

Release Notes: Junos[®] OS Release 17.4R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

17 September 2021

Contents	Introduction 11
	Junos OS Release Notes for ACX Series 11
	New and Changed Features 12
	Release 17.4R2 New and Changed Features 12
	Release 17.4R1 New and Changed Features 12
	Changes in Behavior and Syntax 13
	Management 14
	Security 14
	Subscriber Management and Services 14
	Known Behavior 15
	Known Issues 16
	Interfaces and Chassis 16
	Resolved Issues 17
	Documentation Updates 17
	Migration, Upgrade, and Downgrade Instructions 18
	Upgrade and Downgrade Support Policy for Junos OS Releases 18

Product Compatibility | 19

Hardware Compatibility | 19

Junos OS Release Notes for EX Series Switches | 20

New and Changed Features | 20

Release 17.4R2 New and Changed Features | 21

Release 17.4R1 New and Changed Features | 22

Changes in Behavior and Syntax | 27

EVPNs | 27

Management | 28

Multicast | 28

Network Management and Monitoring | 28

Security | 29

Software Licensing | 29

Subscriber Management and Services | 29

Virtual Chassis | 30

Known Behavior | 31

High Availability (HA) and Resiliency | 32

Infrastructure | 32

Interfaces and Chassis | 32

Platform and Infrastructure | 32

Virtual Chassis | 33

Known Issues | 33

Infrastructure | 33

Platform and Infrastructure | 34

Resolved Issues | 35

Resolved Issues: 17.4R2 | 35

Resolved Issues: 17.4R1 | 39

Documentation Updates | 41

Migration, Upgrade, and Downgrade Instructions | 42

Upgrade and Downgrade Support Policy for Junos OS Releases | 42

Product Compatibility | 43

Hardware Compatibility | 43

Junos OS Release Notes for Junos Fusion Data Center | 44

New and Changed Features | 44

Changes in Behavior and Syntax | 45

Known Behavior | 45

Junos Fusion Data Center | 46

Known Issues | 46

Resolved Issues | 47

Resolved Issues: Junos OS Release 17.4R2 | 47

Resolved Issues: Junos OS Release 17.4R1 | 47

Documentation Updates | 48

Migration, Upgrade, and Downgrade Instructions | 48

Basic Procedure for Upgrading an Aggregation Device | 49

Preparing the Switch for Satellite Device Conversion | 51

Autoconverting a Switch into a Satellite Device | 53

Manually Converting a Switch into a Satellite Device | 56

Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 58

Configuring Satellite Device Upgrade Groups | 59

Converting a Satellite Device to a Standalone Device | 61

Upgrade and Downgrade Support Policy for Junos OS Releases | 61

Downgrading from Release 17.4 | 61

Product Compatibility | 62

Hardware Compatibility | 62

Junos OS Release Notes for Junos Fusion Enterprise | 63

New and Changed Features | 63

Release 17.4R2 New and Changed Features | 64

Release 17.4R1 New and Changed Features | 64

Changes in Behavior and Syntax | 65

Junos Fusion Enterprise | 66

Known Behavior | 66

Junos Fusion Enterprise | 66

Known Issues | 67

Resolved Issues | 68**Resolved Issues: 17.4R2 | 68****Resolved Issues: 17.4R1 | 68****Documentation Updates | 69****Migration, Upgrade, and Downgrade Instructions | 70****Basic Procedure for Upgrading Junos OS on an Aggregation Device | 70****Upgrading an Aggregation Device with Redundant Routing Engines | 72****Preparing the Switch for Satellite Device Conversion | 73****Converting a Satellite Device to a Standalone Switch | 74****Upgrade and Downgrade Support Policy for Junos OS Releases | 74****Downgrading from Release 17.4 | 75****Product Compatibility | 75****Hardware and Software Compatibility | 76****Hardware Compatibility Tool | 76****Junos OS Release Notes for Junos Fusion Provider Edge | 77****New and Changed Features | 77****Release 17.4R2 New and Changed Features | 78****Release 17.4R1 New and Changed Features | 78****Changes in Behavior and Syntax | 78****Security | 79****Known Behavior | 79****Junos Fusion Provider Edge | 80****Known Issues | 80****Junos Fusion Provider Edge | 81****Resolved Issues | 82****Resolved Issues: 17.4R2 | 82****Resolved Issues: 17.4R1 | 83****Documentation Updates | 83****Migration, Upgrade, and Downgrade Instructions | 84****Basic Procedure for Upgrading an Aggregation Device | 84****Upgrading an Aggregation Device with Redundant Routing Engines | 86****Preparing the Switch for Satellite Device Conversion | 87****Converting a Satellite Device to a Standalone Device | 88****Upgrading an Aggregation Device | 90**

Upgrade and Downgrade Support Policy for Junos OS Releases	91
Downgrading from Release 17.4	91
Product Compatibility	92
Hardware Compatibility	92
Junos OS Release Notes for MX Series 5G Universal Routing Platforms	93
New and Changed Features	93
Release 17.4R2-S2 New and Changed Features	94
Release 17.4R2 New and Changed Features	94
Subscriber Management and Services	96
Release 17.4R1 New and Changed Features	96
Changes in Behavior and Syntax	126
EVPNs	126
High Availability (HA) and Resiliency	127
Interfaces and Chassis	127
Management	128
MPLS	128
Multicast	130
Network Management and Monitoring	130
Routing Protocols	131
Security	131
Services Applications	132
Software Defined Networking	133
Software Installation and Upgrade	133
Software Licensing	133
Subscriber Management and Services	133
User Interface and Configuration	136
Known Behavior	136
General Routing	137
EVPN	139
Interfaces and Chassis	139
Layer 2 Ethernet Services	140
MPLS	140
Routing Protocols	140
Services Applications	140

Software Installation and Upgrade	141
Subscriber Management and Services	141
Known Issues	142
General Routing	143
Class of Service (CoS)	148
EVPN	148
Forwarding and Sampling	149
High Availability (HA) and Resiliency	150
Infrastructure	150
Interfaces and Chassis	150
Layer 2 Features	151
Layer 2 Ethernet Services	151
Multiprotocol Label Switching (MPLS)	151
Platform and Infrastructure	152
Routing Protocols	155
Services Applications	157
VPNs	157
Resolved Issues	157
Resolved Issues: 17.4R2	158
Resolved Issues: 17.4R1	184
Documentation Updates	199
Subscriber Management Provisioning guide	199
Migration, Upgrade, and Downgrade Instructions	200
Basic Procedure for Upgrading to Release 17.4	201
Procedure to Upgrade to FreeBSD 11.x-Based Junos OS	201
Procedure to Upgrade to FreeBSD 6.x-Based Junos OS	203
Upgrade and Downgrade Support Policy for Junos OS Releases	205
Upgrading a Router with Redundant Routing Engines	206
Downgrading from Release 17.4	206
Product Compatibility	207
Hardware Compatibility	207

Junos OS Release Notes for NFX Series | 208

New and Changed Features | 208

Release 17.4R2 New and Changed Features | 209

Release 17.4R1 New and Changed Features | 209

Changes in Behavior and Syntax | 209

Known Behavior | 210

Known Issues | 210

Known Issues: 17.4R2 | 211

Known Issues: 17.4R1 | 211

Resolved Issues | 211

Documentation Updates | 212

Migration, Upgrade, and Downgrade Instructions | 212

Upgrade and Downgrade Support Policy for Junos OS Releases | 212

Basic Procedure for Upgrading to Release 17.4 | 213

Product Compatibility | 214

Hardware Compatibility | 215

Junos OS Release Notes for PTX Series Packet Transport Routers | 216

New and Changed Features | 217

Release 17.4R2 New and Changed Features | 217

Release 17.4R1 New and Changed Features | 217

Changes in Behavior and Syntax | 229

Class of Service (CoS) | 230

Interfaces and Chassis | 230

Management | 232

MPLS | 232

Multicast | 233

Network Management and Monitoring | 233

Security | 235

Software Licensing | 235

Subscriber Management and Services | 235

Known Behavior | 236

General Routing | 236

Interfaces and Chassis | 237

Known Issues | 237**General Routing | 238****Interfaces and Chassis | 240****MPLS | 240****Platform and Infrastructure | 240****Resolved Issues | 240****Resolved Issues: 17.4R2 | 241****Resolved Issues: 17.4R1 | 244****Documentation Updates | 247****Migration, Upgrade, and Downgrade Instructions | 247****Upgrade and Downgrade Support Policy for Junos OS Releases | 247****Upgrading a Router with Redundant Routing Engines | 248****Basic Procedure for Upgrading to Release 17.4 | 248****Product Compatibility | 252****Hardware Compatibility | 252****Junos OS Release Notes for the QFX Series | 253****New and Changed Features | 253****Release 17.4R2 New and Changed Features | 255****Release 17.4R1 New and Changed Features | 255****Changes in Behavior and Syntax | 266****Class of Service (CoS) | 267****EVPNs | 267****General Routing | 267****Management | 267****MPLS | 267****Network Management and Monitoring | 269****Routing Policy and Firewall Filters | 270****Security | 270****Software Licensing | 270****Virtual Chassis | 270****Known Behavior | 272****Class of Service (CoS) | 272****EVPN | 272****Interfaces and Chassis | 274**

Layer 2 Features	274
MPLS	274
Routing Protocols	274
Platform and Infrastructure	275
Virtual Chassis	275
Known Issues	276
EVPN	276
Layer 2 Features	277
MPLS	277
Platform and Infrastructure	277
Routing Protocols	279
Resolved Issues	280
Resolved Issues: 17.4R2	280
Resolved Issues: 17.4R1	287
Documentation Updates	291
Migration, Upgrade, and Downgrade Instructions	292
Upgrading Software on QFX Series Switches	292
Installing the Software on QFX10002 Switches	295
Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches	295
Installing the Software on QFX10008 and QFX10016 Switches	297
Performing a Unified ISSU	301
Preparing the Switch for Software Installation	302
Upgrading the Software Using Unified ISSU	302
Upgrade and Downgrade Support Policy for Junos OS Releases	305
Product Compatibility	305
Hardware Compatibility	306
Junos OS Release Notes for SRX Series	307
New and Changed Features	307
Release 17.4R2 New and Changed Features	308
Release 17.4R1-S1 New and Changed Features	308

Release 17.4R1 New and Changed Features	310
Changes in Behavior and Syntax	319
Chassis Cluster	319
IDP	320
Forwarding and Sampling	320
System Logging	321
User Interface and Configuration	321
Known Behavior	321
Authentication and Access	322
Chassis Clustering	322
Install and Upgrade	323
Interfaces and Chassis	323
J-Web	324
Layer 2 Ethernet Services	325
User Interface and Configuration	325
VPNs	325
Known Issues	325
Outstanding Issues	326
Resolved Issues	328
Resolved Issues: 17.4R2	329
Resolved Issues: 17.4R1	338
Documentation Updates	342
Migration, Upgrade, and Downgrade Instructions	342
Upgrade and Downgrade Scripts for Address Book Configuration	343
Product Compatibility	346
Hardware Compatibility	346
Upgrading Using ISSU	347
Compliance Advisor	347
Finding More Information	347
Documentation Feedback	348
Requesting Technical Support	349
Self-Help Online Tools and Resources	349
Creating a Service Request with JTAC	350
Revision History	350

Introduction

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

These release notes accompany Junos OS Release 17.4R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- New and Changed Features | 12
- Changes in Behavior and Syntax | 13
- Known Behavior | 15
- Known Issues | 16
- Resolved Issues | 17
- Documentation Updates | 17
- Migration, Upgrade, and Downgrade Instructions | 18
- Product Compatibility | 19

These release notes accompany Junos OS Release 17.4R2 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R2 New and Changed Features | 12](#)
- [Release 17.4R1 New and Changed Features | 12](#)

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

Release 17.4R2 New and Changed Features

There are no new features or enhancements to existing features for ACX Series in Junos OS Release 17.4R2.

Release 17.4R1 New and Changed Features

Management

- **Support for multiple, smaller configuration YANG modules (ACX Series)**—Starting in Junos OS Release 17.4R1, the YANG module for the Junos OS configuration schema is split into a root configuration module that is augmented by multiple, smaller modules. The root configuration module comprises the top-level configuration node and any nodes that are not emitted as separate modules. Separate, smaller modules augment the root configuration module for the different configuration statement hierarchies. Smaller configuration modules enable YANG tools and utilities to more quickly and efficiently compile and work with the modules, because they only need to import the modules required for the current operation.

[See [Understanding the YANG Modules That Define the Junos OS Configuration.](#)]

Timing and Synchronization

- **Enterprise profile for Precision Time Protocol (PTP) (ACX1100 Router)**—Starting with Junos OS Release 17.4R1, the enterprise profile, which is based on PTPv2, provides the ability for enterprise and financial markets to timestamp on different systems and to handle a range of latency and delays. The enterprise profile supports the following options:
 - IPv4 multicast transport
 - Boundary clocks
 - 512 downstream slave clocks

You can enable the enterprise profile at the [edit protocols ptp profile-type] hierarchy.

NOTE: On ACX Series, the enterprise profile for PTP is supported only on ACX1100 AC router.

SEE ALSO

Changes in Behavior and Syntax 13
Known Behavior 15
Documentation Updates 17
Known Issues 16
Resolved Issues 17
Migration, Upgrade, and Downgrade Instructions 18
Product Compatibility 19

Changes in Behavior and Syntax

IN THIS SECTION

- [Management | 14](#)
- [Security | 14](#)
- [Subscriber Management and Services | 14](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 17.4R2 for the ACX Series Universal Metro Routers.

Management

- **Changes to Junos OS YANG module naming conventions (ACX Series)**—Starting in Junos OS Release 17.4R1, the native Junos OS YANG modules use a new naming convention for the module's name, filename, and namespace. The module name and filename include the device family and the area of the configuration or command hierarchy to which the schema in the module belongs. In addition, the module filename includes a revision date. The module namespace is simplified to include the device family, the module type, and an identifier that is unique to each module and that differentiates the namespace of the module from that of other modules.

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

Security

- **Support to log the SSH key changes**—Starting with Junos OS 17.4R1, the configuration statement **log-key-changes** is introduced at the `[edit system services ssh]` hierarchy level. When the **log-key-changes** configuration statement is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time the **log-key-changes** configuration statement was enabled. If the **log-key-changes** configuration statement was never enabled, then Junos OS logs all the authorized SSH keys.

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (ACX Series)**—Starting in Junos OS Release 17.4R2, the `jdhcpd` process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

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

In earlier releases:

1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.

- 2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

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

SEE ALSO

New and Changed Features	 12
Known Behavior	 15
Documentation Updates	 17
Known Issues	 16
Resolved Issues	 17
Migration, Upgrade, and Downgrade Instructions	 18
Product Compatibility	 19

Known Behavior

There are no known limitations in Junos OS Release 17.4R2 for the ACX Series.

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

SEE ALSO

New and Changed Features	 12
Changes in Behavior and Syntax	 13
Documentation Updates	 17
Known Issues	 16
Resolved Issues	 17
Migration, Upgrade, and Downgrade Instructions	 18
Product Compatibility	 19

Known Issues

IN THIS SECTION

- [Interfaces and Chassis | 16](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R2 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.

Interfaces and Chassis

- On ACX Series router, when link-speed is configured explicitly, aggregate interface goes down permanently after reboot. [PR1022248](#)
- In a normal software MAC learning mode, when incremental MAC traffic of range higher than the profile is received, then after febr restart the MAC entries is not seen in the software table, although present in the hardware table. As a workaround, in the hardware MAC learning mode, delete the routing instance and reconfigure it again to make the MAC entries seen in the software table. In the software MAC learning mode, deactivate the routing instance, clear the pending entries or allow the pending entries to be aged out and then activate the routing instance to solve this issue. [PR1277436](#)
- On ACX Series PE routers with Layer 3 VPN configured, when running traceroute on ingress PE to CE, only P hop and CE are displayed. The PE information is not displayed. This is a hardware limitation and a workaround is not available. [PR1313013](#)

SEE ALSO

[New and Changed Features | 12](#)

[Changes in Behavior and Syntax | 13](#)

[Known Behavior | 15](#)

[Documentation Updates | 17](#)

[Resolved Issues | 17](#)

[Migration, Upgrade, and Downgrade Instructions | 18](#)

[Product Compatibility | 19](#)

Resolved Issues

There are no fixed issues in Junos OS 17.4R2 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	 12
Changes in Behavior and Syntax	 13
Known Behavior	 15
Documentation Updates	 17
Known Issues	 16
Migration, Upgrade, and Downgrade Instructions	 18
Product Compatibility	 19

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 for the ACX Series documentation.

SEE ALSO

New and Changed Features	 12
Changes in Behavior and Syntax	 13
Known Behavior	 15
Known Issues	 16
Resolved Issues	 17
Migration, Upgrade, and Downgrade Instructions	 18
Product Compatibility	 19

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 18](#)

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 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

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

SEE ALSO

[New and Changed Features | 12](#)

[Changes in Behavior and Syntax | 13](#)

[Known Behavior | 15](#)

[Documentation Updates | 17](#)

[Known Issues | 16](#)

[Resolved Issues | 17](#)

[Product Compatibility | 19](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 19](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

[New and Changed Features | 12](#)

[Changes in Behavior and Syntax | 13](#)

[Known Behavior | 15](#)

[Documentation Updates | 17](#)

[Known Issues | 16](#)

[Resolved Issues | 17](#)

[Migration, Upgrade, and Downgrade Instructions | 18](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- New and Changed Features | 20
- Changes in Behavior and Syntax | 27
- Known Behavior | 31
- Known Issues | 33
- Resolved Issues | 35
- Documentation Updates | 41
- Migration, Upgrade, and Downgrade Instructions | 42
- Product Compatibility | 43

These release notes accompany Junos OS Release 17.4R2 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R2 New and Changed Features | 21
- Release 17.4R1 New and Changed Features | 22

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R2 for the EX Series.

