5110-32Q**—The QFX5110 line of switches is Juniper Network’s versatile fixed-configuration solution for hybrid cloud deployments. The model QFX5110-32Q is a flexible configuration switch allowing either 32 ports of 40-Gigabit Ethernet quad small form-factor pluggable plus (QSFP+) or 20 ports of QSFP+ and 4 port of high-density 100-Gigabit Ethernet quad small form-factor pluggable solution (QSFP28). Each QSFP+ port can operate as a native 40-Gigabit Ethernet port, or as four independent 10-Gigabit ports when using breakout cables. The four QSFP28 ports are available either as access ports or as uplinks. The QFX5110-32Q provides full duplex throughput of 960 Gbps. The QFX5110-32Q has a 1 U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow. The model is available with either AC or DC power supplies.

- **QFX10000-60S-6Q Line Card (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.2R1, the QFX10000-60S-6Q line card support 1 Gbps speeds on the 10 Gigabit Ethernet SFP+ ports. See [QFX10000-60S-6Q Line Card](#).
- **QFX10K-12C-DWDM Coherent Line Card (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.2R1, QFX10008 and QFX10016 modular switch chassis support the QFX10K-12C-DWDM Coherent Line Card. The QFX10K-12C-DWDM Coherent Line Card provides up to 1.2 Tbps packet forwarding for cloud providers, service providers, and enterprises that need coherent dense wavelength-division multiplexing (DWDM) with MACsec security features. The six-port line card, with built-in optics, supports flexible rate modulation at 100 Gbps, 150 Gbps, and 200 Gbps speeds. A maximum of four QFX10K-12C-DWDM Coherent Line Cards are supported in either the QFX10008 switch chassis or the QFX10016 switch chassis. [See [QFX10K-12C-DWDM Coherent Line Card](#).]

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

- **Access control and authentication (QFX5100 switches)**—Starting in Junos OS Release 17.2R1, QFX5100 switches support controlling access to your network using 802.1X authentication and MAC RADIUS authentication. 802.1X authentication provides port-based network access control (PNAC) as defined in the IEEE 802.1X standard. QFX5100 switches support 802.1X features including guest VLAN, private VLAN, server fail fallback, dynamic changes to a user session, RADIUS accounting, and configuration of port-filtering attributes on the RADIUS server using vendor-specific attributes (VSAs). MAC RADIUS authentication is used to authenticate end devices independently of whether they are enabled for 802.1X authentication. You can permit end devices that are not 802.1X-enabled to access the LAN by configuring MAC RADIUS authentication on the switch interfaces to which the end devices are connected. You configure access control and authentication features at the the `[edit protocols dot1x]` hierarchy level. This feature was previously supported in an “X” release of Junos OS.

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

Class of Service (CoS)

- **Support for Class of Service on QFX5200 Switches**—Starting in Junos OS Release 17.2R1, the QFX5200 supports class-of-service (CoS). When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default class-of-service(CoS) settings or CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and any rewrite rules to re-mark packets.

[See [Traffic Management User Guide for the QFX Series](#)]

- **Support for FIP snooping and DCBX on QFX5200 Switches**—Starting in Junos OS Release 17.2R1, the QFX5200 supports both FIP snooping and DCBX. FIP snooping filters prevent an FCoE device from gaining unauthorized access to a Fibre Channel (FC) storage device or to another FCoE device. Data

Center Bridging Capability Exchange Protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX advertises the capabilities of applications on interfaces by exchanging application protocol information through application time-length-values (TLVs).

[See [Traffic Management User Guide for the QFX Series.](#)]

Dynamic Host Configuration Protocol (DHCP)

- **User-defined interface description for DHCP relay (QFX5100, QFX5110, and QFX5200 switches)**--Starting in Junos OS Release 17.2R1, you can define an interface description to be included in DHCP relay option 82 that is independent of the textual interface description configured at the **[edit interfaces interface-name]** hierarchy level.

[See [user-defined.](#)]

EVPNs

- **Support for IGMP snooping for EVPN-VXLAN in a multihomed environment (QFX10000 switches)**--Starting in Junos OS Release 17.2R1, QFX10000 switches support IGMP snooping with Ethernet EVPN (EVPN). This feature is useful in an EVPN-VXLAN environment with significant multicast traffic. IGMP snooping enables PE devices to send multicast traffic to CE devices only as needed. To configure IGMP snooping, include the **igmp-snooping (all | vlan-number)** set of statements at the **[edit protocols]** hierarchy level. You must also include the **proxy** statement in the IGMP snooping configuration. All multihomed interfaces must have the same configuration. The following new operational commands are also supported: **show evpn igmp snooping database extensive**, **show igmp snooping evpn database**, **show igmp snooping evpn membership**, and **show evpn multicast-snooping next-hops**.

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

- **Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network (QFX5100 switches)**--Starting in Junos OS Release 17.2R1, QFX5100 switches that function as Layer 2 VXLAN tunnel endpoints (VTEPs) can tunnel single- and double-tagged Q-in-Q packets through an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network. In addition to tunneling Q-in-Q packets, the ingress and egress VTEPs can perform the following Q-in-Q actions:
 - Delete, or pop, an outer service provider VLAN (S-VLAN) tag from an incoming packet.
 - Add, or push, an outer S-VLAN tag onto an outgoing packet.
 - Map a configured range of customer VLAN (C-VLAN) IDs to an S-VLAN.

NOTE: The QFX5100 switch does not support the pop and push actions with a configured range of VLANs.

The ingress and egress VTEPs support the tunneling of Q-in-Q packets and the Q-in-Q actions in the context of specific traffic patterns.

To enable the tunneling of the Q-in-Q packets on the VTEPs, you must configure a flexible VLAN tagging interface, which can transmit 802.1Q VLAN single- and double-tagged packets, on ingress and egress VTEPs. It is also important to configure the interface to retain the inner C-VLAN tag while a packet is tunneled.

[See [Examples: Configuring QFX5100 Switches to Tunnel Q-in-Q Traffic Through an EVPN-VXLAN Overlay Network](#).]