NOTE: The following EX Series switches are supported in Release 17.4R2: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.4R2, 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 17.4A1 for EX4300 and EX4600 Switches](#).

Release 17.4R2 New and Changed Features

EVPNs

- **EVPN proxy ARP and ARP suppression without IRB interfaces (MX Series routers with MPCs, EX9200 switches)**—MX Series routers and EX9200 switches that function as provider edge (PE) devices in an Ethernet VPN-MPLS (EVPN-MPLS) or EVPN-Virtual Extensible LAN (EVPN-VXLAN) environment support the proxy Address Resolution Protocol (ARP) and ARP suppression. Both ARP capabilities are enabled by default.

Starting with Junos OS Release 17.4R2, these features no longer require the configuration of an IRB interface on the PE device. Any interface configured on a PE device can now deliver ARP requests from both local customer edge (CE) devices only. Proxy ARP and ARP suppression are not supported on remote CE devices.

Also, you can now control the following aspects of the MAC-IP address bindings database on a PE device:

- The maximum number of MAC-IP address entries in the database.
- The amount of time a locally learned MAC-IP address binding remains in the database.

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

Restoration Procedures and Failure Handling

- **Device recovery mode support introduced in Junos OS with upgraded FreeBSD (EX Series)**—Starting in Junos OS Release 17.4R2, devices running Junos OS with an upgraded FreeBSD and a saved rescue configuration have an automatic device recovery mode should the system go into amnesiac mode. The new process has the system automatically reboot with the saved rescue configuration. Then the system displays "Device is in recovery mode" in the CLI (in both operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

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

Release 17.4R1 New and Changed Features

Hardware

- **Aggregation device support on EX9200 with EX9200-RE2 routing engine (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.4, EX9200 switches with the EX9200-RE2 Routing Engine module are supported as aggregation devices in a Junos Fusion Enterprise. The EX9200-RE2 module supports virtual machine (VM) architecture in an EX9200 switch.

[See [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#).]

Authentication, Authorization and Accounting (AAA)

- **Periodic refresh of authorization profile on TACACS server (EX Series)**—Starting with Junos OS Release 17.4R1, periodic refresh of the authorization profile that is received from the TACACS server is supported. The authorization profile that is configured for the user on the TACACS server is sent to the Junos OS device after the user is successfully authenticated. The authorization profile is stored locally on the Junos OS device. With the periodic refresh feature, the authorization profile is periodically fetched from the TACACS server to refresh the authorization profile that is stored locally. User authorization is reevaluated using the refreshed authorization profile.

[See [Configuring Periodic Refresh of the TACACS+ Authorization Profile](#).]

EVPNs

- **EVPN-MPLS interworking with Junos Fusion Enterprise and MC-LAG (EX9200 switches)**—Starting with Junos OS Release 17.4R1, you can use Ethernet VPN (EVPN) to extend your Junos Fusion Enterprise or MC-LAG network over an MPLS network. Typically, Junos Fusion Enterprise is extended to a geographically distributed campus or enterprise network, while an MC-LAG network is extended to a data center network or geographically distributed campus or enterprise network.

The EVPN-MPLS interworking feature offers the following benefits:

- Ability to use separate virtual routing and forwarding (VRF) instances to control inter-VLAN routing.
- VLAN translation.
- Default Layer 3 virtual gateway support, which eliminates the need to run such protocols as Virtual Router Redundancy Protocol (VRRP).
- Load balancing to better utilize both links when using EVPN multihoming.
- The use of EVPN type 2 advertisement routes (MAC+IP) reduces the need for flooding domains with ARP packets.

[See [Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG](#).]

- **Support for duplicate MAC address detection and suppression (EX9200 switches)**— When a MAC address relocates, PE devices can converge on the latest location by using sequence numbers in the extended community field. Misconfigurations in the network can lead to duplicate MAC addresses. Starting in Junos OS Release 17.4R1, Juniper supports duplicate MAC address detection and suppression.

You can modify the duplicate MAC address detection settings on the switch by configuring the detection window for identifying duplicate MAC address and the number of MAC address moves detected within the detection window before duplicate MAC detection is triggered and the MAC address is suppressed. In addition, you can also configure an optional recovery time that the switch waits before the duplicate MAC address is automatically unsuppressed.

To configure duplicate MAC detection parameters, use the **detection-window**, **detection-threshold**, and **auto-recovery-time** statements at the `[edit routing instance routing-instance-name protocols evpn duplicate-mac-detection]` hierarchy level.

To clear duplicate MAC suppression manually, use the **clear evpn duplicate-mac-suppression** command.

[See [Overview of MAC Mobility](#).]

Junos OS XML API and Scripting

- **Automation script library additions and upgrades (EX Series)**—Starting in Junos OS Release 17.4R1, devices running Junos OS include new and upgraded Python modules as well as upgraded versions of Junos PyEZ and libslax. On-box Python automation scripts can use features supported in Junos PyEZ Release 2.1.4 and earlier releases to perform operational and configuration tasks on devices running Junos OS. Python automation scripts can also leverage new on-box Python modules including **ipaddress**, **jxmlease**, **pyang**, **serial**, and **six**, as well as upgraded versions of existing modules. In addition, SLAX automation scripts can include features supported in libslax release 0.22.0 and earlier releases.

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

Layer 2 Features

- **Layer 2 protocol tunneling support (EX4600 switches and Virtual Chassis)**—Starting with Junos OS Release 17.4R1, Layer 2 protocol tunneling (L2PT) is supported on EX4600 switches and EX4600 Virtual

Chassis. You can configure the switch to tunnel any of the following Layer 2 protocols: CDP, E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, STP (including RSTP and MSTP), UDLD, VSTP, and VTP.

[See [Layer 2 Protocol Tunneling](#).]

- **Q-in-Q support on redundant trunk links using LAGs with link protection (EX4300 switches and Virtual Chassis)**—Starting in Junos OS Release 17.4R1, Q-in-Q is supported on redundant trunk links (also called “RTGs”) using LAGs with link protection. Redundant trunk links provide a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.

Q-in-Q support on redundant trunk links on a LAG with link protection also includes support for the following items:

- Configuration of flexible VLAN tagging on the same LAG that supports the redundant links configurations
- Multiple redundant-link configurations on one physical interface
- Multicast convergence

[See [Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection](#).]

Management

- **Enhancements to LSP events sensor for Junos Telemetry Interface (EX4600 and EX9200 switches)**—Starting with Junos OS Release 17.4R1, telemetry data streamed through gRPC for LSP events and properties is reported separately for each routing instance. To export data for LSP events and properties, you must now include `/network-instances/network-instance/[name_'instance-name']/` in front of all supported paths. For example, to export LSP events for RSVP Signaling protocol attributes, use the following path: `/network-instances/network-instance[name_'instance-name']/mpls/signaling-protocols/rsvp-te/`. Use the `telemetrySubscribe` RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Support for multiple, smaller configuration YANG modules (EX Series)**—Starting in Junos OS Release 17.4R1, the YANG module for the Junos OS configuration schema is split into a root configuration module that is augmented by multiple, smaller modules. The root configuration module comprises the top-level configuration node and any nodes that are not emitted as separate modules. Separate, smaller modules augment the root configuration module for the different configuration statement hierarchies. Smaller configuration modules enable YANG tools and utilities to more quickly and efficiently compile and work with the modules, because they only need to import the modules required for the current operation.

[See [Understanding the YANG Modules That Define the Junos OS Configuration](#).]

- **Enhancement to BGP sensor for Junos Telemetry Interface (EX4600 and E9200 switches)**—Starting with Junos OS Release 17.4R1, you can specify to export the number of BGP peers in a BGP group for telemetry data exported through gRPC. To export the number of BGP peers for a group, use the following OpenConfig path: `/network-instances/network-instance[name_'instance-name']/protocols/protocol/bgp/peer-groups/peer-group[name_'peer-group-name']/state/peer-count/`. The BGP peer count value exported reflects the number of peering sessions in a group. For example, for a BGP group with two devices, the peer count reported is 1 (one) because each group member has one peer. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

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

Multicast

- **MLD snooping versions 1 and 2 (EX4600 switches and Virtual Chassis)**—Starting with Junos OS Release 17.4R1, EX4600 switches and EX4600 Virtual Chassis support Multicast Listener Discovery (MLD) snooping version 1 (MLDv1) and version 2 (MLDv2). MLD snooping constrains the flooding of IPv6 multicast traffic on VLANs. When MLD snooping is enabled on a VLAN, the switch examines MLD messages encapsulated within ICMPv6 packets transferred between hosts and multicast routers. The switch learns which hosts are interested in receiving traffic for a multicast group and forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces. You configure MLD snooping parameters and enable MLD snooping using configuration statements at the **[edit protocols] mld-snooping vlan *vlan-name*** hierarchy.

[See [Understanding MLD Snooping on Switches](#).]

Routing Protocols

- **Support for EBGp route server (EX Series)**—Starting in Junos OS Release 17.4R1, BGP feature is enhanced to support EBGp route server functionality. A BGP route server is the external BGP (EBGP) equivalent of an internal IBGP (IBGP) route reflector that simplifies the number of direct point-to-point EBGp sessions required in a network. EBGp route server propagates unmodified BGP routing information between external BGP peers to facilitate high scale exchange of routes in peering points such as Internet Exchange Points (IXPs). When BGP is configured as a route server, EBGp routes are propagated between peers unmodified, with full attribute transparency (NEXT_HOP, AS_PATH, MULTI_EXIT_DISC, AIGP, and Communities).

The BGP JET **bgp_route_service.proto** API has been enhanced to support route server functionality as follows:

- Program the EBGp route server.
- Inject routes to the specific route server RIB for selectively advertising it to the client groups in client-specific RIBs.

The BGP JET **bgp_route_service.proto** API includes a peer-type object that identifies individual routes as either EBGp or IBGP (default).

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

- **Support for importing IGP topologies into BGP-LS (EX Series)**—Starting in Junos OS Release 17.4R1, you can import IGP, that is IS-IS and OSPF topologies into BGP-LS. Prior to Junos OS Release 17.4R1, Junos OS BGP-LS implementation exports only Traffic Engineering enabled (RSVP-enabled) links. This feature allows you to export IGP links (that do not have RSVP enabled) and Traffic Engineering enabled links into BGP-LS.

Software Installation and Upgrade

- **Configuration validation for image upgrade or downgrade (EX4300)**—Starting in Junos OS Release 17.4R1, when you install a new version of Junos OS on the switch, the system validates that the existing configuration is compatible with the new image. Without the validation feature, configuration incompatibilities or insufficient memory to load the new image might cause the system to lose its current configuration or go offline. With the validation feature, if validation fails, the new image is not loaded, and an error message provides information about the failure.

Image validation is supported only on the **jinstall** package.

If you invoke validation from an image that does not support validation, the new image is loaded but validation does not occur.

Invoke validation by issuing either **request system software add** or **request system software nonstop-upgrade**. You can also issue **request system software validate** to run just configuration validation.

Image validation does not work in a downgrade from Release 17.4 to 17.2 or earlier if graceful switchover is enabled and image loading is done without NSSU. Use one of the following options:

- To downgrade with graceful switchover but without image validation—Issue the **request system software add image-name reboot no-validate** command.
- To downgrade with image validation but without graceful switchover—Remove the graceful-switchover configuration and then issue the **request system software add image-name reboot** command.
- To downgrade with image validation and graceful switchover—Use NSSU by issuing the **request system software nonstop-upgrade image-name** command.

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

SEE ALSO

Changes in Behavior and Syntax	27
Known Behavior	31
Known Issues	33
Resolved Issues	35
Documentation Updates	41
Migration, Upgrade, and Downgrade Instructions	42

Changes in Behavior and Syntax

IN THIS SECTION

- [EVPNs | 27](#)
- [Management | 28](#)
- [Multicast | 28](#)
- [Network Management and Monitoring | 28](#)
- [Security | 29](#)
- [Software Licensing | 29](#)
- [Subscriber Management and Services | 29](#)
- [Virtual Chassis | 30](#)

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 17.4R2 for the EX Series.

EVPNs

- **Change to show vlans evpn command (EX9200 switches)**—Starting with Junos OS Release 17.4R2, the **show vlans evpn** command is replaced by the **show ethernet-switching evpn** command.
- On EX9200 switches, you can configure EVPN to extend a Junos Fusion Enterprise or multichassis link aggregation group (MC-LAG) network over an MPLS network to a data center or campus network. For both Junos Fusion Enterprise and MC-LAG use cases, you must include the **bgp-peer** configuration statement in the **[edit routing-instances name protocols evpn mclag]** hierarchy level. This configuration enables the interworking of EVPN-MPLS with Junos Fusion Enterprise or MC-LAG. If you do not include the **bgp-peer** configuration statement in your configuration, unexpected behavior and a core dump could result. To enforce this configuration, we now check for this configuration during the commit. If the configuration is not present, an error occurs.

See [[Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG .](#)]

Management

- **Changes to Junos OS YANG module naming conventions (EX Series)**—Starting in Junos OS Release 17.4R1, the native Junos OS YANG modules use a new naming convention for the module's name, filename, and namespace. The module name and filename include the device family and the area of the configuration or command hierarchy to which the schema in the module belongs. In addition, the module filename includes a revision date. The module namespace is simplified to include the device family, the module type, and an identifier that is unique to each module and that differentiates the namespace of the module from that of other modules.

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

Multicast

- **Support for per-source multicast traffic forwarding with IGMPv3 (EX4300)**—Starting in Junos OS Release 17.4R2, EX4300 switches forward multicast traffic on a per-source basis according to received IGMPv3 INCLUDE and EXCLUDE reports. In releases prior to this release, EX4300 switches process IGMPv3 reports, but instead of source-specific multicast (SSM) forwarding, they consolidate IGMPv3 INCLUDE and EXCLUDE mode reports for a group into one route for all sources sending to the group. As a result, with the prior behavior, receivers might get traffic from sources they didn't specify.

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

Network Management and Monitoring

- **Change in default log level setting (EX Series)**—In Junos OS Release, 17.4R1, the following changes were made in default logging levels:

Before this change:

- SNMP_TRAP_LINK_UP was LOG_INFO for both the physical (IFD) and logical (IFL) interfaces.
- SNMP_TRAP_LINK_DOWN was LOG_WARNING for both the physical (IFD) and logical (IFL) interfaces.

After this change:

- IFD LinkUp -> LOG_NOTICE (because this is an important message but less frequent)
- IFL LinkUp -> LOG_INFO (no change)
- IFD and IFL LinkDown -> LOG_WARNING (no change)

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

- **SNMP syslog messages changed (EX Series)**—Starting in Junos OS Release 17.4R1, 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](#).]

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a non-default routing instance with a non-default logical system (EX Series)**—Starting in Junos OS Release 17.4R2, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

Security

- **Support for logging SSH key changes**—Starting with Junos OS Release 17.4R1, the configuration statement **log-key-changes** is introduced at the **[edit system services ssh]** hierarchy level. When **log-key-changes** configuration statement is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time **log-key-changes** was enabled. If **log-key-changes** was never enabled, then Junos OS logs all the authorized SSH keys.

Software Licensing

- **Key generator adds one day to make the duration of license show as 365 days (EX Series)**—Starting in Junos OS Release 17.4R1, the duration of subscription licenses as generated by the **show system license** command and shown in the output is correct to the numbers of days. Before this fix, for example, for a 1-year subscription license, the duration was generated as 364 days. After the fix, the duration of the 1-year subscription now shows as 365 days.

[See [show system license](#).]

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (EX Series)**—Starting in Junos OS Release 17.4R2, the **jdhcpd** process handles the second renew request differently if the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

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

In earlier releases:

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

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

Virtual Chassis

- **New configuration option to disable automatic Virtual Chassis port conversion (EX4300 and EX4600 Virtual Chassis)**—Starting in Junos OS Release 17.4R2, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in an EX4300 or EX4600 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

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

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

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

SEE ALSO

New and Changed Features	 20
Known Behavior	 31
Known Issues	 33
Resolved Issues	 35
Documentation Updates	 41
Migration, Upgrade, and Downgrade Instructions	 42
Product Compatibility	 43

Known Behavior

IN THIS SECTION

- [High Availability \(HA\) and Resiliency](#) | [32](#)
- [Infrastructure](#) | [32](#)
- [Interfaces and Chassis](#) | [32](#)
- [Platform and Infrastructure](#) | [32](#)
- [Virtual Chassis](#) | [33](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 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

- 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](#)

Infrastructure

- The issue is specific to a downgrade(17.4T) and a core is seen only once during the downgrade because of a timing issue in the sdk toolkit upgradation, after which dcpfe recovers on its own and no issues are seen after that. [PR1337008](#)

Interfaces and Chassis

- Configuring link aggregation group (LAG) hashing with the **[edit forwarding-options enhanced-hash-key] inet vlan-id** statement uses the VLAN ID in the hashing algorithm calculation. On some switching platforms, when this option is configured for a LAG that spans FPCs, such as in a Virtual Chassis or Virtual Chassis Fabric (VCF), packets are dropped due to an issue with using an incorrect VLAN ID in the hashing algorithm. As a result, the **vlan-id** hashing option is not supported in a Virtual Chassis or VCF containing any of the following members: EX4300, EX4600, QFX5100, or QFX5110 switches. Under these conditions, use any of the other supported **enhanced-hash-key** hashing configuration options instead. [PR1293920](#)

Platform and Infrastructure

- On EX4300 and EX4600 switches, if a remote analyzer has an output IP address that is reachable through a route learned by BGP, the analyzer might be in a down state. [PR1007963](#)
- On an EX4300 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On EX4300 switches, when 802.1X single-supplicant authentication is initiated, multiple "EAP Request Id Frame Sent" packets might be sent. [PR1163966](#)
- On EX4300 10G links, preexisting MACsec sessions might not come up after the following events: Process (pfex, dot1x) restart or system restart link flaps [PR1294526](#)
- mcsnoopd might crash when all the core facing interfaces that are part of the L2 domain have flapped and it is attempting to flood a packet received over a CE interface, over the core-facing interfaces. [PR1329694](#)

Virtual Chassis

- Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on an QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur and its considered to be known behavior. Release note reference:
https://www.juniper.net/documentation/en_US/junos/information-products/topic-collections/release-notes/172/topic-118735.html#PR1347902

SEE ALSO

New and Changed Features 20
Changes in Behavior and Syntax 27
Known Issues 33
Resolved Issues 35
Documentation Updates 41
Migration, Upgrade, and Downgrade Instructions 42
Product Compatibility 43

Known Issues

IN THIS SECTION

- Infrastructure | 33
- Platform and Infrastructure | 34

This section lists the known issues in hardware and software in Junos OS Release 17.4R2 for the EX Series.

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

Infrastructure

- The **set system ports console log-out-on-disconnect** configuration statement does not work. [PR1146891](#)

- When ML license is installed and if the master Routing Engine is operating at scale beyond the default limit, a ksyncd core file and a vmcore can be seen if gres configuration is enabled [while the master is scaling beyond the default limit]. [PR1376362](#)

Platform and Infrastructure

- On an EX9200-12QS line card, interfaces with the default speed of 10-Gigabit Ethernet are not brought down even when the remote end of a connection is misconfigured as 40-Gigabit Ethernet. [PR1175918](#)
- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, **fpc7 KERNEL/PFE APP=NH OUT OF SYNC: error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none**. There is no service impact observed in MPC2 and MPC3 type cards. [PR1205593](#)
- On EX4300 switches, when a policer with the action of loss of priority is applied to the lo0 interface, all ICMP packets might be dropped. [PR1243666](#)
- On EX4300 10G links, preexisting MACsec sessions might not come up after the following events: process (pfex, dot1x) restart or system restart link flaps. [PR1294526](#)
- In Streaming Telemetry scenario, if **commit full** is performed, na-grpd daemon might restart causing disconnection of streaming telemetry. [PR1326366](#)
- MPC5 - inline-ka PPP echo requests are not transmitted when anchor-point is lt-x/2/x or lt-x/3/x in a pseudowire deployment. [PR1345727](#)

SEE ALSO

[New and Changed Features | 20](#)

[Changes in Behavior and Syntax | 27](#)

[Known Behavior | 31](#)

[Resolved Issues | 35](#)

[Documentation Updates | 41](#)

[Migration, Upgrade, and Downgrade Instructions | 42](#)

[Product Compatibility | 43](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R2 | 35](#)
- [Resolved Issues: 17.4R1 | 39](#)

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

Resolved Issues: 17.4R2

Authentication and Access Control

- Macsec statistics display output is not proper. [PR1355339](#)

EVPN

- The traffic might get dropped as the core-facing interface is down. [PR1343515](#)
- Proxy ARP might not work as expected in an EVPN environment. [PR1368911](#)

High Availability (HA) and Resiliency

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

Infrastructure

- Unable to provide management when em0 interface of FPC is connected to another FPC L2 interface of the same Virtual Chassis. [PR1299385](#)
- The file system might be corrupted multiple times during an image upgrade or a commit operation. [PR1317250](#)
- The upgrade might fail if bad blocks are in the flash/filesystem and corruption occurs. [PR1317628](#)
- PFC feature might not work on an EX4600. [PR1322439](#)
- ifinfo core files can be created on an EX4600 Virtual Chassis. [PR1324326](#)
- There is support for archiving dmesg file `/var/run/dmesg.boot`. [PR1327021](#)

- Enabling mac-move-limit stops ping on flexible-vlan-tagging enabled interface. [PR1357742](#)
- The dot1x filter might be removed from the Packet Forwarding Engine when **static-mac-address** ages out or is learned by eswd. [PR1335125](#)

Interfaces and Chassis

- An identical IP address can be configured on different logical interfaces from different physical interfaces in the same routing instance (including the master routing instance). [PR1221993](#)
- An EX4300 Virtual Chassis LACP flap is observed after rebooting a master FPC with PDT configurations [PR1301338](#)
- The interface might not work properly after FPC restarts. [PR1329896](#)
- The MAC address assigned to an aggregated Ethernet member interface is not the same as that of its parent aggregated Ethernet interface upon master node removal. [PR1333734](#)
- An EX4600 MC-LAG is observed after the reboot of a VRRP master and backup There are also black holes in traffic to downstream switches. [PR1345316](#)

Platform and Infrastructure

- After access is rejected, the dot1x process might crash due to a memory leak. [PR1160059](#)
- The mismatch of VLAN-ID between an interface IFL and VLAN configuration might result in a traffic black hole. [PR1259310](#)
- MACsec session cannot be recovered after physically flapping one link of an aggregated Ethernet. [PR1283314](#)
- Performing load replace terminal and attempting to replace the interface stanza might terminate the current CLI session and leave the user session hanging. [PR1293587](#)
- You might observe some eswd core files if **apply-groups** is configured under **interface-range**. [PR1300709](#)
- Multicast receiver connected to EX4300 might not be able to get the multicast streaming. [PR1308269](#)
- Traceroute is not working in an EX9200 device for routing instances running on Junos OS Release 17.1R3. [PR1310615](#)
- Autonegotiation is not working as expected between an EX4300 and an SRX5800. [PR1311458](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- IGMP snooping might not learn a multicast router interface dynamically. [PR1312128](#)
- PEM alarms and I2C failures are observed on EX9200 Series. [PR1312336](#)
- The DHCP-security binding table might not get updated. [PR1312670](#)
- Traffic going through an aggregated Ethernet interface might be dropped if there is a mastership change. [PR1327578](#)
- A memory leak is seen for dot1xd. [PR1313578](#)

- The Fan speed might frequently fluctuate between normal and full for MX Series platform. [PR1316192](#)
- The interface with 1G SFP might go down if no-auto-negotiation is configured. [PR1315668](#)
- Replace the **show vlans evpn** command to the **show ethernet-switching evpn** command for the EX9200 line of switches.. [PR1316272](#)
- IGMPv3 on EX4300 does not have the correct outgoing interfaces in the Packet Forwarding Engine that are listed in the kernel. [PR1317141](#)
- The L2cpd core files might be seen if the interface is disabled under VSTP and enabled under RSTP. [PR1317908](#)
- The vmcore might be seen and the device might reboot after the ICL is changed from an aggregated Ethernet to a physical interface. [PR1318929](#)
- High latency might be observed between a master Routing Engine and another FPC. [PR1319795](#)
- VLAN might not be processed, which leads to improper STP convergence. [PR1320719](#)
- Multicast traffic might not be forwarded to one of the receivers. [PR1323499](#)
- MAC learning issue and new VLANs creation failure might happen for some VLANs on an EX4300 platform. [PR1325816](#)
- The L2cpd might create a core file. [PR1325917](#)
- Extra EAP request packets might be sent unnecessarily. [PR1328390](#)
- EX4300 crashes when it receives more than 120kpps ARPs on me0 interface. [PR1329430](#)
- EX Series switches do not send RADIUS request after modifying the interface-range configuration. [PR1326442](#)
- The major alarm **Fan & PSU Airflow direction mismatch** might be seen by removing the management cable. [PR1327561](#)
- The SNMP trap message is always sent out with log about **Fan/Blower OK** on an EX4300 Virtual Chassis switch. [PR1329507](#)
- When exhausting a TCAM table, the filter might be incorrectly programmed. [PR1330148](#)
- The Rpd process crashed and generated core files on the new backup Routing Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler** after disabling NSR and GRES. [PR1330750](#)
- The dot1xd might crash if ports in multi-supplciant mode flaps. [PR1332957](#)
- The interface on which the VSTP is disabled by CLI might stay in the **Discarding** state after rebooting the device. [PR1333684](#)
- STP BPDUs are not sent out on the other active child when the anchor FPC has no active child. [PR1333872](#)
- MQSS errors and alarms might occur when the interface goes down. [PR1334928](#)

- EX9208: vstp vlan all statement has created L2CPD core files are generated during Routing Engine switchover or commit. [PR1341246](#)
- EX4300 storm control does not generate any action log after adding an RTG configuration. [PR1335256](#)
- IGMP packets are forwarded out of an RTG backup interface. [PR1335733](#)
- An L2cpd memory leak appears on EX Series platforms with VoIP configured. [PR1337347](#)
- The **show spanning-tree statistics bridge** command output gives 0 for all VLAN instance IDs. [PR1337891](#)
- MAC source address filter with the configuration statement **accept-source-mac**. does not work if MAC move limit is configured. [PR1341520](#)
- MSTP might not work normally after permitting a commit. [PR1342900](#)
- The filter might not be programmed in the Packet Forwarding Engine even though TCAM entries are available. [PR1345296](#)
- Statistics daemon PFED might generate core files on an upgrade between certain releases. [PR1346925](#)
- After the EX9200 FPC comes online, the other FPC CPU might use 100 percent and has traffic loss for about 30 seconds. [PR1346949](#)
- On EX4300 or EX4600 switches the VLAN translation feature does not work for the control plane traffic. [PR1348094](#)
- On EX4300 platforms, traffic drop might happen if LLC packets are received with DSAP and SSAP as 0x88 and 0x8e. [PR1348618](#)
- Running RSI via console port might cause system crash and reboot. [PR1349332](#)
- EX4600 detects a **LATENCY OVER-THRESHOLD** event with the incorrect value. [PR1348749](#)
- Commit error observed if box is downgraded from Junos OS 18.2/18.3 release to Junos OS Release 17.3R3. [PR1355542](#)
- On EX4300 platforms (Virtual Chassis and standalone) running Junos OS Release 16.1R5 or Junos OS Release 16.1R6, a firewall filter with a syslog option is unable to send syslog messages to the syslog server. [PR1351548](#)
- A high usage chassis alarm in "/var" does not clear from the EX4300 Virtual Chassis when a file is copied from fpc1 (master) to fpc0 (backup). [PR1354007](#)
- The ports using an SFP-T transceiver might still be up after system halt. [PR1354857](#)
- The FPC might crash due to the memory leak caused by the VTEP traffic. [PR1356279](#)
- Some interfaces cannot be added under STP configuration. [PR1363625](#)
- On EX4300/EX4600 platforms, the l2ald process might crash in dot1x scenario. [PR1363964](#)
- Packet Forwarding Engine might crash if encountering frequent MAC move. [PR1367141](#)
- The **request system zeroize** non-interactively might not erase the configuration on EX4300. [PR1368452](#)

Routing Protocols

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

Resolved Issues: 17.4R1

Authentication, Authorization, and Accounting (AAA)

- Dot1x crash on EX4300 can occur when traffic is flooded while a VLAN configuration commit is in progress [PR1293011](#)

Class of Service (CoS)

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

EVPNs

- An l2ald crash occurs with no apparent trigger. [PR1302344](#)

Infrastructure

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

Junos Fusion Enterprise

- CoS shaping is not happening properly according to the configured shaping rate. [PR1268084](#)
- Request chassis satellite beacon functionality to specific SD is not working, causing all the SDs to enable the beacon LED. [PR1272956](#)
- On Dual-AD JFE setup, while applying Routing Engine lo0 filters and setting the cascade port down on AD2, the SD goes to "ProvSessionDown" on that AD2 while it stays online on AD1. [PR1275290](#)
- Issues are seen during conversion from Junos OS release to SNOS. [PR1289809](#)
- VRRP has a split-brain in dual autodiscovery Junos Fusion. [PR1293030](#)
- AD without cascade port cannot reach hosts over ICL link if they are authenticated by dot1x in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)
- The dot1x authentication might fail in a Junos Fusion setup. [PR1299532](#)

- IPv6 multicast is not forwarded over MC-LAG ICL interface until interface toggle. [PR1301698](#)
- Dot1x might crash in a Junos Fusion setup with dual AD. [PR1303909](#)
- All the dot1x sessions are removed when AUTO ICCP link is disabled. [PR1307588](#)
- LACP aggregated Ethernet interfaces go to a down state when performing **commit synchronize**. [PR1314561](#)

Layer 2 Features

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

Network Management and Monitoring

- The **show snmp mib walk** command used for jnxMIMstMstiPortState does not display anything in Junos OS Release 17.1R2 on the EX4600 platform. [PR1305281](#)

Platform and Infrastructure

- Layer 3 protocol packets are not being sent out from the switch. [PR1226976](#)
- PXE unicast ACK packets are dropped on EX4300. [PR1230096](#)
- The EOAM LFM adjacency on EX9200 might flap when the unrelated MIC that is in the same MPC slot is brought online. [PR1253102](#)
- The **interface-range** command cannot be used to set speed and autonegotiation properties for a group of interfaces. [PR1258851](#)
- On EX4300 Virtual Chassis, a 10-Gigabit Ethernet VCP might not get a neighbor after a system reboot. [PR1261363](#)
- CPU utilization for pfex_junos usage might go high if DHCP relay packets are coming continually. [PR1276995](#)
- On EX4300 some functions of IPv6 Router Advertisement Guard do not work. [PR1294260](#)
- **ERROR: /dev/da0s1a is not a JUNOS snapshot** is seen during system startup. [PR1297888](#)
- On EX4300 switches, when unknown unicast ICMP packets are received by an interface, packets are routed, so TTL is decremented. [PR1302070](#)
- On EX4300 Virtual Chassis, the FRU PSU removal and insertion traps are not generated for master or backup FPCs. [PR1302729](#)

Port Security

- MACsec might not work on a 10-Gigabit Ethernet interface after the switch is rebooted. [PR1276730](#)

User Interface and Configuration

- On EX4300, J-Web allows configuration of source-address-filter. [PR1281290](#)

Virtual Chassis

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

VLAN Infrastructure

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

SEE ALSO

New and Changed Features 20
Changes in Behavior and Syntax 27
Known Behavior 31
Known Issues 33
Documentation Updates 41
Migration, Upgrade, and Downgrade Instructions 42
Product Compatibility 43

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 for the EX Series switches documentation.

SEE ALSO

New and Changed Features 20
Changes in Behavior and Syntax 27
Known Behavior 31
Known Issues 33
Resolved Issues 35

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 42](#)

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. 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>.

SEE ALSO

Known Behavior 31
Known Issues 33
Resolved Issues 35
Documentation Updates 41
Product Compatibility 43

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 43

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 20
Changes in Behavior and Syntax 27
Known Behavior 31
Known Issues 33
Resolved Issues 35
Documentation Updates 41

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

- [New and Changed Features | 44](#)
- [Changes in Behavior and Syntax | 45](#)
- [Known Behavior | 45](#)
- [Known Issues | 46](#)
- [Resolved Issues | 47](#)
- [Documentation Updates | 48](#)
- [Migration, Upgrade, and Downgrade Instructions | 48](#)
- [Product Compatibility | 62](#)

These release notes accompany Junos OS Release 17.4R2 for the Junos Fusion Data Center. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

There are no new features in Junos OS Release 17.4R2 for Junos Fusion Data Center.

SEE ALSO

Changes in Behavior and Syntax 45
Known Behavior 45
Known Issues 46
Resolved Issues 47

[Documentation Updates | 48](#)

[Migration, Upgrade, and Downgrade Instructions | 48](#)

[Product Compatibility | 62](#)

Changes in Behavior and Syntax

There are no changes in behavior and syntax for Junos Fusion Data Center in Junos OS Release 17.4R2.

SEE ALSO

[New and Changed Features | 44](#)

[Known Behavior | 45](#)

[Known Issues | 46](#)

[Resolved Issues | 47](#)

[Documentation Updates | 48](#)

[Migration, Upgrade, and Downgrade Instructions | 48](#)

[Product Compatibility | 62](#)

Known Behavior

IN THIS SECTION

- [Junos Fusion Data Center | 46](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for Junos Fusion Data Center.

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

Junos Fusion Data Center

- The license installed will not be deleted, unless it is explicitly deleted using the **request** command. After disabling the cascade port, the license count will be marked as zero only after the satellite information is purged from the neighbor database. Previously, this satellite neighbor information persisted only for 8 minutes; now, neighbor information is being held for 8 hours. This time delay is introduced to avoid repeating the initial recognition of the satellite device for interface-down events. As a workaround, delete the FPC instance for the satellite device to see the license removed for the corresponding satellite device. [PR1294951](#)

SEE ALSO

New and Changed Features 44
Changes in Behavior and Syntax 45
Known Issues 46
Resolved Issues 47
Documentation Updates 48
Migration, Upgrade, and Downgrade Instructions 48
Product Compatibility 62

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.4R2 for Junos Fusion Data Center.