- **EVPN-VXLAN support of Virtual Chassis and Virtual Chassis Fabric (QFX5100, QFX5100 Virtual Chassis, and Virtual Chassis Fabric)**—Ethernet VPN (EVPN) supports multihoming active-active mode, which enables a host to be connected to two leaf devices through a Layer 2 link aggregation group (LAG) interface. In previous Junos OS releases, the two leaf devices had to be QFX5100 standalone switches. Starting in Junos OS Release 17.2R1, the two leaf devices can be QFX5100 standalone switches, QFX5100 switches configured as a Virtual Chassis (VC), QFX5100 switches configured as a Virtual Chassis Fabric (VCF), or a mix of these options.

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

[See [EVPN-VXLAN Support of Virtual Chassis and Virtual Chassis Fabric](#).]

- **EVPN pure type-5 route support (QFX10000 switches)**—Starting in Junos OS Release 17.2R1, you can configure pure type-5 routing in an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which is used to carry the MAC address of the sending switch and to provide next hop reachability for the prefix. This router MAC extended community provides next-hop reachability without requiring an overlay next hop or supporting type-2 route. To configure pure type-5 routing, include the **ip-prefix-routes advertise direct-nexthop** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. Pure type-5 routing was previously introduced in Junos OS Release 15.1x53-D60.

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

Infrastructure

- **Secure Boot (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement Secure Boot.

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

Interfaces and Chassis

- **Resilient hashing support for link aggregation groups and equal cost multipath routes (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, resilient hashing is now supported by link aggregation groups (LAGs) and equal cost multipath (ECMP) sets.

Resilient hashing enhances LAGs by minimizing destination remapping when a new member is added to or deleted from the LAG.

Resilient hashing works in conjunction with the default static hashing algorithm. It distributes traffic across all members of a LAG by tracking the flow's LAG member utilization. When a flow is affected by a LAG member change, the packet forwarding engine (PFE) rebalances the flow by reprogramming the flow set table. Destination paths are remapped when a new member is added to or existing members are deleted from a LAG.

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

[See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups.](#)]

- **Multichassis link aggregation groups (MC-LAG) (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG is an MC-LAG client that has one or more physical links in a LAG. This client does not need to detect the MC-LAG. On the other side of the MC-LAG are two switches. Each of these switches has one or more physical links connected to a single client. The switches coordinate with each other to ensure that data traffic is forwarded properly.

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

[See [Multichassis Link Aggregation Features, Terms, and Best Practices.](#)]

- **Channelizing 100-Gigabit Ethernet QSFP28 interfaces (QFX5200 switch)**—This feature enables you to channelize the 100-Gigabit Ethernet interfaces to two independent 50-Gigabit Ethernet or to four independent 25-Gigabit Ethernet interfaces. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated

40-Gigabit Ethernet interfaces or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables.

To channelize the ports, manually configure the port speed using the **set chassis fpc slot-number port port-number channel-speed speed** command, where the speed can be set to 10G, 25G, or 50G. The ports do not support autochannelization.

NOTE: If a 100G transceiver is connected to the switch, channelize the port only to 25G or 50G. If a 40G transceiver is connected, channelize the port only to 10G. Note that there is no commit check for these options.

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

[See [Channelizing Interfaces on QFX5200 Switches](#).]

- **IRB interface in a PVLAN (QFX5110 switches)**—Starting with Junos OS Release 17.2R1, you can configure an integrated routing and bridging (IRB) interface in a private VLAN (PVLAN) on QFX5110 switches so that devices within community VLANs and isolated VLANs can communicate with each other and with devices outside the PVLAN at Layer 3 without requiring you to install a router. This feature was previously supported in an "X" release of Junos OS.

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

IPv4

- **Generic routing encapsulation (GRE) support (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, you can use GRE tunneling services on QFX5110 switches to encapsulate any network layer protocol over an IP network. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first adds a GRE header and then adds an outer IP header that is used to route the packet. When it receives the packet, a switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect noncontiguous networks and to provide options for networks that contain protocols with limited hop counts.

IPv6

- **IPv6 feature support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can configure the Neighbor Discovery Protocol, the Virtual Router Redundancy Protocol (VRRP) for IPv6, and Protocol Independent Multicast (PIM) for IPv6. You can also configure BGP and IS-IS for IPv6 as well as OSPFv3. Additionally, unicast IPv6 is supported for virtual router instances. DHCPv6 is also supported. IPv6 feature support for QFX5110 and QFX5200 switches was previously introduced in "X" releases of Junos OS.

[See [Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery](#) and [Verifying and Managing DHCPv6 Local Server Configuration](#).]

Layer 2 Features

- **Layer 2 Features (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, the following features are supported:
 - VLAN support—Enables you to divide one physical broadcast domain into multiple virtual domains.
 - LLDP—Enables a switch to advertise its identity and capabilities on a LAN as well as receive information about other network devices.
 - Q-in-Q tunneling support—Allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.
 - Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) provide Layer 2 loop prevention.

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

- **Q-in-Q tunneling support (QFX5200 switches)**—Starting in Junos OS Release 17.2R1, QFX5200 switches support Q-in-Q tunneling, which enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag. This feature was previously supported in an "X" release of Junos OS.

[See [Understanding Q-in-Q Tunneling](#).]

Layer 3 Features

- **Support for hierarchical ECMP groups (QFX5200 switches)**—Starting in Junos OS Release 17.2R1, hierarchical equal-cost multipath (ECMP) groups are enabled by default at system start. Hierarchical ECMP provides for two-level route resolution automatically through the Packet Forwarding Engine. Two-level route resolution through ECMP groups enhances load-balancing of traffic. This feature was previously introduced Junos OS Release 15.1X53-D30.

[See [Overview of Hierarchical ECMP Groups](#).]

- **Support for 64 next-hop gateways for ECMP (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, you can configure as many as 64 equal-cost-multipath (ECMP) next hops for RSVP and LDP LSPs or external BGP peers. The following Layer 3 protocols are supported as ECMP gateways for both IPv4 and IPv6 traffic: OSPF, ISIS, EBGp, and IBGP (resolving over IGP routes). Include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level. This feature was previously introduced on QFX5110 switches in Junos OS Release 15.1X53-D210.

[See [maximum-ecmp](#).]