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

SEE ALSO

New and Changed Features 44
Changes in Behavior and Syntax 45
Known Behavior 45
Resolved Issues 47
Documentation Updates 48

[Migration, Upgrade, and Downgrade Instructions | 48](#)[Product Compatibility | 62](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Junos OS Release 17.4R2 | 47](#)
- [Resolved Issues: Junos OS Release 17.4R1 | 47](#)

This section lists the issues fixed in the Junos OS Release 17.4R2 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: Junos OS Release 17.4R2

- The LAG interface might flap if rebooting aggregation device. [PR1315879](#)
- Duplicated packets might be received on the multicast downstream devices and multicast receivers. [PR1316499](#)
- The aggregate device might show a plus sign on the ICL link for a satellite device. [PR1335373](#)
- Aggregation device failure (power off) in a Junos Fusion Data Center causes complete or partial traffic loss for an extended period. [PR1352167](#)

Resolved Issues: Junos OS Release 17.4R1

There are no resolved issues in Junos OS Release 17.4R1 for Junos Fusion Data Center.

SEE ALSO

[New and Changed Features | 44](#)[Changes in Behavior and Syntax | 45](#)[Known Behavior | 45](#)

[Known Issues | 46](#)

[Documentation Updates | 48](#)

[Migration, Upgrade, and Downgrade Instructions | 48](#)

[Product Compatibility | 62](#)

Documentation Updates

This section lists the errata or changes in Junos OS Release 17.4R2 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 | 44](#)

[Changes in Behavior and Syntax | 45](#)

[Known Behavior | 45](#)

[Known Issues | 46](#)

[Resolved Issues | 47](#)

[Migration, Upgrade, and Downgrade Instructions | 48](#)

[Product Compatibility | 62](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 49](#)
- [Preparing the Switch for Satellite Device Conversion | 51](#)
- [Autoconverting a Switch into a Satellite Device | 53](#)
- [Manually Converting a Switch into a Satellite Device | 56](#)
- [Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 58](#)
- [Configuring Satellite Device Upgrade Groups | 59](#)
- [Converting a Satellite Device to a Standalone Device | 61](#)

- Upgrade and Downgrade Support Policy for Junos OS Releases | 61
- Downgrading from Release 17.4 | 61

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

To download and install Junos OS:

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

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

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

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

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

All other customers, use the following command:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Preparing the Switch for Satellite Device Conversion

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

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

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

Customers with EX4300 switches, use the following command:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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 17.4R1 or later, and that the satellite devices are running a compatible conversion release of Junos OS. See [Junos Fusion Hardware and Software Compatibility Matrices](#).

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 to the upgrade group:

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

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

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 17.4R1 or later, and that the switches that will become satellite devices are running a compatible conversion release of Junos OS. See [Junos Fusion Hardware and Software Compatibility Matrices](#).

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 group-name 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 to the upgrade group:

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

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

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 a compatible conversion release of Junos OS. See [Junos Fusion Hardware and Software Compatibility Matrices](#).
- 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.

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.n-signed.tgz** stored in the `/var/tmp/` directory on the switch, where *n* is the spin number:

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

In the event that you need to convert a satellite device to a standalone device, you will need to install a new Junos OS software package on the satellite device and remove it from the Junos Fusion topology. For more information, see [Converting a Satellite Device to a Standalone Device](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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 17.4

To downgrade from Release 17.4 to another supported release, follow the procedure for upgrading, but replace the 17.4 **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 | 44](#)

[Changes in Behavior and Syntax | 45](#)

[Known Behavior | 45](#)

Known Issues 46
Resolved Issues 47
Documentation Updates 48
Product Compatibility 62

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 62

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 44
Changes in Behavior and Syntax 45
Known Behavior 45
Known Issues 46
Resolved Issues 47
Documentation Updates 48
Migration, Upgrade, and Downgrade Instructions 48

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- New and Changed Features | 63
- Changes in Behavior and Syntax | 65
- Known Behavior | 66
- Known Issues | 67
- Resolved Issues | 68
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 70
- Product Compatibility | 75

These release notes accompany Junos OS Release 17.4R2 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

New and Changed Features

IN THIS SECTION

- Release 17.4R2 New and Changed Features | 64
- Release 17.4R1 New and Changed Features | 64

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

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

Release 17.4R2 New and Changed Features

There are no new features in Junos OS Release 17.4R2 for Junos Fusion Enterprise.

Release 17.4R1 New and Changed Features

Junos Fusion Enterprise

- **Cascade port support on EX9200 line cards (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.4R1, interfaces on the following EX9200 line cards can be converted to cascade ports in a Junos Fusion Enterprise topology:

- EX9200-12QS
- EX9200-40XS
- EX9200-40F
- EX9200-40F-M

In a Junos Fusion Enterprise topology, the EX9200 switch acts as the aggregation device. A cascade port is a port on the aggregation device that sends and receives control and network traffic from an attached satellite device.

[See [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#).]

- **Aggregation device support on EX9200 with EX9200-RE2 Routing Engine (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.4, EX9200 switches with the EX9200-RE2 Routing Engine module are supported as aggregation devices in a Junos Fusion Enterprise. The EX9200-RE2 module supports virtual machine (VM) architecture in an EX9200 switch.

[See [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#).]

- **EVPN-MPLS interworking with Junos Fusion Enterprise (EX9200 switches)**—Starting with Junos OS Release 17.4R1, you can use Ethernet VPN (EVPN) to extend your Junos Fusion Enterprise over an MPLS network to a geographically distributed campus or enterprise network.

The EVPN-MPLS interworking feature offers the following benefits:

- Ability to use separate virtual routing and forwarding (VRF) instances to control inter-VLAN routing.
- VLAN translation.

- Default Layer 3 virtual gateway support, which eliminates the need to run such protocols as Virtual Router Redundancy Protocol (VRRP).
- Load balancing to better utilize both links when using EVPN multihoming.
- The use of EVPN type 2 advertisement routes (MAC+IP) reduces the need for flooding domains with ARP packets.

[See [Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG.](#)]

SEE ALSO

Changes in Behavior and Syntax	 65
Known Behavior	 66
Known Issues	 67
Resolved Issues	 68
Documentation Updates	 69
Migration, Upgrade, and Downgrade Instructions	 70
Product Compatibility	 75

Changes in Behavior and Syntax

IN THIS SECTION

- [Junos Fusion Enterprise](#) | [66](#)

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 17.4R2 for Junos Fusion Enterprise.

Junos Fusion Enterprise

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

SEE ALSO

New and Changed Features 63
Known Behavior 66
Known Issues 67
Resolved Issues 68
Documentation Updates 69
Migration, Upgrade, and Downgrade Instructions 70
Product Compatibility 75

Known Behavior

IN THIS SECTION

- [Junos Fusion Enterprise | 66](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- On a Junos Fusion, when using LLDP, the **Power via MDI** and **Extended Power via MDI** TLVs are not transmitted. [PR1105217](#)
- On a Junos Fusion Enterprise, when you issue the **show ethernet-switching table** CLI command, it takes a few minutes to show entries when an extended port receives with the MAC count set to 150K. [PR1117567](#)

- On a Junos Fusion Enterprise, when the satellite devices of a cluster are rebooted, the output of the CLI command **show chassis satellite** shows the Port State of the cascade ports as **Present**. [PR1175834](#)
- On a Junos Fusion Enterprise, in order to use a non-default port as a clustering port in a clustering port policy, the policy must include at least one port that is a default uplink/clustering port for that platform. [PR1241808](#)

SEE ALSO

[New and Changed Features | 63](#)

[Changes in Behavior and Syntax | 65](#)

[Known Issues | 67](#)

[Resolved Issues | 68](#)

[Documentation Updates | 69](#)

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

[Product Compatibility | 75](#)

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.4R2 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 | 63](#)

[Changes in Behavior and Syntax | 65](#)

[Known Behavior | 66](#)

[Resolved Issues | 68](#)

[Documentation Updates | 69](#)

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

[Product Compatibility | 75](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R2 | 68](#)
- [Resolved Issues: 17.4R1 | 68](#)

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

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

Resolved Issues: 17.4R2

- Mirrored packets are dropped if analyzer output extended port is reachable via the ICL link. [PR1211123](#)
- In a Junos Fusion environment the satellite device displays **U-boot** on the LCD screen. [PR1304784](#)
- In a Junos Fusion, a packet loss of 2-3 seconds is seen every 5 minutes. [PR1320254](#)
- On Junos Fusion Enterprise, an SCPD core might be seen on an aggregation device when DACL on 802.1X-enabled port is installed on a single-homed satellite device. [PR1328247](#)
- On Junos Fusion Enterprise, DHCP security binding entries are not synced after the FPC goes offline and comes back online. [PR1332828](#)
- Issue with 802.1X re-authentication. [PR1345365](#)
- A satellite device does not recover PoE after the device is offline for more than 10 minutes and rejoins the aggregation device. [PR1356478](#)
- The Junos Fusion satellite device reboots after an automatic PoE firmware upgrade. [PR1359065](#)
- The ppm-lite process might generate a core file on the Junos Fusion satellite devices. [PR1364265](#)

Resolved Issues: 17.4R1

- On Junos Fusion Enterprise, traffic shaping is not supported on the extended ports. [PR1268084](#)
- On a Junos Fusion Enterprise with dual aggregation devices (ADs), if you apply Routing Engine loopback filters and bring down the cascade port on one of the ADs, the satellite device (SD) on the AD where the cascade port is down goes to ProvSessDown due to a TCP session drop over the ICL interface. [PR1275290](#)

- VRRP has a split-brain state in dual autodiscovery Junos Fusion. [PR1293030](#)
- An aggregation device without a cascade port cannot reach hosts over ICL link if they are authenticated by 802.1X in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)
- The 802.1X authentication might fail in a Junos Fusion setup. [PR1299532](#)
- IPv6 multicast is not forwarded over an MC-LAG ICL interface until the interface is toggled. [PR1301698](#)
- The l2ald process generates a core file with no apparent trigger. [PR1302344](#)
- All 802.1X authentication sessions are removed when the AUTO ICCP link is disabled. [PR1307588](#)
- The dot1x process might generate a core file in a Junos Fusion setup with dual aggregation devices. [PR1303909](#)
- LACP aggregated Ethernet interfaces go to down state when performing **commit synchronize**. [PR1314561](#)

SEE ALSO

New and Changed Features 63
Changes in Behavior and Syntax 65
Known Behavior 66
Known Issues 67
Documentation Updates 69
Migration, Upgrade, and Downgrade Instructions 70
Product Compatibility 75

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 for Junos Fusion Enterprise documentation.

SEE ALSO

New and Changed Features 63
Changes in Behavior and Syntax 65
Known Behavior 66
Known Issues 67
Resolved Issues 68

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading Junos OS on an Aggregation Device | 70
- Upgrading an Aggregation Device with Redundant Routing Engines | 72
- Preparing the Switch for Satellite Device Conversion | 73
- Converting a Satellite Device to a Standalone Switch | 74
- Upgrade and Downgrade Support Policy for Junos OS Releases | 74
- Downgrading from Release 17.4 | 75

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 17.4R2:

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

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

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

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

All other customers, use the following commands, where *n* is the spin number:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.

3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

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

Preparing the Switch for Satellite Device Conversion

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

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

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

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

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

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

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 it from the Junos Fusion topology. For more information, see [Converting a Satellite Device to a Standalone Device](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Downgrading from Release 17.4

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

NOTE: You cannot downgrade more than three releases.

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

To downgrade a Junos Fusion Enterprise from Junos OS Release 17.4R1, follow the procedure for upgrading, but replace the 17.4 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

- [New and Changed Features | 63](#)
- [Changes in Behavior and Syntax | 65](#)
- [Known Behavior | 66](#)
- [Known Issues | 67](#)
- [Resolved Issues | 68](#)
- [Documentation Updates | 69](#)
- [Product Compatibility | 75](#)

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 76](#)
- [Hardware Compatibility Tool | 76](#)

Hardware and Software Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features	 63
Changes in Behavior and Syntax	 65
Known Behavior	 66
Known Issues	 67
Resolved Issues	 68
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 | 77
- Changes in Behavior and Syntax | 78
- Known Behavior | 79
- Known Issues | 80
- Resolved Issues | 82
- Documentation Updates | 83
- Migration, Upgrade, and Downgrade Instructions | 84
- Product Compatibility | 92

These release notes accompany Junos OS Release 17.4R2 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R2 New and Changed Features | 78
- Release 17.4R1 New and Changed Features | 78

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R2 for Junos Fusion Provider Edge.

Release 17.4R2 New and Changed Features

There are no new features in Junos OS Release 17.4R2 for Junos Fusion Provider Edge.

Release 17.4R1 New and Changed Features

Hardware

- **Support for MX204 routers (Junos Fusion Provider Edge)**—Starting in Junos OS Release 17.4R1, you can configure MX204 Universal Routing Platforms as aggregation devices in a Junos Fusion Provider Edge topology. Junos Fusion Provider Edge brings the Junos Fusion technology to the service provider edge. In a Junos Fusion Provider Edge, MX Series routers act as aggregation devices, while EX4300, QFX5100, QFX5110, or QFX5200 switches act as satellite devices.

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

SEE ALSO

Changes in Behavior and Syntax	78
Known Behavior	79
Known Issues	80
Resolved Issues	82
Documentation Updates	83
Migration, Upgrade, and Downgrade Instructions	84
Product Compatibility	92

Changes in Behavior and Syntax

IN THIS SECTION

- [Security](#) | [79](#)

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 17.4R2 for Junos Fusion Fabrics.

Security

- **Support for logging the SSH key changes**—Starting with Junos OS Release 17.4R1, the configuration statement **log-key-changes** is introduced at the `[edit system services ssh]` hierarchy level. When **log-key-changes** is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time **log-key-changes** was enabled. If the **log-key-changes** was never enabled, then Junos OS logs all the authorized SSH keys.

SEE ALSO

[New and Changed Features | 77](#)

[Known Behavior | 79](#)

[Known Issues | 80](#)

[Resolved Issues | 82](#)

[Documentation Updates | 83](#)

[Migration, Upgrade, and Downgrade Instructions | 84](#)

[Product Compatibility | 92](#)

Known Behavior

IN THIS SECTION

- [Junos Fusion Provider Edge | 80](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for Junos Fusion Provider Edge.

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

Junos Fusion Provider Edge

- An L2 filter with multiple terms containing mixed L2 and L3/L4 match conditions do not get programmed a QFX10000 switch as an aggregation device. This is due to an ASIC limitation. [PR1286708](#)
- The FPCs are not online after an image upgrade due to lack of space in the `/var/tmp` directory. [PR1296082](#)
- The `no-mac-learning` and `interface-mac-limit` statements are not supported on extended ports or LAGs of extended ports. [PR1296731](#)
- The CLI `interface-set` command in the firewall filter match condition is not supported on a QFX10000 switch as an aggregation device. [PR1298633](#)
- The policy route action is not supported on interfaces with a `vxlan-vni` configuration along with routing instances. [PR1298683](#)
- The next-ip action for the firewall filters is not supported with an EVPN-VXLAN VNI configuration. [PR1298688](#)
- Configuration synchronization is not triggered when you issue the rollback command on the local aggregation device (AD). [PR1298747](#)

SEE ALSO

New and Changed Features 77
Changes in Behavior and Syntax 78
Known Issues 80
Resolved Issues 82
Documentation Updates 83
Migration, Upgrade, and Downgrade Instructions 84
Product Compatibility 92

Known Issues

IN THIS SECTION

- [Junos Fusion Provider Edge | 81](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R2 for Junos Fusion Provider Edge.

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

Junos Fusion Provider Edge

- The license installed will not be deleted, unless it is explicitly deleted using the **request** command. After disabling the cascade port, the license count will be marked as zero only after the satellite information is purged from the neighbor database. Previously this satellite neighbor information persisted for only 8 minutes; now neighbor information is being held for 8 hours. This time delay is introduced to avoid repeating the initial recognition of the satellite device for interface-down events. **user@host> show configuration | display set | grep et-0/0/30 set groups user-host-grp interfaces et-0/0/30 cascade-port set chassis satellite-management fpc 101 cascade-ports et-0/0/30 set interfaces et-0/0/30 disable {master:0} user@host> show chassis satellite terse**

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

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

Interface	State	Port	Info	System	Name	Model
SW Version et-0/0/30 Dn et-0/0/18 Two-Way et-0/0/18	sd102	QFX5200-32C-32Q	17.4-20170726_common_xxx.0	et-0/0/12 Two-Way et-0/0/50	sd103	QFX5110-48S-4C
17.4-20170726_common_xxx.0 et-0/0/6 Two-Way et-0/1/3	sd100	EX4300-48T	17.4-20170726_common_xxx.0	{master:0} user@host> show system license	License usage:	Licenses

Licenses	Licenses Expiry	Feature name	used	installed	needed	bgp	1	0	1	invalid	SD-QFX5100-48SH-48TH
0	4	0	permanent	Licenses installed:	License identifier:	JUNOSxxxxxx	License version:	4	Software Serial Number:	99999B999999999	Customer ID:

USER-SWITCH Features: SD-QFX5100-48SH-48TH-4PK - SD 4 pack QFX5000-10-JFD permanent {master:0} user@host> show system license usage

Licenses	Licenses Expiry	Feature name	used	installed	needed	bgp	1	0	1	invalid	SD-QFX5100-48SH-48TH
0	4	0	permanent	{master:0} user@host> show system alarms	4 alarms currently active	Alarm time	Class	Description	2017-08-29 13:14:27 UTC	Minor BGP Routing Protocol usage requires a license	2017-08-28 17:25:27 UTC

Major FPC0: PEM 1 Not Powered 2017-08-28 17:25:27 UTC Major FPC Management1 Ethernet Link Down[PR1294951](#)
- Configuration synchronization is not triggered when you issue the rollback command on the local aggregation device (AD). [PR1298747](#)
- When changing fpc slot-id, always delete the old configuration, commit, and then apply the new configuration. Otherwise, sdpd and mib2d might generate core files. Example: (1) Delete chassis satellite-management fpc 101 cascade-ports et-0/0/11 (2) commit (3) Set chassis satellite-management fpc 102 cascade-ports et-0/0/11 (4) commit. [PR1309080](#)

SEE ALSO

New and Changed Features	77
Changes in Behavior and Syntax	78
Known Behavior	79
Resolved Issues	82
Documentation Updates	83
Migration, Upgrade, and Downgrade Instructions	84
Product Compatibility	92

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.4R2 | 82
- Resolved Issues: 17.4R1 | 83

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: 17.4R2

Junos Fusion Provider Edge

- The **show interfaces diagnostics optics satellite** command does not display any outputs. [PR1327876](#)
- High IGMP leave latency with IGMP snooping in an EVPN. [PR1327980](#)
- SSH key-based authentication fails after a reboot if **chassis satellite-management** is configured. [PR1344392](#)

Resolved Issues: 17.4R1

Junos Fusion Provider Edge

- Chassis alarms are not generated after the uplinks are made down from the satellite device. [PR1275480](#)

SEE ALSO

New and Changed Features 77
Changes in Behavior and Syntax 78
Known Behavior 79
Known Issues 80
Documentation Updates 83
Migration, Upgrade, and Downgrade Instructions 84
Product Compatibility 92

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 for Junos Fusion Provider Edge documentation.

SEE ALSO

New and Changed Features 77
Changes in Behavior and Syntax 78
Known Behavior 79
Known Issues 80
Resolved Issues 82
Migration, Upgrade, and Downgrade Instructions 84
Product Compatibility 92

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 84
- Upgrading an Aggregation Device with Redundant Routing Engines | 86
- Preparing the Switch for Satellite Device Conversion | 87
- Converting a Satellite Device to a Standalone Device | 88
- Upgrading an Aggregation Device | 90
- Upgrade and Downgrade Support Policy for Junos OS Releases | 91
- Downgrading from Release 17.4 | 91

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 17.4R2 is different that for earlier Junos OS releases.

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

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

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

For upgrades from Junos OS Release 14.2 and earlier:

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

All other upgrades:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.4R2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

- The switch can only be converted to SNOS 3.1 and higher.
- The switch can be converted to a satellite device if it is in factory-default or it has the **set chassis auto-satellite-conversion** statement in its configuration.

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 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-D43 is named `install-media-pxe-qfx-5-14.1X53-D43.7-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-D43 software image for your platform.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.
7. Copy the software to the routing platform or to your internal software distribution site.
8. Remove the satellite device from the automatic satellite conversion configuration.

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

```
[edit]
```

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

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

```
[edit]
```

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

```
[edit]
```

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

```
[edit]
user@aggregation-device> request chassis satellite install  
/var/tmp/install-media-pxe-qfx-5-14.1X53-D43.7-signed.tgz fpc-slot 101
```