- **Support to disable hierarchical ECMP (QFX5200 switches)**—Starting with Junos OS Release 17.2R1, you can disable hierarchical equal-cost multipath (ECMP) groups at system start time. Hierarchical ECMP is enabled by default. Disabling this feature effectively increases the number of ECMP groups. Include the **no-hierarchical-ecmp** statement at the **[edit forwarding-options]** hierarchical level. Disabling hierarchical ECMP causes the Packet Forwarding Engine to restart. To reenab le hierarchical ECMP, issue the following command: **delete forwarding-options no-hierarchical-ecmp**. This feature was previously introduced in Junos OS Release 15.1X53-D210.

[See [no-hierarchical-ecmp](#).]

Management

- **Support for device family and release in Junos OS YANG modules (QFX Series)**—Starting in Junos OS Release 17.2, Junos OS YANG modules are specific to a device family, and each module's namespace includes the module name, device family, and Junos OS release string. Furthermore, each **juniper-command**

module uses its own unique module name as the module's prefix. Device families include **junos**, **junos-es**, **junos-ex**, and **junos-qfx**.

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

- **Support for the Junos Telemetry Interface (QFX10000 switches)**—Starting with Junos OS Release 17.2R1, the Junos Telemetry Interface is supported on QFX10000 switches. Both UDP and gRPC streaming of statistics are supported. Junos Telemetry Interface enables you to provision sensors to export telemetry data for various network elements without involving polling.

The following sensors are supported on QFX10000 switches:

- Logical interfaces (UDP and gRPC streaming)
- Physical interfaces (UDP and gRPC streaming)
- Firewall filters, including traffic-class counters (UDP and gRPC streaming)
- LSP statistics (UDP and gRPC streaming)
- LSP events and properties (gRPC streaming)
- Optical interfaces (UDP and gRPC streaming)
- Network processing unit (NPU) memory (UDP and gRPC streaming)
- NPU memory utilization (UDP and gRPC streaming)
- CPU memory (UDP and gRPC streaming)
- Chassis components (gRPC streaming only)
- RSVP interface events (gRPC streaming only)
- BGP peers (gRPC streaming only)
- Memory utilization for routing protocol tasks (gRPC streaming only)
- Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol (gRPC streaming only)
- Ethernet interfaces enabled configured with the Link Layer Discovery Protocol (gRPC streaming only)
- Network Discovery Protocol table state (gRPC streaming only)
- Address Resolution Protocol table state (gRPC streaming only)

To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. To provision sensors to stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters for a specified list of OpenConfig command paths. Because QFX10000 switches run a version Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models.

The LSP events and properties sensor is supported in Junos OS Release 17.2R1 for the first time. You can export statistics for ingress point-to-point LSPs, point-to-multipoint LSPs, bypass LSPs, and dynamically created LSPs. To export data through gRPC, use the `/mpls/lsp/` or `/mpls/signal-protocols/` set of OpenConfig subscription paths. To enable the LSP events and properties sensor, you must also include the **lsp-telemetry** statement at the **[edit routing-options]** hierarchy level.

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

- **Support for the Junos Telemetry Interface (QFX5200 switches)**—Starting with Junos OS Release 17.2R1, you can provision sensors through the Junos Telemetry Interface to export telemetry data for various network elements without involving polling. On QFX5200 switches, only gRPC streaming of statistics is supported. UDP streaming is not supported.

The following sensors are supported:

- Chassis components
- Aggregated Ethernet interfaces configured with the Link Aggregation Control Protocol
- Ethernet interfaces enabled configured with the Link Layer Discovery Protocol
- BGP peers
- RSVP interface events
- Memory utilization for routing protocol tasks
- Address Resolution Protocol table state
- Network Discovery Protocol table state

To provision sensors to stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters for a specified list of OpenConfig commands paths. You must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models.

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

MPLS

- **TE++ Dynamic bandwidth management using container LSPs (QFX5100)**—Starting with Junos OS Release 17.2R1, a new type of label-switched path (LSP), called a container LSP, is introduced to enable load balancing across multiple point-to-point member LSPs between the same ingress and egress routers. Each member LSP takes a different path to the same destination and can be routed along a different interior gateway protocol (IGP) cost path. Based on the configuration and aggregate traffic, a container LSP provides support for dynamic bandwidth management by enabling the ingress router to dynamically add and remove member LSPs through a process called LSP splitting and LSP merging, respectively. Member LSPs can also be re-optimized with different bandwidth values in a make-before-break way. The feature was previously supported in a "X" release of Junos OS.

[See [Dynamic Bandwidth Management Using Container LSP Overview](#).]

- **Entropy labels for LSPs (QFX10000 switches)**—Starting with Junos OS Release 17.2R1, you can configure entropy labels for label-switched paths (LSPs). An entropy label is a special load-balancing label that enhances the ability of the switch to load-balance traffic across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs). The entropy label allows the switch to efficiently load-balance traffic using just the label stack rather than deep packet inspection (DPI). To configure entropy labels, include the `entropy-label` statement at the `[edit protocols mpls labeled-switched-path labeled-switched-path-name]` hierarchy level.

[See [Understanding Entropy Label for BGP Labeled Unicast LSPs](#).]

[See [Automatic Bandwidth Allocation for LSPs](#).]

- **Support for a label stack for BGP label unicast for MPLS advertisements (QFX10000 switches)**—Starting with Junos OS 17.2R1, QFX10000 switches implement RFC 3701, which supports a stack of labels in BGP label unicast for both IPv4 and IPv6 traffic. Previously, only one label per prefix was supported in the BGP unicast label. You can now specify to include up to five labels per prefix in the BGP labeled unicast updates. This feature enables the use of the BGP label unicast stack to program a stack of labels to control packet forwarding in a network configured with hierarchical MPLS label-switched paths. To configure as many as five labels to advertise through MPLS, include the `maximum-labels number` statement at the `[edit interfaces interface-name unit logical-unit-number family mpls]` hierarchy level. The `show route receive-protocol bgp neighbor-address detail` and `show route advertising-protocol neighbor-address detail` operational commands are enhanced to display multiple labels for one prefix in the Labels field.

[See [Configuring the Maximum Number of MPLS Labels](#).]