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

```
[edit]
user@aggregation-device> request chassis satellite install  
/var/tmp/jinstall-ex-4300-14.1X53-D43.7domestic-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 17.4R2, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. 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 17.4

To downgrade from Release 17.4 to another supported release, follow the procedure for upgrading, but replace the 17.4 **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 | 77](#)

[Changes in Behavior and Syntax | 78](#)

[Known Behavior | 79](#)

[Known Issues | 80](#)

[Resolved Issues | 82](#)

[Documentation Updates | 83](#)

[Product Compatibility | 92](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 92](#)

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 77
Changes in Behavior and Syntax 78
Known Behavior 79
Known Issues 80
Resolved Issues 82
Documentation Updates 83
Migration, Upgrade, and Downgrade Instructions 84

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

IN THIS SECTION

- New and Changed Features | 93
- Changes in Behavior and Syntax | 126
- Known Behavior | 136
- Known Issues | 142
- Resolved Issues | 157
- Documentation Updates | 199
- Migration, Upgrade, and Downgrade Instructions | 200
- Product Compatibility | 207

These release notes accompany Junos OS Release 17.4R2 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R2-S2 New and Changed Features | 94
- Release 17.4R2 New and Changed Features | 94
- Subscriber Management and Services | 96
- Release 17.4R1 New and Changed Features | 96

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R2 for the MX Series routers.

Release 17.4R2-S2 New and Changed Features

Routing Protocols

- **Support for creating IS-IS topology independent LFA for prefix-SIDs learned from LDP mapping server**
—Starting in Junos OS Release 17.4R2-S2, you can configure a point of local repair to create a topology independent loop-free alternate backup path for prefix-SIDs derived from LDP mapping server advertisements in an IS-IS network. In a network configured with segment routing, IS-IS uses the LDP mapping server advertisements to derive prefix-SIDs. LDP Mapping server advertisements for IPv6 are currently not supported.

To attach flags to LDP mapping server advertisements, include the **attached** statement at the **[edit routing-options source-packet-routing mapping-server-entry *mapping-server-name*]** hierarchy level.

Release 17.4R2 New and Changed Features

EVPNs

- **EVPN proxy ARP and ARP suppression without IRB interfaces (MX Series routers with MPCs, EX9200 switches)**—MX Series routers and 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.

Starting with Junos OS Release 17.4R2, these features no longer require the configuration of an integrated routing and bridging (IRB) interface on the PE device. Now, any interface configured on a PE device can deliver ARP requests from local remote customer edge (CE) devices. ARP proxy and ARP suppression are not supported on remote CE's.

In addition, you can now control the following aspects of the media access control (MAC)-IP address bindings database on a PE device:

- The maximum number of MAC-IP address entries in the database
- The amount of time a locally learned MAC-IP address binding remains in the database

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

Interfaces and Chassis

- **Enhancement to increase the threshold of corrected single-bit errors (MPC7E, MPC8E, MPC9E on MX Series)**—In Junos OS Release 17.4R2, the threshold of corrected single-bit error is increased from 32 to 1024, and the alarm severity is changed from Major to Minor for those error messages. There is no operational impact upon corrected single bit errors. Also, a log message is added to display how many single-bit errors have been corrected between the reported events as follows:

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

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

[See [Alarm Overview](#).]

MPLS

- **Interoperability of segment routing with LDP (MX Series)**—In an LDP network with gradual deployment of segment routing, some devices may not support segment routing, which can cause interoperability issues in the network. Starting in Junos OS Release 18.2R1, and 17.4R2, you can use OSPF or ISIS to enable segment routing devices to operate with the LDP devices that are not segment routing capable.

To implement this feature using OSPF, an extended prefix link-state advertisement (LSA) with Range type, length, and value (TLV) for all the LDP prefixes is generated, and mapping routes corresponding to the prefix is installed in the inet.3 and mpls.0 routing tables.

To implement this feature using ISIS, a server-client configuration is required under protocols ISIS and LDP, respectively, and routes from the inet.3 or inet.0 routing tables are used for stitching of segment routing LSP with an LDP LSP and vice-versa.

[See [LDP Mapping Server for Interoperability of Segment Routing with LDP Overview](#) .]

Restoration Procedures and Failure Handling

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (MX Series)**—In Junos OS Release 17.4R2, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode. The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

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

Software Installation and Upgrade

- **ZTP support is added for MX VM host platforms (MX Series)**—In Junos OS Release 17.4R2, ZTP, which automates the provisioning of the device configuration and software image with minimal manual intervention, is supported on MX Series VM hosts. When you physically connect a supported device to the network and boot it with a factory configuration, the device attempts to upgrade the Junos OS software image automatically and autoinstall a configuration provided on the DHCP server.

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

Subscriber Management and Services

- **Controlling search behavior for address allocation from linked pools (MX Series)**—Starting in Junos OS Release 17.4R2, you can use the **linked-pool-aggregation** statement at the **[edit access]** hierarchy level to change how addresses are allocated from linked IP address pools. When you configure the statement, addresses can be assigned from a later pool in the chain before an earlier pool is depleted. When the statement is not configured, IP addresses are assigned contiguously, so that all addresses are allocated from the matching pool and then the first pool in the chain before addresses are assigned from a linked pool.

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

Release 17.4R1 New and Changed Features

Hardware

- **Support for the CFP2-DCO-T-WDM-1 transceiver on the MPC5E-100G10G MPC and the MIC6-100G-CFP2 MIC (MX Series)**—Starting in Junos OS Release 17.4R1, you can install the CFP2-DCO-T-WDM-1 transceiver on the MPC5E-100G10G MPC and the MIC6-100G-CFP2 MIC (installed on the MX2K-MPC6E MPC). The CFP2-DCO-T-WDM-1 transceiver is a 100-Gigabit digital pluggable CFP2 digital coherent optical module.

The CFP2-DCO-T-WDM-1 transceiver supports the following:

- International Telecommunication Union (ITU)-standard OTN performance monitoring and alarm management
- 100-Gigabit quadrature phase shift keying (QPSK) with differential encoding mode and soft-decision forward error correction (SD-FEC)
- proNX Service Manager (PSM)
- Junos OS YANG extensions
- Firmware upgrade

[See [2x100GE + 4x10GE MPC5E](#) and [100-Gigabit Ethernet MIC with CFP2](#).]

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

- **Periodic refresh of authorization profile on TACACS+ server (MX Series)**—Starting with Junos OS Release 17.4R1, periodic refresh of the authorization profile that is received from the TACACS server is supported. The authorization profile that is configured for the user on the TACACS server is sent to the Junos OS device after the user is successfully authenticated. The authorization profile is stored locally on the Junos OS device. With the periodic refresh feature, the authorization profile is periodically fetched from the TACACS server to refresh the authorization profile that is stored locally. User authorization is reevaluated using the refreshed authorization profile.

[See [Configuring Periodic Refresh of the TACACS+ Authorization Profile](#).]

- **Enhanced TACACS+ support for the dedicated management instance (MX Series and vMX)**—Starting in Junos OS Release 17.4R1, TACACS+ behavior is enhanced to support the management interface in a non-default virtual routing and forwarding (VRF) instance. For supported platforms, TACACS+ packets can now be sent to the server successfully even with the **management-instance** configuration statement enabled. The dedicated management instance was released in Junos OS Release 17.3R1.

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

Class of Service (CoS)

- **New criteria introduced for when to throttle logins based on CoS queues (MX Series)**—Starting in Junos OS Release 17.4R1, new criteria are incorporated into the throttling decision for subscriber access. CoS resources (queues) are taken into account when deciding whether to avoid accepting new subscriber logins when there are insufficient CoS resources. To support this behavior, a new CLI configuration statement (**high-cos-queue-threshold**) is introduced to enable usage of CoS resource monitoring in throttling decisions and to set the threshold of CoS resource usage above which new logins are not permitted. A new show command (**show system resource-monitor ifd-cos-queue-mapping fpc**) is also introduced.

[See “Throttling Subscriber Load Based on CoS Resource Capacity” in [Resource Monitoring for Subscriber Management and Services Overview](#), [high-cos-queue-threshold](#), and [show system resource-monitor ifd-cos-queue-mapping fpc](#).]

- **Support for static Type of Service (ToS)/Traffic Class on GRE tunnels (MX Series)**—Starting in Junos OS Release 17.4R1, MPCs on MX Series routers support the setting of a static ToS/Traffic Class value in the IPv4/IPv6 header, respectively, of a GRE tunnel. You can set a **traffic-class** value at the **interfaces gre-interface-name unit logical-unit-number tunnel** hierarchy level. The value represents the entire 8-bit differentiated services (DS) field in the IP header, ranging from **0-255**, and should be chosen based on the desired DSCP/IP precedence value. For example, if a DSCP value of **111000** is desired, then configure the **traffic-class** value to be **224** (corresponding to **111000 00**).

[See [traffic-class \(Tunnels\)](#).]

Dynamic Host Configuration Protocol (DHCP)

- **Support for RADIUS reauthentication of DHCPv4 and DHCPv6 clients (MX Series)**—Starting in Junos OS Release 17.4R1, reissue of the RADIUS authentication request [**access-request**] is supported as an alternative to RADIUS Change of Authorization (CoA) to change subscriber session characteristics.

Reauthentication is enabled by the following triggers:

- The **reauthenticate remote-id-mismatch** command specifies reauthentication when there is a remote-id change in the option of the control packet (for example, RENEW, REBIND, DISCOVER, or SOLICIT) for the DHCPv4 or DHCPv6 client.
- The **reauthenticate lease-renewal** command specifies reauthentication for a renew or rebind.
- The **reauthentication-on-renew** command indicates to reauthentication on every renew or rebind from the DHCPv4 or DHCPv6 client.
- If both **reauthenticate lease-renewal** and the **Reauthentication-on-renew** are specified for a given subscriber, the Junos DHCPD (DHCP daemon) requests reauthentication from the RADIUS server every time the DHCP client sends a DHCP renew request. If the **reauthentication-on-renew** vendor-specific attribute (VSA) is disabled, then behavior reverts to **reauthenticate lease-renewal** configuration.
- If both **reauthenticate lease-renewal** and the **reauthentication-on-renew** VSA are enabled for a given subscriber
 - Junos OS DHCPD requests reauthentication from the RADIUS server every time the DHCP client sends a DHCP renew request (as **reauthentication-on-renew** VSA is enabled).
 - If the client sends a discover or solicit with DHCP options indicating a service plan change (different remote-id), Junos DHCPD will request reauthentication (as Junos OS DHCPD configuration reauthenticates on remote-id mismatch).
 - If the client sends a discover or solicit with DHCP options indicating No service plan change (same remote-id), Junos OS DHCPD will not request reauthentication (as the discover or solicit are not renews, and there is no remote-id mismatch).
 - If the reauthentication-on-renew VSA is disabled, then Junos OS DHCPD only reauthenticates when there is a renew, discover or solicit with a remote-id change (service plan change).

[See [RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCPv4 and DHCPv6 Subscribers Overview](#).]

- **Support for forward-only action for DHCP relayed traffic with unknown DHCP server address (MX Series)**—Starting in Junos OS Release 17.4R1, forward-only action for DHCP relayed traffic is supported with unknown DHCP server address. Administrator is able to configure for which servers (clients are binding) they need to have relay subscriber entry, apply dynamic profile, policies and more, and for whom they want to forward only. This feature also introduces configuration for processing destination address, **option-54** and **option-2** on DHCP relay.

DHCP relay agent entry will be useful for authentication, authorization, accounting, applying filtering, QoS to client, processing of options specified in the packet. Customer networks can contain non-customer controlled bindings for which the customer does not want these relay agent entry functionalities. Hence relay agent subscriber entries are not created for non-customer controlled bindings.

Prior to 17.4R1 Release, subscriber entry creation constituted of Junos OS DHCPD (DHCP daemon) memory resources, session database resources, authentication procedure, accounting, dynamic profile instantiation, dynamic interface creation, firewall, CoS association, and more. If a customer network has some non-customer controlled traffic for which a relay agent entry is created then it would be an unnecessary utilization of resources, and an incorrect association of profiles.

[See [Forward-only Action for DHCPv4 and DHCPv6 Relay Traffic with Unknown DHCP Server Address Overview](#).]

EVPNs

- **Support for duplicate MAC address detection and suppression (MX Series)**—When a MAC address relocates, PE devices can converge on the latest location by using sequence numbers in the extended community field. Misconfigurations in the network can lead to duplicate MAC addresses. Starting in Junos OS Release 17.4R1, Juniper supports duplicate MAC address detection and suppression.

You can modify the duplicate MAC address detection settings on the router by configuring the detection window for identifying duplicate MAC address and the number of MAC address moves detected within the detection window before duplicate MAC detection is triggered and the MAC address is suppressed. In addition, you can also configure an optional recovery time that the router waits before the duplicate MAC address is automatically unsuppressed.

To configure duplicate MAC detection parameters, use the **detection-window**, **detection-threshold**, and **auto-recovery-time** statements at the `[edit routing instance routing-instance-name protocols evpn duplicate-mac-detection]` hierarchy level.

To clear duplicate MAC suppression manually, use the **clear evpn duplicate-mac-suppression** command.

[See [Overview of MAC Mobility](#).]

- **Enhancements to composite next hops (MX Series)**—Starting in Junos OS Release 17.4R1, you can enable dynamic list next hop. By enabling this feature, when the link fails between the CE device and a multihomed PE device in EVPN active-active multihoming, the routing process daemon (rpd) dynamically modifies the next-hop list without first removing the next-hop entry and creating a new entry. This reduces mass MAC route withdrawals and improves convergence and performance.

To enable dynamic list next hop, include the **dynamic-list-next-hop** statement at the `[edit routing-options forwarding-table]` hierarchy level. If you perform a unified ISSU to upgrade your device from an OS release prior to Junos OS Release 17.4R1, you must upgrade both the Routing engine and the backup Routing Engine before enabling dynamic list next hop.

[See [Configuring Dynamic List Next Hop](#).]

- **EVPN active standby multihoming to a single PE device (MX Series)**—Starting in Junos OS Release 17.4R1, Juniper supports EVPN active-standby multihoming. When you configure a protect (backup)

interface for a primary interface on the same PE router, the protect interface becomes active when the primary interface fails and network traffic is switched to the protect interface.

To configure a protect interface, include the **protect-interface** statement at the **[edit interfaces]** hierarchy level for a routing instance, EVPN bridge domain, and the EVPN protocol under EVPN VPWS routing instance.

[See [Configuring EVPN Active-Standby Multihoming to a Single PE.](#)]

- **SPRING support for EVPN (MX Series)**—Starting in Junos OS Release 17.4R1, Junos OS supports using Source Packet Routing in Networking (SPRING) as the underlay transport in EVPN. SPRING tunnels enable routers to steer a packet through a specific set of nodes and links in the network.

To configure SPRING, use the **source-packet-routing** statement at the **[edit protocols isis]** hierarchy level.

[See [Understanding Source Packet Routing in Networking \(SPRING\).](#)]

- **EVPN-MPLS interworking with MC-LAG (MX Series routers)**—Starting with Junos OS Release 17.4R1, you can use Ethernet VPN (EVPN) to extend your MC-LAG network over an MPLS network. Typically, an MC-LAG network is extended to a data center network or geographically distributed campus or enterprise network.

The EVPN-MPLS interworking feature offers the following benefits:

- Ability to use separate virtual routing and forwarding (VRF) instances to control inter-VLAN routing.
- VLAN translation.
- Default Layer 3 virtual gateway support, which eliminates the need to run such protocols as Virtual Router Redundancy Protocol (VRRP).
- Load balancing to better utilize both links when using EVPN multihoming.
- The use of EVPN type 2 advertisement routes (MAC+IP) reduces the need for flooding domains with ARP packets.

[See [Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG.](#)]

General Routing

- **Support for PTP over IPv4 and hybrid mode on 10GE, 40G, and 100GE WAN ports (MX10003, MX204)**—Starting in Junos OS Release 17.4R1, the 10GE, 40G, and 100GE WAN ports support the following features:
 - **PTP over IPV4 Encapsulation**—In PTP over IPv4, the nodes (master and slave devices) participate in unicast negotiation in which the slave node is provisioned with the IP address of the master node and requests unicast messages to be sent to it from the master node.
 - **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.

[See [Understanding Hybrid Mode](#)]

- **PHY timestamping support**—PHY timestamping is the timestamping of the 1588 event packets at the PHY. Timestamping the packet in the PHY eliminates the noise or the Packet Delay Variation (PDV) that is introduced by the Packet Forwarding Engine (PFE).

[See [phy-timestamping](#)]

- **Support for PTP over Ethernet, hybrid mode, and G.8275.1 profile (MPC7E-10G, MPC7E-MRATE, MPC8E, MPC9E)**—Starting in Junos OS Release 17.4R1, MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E 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. The IEEE 1588 standard defines two types of multicast MAC addresses 01-80-C2-00-00-0E (link local multicast) and 01-1B-19-00-00-00 (standard Ethernet multicast) for PTP over Ethernet operations.
 - **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.

[See [Understanding Hybrid Mode](#)]

- **G.8275.1 profile**— The G.8275.1 is a PTP profile for applications requiring accurate phase and time synchronization. It supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support and is based on the second version of PTP defined in (IEEE 1588). You can configure the G.8275.1 profile by including the **profile-type g.8275.1** statement at the **[edit protocols ptp]** hierarchy level.

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

High Availability (HA) and Resiliency

- **Hardware resiliency support (MX204)**—Starting in Junos OS Release 17.4R1, MX204 routers support the resiliency feature, which includes hardware failure and fault handling. Resiliency on an MX204 enhances its debugging capability in the case of hardware failure of any of its components. For example, the resiliency feature enables the router to recover from inter-integrated circuit (I2C) failure, and improves its voltage monitoring, temperature monitoring, PCI Express error handling and reporting, DRAM single-bit and multibit error checking and correction (ECC), and SSD SMART attribute monitoring capabilities.
- **L2VPN connection last uptime preserved after switchover (MX Series)**—Starting in Junos OS Release 17.4R1, the **show l2vpn connections** command displays the last time that the L2VPN connection was in the **Up** condition, and this value persists after a switchover or unified ISSU.

[See [show l2vpn connections](#)]

Interfaces and Chassis

- **Support for JNP-MIC-100G MIC with MACsec support on MPC8E and MPC9E (MX2000 line of routers)**—Starting in Junos OS Release 17.4R1, the JNP-MIC-100G MIC extends Media Access Control Security (MACsec) capabilities on MPC8E and MPC9E MPCs installed in MX2010, MX2020, and MX2008 routers. Each MPC supports two JNP-MIC-100G MICs. On an MPC8E, each MIC supports 48 10-Gigabit Ethernet, 12 40-Gigabit Ethernet, or 4 100-Gigabit Ethernet MACsec-capable interfaces, or a combination. On an MPC9E, each MIC supports 48 10-Gigabit Ethernet, 12 40-Gigabit Ethernet, or 8 100-Gigabit Ethernet MACsec-capable interfaces, or a combination. Support for MACsec increases security within a data center and also provides secured connectivity between data centers.

[See [Understanding Media Access Control Security \(MACsec\) on MX Series Routers](#) on basic information about MACsec.]

- **MX204 Universal Routing Platform**—Starting in Junos OS Release 17.4R1, the MX204 Universal Routing Platform is added to the MX Series family of routers. The MX204 is a highly dense 1 rack unit (1 U) chassis that offers speeds of up to 400 Gbps and can be used as a preaggregation chassis and in mobile backhaul scenarios.

The MX204 router is a fixed-configuration router, and supports one fixed Routing Engine. The MX204 has four rate-selectable ports that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet ports, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports. The four rate-selectable ports support QSFP28 and QSFP+ transceivers, whereas the eight 10-Gigabit Ethernet ports support SFP+ transceivers.

[See [MX204 Router Rate-Selectability Overview](#) and [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router](#).]

- **MX204 router supports port LED for 4xQSFP ports**—Starting in Junos OS Release 17.4R1, port LED is supported on MX204 routers. LEDs on the interface cards display the status of the ports. In MX204 router, there are four port LEDs per port. Each port provides an individual status LED with four states signaled by the color/LED state: OFF, GREEN, AMBER, RED

[See [MX204 LED Scheme Overview](#).]

- **Support for power management and environmental monitoring in MX204 routers**—Starting with Junos OS Release 17.4R1, Junos OS chassis management software for the MX204 routers provides enhanced environmental monitoring and power management. MX204 routers have one Routing Engine and MPC. The MPC has one Packet Forwarding Engine that supports a bandwidth up to 400 Gbps. The MPC supports two fixed Physical Interface Card (PIC) where PIC0 comprises four QFP28 ports and PIC1 comprises 8 XSFPP ports. The power supply and the fan trays are upgradable. The cooling system contains three fan assemblies with two fans in each assembly. The chassis has two redundant power supply modules (PSM): DC PSM and AC PSM. Each of these PSMs deliver 650 W of power.
- **Software feature support on MX204 routers**— Starting with Junos OS Release 17.4R1, Junos OS supports the MX204 Universal Routing Platform (model number: JNP204 [MX204]). The MX204 chassis is a monolithic system containing in-built MPC with one EA ASICs (operating in 400G mode) and supports 2 fixed port PICs (4xQSFP28 PIC and 8xSFPP PIC). All the devices including Packet Forwarding Engines,

WAN interfaces are managed by the CPU subsystem (8 core Broadwell CPU). There are no fabric ASICs in the MX204 router.

The MX204 router is a 400G capable monolithic platform having a single board with 8 Core Intel Broadwell CPU with 1 EA Packet Forwarding Engine ASICs connected to each other back to back.

The following features are supported on MX204 platform:

- 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)
- Layer 2 protocols
- Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs
- Layer 3 routing protocols and MPLS
- Layer 3 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
- **Support for MACsec PSK keychain (MX2010, MX2020)**—Starting in Junos OS Release 17.4R1, MX2020 and MX2010 supports Key Agreement Protocol Fail Open mode. The MACsec PSK chains hitless rollover feature is documented in Junos OS Release 17.4R1, but not supported.
- **Strong encryption for configuration secrets (MX2020, MX2010, and MX2008 routers)**—Starting in Junos OS Release 17.4R1, the MX2020, MX2010 and MX2008 routers support strong encryption for configuration secrets. To use strong encryption for your configuration secrets, you need to configure a master password. The master password enables you to derive an encryption key that you use with the AES256-GCM standard to encrypt configuration secrets. This new encryption method uses the \$8\$ formatted strings.