- **MPLS support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, MPLS is supported on the QFX5110 and QFX5200 switches. MPLS supports both label edge routers (LER) and label switch routers (LSR) and provides the following capabilities:
 - Support for both MPLS major protocols, LDP and RSVP
 - IS-IS interior gateway protocol (IGP) traffic engineering
 - Class of service (CoS)

- Object access method, including ping, traceroute, and Bidirectional Forwarding Detection (BFD)
- Fast reroute (FRR), a component of MPLS local protection (Both one-to-one local protection and many-to-one local protection are supported.)
- Loop-free alternate (LFA)
- 6PE devices
- Layer 3 VPNs for both IPv4 and IPv6
- LDP tunneling over RSVP

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

[See [MPLS Overview for Switches](#).]

- **Support for equal cost multipath (ECMP) routing on label-switching routers (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can configure ECMP on MPLS label-switched routers (LSRs). ECMP is a layer 3 mechanism for load-balancing traffic to a destination over multiple equal-cost next hops. When a link goes down, ECMP uses fast reroute protection to shift packet forwarding to use operational links, thereby decreasing packet loss. This feature was previously supported in an “X” release of Junos OS.

[See [Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing](#).]

- **Ethernet over MPLS (Layer 2 circuit) support (QFX5100 Virtual Chassis and Virtual Chassis Fabric)**—Starting in Junos OS Release 17.2R1, a QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF) supports Ethernet over MPLS (Layer 2 circuit). The Virtual Chassis or VCF can act as a provider edge switch on which you configure MPLS and LDP for the interfaces that will carry the Layer 2 circuit traffic. The Layer 2 circuit can be port-based (pseudo-wire) or VLAN-based. These features were previously supported for a QFX5100 Virtual Chassis or VCF in an “X” release of Junos OS.

[See [Understanding Ethernet-over-MPLS \(L2 Circuit\)](#) and [Configuring Ethernet over MPLS \(L2 Circuit\)](#).]

Multicast

- **Layer 3 multicast support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, IGMP, including versions 1,2, and 3, IGMP snooping, PIM sparse mode and PIM source-specific multicast are supported. You can also configure IGMP, IGMP snooping and PIM in virtual router instances. Multicast Source Discovery Protocol (MSDP) is also supported. Configure IGMP at the **[edit protocols igmp]** hierarchy level. Configure IGMP snooping at **[edit protocols igmp-snooping]** hierarchy level. Configure PIM at the **[edit protocols pim]** hierarchy level. Configure MSDP at the **[edit protocols msdp]** hierarchy level. Layer 3 multicast support was previously introduced in “X” releases of Junos OS.

[See [Multicast Overview](#).]

- **Support for static multicast route leaking for VRF and virtual-router instances (QFX5100 and EX4300 switches)**—Starting in Junos OS Release 17.2R1, you can configure your switch to share IPv4 multicast routes among different virtual routing and forwarding (VRF) instances or different virtual-router instances.

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

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

Network Management and Monitoring

- **sFlow Enhancements (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 17.2R1, sFlow IPv4 and IPv6 packets support extended router information, including the IP address of the next-hop router, the outgoing VLAN ID, the source IP address prefix length, and the destination IP address prefix length. This information is collected only if BGP is configured on the switch.

In addition, a configuration statement was introduced that allows the sFlow sampling rate to stay within the maximum sampling rate of 1 out of 64,000 packets. Packet-based sampling is implemented in the hardware, so all of the interfaces can be monitored with very little overhead. However, if traffic levels are unusually high, the hardware generates more samples than it can handle. The extra samples are dropped by the software rate-limiting algorithm and can cause inaccurate results. You can include the **disable-sw-rate-limiter** statement at the **[edit protocols sFlow]** hierarchy to disable the software, allowing the hardware sampling rate to stay within the maximum sampling rate for sFlow.

[See [Understanding How to Use sFlow Technology for Network Monitoring on a Switch.](#)]

- **sFlow technology support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, the QFX5110 and QFX5200 switches support sFlow technology. sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring randomly samples network packets and sends the samples to a monitoring station called a collector. You can configure sFlow monitoring on the switch to continuously monitor traffic at wire speed on all interfaces simultaneously. sFlow monitoring also collects samples of network packets, providing you with visibility into network traffic information. You configure sFlow monitoring at the **[edit protocols sflow]** hierarchy level. sFlow operational commands include **show sflow** and **clear sflow collector statistics**. This feature was previously supported in an "X" release of Junos OS

[See [Understanding How to Use sFlow Technology for Network Monitoring on a Switch.](#)]

- **Port mirroring (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can use port mirroring on QFX5110 and QFX5200 switches to copy packets entering or exiting a port or entering a VLAN and send the copies to a local interface for local monitoring or to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring

compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. This feature was previously supported in an "X" release of Junos OS.

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

- **SNMP support for monitoring tunnel statistics (QFX Series)**—Starting in Junos OS Release 17.2R1, SNMP MIB `jnxTunnelStat` supports monitoring of tunnel statistics for IPv4 over IPv6 tunnels. This is a new enterprise-specific MIB, Tunnel Stats MIB, that currently displays three counters: tunnel count in `rp`, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine. This MIB can be extended to support other tunnel statistics. The MIB is defined in `jnx-tunnel-stats.txt`. This MIB is attached to `jnxMibs`.

Port Security

- **Media Access Control Security (MACsec) support (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 17.2R1, MACsec is supported on all six interfaces of the QFX10K-12C-DWDM line card when it is installed in a QFX10008 or QFX10016 switch. MACsec is an 802.1AE IEEE industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is capable of identifying and preventing most security threats, and can be used in combination with other security protocols to provide end-to-end network security. MACsec can be enabled only on domestic versions of Junos OS software.

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

- **Access security support (QFX5110 switches)**—Starting in Junos OS Release 17.2R1, the following access security features are supported on QFX5110 switches:
 - **DHCP snooping**—DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information, which it uses to build and maintain a database of valid IP-address-to-MAC-address (IP-MAC) bindings called the DHCP snooping database. Clients on untrusted ports are only allowed to access the network if they can be validated against the database.
 - **DHCPv6 snooping**—DHCP snooping for DHCPv6.
 - **DHCP option 82**—You can use DHCP option 82, also known as the *DHCP relay agent information* option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
 - **DHCPv6 option 37**—Option 37 is the DHCPv6 equivalent of the remote ID suboption of DHCP option 82. It is used to insert information about the network location of the remote host into DHCPv6 packets.
 - **Dynamic ARP inspection (DAI)**—DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as *ARP poisoning* or *ARP cache poisoning*). ARP requests

and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons.

- IPv6 neighbor discovery (ND) inspection—IPv6 ND inspection mitigates attacks based on the Neighbor Discovery Protocol by inspecting neighbor discovery messages and verifying them against the DHCPv6 snooping table.
- MAC limiting—You can configure a MAC limit per interface and per VLAN, and set an action to take on the next packet the interface or VLAN receives after the limit is reached.
- MAC move limiting—You can configure MAC move limiting to track MAC address movements on the switch, so that if a MAC address changes more than the configured number of times within one second, the changes to MAC addresses are dropped, logged, or ignored, or the interface is shut down.
- Persistent MAC learning—Persistent (also called *sticky*) MAC addresses help restrict access to an access port by identifying the MAC addresses of workstations that are allowed access to a given port. Secure access to these workstations is retained even if the switch is restarted.

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

Routing Protocols

- **Support for BGP Monitoring Protocol (BMP) Version 3 (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, you can configure BMP, which sends BGP route information from the switch to a monitoring application, or station, on a separate device. To deploy BMP in your network, you need to configure BMP on each switch and at least one BMP monitoring station. Only version 3 is supported. To configure BMP, include the **bmp** set of statements at the **[edit routing-options]** hierarchy level. To configure a BMP monitoring station, include the **station-address ip-address** and the **station-port number** statements at the **[edit routing-options bmp]** hierarchy level.

[See [Configuring BGP Monitoring Protocol Version 3](#).]

- **Support for segment routing for IS-IS (QFX5100 switches and QFX10000 switches)**—Starting with Junos OS Release 17.2R1, you can advertise MPLS labels through IS-IS to support segment routing. IS-IS advertises a set of segments, which enables an ingress device to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the path to take. Two types of segments are supported: node and adjacency. A node segment represents a shortest-path link to a node. An adjacency segment represents a specific adjacency to a node. To enable segment routing, include the **source-packet-routing** statement at the **[edit protocols isis]** hierarchy level. By default, segment routing is enabled on all IS-IS levels. To disable advertising of the adjacency segment for a specified interface, include the **no-advertise-adjacency-segment** statement. You can also specify an interval for maintaining adjacency segments by including the **adjacency-segment hold-time milliseconds** statement.

To enable node segments, include the **node-segment** statement at the **[edit protocols isis source-packet-routing]** hierarchy level. You have two options for advertising a range of indices for IPv4

or IPv6 addresses. Use the **index-range** statement to specify a dynamic label range managed by MPLS. To specify a specific block of indices, also known as a segment routing global block, include the **start-label <number> index-range <number>** statements at the **[edit protocols isis source-packet-routing srgb]** hierarchy level. This configuration enables MPLS to reserve the specified label range.

Segment routing in IS-IS also supports provisioning prefix segment indices (SIDs) and anycast SIDs for both IPv4 and IPv6 prefixes. These SIDs are provisioned through a routing policy for each prefix. Include the **then prefix-segment index number** statement at the **[edit policy options policy-statement policy-name]** hierarchy level. You can also enable IPG shortcuts for prefix segment routes. Include the **shortcuts** statement at the **[edit protocols isis traffic-engineering family (inet-mpls | inet6-mpls)]** hierarchy level.

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

- **Support for segment routing for OSPF (QFX5100 switches and QFX10000 switches)**—Starting with Junos OS Release 17.2R1, you can advertise MPLS labels through OSPF to support segment routing. Only IPv4 is supported. OSPFv3 is not supported. OSPF advertises a set of segments, which enables an ingress device to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the path to take. Two types of segments are supported: node and adjacency. A node segment represents a shortest-path link to a node. An adjacency segment represents a specific adjacency to a node. To enable segment routing, include the **source-packet-routing** statement at the **[edit protocols ospf]** hierarchy level. By default, segment routing is enabled for all OSPF areas. To disable for a specific area, include the **no-source-packet-routing** statement at the **[edit protocols ospf area area-id]** hierarchy level. To enable node segments, include the **node-segment** statement. You can specify a range for IPv4 addresses to advertise, which MPLS manages dynamically. To disable advertising of the adjacency segment for a specified interface, include the **no-advertise-adjacency-segment** statement.

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

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

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

Security

- **Support for filter-based decapsulation over an IP-IP interface (QFX10000 switches)**—Starting in Junos OS Release 17.2R1, you can use a firewall filter over an IP-IP interface to de-encapsulate traffic on the switch, without the need to create any tunnel interfaces. IP-in-IP packets are special IP tunneling packets with no GRE header. With this feature, you can define a filter with filtering terms to classify packets based on packet fields such as destination IP address and protocol type. This provides significant benefits in terms of scalability, performance, and flexibility.

[See [Configuring a Firewall Filter to De-Encapsulate IP-in-IP Traffic.](#)]

- **Policing support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can use policing (or rate-limiting) to apply limits to traffic flow and to set consequences for packets that exceed those limits. A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service. This feature was previously supported in an “X” release of Junos OS.

[See [Overview of Policers.](#)]

- **Storm control support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can monitor traffic levels and take a specified action when a defined traffic level (called the storm control level) is exceeded, preventing packets from proliferating and degrading service. You can configure the switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs. This feature was previously supported in an “X” release of Junos OS.

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

- **Firewall filters support (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can provide rules that define whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces. This feature was previously supported in an “X” release of Junos OS.

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

- **Generic routing encapsulation (GRE) support (QFX5100 and QFX5200 switches)**—You can use GRE tunneling services to encapsulate any network layer protocol over an IP network. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first adds a GRE header and then adds an outer IP header that is used to route the packet. When it receives the packet, a switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect non-contiguous networks and to provide options for networks that contain protocols with limited hop counts. This feature was previously supported in an “X” release of Junos OS.

[See [Configuring a Firewall Filter to De-Encapsulate GRE Traffic.](#)]

Services Applications

- **Support for IPFIX templates for flow aggregation (QFX10002 switches)**—Starting with Junos OS Release 17.2R1, you can define a flow record template for unicast IPv4 and IPv6 traffic in IP Flow Information Export (IPFIX) format. Templates are transmitted to the collector periodically. To define an IPFIX template, include the **version-ipfix template *template-name*** set of statements at the **[edit services flow-monitoring]** hierarchy level.