[See [Hardening Shared Secrets in Junos OS.](#)]
- **Support for pre-FEC BER monitoring when using the CFP2-DCO-T-WDM-1 transceiver (MX Series)**—Starting in Junos OS Release 17.4R1, you can monitor the condition of an OTN link by using the pre-forward error correction (pre-FEC) bit error rate (BER) when using the CFP2-DCO-T-WDM-1 transceiver.
[See [Understanding Pre-FEC BER Monitoring and BER Thresholds.](#)]

Junos OS XML API and Scripting

- **Automation script library additions and upgrades (MX Series)**—Starting in Junos OS Release 17.4R1, devices running Junos OS include new and upgraded Python modules as well as upgraded versions of Junos PyEZ and libslax. On-box Python automation scripts can use features supported in Junos PyEZ Release 2.1.4 and earlier releases to perform operational and configuration tasks on devices running Junos OS. Python automation scripts can also leverage new on-box Python modules including **ipaddress**, **jxmlease**, **pyang**, **serial**, and **six**, as well as upgraded versions of existing modules. In addition, SLAX automation scripts can include features supported in libslax release 0.22.0 and earlier releases.

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

Layer 2 Features

- **Support for new configuration statements to perform qualified MAC learning on inner VLAN tags (MX Series)** —Starting with Junos OS Release 17.4R1, MX series routers support the following new configuration statements:
 - **deep-vlan-qualified-learning *vlan_tag_number*** at the **[edit interfaces unit *logical_unit_number*]** hierarchy level to enable qualified mac-learning on the third VLAN tag (innermost) of an ingress 3-tagged packet, without any kind of implicit VLAN manipulation. If the packet has two tags, MAC learning happens on the second VLAN. If the ingress packet has more than three tags, all tags beyond the third tag are treated as part of data. For bidirectional traffic flow, **input-vlan-map pop** has to be configured.
 - **vlan-id inner-all** at the **[edit routing instances *instance_name*]** to enable qualified MAC learning on the second (inner) VLAN tag of an ingress double tagged packet, without removing the first (outer) tag implicitly. For a single-tagged packet, qualified MAC learning happens on VLAN 4096. If the ingress packet has more than two tags, all tags beyond the second tag are treated as part of data.

Logical Systems

- **Storm control In logical systems (MX Series)**—Starting in Junos OS Release 17.4R1, support for storm control has been added for logical systems running on MX Series devices. With storm control, you can set a traffic threshold and enable traffic monitoring so that whenever the threshold is reached, the router automatically starts dropping broadcast, unknown unicast, and/or multicast (BUM) packets in order to prevent a “storm” of packets from proliferating on the network.

To use this feature with a given logical system, create a storm control profile at the **[edit logical-systems *name* forwarding-options storm-control-profiles *name*]** hierarchy level.

[See [Understanding Storm Control for Managing Traffic Levels](#).]

- **EVPNs on logical systems (MX Series)**—Starting with Junos OS Release 17.4R1, support for Ethernet Virtual Private Network (EVPN) has been added for logical systems running on MX Series devices. Running EVPN in a logical system provides the same options and performance as running EVPN on a physical system, which adheres to the standards described in RFC 7432. Note that Graceful Restart, Graceful Routing Engine switchover (GRES), and nonstop active routing (NSR) are not supported.

Configure EVPN on a logical system at the `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols evpn]` level.

[See [EVPN Overview](#) .]

Management

- **Support for IS-IS sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can export data for the IS-IS routing protocol through the Junos Telemetry Interface. Only gRPC streaming is supported. To export statistics for IS-IS, include the `/network-instances/network-instance[name_'instance-name']/protocols/protocol/isis/levels/level/` and `/network-instances/network-instance[name_'instance-name']/protocols/protocol/isis/interfaces/interface/levels/level/` set of paths. Use the `telemetrySubscribe` RPC to specify telemetry parameters and provision the sensor. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

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

- **Support for Packet Forwarding Engine traffic sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can export Packet Forwarding Engine traffic statistics through the Junos Telemetry Interface. Both UDP and gRPC are supported. This sensor tracks reporting of Packet Forwarding Engine statistics counters and provides visibility into Packet Forwarding Engine error and drop statistics. The resource name for the sensor is `/junos/system/linecard/packet/usage/`. The OpenConfig path is `/components/component/subcomponents/subcomponent[name='FPC<id>:NPU<id>']/properties/property/`, where NPU refers to the Packet Forwarding Engine. To provision the sensor to export data through gRPC, use the `telemetrySubscribe` RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the `[edit services analytics]` hierarchy level.

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

- **Enhancements to LSP events sensor for Junos Telemetry Interface (MX Series)** —Starting with Junos OS Release 17.4R1, telemetry data streamed through gRPC for LSP events and properties is reported separately for each routing instance. To export data for LSP events and properties, you must now include `/network-instances/network-instance[name_'instance-name']/` in front of all supported paths. For example, to export LSP events for RSVP signaling protocol attributes, use the following path: `/network-instances/network-instance[name_'instance-name']/mpls/signaling-protocols/rsvp-te/`. Use the `telemetrySubscribe` RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Enhancement to BGP sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can specify to export the number of BGP peers in a BGP group for telemetry data exported through gRPC. To export the number of BGP peers for a group, use the following OpenConfig path: `/network-instances/network-instance[name_'instance-name']/protocols/protocol/`

`bgp/peer-groups/peer-group[name_ 'peer-group-name']/state/peer-count/`. The BGP peer count value exported reflects the number of peering sessions in a group. For example, for a BGP group with two devices, the peer count reported is 1 (one) because each group member has one peer. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

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

- **Broadband edge (BBE) telemetry sensors (MX Series routers)**—In Junos OS Release 17.4R1, support is expanded for BBE telemetry sensors. These sensors are used to proactively manage a broadband network gateway (BNG) and are configured using both Junos Telemetry Interface (JTI) and gRPC streaming. The new sensors are grouped in the following functional areas:

- Chassis and system extensions
- Authentication, authorization, and accounting (AAA)
- Dynamic Host Configuration Protocol (DHCP)
- Packet Forwarding Engine resource monitoring

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

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

- **Enhancements to MPLS sensor for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can export statistics for MPLS through the Junos Telemetry Interface in the following categories:

- Shared Risk Link Groups (SRLGs)
- Traffic engineering global attributes
- Traffic engineering interface attributes

Additional RSVP signaling protocol attributes, such as counters and interfaces, that were not previously available are also supported. Only gRPC streaming is supported.

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

- **Support for bidirectional authentication for gRPC for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can configure gRPC to require client authentication as well as server authentication. Previously, only the client initiating an RPC request was able to authenticate the server; that is, a Juniper device using SSL certificates. To enable bidirectional authentication, include the **mutual-authentication** statement at the `[edit system-services extension-service request-response grpc ssl]` hierarchy level. You must also configure and reference a certificate-authority profile. Include the **certificate-authority profile name** statement at the `[edit system services extension-service request-response grpc ssl]` hierarchy level. For **profile-name**, include the name of **certificate-authority** profile configured at the `[edit security pki ca-profile]` hierarchy level. This profile is used to validate the certificate provided by the client.

NOTE: MX80 and M104 routers do not support gRPC.

[See [gRPC Services for Junos Telemetry Interface](#).]

- **Support for BGP routing table sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can provision Junos Telemetry Interface sensors to export data for BGP routing tables (RIBs) for IPv4 and IPv6 routes. Each address family supports exporting data for five different tables. Only gRPC streaming is supported.

The tables are:

- **local-rib**—Main BGP routing table for the main routing instance.
- **adj-rib-in-pre**—NLRI updates received from the neighbor before any local input policy filters have been applied.
- **adj-rib-in-post**—Routes received from the neighbor eligible for best-path selection after local input policy filters have been applied.
- **adj-rib-out-pre**—Routes eligible for advertising to the neighbor before output policy filters have been applied.
- **adj-rib-out-post**—Routes eligible for advertising to the neighbor after output policy filters have been applied.

To stream data for the main BGP routing table for IPv4 routes, include the **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/loc-rib/** set of paths. To stream data for the main BGP routing table for IPv6 routes, include the **/bgp-rib/afi-safis/afi-safi/ipv6-unicast/loc-rib/** set of paths.

For the neighbor BGP routing tables for IPv4 routes, include the following sets of paths:

- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-in-pre/**
- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-in-post/**
- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-out-pre/**
- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-out-post/**

To stream data for IPv6 routes, change **ipv4-unicast** to **ipv6-unicast** in any of the paths.

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

- **Junos Telemetry Interface support for virtual MX Series routers (vMX)**—Starting with Junos OS Release 17.4R1, the Junos Telemetry Interface is supported on vMX routers. The Junos Telemetry Interface enables you to provision sensors to stream telemetry data for network elements without involving polling. All sensors supported on MX Series routers are supported on vMX routers, except for the following: fabric statistics and high queue-scale statistics. To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. For UDP streaming, all parameters

are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

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

- **Multiservices MPC (MS-MPC) support for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, interfaces configured on MS-MPCs support the Junos Telemetry Interface, which enables you to provision sensors to stream telemetry data for network elements without involving polling. Only streaming through UDP is supported. gRPC streaming is not supported. To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level.

Only the following sensors are supported on MS-MPCs:

- Firewall filters
- CPU memory
- NPU memory
- NPU memory utilization
- Physical interfaces

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

- **Junos Telemetry Interface support on MX2008 routers (MX Series)**—Starting with Junos OS Release 17.4R1, the Junos Telemetry Interface, which enables you to provision sensors to stream telemetry data for network elements without involving polling, is supported on MX2008 routers. Both UDP and gRPC streaming are supported. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

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

- **Support for dynamic tunnel statistics for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can export counter statistics for Packet Forwarding Engine dynamic tunnels. Both UDP and gRPC streaming are supported. The resource string to export statistics is `/junos/services/ip-tunnel/usage/`. The OpenConfig path is `/junos/services/ip-tunnel[name='tunnel-name']/usage/counters[name='counter-name']`. All parameters for UDP sensors are configured at the **[edit services analytics]** hierarchy level. To export data through gRPC, use the **telemetrySubscribe** RPC. To stream data through gRPC, you must also download the OpenConfig for Junos OS module. MX80 and MX104 routers only support UDP streaming. They do not support gRPC.

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

- **Support for bypass LSP statistics for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.4R1, you can export statistics for bypass label-switched paths (LSPs). Previously, only statistics

for the primary LSP path were exported. The ability to export bypass LSP statistics helps to monitor the efficiency of global convergence when the bypass LSP is used to carry traffic during a link or node failure.

Statistics are exported for the following:

- Bypass LSP originating at the ingress router of the protected LSP
- Bypass LSP originating at the transit router of the protected LSP
- Bypass LSP protecting the transit LSP as well as the locally originated LSP

When the bypass LSP is active, traffic is exported both on the bypass LSP and the ingress (protected) LSP. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module. You must also include the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level.

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

- **Support for multiple, smaller configuration YANG modules (MX Series)**—Starting in Junos OS Release 17.4R1, the YANG module for the Junos OS configuration schema is split into a root configuration module that is augmented by multiple, smaller modules. The root configuration module comprises the top-level configuration node and any nodes that are not emitted as separate modules. Separate, smaller modules augment the root configuration module for the different configuration statement hierarchies. Smaller configuration modules enable YANG tools and utilities to more quickly and efficiently compile and work with the modules, because they only need to import the modules required for the current operation.

[See [Understanding the YANG Modules That Define the Junos OS Configuration](#).]

MPLS

- **Support for Ethernet CCC encapsulation on pseudowire subscriber transport and services logical interfaces (MX Series)**—Starting in Junos OS Release 17.4R1, you can configure the same Ethernet circuit cross-connect (CCC) encapsulation (also known as VLAN-ID) on pseudowire subscriber transport and service logical interface. The primary reason for Ethernet CCC encapsulation on the pseudowire subscriber transport is for interoperability between the existing access node and aggregation node in the network.

Prior to Release 17.4R1, Junos OS does not allow the same VLAN-ID to be configured on more than one logical interface under the same pseudowire subscriber physical interface. To establish a pseudowire connection from an access node or aggregation node to a Multi-Service Edge (MSE) node, **ignore-encapsulation-mismatch** configuration statement is used. This statement is a Junos OS feature and the access or aggregation device may not support this feature. To overcome this restriction, you can configure same VLAN-ID on transport and service logical interface.

[See [VLAN CCC Encapsulation on Transport Side of Pseudowire Subscriber Logical Interfaces Overview](#).]

- **Support for static adjacency segment identifier for IS-IS (MX Series)**—Starting with Junos OS Release 17.4R1, you can configure static adjacency segment ID (SID) labels for an interface. You can configure

two IPv4 adjacency SIDs (protected and unprotected), IPv6 adjacency SIDs (protected and unprotected) per level per interface. You can use the same adjacent SID for multiple interfaces by grouping a set of interfaces under an interface-group and configuring the adjacency-segment for that interface-group. For static adjacency SIDs, the labels are picked from either a static reserved label pool or from segment routing global block (SRGB).

[See [Static Adjacency Segment Identifier for ISIS](#).]

- **Support for static adjacency segment identifier for aggregate Ethernet member links using single-hop static LSP (MX Series)**—Starting with Junos OS Release 17.4R1, you can configure a transit single-hop static label switched path (LSP) for a specific member link of an aggregated Ethernet (AE) interface. A static labeled route is added with next-hop pointing to the AE member link of an aggregate interface. Label for these routes is picked from the segment routing local block (SRLB) pool of the configured static label range. This feature is supported for AE interfaces only.

A new **member-interface** CLI command is added under the **next-hop** configuration at the **[edit protocols mpls static-label-switched-path lsp-name transit]** hierarchy to configure the AE member interface name. The static LSP label is configured from a defined static label range.

[See [Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-hop Static LSP](#).]

- **Support for segment routing statistics (MX Series Routers with MPCs and MICs)**—Starting in Junos OS Release 17.4R1, the traffic statistics in a segment routing (SR) network can be recorded in an OpenConfig compliant format for Layer 3 interfaces. The statistics is recorded for the Source Packet Routing in Networking (SPRING) traffic only, excluding RSVP and LDP-signaled traffic, and the family MPLS statistics per interface is accounted for separately. The SR statistics also includes SPRING traffic statistics per link aggregation group (LAG) member, and per service identifier (SID).

To enable recording of SR statistics, include the **sensor-based-stats (per-interface-per-member-link <ingress | egress> | per-sid ingress)** statement at the **[edit protocol isis source-packet-routing]** hierarchy level.

[See [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

- **IPv6 next-hop support for static egress LSPs (MX Series)**—Starting in Junos OS Release 17.4R1, 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 \(Protocols MPLS\)](#).]

Operation, Administration, and Maintenance (OAM)

- **Support for Inline performance monitoring (MX Series Routers)**—Starting in Junos OS Release 17.4R1, Junos OS supports inline mode for MEF 35 compliant service OAM performance monitoring on MX Series routers. Performance monitoring functions include measurement of Ethernet frame delay, frame

delay variations, frame loss, and availability of service. By default, performance monitoring packets are handled by the CPU of a line-card, such as Modular Port Concentrator (MPC). Enabling inline mode of performance monitoring delegates the processing of the protocol data units (PDUs) to the forwarding ASIC (that is, to the hardware). By enabling inline mode of performance monitoring, the load on the CPU of the line-card is reduced and you can configure an increased number of performance monitoring sessions and achieve maximum scaling for service OAM performance monitoring sessions.

Inline mode of performance monitoring is supported only for proactive mode of frame delay measurement (Two-way Delay Measurements) and synthetic loss measurements (SLM) sessions. Performance monitoring functions configured using the iterator profile (CFM) are referred to as proactive performance monitoring. Inline mode of performance monitoring for frame loss measurement using service frames (LM) is not supported.

NOTE: MPC3E (MX-MPC3E-3D) and MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE) do not support inline performance monitoring. User-defined Data TLV is not supported if you have configured inline mode of performance monitoring. Also, only 12 history records per PM sessions are supported.

- **Support for CFM monitoring on pseudowire services interfaces(MX Series Routers)**—Starting in Junos OS Release 17.4R1, Junos OS supports IEEE 802.1ag connectivity fault management (CFM) on pseudowire service interfaces. Pseudowire service interfaces support configuring of subscriber interfaces over MPLS pseudowire termination. Termination of subscriber interfaces over PW enables network operators to extend their MPLS domain from the Access/Aggregation network to the service edge and use uniform MPLS label provisioning for a larger portion of their network.

To enable support for CFM on pseudowire service interfaces, configure maintenance intermediate points (MIPs) on the pseudowire service interfaces. The CFM MIP session is supported only on the pseudowire services interface and not on the pseudowire services tunnel interface.

Routing Protocols

- **Support for timing and synchronization on MX204 Routers**—Starting in Junos OS Release 17.4R1, MX204 routers support the following timing and synchronization features:
 - **SyncE support with ESMC**—Synchronized Ethernet with Ethernet Synchronization Message Channel (ESMC) is supported as per the ITU G.8264 specification. ESMC is a logical communication channel. It transmits synchronization status message information, which is the quality level of the transmitting Synchronous Ethernet equipment clock, by using ESMC protocol data units.
 - **PTP support**—Precision Time Protocol (PTP), also known as IEEE 1588v2, is a packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks. IEEE 1588 PTP (Version 2) clock synchronization standard is a highly precise protocol for time synchronization that synchronizes clocks in a distributed system. The time synchronization is achieved through packets that are transmitted and received in a session between a master clock and a slave clock. One-step clock mode operation for the master clock is supported.

- **BITS (T1/E1) Interface support**—BITS support for input and output on T1/E1 framed and 2.048MHz unframed clock input.
- **GPS external clock interface and TOD support**—GPS input and output support for 1 MHz/5 MHz/10 MHz and PPS signal
- **Support for importing IGP topology information into BGP-LS (MX Series)**—Starting in Junos OS Release 17.4R1, you can import interior gateway protocol (IGP) topology information into BGP-Link State (BGP-LS) in addition to RSVP-traffic engineering (RSVP-TE) topology information through the `Isdist.0` routing table. This allows you to monitor both IGP and traffic engineering topology information.

To install IGP topology information into the traffic engineering database, use the **set igp-topology** configuration statement at the **[edit protocols isis traffic-engineering]** and **[edit protocols ospf traffic-engineering]** hierarchy levels. To import IGP topology information into BGP-LS from `Isdist.0`, use the **set bgp-ls** configuration statement at the **[edit protocols mpls traffic-engineering database import igp-topology]** hierarchy level.

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

- **BGP supports segment routing policy for traffic engineering (MX Series)**—Starting in Junos OS Release 17.4R1, a BGP speaker supports traffic steering based on a segment routing policy at ingress routers. The controller can specify a segment routing policy consisting of multiple paths to steer labeled or IP traffic. The segment routing policy adds an ordered list of segments to the header of a packet for traffic steering. Static policies can be configured at ingress routers to allow routing of traffic even when the link to the controller fails.

To enable BGP IPv4 segment routing traffic engineering capability for an address family, include the **segment-routing-te** statement at the **[edit protocols bgp family inet]** hierarchy level.

[See [Understanding Ingress Peer Traffic Engineering for BGP SPRING.](#)]

- **Support for EVPN control plane with VXLAN data plane encapsulation (MX150)**—Starting in Junos OS Release 17.4R1, MX150 routers, powered with vMX, decouples an underlay network from the tenant overlay network with VXLAN. By using a Layer 3 IP-based underlay coupled with a VXLAN-EVPN overlay, you can deploy larger networks than those possible with traditional Layer 2-based networks. With overlays, end-points (servers and virtual machines) can be placed anywhere in the network and remain connected to the same logical Layer 2 network. One of the key benefits is that virtual topology can be decoupled from the physical topology.
- **Support for Layer 2 VXLAN gateway (MX150)**—Starting in Junos OS Release 17.4R1, MX150 routers, powered with vMX, that support a Virtual Extensible LAN (VXLAN) can function as a hardware virtual tunnel endpoint (VTEP). In this role, the Juniper Networks device encapsulates in VXLAN packets Layer 2 Ethernet frames received from software applications that run directly on a physical server. The VXLAN packets are tunneled over a Layer 3 fabric. Upon receipt of the VXLAN packets, software VTEPs in the virtual network de-encapsulate the packets and forward the packets to virtual machines (VMs).
- **Support for BGP advertising aggregate bandwidth across external BGP links for load balancing (MX Series)**—Starting in Junos OS Release 17.4R1, BGP uses a new link bandwidth extended community, **aggregate-bandwidth**, to advertise aggregated bandwidth of multipath routes across external links. BGP