You must also perform the following configuration:

- Sampling instance at the **[edit forwarding-options]** hierarchy level.
- Associate the sampling instance with the FPC at the **[edit chassis]** hierarchy level and with a template configured at the **[edit services flow-monitoring]** hierarchy level.
- Firewall filter for the family of traffic to be sampled at the **[edit firewall]** hierarchy level.

[See [Configuring Flow Aggregation to Use IPFIX Flow Templates.](#)]

Software Defined Networking (SDN)

- **OVSDB-VXLAN support with Contrail (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, the Open vSwitch Database (OVSDB) management protocol provides a means through which a Contrail controller can communicate with QFX5110 and QFX5200 switches to provision them as Layer 2 VXLAN gateways. In an environment in which Contrail Release 2.22 or later is deployed, a Contrail controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and vice versa.

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

[See [Understanding the OVSDB Protocol Running on Juniper Networks Devices.](#)]

- **Layer 2 VXLAN gateway (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, you can implement a QFX5110 or QFX5200 switch as a Layer 2 Virtual Extensible LAN (VXLAN) gateway. VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. You can use VXLAN tunnels to enable migration of virtual machines (VMs) between servers that exist in separate Layer 2 domains by tunneling the traffic through Layer 3 networks. This functionality allows you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

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

[See [Understanding VXLANs.](#)]

- **BFD in a VMware NSX environment with OVSDB and VXLAN (QFX5100 switches, QFX5100 Virtual Chassis)**

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

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

By exchanging BFD control messages with replicators at regular intervals, the hardware VTEP can monitor the replicators to ensure that they are functioning and reachable.

Software Installation and Upgrade

- **Support for FreeBSD 10 kernel for Junos OS (QFX5200 and QFX5110 switches)**—Starting with Junos OS Release 17.2R1, FreeBSD 10 is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. Because installation restructures the file system, logs and configurations are lost unless precautions are taken. Now there are Junos OS and OAM volumes, that provide the ability to boot from the OAM volume upon failures. Some system commands display different output and a few others are deprecated.

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

[See [Understanding Junos OS with Upgraded FreeBSD.](#)]

Software Licensing

- **Integrated software feature licenses (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, the standard QFX Series premium feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB) software license and the standard QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Multi-protocol Label Switching (MPLS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB) license are supported.

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

System Management

- **Support for Precision Time Protocol (PTP) transparent clock (QFX5100 and QFX5110 switches)**—Starting in Junos OS Release 17.2R1, PTP synchronizes clocks throughout a packet-switched network. With a

transparent clock, the PTP packets are updated with residence time as the packets pass through the switch. There is no master/slave designation. End-to-end transparent clocks are supported. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that the timestamps are sent in one packet. In a two-step process, estimated timestamps are sent in one packet, and additional packets contain updated timestamps. In addition, User UDP over IPv4 and IPv6 and unicast and multicast transparent clock are supported.

You can configure the transparent clock at the **[edit protocols ptp]** hierarchy.

[See [Understanding Transparent Clocks in Precision Time Protocol](#).]

- **Zero Touch Provisioning (QFX5100, QFX5110, and QFX5200 switches)**—Starting with Junos OS Release 17.2R1, Zero Touch Provisioning allows you to provision new Juniper Networks switches in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version.

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

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

VLAN Infrastructure

- **Double VLAN tags on Layer 3 subinterfaces (QFX10000 switches)**—Starting in Junos OS Release 17.2R1, you can configure double VLAN tags on Layer 3 subinterfaces (also called "Layer 3 logical interfaces") on QFX10000 switches. Layer 3 double-tagged logical interfaces support **inet**, **inet6**, and **mpls** families.

Support for double-tagging VLANs on Layer 3 logical interfaces includes:

- Configuration of an IPv4, an IPv6, or an **mpls** family on the logical interface
- Configuration over an aggregated Ethernet interface
- Configuration of multiple logical interfaces on a single physical interface

[See [Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces](#).]

SEE ALSO

[Changes in Behavior and Syntax | 247](#)

[Known Behavior | 249](#)

[Known Issues | 250](#)

[Resolved Issues | 254](#)[Documentation Updates | 256](#)[Migration, Upgrade, and Downgrade Instructions | 257](#)[Product Compatibility | 263](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 248](#)
- [Management | 248](#)

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 17.2R1 for the QFX Series.

Class of Service (CoS)

- Following options are hidden under Traffic control profiles for QFX10002 and QFX10008 products:
 - > delay-buffer-rate
 - > excess-rate
 - > excess-rate-high
 - > excess-rate-low
 - > excess-rate-medium-high
 - > excess-rate-medium-low

Shaping-rate option is hidden on QFX10008 but not on QFX10002 as shaping rate configurations are used in QFX10002 satellite solution setup. [PR1261988](#)

Management

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

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

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

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

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

- **Enhancement to the Junos Telemetry Interface (QFX10000 and QFX5200 switches)**—Starting in Junos OS Release 17.2R1, the values displayed in the **oper-status** key-value field of data streamed through gRPC for the physical interfaces sensor have changed.

The following values are now displayed to indicate the operational status of an interface:

- operational status up—**UP**
- operational status down—**DOWN**

- operational status unknown—**UNKNOWN**
- **Enhancement to NPU memory sensors for Junos Telemetry Interface (QFX10000 switches)**—Starting with Junos OS Release 17.2R1, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is `/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/`
[See [Guidelines for gRPC Sensors.](#)]

SEE ALSO

New and Changed Features	 226
Known Behavior	 249
Known Issues	 250
Resolved Issues	 254
Documentation Updates	 256
Migration, Upgrade, and Downgrade Instructions	 257
Product Compatibility	 263

Known Behavior

IN THIS SECTION

- [High Availability \(HA\) and Resiliency](#) | 250
- [Virtual Chassis](#) | 250

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.2R1 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.

High Availability (HA) and Resiliency

- **On QFX5100 switches, residual and baseline statistics loss from unified ISSU**—Using unified ISSU to upgrade to Junos OS Release 17.2R1 or later will result in a loss of residual and baseline statistics for interfaces, interface set specific statistics, and BBE subscriber service statistics because of an update to the statistics database.

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

Virtual Chassis

- If a QFX5100 switch running Junos OS Release 17.2R1 or later is in the same Virtual Chassis or Virtual Chassis Fabric (VCF) as a Juniper Networks device that does not support Virtual Extensible LAN (VXLAN), for example, an EX4300 switch, the Junos OS CLI of the EX4300 switch supersedes the Junos OS CLI of the QFX5100. As a result, the `vxlan` configuration statement at the `[edit vlans vlan-name]` hierarchy level does not appear. [PR1176054](#)

SEE ALSO

New and Changed Features 226
Changes in Behavior and Syntax 247
Known Issues 250
Resolved Issues 254
Documentation Updates 256
Migration, Upgrade, and Downgrade Instructions 257
Product Compatibility 263

Known Issues

IN THIS SECTION

- [General Routing | 251](#)
- [High Availability \(HA\) and Resiliency | 252](#)
- [Infrastructure | 252](#)
- [Interfaces and Chassis | 252](#)
- [IPsec | 252](#)

- Layer 2 Features | 252
- MPLS | 252
- Open vSwitch Database Management Protocol (OVSDb) | 253
- OpenFlow | 253
- Platform and Infrastructure | 253
- Routing Protocols | 254

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.2R1.

General Routing

- When a QFX5110 switch is acting as a DHCP relay agent, DHCP Reply packets might not be relayed back to the client by the relay agent if the Reply packets enter on a GRE tunneling interface. [PR1198982](#)
- 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 QFX10000 switches, sFlow monitoring technology output might display a negative number of samples after a long run. As a workaround, issue the **clear sflow collector** command to show or reset the count. [PR1244080](#)
- On the 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](#)
- Currently maximum 64 LAG members under a single aggregated Ethernet(ae) bundle is supported for QFX series platform. [PR1259515](#)
- If and when we release an upgrade to the existing line card FPGAs, then that upgrade, released in future, will need to be executed out of RE0 - it will not work out of RE1. When we release the upgrade, it will have to be release noted as such. At this time no action is required. [PR1276300](#)

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an QFX4300 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](#)

Infrastructure

- When there is a high route churn or when there is high rate of route updates being pushed to kernel, show interface command might show delay or not show all statistics due to route updates being prioritized over statistics messages. [PR1250328](#)

Interfaces and Chassis

- With MAC, ARP inside IFA block, DCD control path was taking different route, changing routing instance to default. [PR1239191](#)
- Deactivation followed by activation of both AE and MC-AE interfaces, blocks the flow of multicast traffic after that. [PR1257586](#)
- When optics TCA alarm is raised, we will see OT link flap. However this does not result in any traffic loss nor loss of signal is reported. [PR1279351](#)
- If the link between a QFX5200 switch interface and a NIC on a third-party device remains up irrespective of whether forward error correction (FEC) is enabled or disabled on the switch, then match the FEC mode (by setting FEC on or off) on the QFX5200 with that on the third-party device. [PR1246654](#)

IPsec

- On some QFX-Series platforms including QFX5110, QFX5200 and QFX10k platforms, the OSPFv3 authentication using IPsec SA does not work, which might cause the OSPFv3 neighbors not to be established. [PR1301428](#)

Layer 2 Features

- On QFX5100 VC interfaces on which the flexible-vlan-tagging statement is specified, STP, RSTP, MSTP, and VSTP are not supported. [PR1075230](#)

MPLS

- In Junos OS Release 17.2R1: Next-hops of static LSPs configured for "stitch" will resolve over the better preference resolving route in inet.0 or inet.3. In Junos OS 17.2R1 and later releases: Next-hops of static

LSPs configured for "stitch" will resolve over routes in inet.3 by default. To resolve the next-hop of a static LSP over routes in inet(6).0 as well, the following config needs to be enabled: `routing-options { resolution { rib mpls.0 { inet-resolution-ribs [inet.3 inet.0]; inet6-resolution-ribs [inet6.3 inet6.0]; } } }` Enabling this config will cause, the resolution to happen over the more preferred resolving route (lowest protocol preference) among all the longest-matching-prefix routes in both inet(6).0 and inet(6).3.

[PR1259238](#)

- On the QFX5110-32Q, while doing the 5000 tunnel scale test in 17.2R1 the below mentioned steps are followed, it will lead to dcpfe core. 1)Scale for 5000 Ingress sessions 2)delete the configurations 3)Scale for 5000 Ingress/Egress sessions 4)delete the configurations 5) Repeat step 1.The same core will also be noticed with 3000 and 4000 tunnel scale tests[PR1263201](#)

Open vSwitch Database Management Protocol (OVSDB)

- OVSDB (QFX5110, QFX5200, and QFX10000 switches)-- Use of OVSDB is not recommended on these switches with Junos OS 17.2R1. [PR1288227](#)

OpenFlow

- OpenFlow (QFX5110, QFX5200, and QFX10000 switches)-- Use of OpenFlow is not recommended on these switches with Junos OS 17.2R1. [PR1288227](#)

Platform and Infrastructure

- On QFX5100, QFX5110, and QFX5200 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](#)
- On QFX10002 switches, the **request system snapshot** command does not work. [PR1048182](#)
- While doing an SSH to a VNF, an error with the message **Unrecognized command** is seen. This has no impact on the functionality. [PR1108785](#)
- For QFX5110-32Q, throughput as per RFC 2544 is not 100% 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 10G, 40G and 100G ports separately.[PR1256671](#)

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](#)
- When polling the SNMP MIB jnxFirewallsEntry and if more than one firewall filter was configured and attached to any logical interfaces on QFX series switches, the counters for only one firewall filter would be returned. Now all filters and counters are returned when polling the MIB. [PR1250776](#)
- After doing a '**restart routing**' in the master Routing Engine, PIM join states of VXLAN multicast groups in the backup Routing Engine are not in sync with the master Routing Engine. The PIM join states in the master Routing Engine are proper and the problem happens only in the backup Routing Engine. There is no impact to the master Routing Engine due to this problem. The workaround to this problem is to apply '**restart routing**' in the backup Routing Engine. This will clear the problem in the backup Routing Engine. [PR1255480](#)
- Unified ISSU is not supported with MPLS configuration. [PR1264786](#)

SEE ALSO

New and Changed Features 226
Changes in Behavior and Syntax 247
Known Behavior 249
Resolved Issues 254
Documentation Updates 256
Migration, Upgrade, and Downgrade Instructions 257
Product Compatibility 263

Resolved Issues

IN THIS SECTION

- [General Routing | 255](#)
- [Layer 2 Features | 255](#)
- [Routing Protocols | 256](#)

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.

General Routing

- When a QFX5110 switch is acting as a DHCP relay agent, DHCP Reply packets might not be relayed back to the client by the relay agent if the Reply packets enter on a GRE tunneling interface. [PR1198982](#)
- The backup link in the aggregated link is not forwarding the traffic when the primary link goes down in the following configuration with Junos OS Release 15.1R4. **root# show interfaces ge-0/0/10 ether-options { 802.3ad { ae0; primary; } } {master:0}[edit]** **root# show interfaces ge-0/0/19 ether-options { 802.3ad { ae0; backup; } }** [PR1208614](#)
- 1G mode is not supported in Junos OS Release 17.1R1. It will be available in Junos OS Release 17.2R1. [PR1239091](#)
- On QFX10000 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](#)
- When you activate and deactivate Route Target per bridge domain in EVPN, the RPD process might crash resulting in traffic loss. [PR1244956](#)
- The output of **show interface** might incorrectly show interfaces as Link-mode: Auto and Speed: Auto even though the interface is manually configured at a speed and duplex setting. This issue is cosmetic in nature as the interface is indeed running at the manually configured Speed and Duplex setting. [PR1260986](#)
- QFX5200, 15.1X53-D210.7, Error log "ifd; does not exist" during SNMP query and SNMP delay when using **set chassis fpc 0 pic 0 port-range 0 127** on standalone QFX5200. [PR1263794](#)
- Removing force-up on an active link can cause programming issues on the QFX5100. Traffic returning from the destination will not be forwarded on egress on the QFX5100. [PR1264650](#)
- The description for a 40G-AOC cable in **show chassis hardware** shows UNKNOWN. This issue is seen with the cable length 1M, 3M, 5M, and 7M. [PR1269018](#)

Layer 2 Features

- On QFX5100 with Junos OS Release 15.1R3, if RTG and VSTP are configured on the same VLAN, communication will not work over RTG interfaces. When only RTG is configured on the VLAN, the communication works as expected. This issue is resolved in Junos OS Release 15.1R6. [PR1230750](#)
- S-Link macs addresses are not moving across MC-LAG chassis on the 10K platform. [PR1260316](#)

- Flexible tagged LAG interface might go down when configuring native VLAN. [PR1262529](#)
- In case any STP protocol is running on a peer device, BPDU packets will be dropped on QFX5100 instead of flooding it to the same segment. [PR1262847](#)

Routing Protocols

- VCF does not forward BUM after fabric-tree-root is configured. [PR1257984](#)
- When you change the member interface of a LAG with an IPv4 and IPv6 configuration on it, traffic forwarding might be impacted. As a workaround, delete or deactivate the LAG bundle configuration and reconfigure the same. [PR1270011](#)

SEE ALSO

New and Changed Features 226
Changes in Behavior and Syntax 247
Known Behavior 249
Known Issues 250
Documentation Updates 256
Migration, Upgrade, and Downgrade Instructions 257
Product Compatibility 263

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.2R1.

SEE ALSO

New and Changed Features 226
Changes in Behavior and Syntax 247
Known Behavior 249
Known Issues 250
Resolved Issues 254
Migration, Upgrade, and Downgrade Instructions 257
Product Compatibility 263

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 257
- Installing the Software on QFX10002 Switches | 259
- Performing an In-Service Software Upgrade (ISSU) | 259
- Preparing the Switch for Software Installation | 260
- Upgrading the Software Using ISSU | 260

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **17.2** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 17.2 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.

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

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

```
user@host> request system software add source/jinstall-host-qfx-5-17.2R1.13-signed.tgz reboot
```

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.2 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D43. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D43 or Junos OS Release 17.2R1.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is not compatible from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-17.2R1.13-secure-signed.tgz
reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-qfx-10-f-flex-17.2R1.10-domestic.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Performing an In-Service Software Upgrade (ISSU)

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 260](#)
- [Upgrading the Software Using ISSU on page 260](#)

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 ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, `jinstall-132_x51_vjunos.domestic.tgz`.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-qfx-5-13.2X51-D15.4-domestic ...
Install jinstall-qfx-5-13.2X51-D15.4-domestic completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
```

```

ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. To ensure that the resilient dual-root partitions feature operates correctly, copy 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](#) | 226

[Changes in Behavior and Syntax | 247](#)

[Known Behavior | 249](#)

[Known Issues | 250](#)

[Resolved Issues | 254](#)

[Documentation Updates | 256](#)

[Product Compatibility | 263](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 263](#)

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 | 226](#)

[Changes in Behavior and Syntax | 247](#)

[Known Behavior | 249](#)

[Known Issues | 250](#)

[Resolved Issues | 254](#)

Third-Party Components

This product includes third-party components. To obtain a complete list of third-party components, see [Copyright and Trademark Information](#).

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 the Junos OS, see the Juniper Networks Problem Report Search application at <https://prsearch.juniper.net>.

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

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/>.

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF format. Find Content Explorer at <https://www.juniper.net/documentation/content-applications/content-explorer/>.

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post sales technical support, you can access our tools and resources online or open a case with JTAC.

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <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:
<https://entitlementsearch.juniper.net/entitlementsearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at
<https://www.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/documentation/feedback/>.

Revision History

3 September 2020—Revision 11, Junos OS Release 17.2R1—ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

14 February 2019—Revision 10, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

14 December 2017—Revision 9, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

28 September 2017—Revision 8, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

1 September 2017—Revision 7, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

20 July 2017—Revision 6, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

6 July 2017—Revision 5, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

29 June 2017—Revision 4, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

20 June 2017—Revision 3, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

13 June 2017—Revision 2, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and Junos Fusion.

6 June 2017—Revision 1, Junos OS Release 17.2R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, and Junos Fusion.

Copyright © 2017, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.