calculates the aggregate of multipaths that have unequal bandwidth allocation and advertises the aggregated bandwidth to external BGP peers. A threshold to the aggregate bandwidth can be configured to restrict the bandwidth usage of a BGP group. In earlier Junos OS releases, a BGP speaker receiving multipaths from its internal peers advertised the link bandwidth associated with the active route. To advertise aggregated bandwidth of multipath routes and to set a maximum threshold, configure a policy with **aggregate-bandwidth** and **limit bandwidth** actions at the **[edit policy-options policy-statement name then]** hierarchy level.

[See [Advertising Aggregate Bandwidth Across External BGP Links for Load Balancing Overview](#).]

- **Topology-independent loop-free alternate for IS-IS (MX Series)**—Starting in Junos OS Release 17.4R1, topology-independent loop-free alternate (TI-LFA) with segment routing provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. You can enable TI-LFA for IS-IS by configuring the **use-post-convergence-lfa** statement at the **[edit protocols isis backup-spf-options]** hierarchy level. TI-LFA provides protection against link failure, node failure, and failures of fate-sharing groups.

You can enable the creation of post-convergence backup paths for a given interface by configuring the **post-convergence-lfa** statement at the **[edit protocols isis interface interface-name level level]** hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

You can enable **node-protection** and/or **fate-sharing-protection** mode for a given interface at the **[edit protocols isis interface interface-name level level post-convergence-lfa]** hierarchy level. To use a particular fate-sharing group as a constraint for the fate-sharing-aware post-convergence path, you need to configure the **use-for-post-convergence-lfa** statement at the **[edit routing-options fate-sharing group group-name]** hierarchy level.

[See [Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS](#).]

- **Support for trace route through an interface through the inactive routes (MX Series)**—Starting in Junos OS Release 17.4R1, you can configure traceroute to send out packets through an inactive next hop by specifying the **traceroute next-hop address** to a destination through an inactive next hop.

[See [Traceroute for Inactive Interface](#).]

- **Support for network instance based BGP configuration (MX Series)**—Starting in Junos OS Release 17.4R1, you can configure BGP in a specific network instance. After the network instance is configured, you will be prompted with options for BGP configuration such as global bgp, neighbor bgp, and so on. See [Mapping OpenConfig Network Instance Commands to Junos Operation](#).
- **Support for EBGp route server (MX Series)**—Starting in Junos OS Release 17.4R1, BGP feature is enhanced to support EBGp route server functionality. A BGP route server is the external BGP (EBGP) equivalent of an internal IBGP (IBGP) route reflector that simplifies the number of direct point-to-point EBGp sessions required in a network. EBGp route server propagates unmodified BGP routing information between external BGP peers to facilitate high scale exchange of routes in peering points such as Internet Exchange Points (IXPs). When BGP is configured as a route server, EBGp routes are propagated between peers unmodified, with full attribute transparency (NEXT_HOP, AS_PATH, MULTI_EXIT_DISC, AIGP, and Communities).

The BGP JET **bgp_route_service.proto** API has been enhanced to support route server functionality as follows:

- Program the EBGp route server.
- Inject routes to the specific route server RIB for selectively advertising it to the client groups in client-specific RIBs.

The BGP JET **bgp_route_service.proto** API includes a peer-type object that identifies individual routes as either EBGp or IBGP (default).

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

Services Applications

- **Inline video monitoring for IPv4-over-MPLS flows on M10003 and MX204 routers**—Starting in Junos OS Release 17.4R1, MX10003 and MX204 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](#)]

- **Port Control Protocol support (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.4R1, the Port Control Protocol (PCP) feature is supported on MS-MPCs and MS-MICs. Before Junos OS Release 17.4R1, PCP was supported only on MS-DPC service cards. PCP provides a mechanism to control the forwarding of incoming packets by upstream devices such as NAT44 and firewall devices, and a mechanism to reduce application keepalive traffic. Use PCP in the context of both carrier-grade NATs and small NATs (for example, residential NATs). PCP allows hosts to operate servers for a long time (for example, a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a carrier-grade NAT operated by their Internet service provider. PCP allows applications to create mappings from an external IP address and port to an internal IP address and port.

PCP on the MS-MPC and MS-MIC supports only NAPT44. PCP with DS-Lite is not supported on the MS-MPC and MS-MIC.

[See [Port Control Protocol Overview](#), [Configuring Port Control Protocol](#), and [Example: Configuring Port Control Protocol with NAPT44](#).]

- **Increased sampling rate for inline Junos Traffic Vision (MX Series)**—Starting in Junos OS Release 17.4R1, the sampling rate that you can configure for inline Junos Traffic Vision (inline active flow monitoring) using the **rate number** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet [inet6])** and **[edit forwarding-options sampling input]** hierarchy levels is increased from 65,535 to 16,000,000. This functionality is supported for Inline Active Flow Monitoring on MX Series and vMX routers. This feature is also supported for PIC-based flow monitoring on MX Series routers with certain MPCs. If a line card does not support a sampling rate higher than 65,535, such as an I-chip-based DPC, the maximum sampling rate is limited to 65,535.

[See [Example: Configuring Flow Monitoring on MS-MIC and MS-MPC](#).]

- **Support for Diffie-Hellman group15, group16, and group24 for IKE SAs and IPsec policies (MX Series)**—Starting in Junos OS Release 17.4R1, Diffie-Hellman group15, group16, and group24 for IKE security associations (SAs) and IPsec policies are supported.

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

- **Port forwarding (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.4R1, support for port forwarding is extended to the MS-MPC and MS-MIC. Port forwarding allows the destination address and port of a packet to be changed to reach the correct host in a Network Address Translation (NAT) gateway. The translation facilitates reaching a host within a masqueraded, typically private, network based on the port number on which the packet was received from the originating host. Port forwarding allows remote computers, such as public machines on the Internet, to connect to a nonstandard port (port other than 80) of a specific computer within a private network. An example of this type of destination is the host of a public HTTP server within a private network. You can also configure port forwarding without translating a destination address.

[See [Port Forwarding Overview](#).]

- **Support for 100,000 simultaneous RPM probes from RPM clients for offload RPM (MX Series)**—Starting in Junos OS Release 17.4R1, you can enable the application of optimized CLI configuration in the offload-RPM scale configuration and the existing legacy RPM clients supported on MS-MIC and MS-MPC by entering the **rpm-scale** statement at the **[edit services rpm probe probe-owner]** hierarchy level and at the **[edit groups group-name services rpm]** hierarchy level.

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

- **Support for CoS revert and direction awareness on services interfaces (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.4R1, you can configure a services interface CoS rule to store the DSCP and forwarding class of a packet that is received in the match direction of the rule; this stored DSCP and forwarding class are then applied to packets that are received in the reverse direction of the same session. You can also configure a service set to create a CoS session when a packet is first received in the wrong match direction for a CoS rule; this results in the CoS rule values being applied as soon as a packet in the correct match direction is received.

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

- **DS-Lite support on MS-MPCs and MS-MICs (MX Series routers)**—Starting in Junos OS Release 17.4R1, the MS-MPC and MS-MIC support dual-stack lite (DS-Lite). DS-Lite employs IPv4-over-IPv6 tunnels to cross an IPv6 access network to reach a carrier-grade IPv4-IPv4 NAT. This facilitates the phased introduction of IPv6 on the Internet by providing backward compatibility with IPv4.

Prior to Junos OS Release 17.4R1, DS-Lite was supported on the MX Series only on MS-DPCs.

DS-Lite running on MS-MPCs or MS-MICs does not support the following features, which are supported on MS-DPCs:

- ALGs
- Limitations per subnet

- Clearing NAT mappings and flows for a specific subscriber, for a basic bridging broadband device (B4), or for a specific service set
- Port Control Protocol

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

- **IPsec NAT-T Support (MX Series)**—Starting in Junos OS Release 17.4R1, NAT-T is supported for IKEv1 and IKEv2. Junos OS Release 17.4R1 also supports UDP encapsulation and decapsulation for IKE and ESP packets by specifying **disable-natt** at the `[edit services ipsec-vpn]` hierarchy levels. NAT-T is enabled by default.

[See [disable-natt \(Services IPsec VPN\)](#).]

- **Multiple syslog servers support (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.4R1, you can commit multiple syslog hosts (up to four) under the `[edit services service-set service-set-name]` hierarchy level.

[See [Configuring System Logging for Service Sets](#).]

- **Support for inline NAT and FlowTapLite on MPC7E, MPC8E, and MPC9E (MX Series)**—Starting in Junos OS Release 17.4R1, you can configure inline NAT and FlowTapLite on the following Modular Port Concentrators: MPC7E, MPC8E, and MPC9E.

[See [Inline Network Address Translation Overview for MPCs](#) and [Configuring FlowTapLite](#).]

- **Support for NAT64 with deterministic IP address and port mapping (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.4R1, there is support for deterministic NAT64 mapping on the MS-MPC and MS-MIC. Deterministic NAT mapping ensures that a given internal IP address and port are always mapped to the same external IP address and port range, and the reverse mapping of a given translated external IP address and port are always mapped to the same internal IP address. Deterministic NAT mapping eliminates the need for logging address translations.

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

- **Support for inline video monitoring for IPv6 flows (MX Series)**—Starting in Junos OS Release 17.4R1, MX Series routers support the inline video monitoring of IPv6 flows and IPv6-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](#).]

- **Support for disabling the filtering of HTTP traffic with an embedded IP address belonging to a blacklisted domain (MX Series)**—Starting in Junos OS Release 17.4R1, you can disable the filtering of HTTP traffic that contains an embedded IP address (for example, `http://10.1.1.1`) belonging to a blacklisted domain name in the URL filter database. To disable the filtering, include the **disable-url-filtering** statement at the `[edit services url-filter profile profile-name template template-name]` hierarchy level when you are configuring URL filtering. However, if the embedded IP address is explicitly identified in the blacklisted URL filter database, then the traffic is still filtered.

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

Software Defined Networking (SDN)

- **Support for YANG-based abstraction to orchestrate GNFs (MX480, MX960, MX2010, MX2020)**—Starting with Junos OS Release 17.4R1, Junos supports YANG-based abstraction to orchestrate guest network functions (GNFs), using single touchpoint. In the single touchpoint method, the SDN controller (for example, OpenDaylight or ODL) communicates only with the base system (BSYS). The BSYS receives the RPC requests from the ODL controller, parses the RPC, and then forwards the adequate RPC to the JDM (based on scripts available at the BSYS). After receiving the response from the JDM, the BSYS parses and forwards the response back to the ODL.

NOTE: Junos Node Slicing also supports management of GNF life cycle using the dual touchpoint method. In this method, ODL sends RPCs to, and receive responses from, JDM and BSYS separately. To enable dual touch point, you just need to mount both BSYS and Juniper Device Manager (JDM) on ODL.

[See [Setting Up YANG-Based Abstraction to Orchestrate GNFs.](#)]

- **Unified ISSU support for Junos Node Slicing (MX480, MX960, MX2010, MX2020)**—Starting with Junos OS Release 17.4R1, Junos Node Slicing supports unified ISSU. ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Now, users with administrator rights can perform unified ISSU on the BSYS, (the base system in a Junos Node Slicing setup) and the guest network functions (GNF) separately. Also, users can run unified ISSU on each GNF independently, without affecting other GNFs.

NOTE: The multi-version software support limitations (such as version difference limits) are also applicable to unified ISSU upgrade.

[See [Understanding the Unified ISSU Process.](#)]

- **Multi-Version software support for Junos Node Slicing (MX480, MX960, MX2010, MX2020)**—Starting from Junos OS Release 17.4R1, Junos Node Slicing supports multi-version software compatibility, enabling the BSYS to interoperate with a guest network function (GNF), which runs a Junos OS version that is higher than the software version on the BSYS. This feature supports a deviation of up to two versions between GNF and BSYS. That is, the GNF software can be up to two versions higher than the BSYS software. However, for this feature to work, both BSYS and GNF must meet a minimum version requirement of Junos OS Release 17.4R1.

NOTE: The multi-version software compatibility support is limited to major releases only.

[See [Understanding Multi-Version Software Compatibility.](#)]

- **Improved debugging ability and serviceability for JDM (MX480, MX960, MX2010, MX2020)**—Starting with Junos OS release 17.4R1, improved debugging ability and serviceability are provided for Juniper Device Manager (JDM). The following are the key capabilities supported:
 - JDM-JDM keepalive to monitor reachability of the peer JDM, and to provide failover in case one of the JDM instances (running on server 0 and server 1) goes down.
 - A new **force** option under the CLI command **request virtual-network-functions** to overwrite a VNF image. Example: **request virtual-network-functions vnf-name add-image image-name force**
 - New CLI command, **show version vnf vnf-name**, to show the version details of the guest network functions (GNFs).
 - Dedicated interfaces for JDM and VNF management.

Configuring JDM on the x86 Servers

- **Abstracted Fabric interface for Junos Node Slicing (MX480, MX960, MX2010, MX2020)**—Starting with Junos OS Release 17.4R1, Junos Node Slicing supports Abstracted Fabric (AF) interface, a pseudointerface that represents the behavior of a first class Ethernet interface. An AF interface is created on a GNF to enable it to communicate with the peer GNF when the two GNFs are configured to be connected to each other. The AF interface facilitates routing control and management traffic between GNFs. You can create or delete AF interface from the BSYS. AF interfaces support the following protocol families: inet, inet6, mpls, ccc, and iso.

NOTE: Most of the Layer 1 features and a few of the Layer 2 and Layer 3 features are disabled on AF interfaces.

[See [Abstracted Fabric Interface](#)]

- **Software Support for Junos Node Slicing (MX480, MX960, MX2010, MX2020)**—Starting from Junos OS Release 17.4R1, Junos Node Slicing supports the following software features:
 - BNG
 - Business PE router
 - L2VPN or EVPN PE router
 - Multicast
 - Junos Telemetry Interface—An MX Series router in the BSYS mode provides full-fledged JTI support. However, guest network functions (GNFs) provide limited support for JTI (only physical and logical interfaces statistics for FPCs owned by GNFs are available through gRPC).
- **Support for OpenDaylight (ODL) controller on MX Series routers**—Starting with Junos OS Release 17.4R1, MX Series routers support OpenDaylight (ODL) controller (Carbon release). The ODL controller, or ODL platform, provides a southbound Network Configuration Protocol (NETCONF) connector API, which uses NETCONF and YANG models to interact with a network device. You can use the ODL

controller to carry out configuration changes in MX Series routers, and orchestrate and provision the routers. Also, ODL controller enables you to execute Remote Procedure Calls (RPCs) to MX Series routers to get state information.

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

Software Installation and Upgrade

- **Support for unified ISSU on MX Series routers with MPC7E-MRATE, MPC7E-10G, MX2K-MPC8E, and MX2K-MPC9E (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Release 17.4R1, Junos OS supports unified in-service software upgrade (ISSU) on MX Series routers with MPC7E-MRATE, MPC7E-10G, MX2K-MPC8E, and MX2K-MPC9E.

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

[See [Getting Started with Unified In-Service Software Upgrade](#)]

- **Support for Zero Touch Provisioning (ZTP) (MX150)**—Starting in Junos OS Release 17.4R1, MX150 routers, powered with vMX, support zero touch provisioning. Zero touch provisioning enables you to provision new routers in your network automatically either by executing a script file or by loading a configuration file. In either case, the information is detected in a file on the Dynamic Host Control Protocol (DHCP) server. When you physically connect the MX150 router to the network and boot it with a default configuration, it attempts to upgrade the Junos OS Software automatically using information detected on the DHCP server. If you do not configure the DHCP server to provide this information, the MX150 router boots with the pre-installed software and default configuration.
- **Support for unified ISSU on the CFP2-DCO-T-WDM-1 transceiver (MX Series)**—Starting in Junos OS Release 17.4R1, unified in-service software upgrade (unified ISSU) is supported on the CFP2-DCO-T-WDM-1 transceiver when the transceiver is installed on the MPC5E-100G10G MPC or the MIC6-100G-CFP2 MIC (installed on the MX2K-MPC6E MPC).

[See [Getting Started with Unified In-Service Software Upgrade.](#)]

Subscriber Management and Services

- **Support for static subscriber daemon gaps for Gx/Gy support (MX Series)**—Starting in Junos OS Release 17.4R1, support for usage based billing are added using the Gy interface for static subscribers. The **service-profile** is added to the **static-subscribers** to apply services for all static subscribers at the hierarchy level **[edit system services static-subscribers group group-name]**.

[See [Subscribers on Static Interfaces Overview.](#)]

- **DHCP session liveness detection based on ARP and neighbor discovery packets (MX Series)**—Starting in Junos OS Release 17.4R1, you can configure bidirectional Layer 2 liveness detection for directly connected DHCPv4 and DHCPv6 subscribers using ARP packets for v4 and neighbor discovery (ND) packets for v6. You can configure Layer 2 liveness detection for both DHCP local server and DHCP relay clients. This method of liveness detection enables the host and the broadband network gateway (BNG) separately to determine the validity and state of the DHCP client session and to clean up inactive sessions.

The liveness detection send functionality enables the BNG to determine client session state based on the host response to request packets the BNG sends at a configurable interval. The liveness detection receive functionality enables the client host to determine session state based on the BNG response to ARP or ND packets sent by the client to the BNG.

Layer 2 liveness detection (AR/ND) and Bidirectional Forwarding Detection (BFD) are mutually exclusive.

[See [DHCP Liveness Detection Overview](#).]

- **RADIUS-sourced DHCPv4 and DHCPv6 Options support for single and dual-stack sessions (MX Series)**—Starting in Junos OS Release 17.4R1, for DHCP dual-stack session subscribers, the DHCPv4 option values are saved in the **SDB_DHCP_OPTIONS** session database (SDB) attribute. Likewise, for DHCPv6 subscribers, option values are saved in the **SDB_DHCPV6_OPTIONS** SDB attribute. However, for single-stack sessions (DHCP or DHCPv6), the DHCP option values for both IPv4 and IPv6 subscribers will be saved in **SDB_DHCP_OPTIONS** SDB attribute.

For both single and dual-stack sessions, DHCPv4 header is saved in the **SDB_DHCP_HEADER** and DHCPv6 header in the **SDB_DHCPV6_HEADER** SDB attributes.

The option values and header values received in DHCPv4 discover and DHCPv6 solicit messages are stored in respective SDBs and thus get populated in the new vendor specific attributes (VSAs). These VSAs are then sent to RADIUS server for authentication. The RADIUS server decodes the options, authenticates the client, and sends the RADIUS-sourced DHCP options back to the DHCP server. The DHCP server copies the RADIUS-sourced DHCP options, and also adds the DHCP server-sourced options to the packet and sends the response back to the client.

[See [Dedicated Session Database and Vendor-Specific Attributes for DHCPv4 and DHCPv6 Subscribers Overview](#).]

- **Appending subscriber information to redirect URLs (MX Series)**—Starting in Junos OS Release 17.4R1, you can append information about the subscriber retrieved from the subscriber session database when the redirect URL is returned to the HTTP client. You can configure the attributes at the **[edit services captive-portal-content-delivery]** hierarchy. Only the following attributes are supported: subscriber IP or IPv6 address, NAS IP address, requested URL, NAS port ID, MAC address, subscriber session ID, and username.

NOTE: This feature is already supported for Routing Engine based and Multiservices Modular PIC Concentrator (MS-MPC) based converged captive-portal-content-delivery (CPCD). From 17.4R1 onward, it is supported for Routing Engine based and MS-MPC based static CPCD.

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

- **Enhancements to share CPE parameters between broadband network gateway (BNG) and RADIUS server (MX Series)**—Starting in Junos OS Release 17.4R1, the following enhancements are made to facilitate better communication between the broadband network gateway (BNG) and the RADIUS server:

- CPE parameters such as DHCPv4 (VSA 26-208) and DHCPv6 (VSA 26-209) packet headers are shared between the broadband network gateway (BNG) and the RADIUS server.
- A new VSA 26-207 is introduced that facilitates the exchange of DHCPv6 options with the RADIUS server, thereby ensuring that VSA 26-55 is dedicated to the exchange of DHCPv4 options.
- A new statement, **family-state-change-immediate-update**. When configured at the **[edit access profile]** hierarchy level, the DHCP (both DHCPv4 and DHCPv6) server sends an immediate interim accounting report to the RADIUS server when the second family (IPv4 or IPv6) is activated or the first family gets deactivated.
- A new VSA 26-210 is added to convey the reason for the accounting-request message in the start and interim accounting request packets sent to the RADIUS server. This helps the RADIUS server to determine the reason of the start and interim accounting that is being sent.

[See [Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview](#).]

- **Virtual broadband network gateway support (MX150)**—Starting in Junos OS Release 17.4R1, MX150 routers, powered with vMX, support most of the subscriber management features available with Junos OS Release 17.4 on vMX to provide a virtual broadband network gateway on MX150 routers. vBNG runs on vMX, so it has similar exceptions; the following subscriber management features available on vMX are not supported for vBNG:

- High availability features such as hot-standby backup for enhanced subscriber management and MX Series Virtual Chassis.

To deploy a vBNG instance, you must purchase the following vBNG license:

- vBNG subscriber scale license for one of these tiers: Introductory, Preferred, or Elite.

- **Support for Broadband Edge on MX204 router**—Starting in Junos OS Release 17.4R1, MX204 supports the next-generation broadband edge software architecture for wireline subscriber management. With enhanced subscriber management, you can take advantage of optimized scaling and performance for configuration and management of dynamic interfaces and services for subscriber management.
- **New criteria introduced for when to throttle logins based on CoS queues (MX Series)**—Starting in Junos OS Release 17.4R1, new criteria are incorporated into the throttling decision for subscriber access. CoS resources (queues) are taken into account when deciding whether to avoid accepting new subscriber logins when there are insufficient CoS resources. To support this behavior, a new CLI configuration statement (**high-cos-queue-threshold**) is introduced to enable usage of CoS resource monitoring in throttling decisions and to set the threshold of CoS resource usage above which new logins are not permitted. A new show command (**show system resource-monitor ifd-cos-queue-mapping fpc**) is also introduced.
- **Improved multicast performance with distributed IGMP (MX Series)**—Starting in Junos OS Release 17.4R1, both dynamic and static interfaces support distributed Internet Group Management Protocol (IGMP). Distributed IGMP moves IGMP processing from the Routing Engine and distributes it across multiple Modular Port Concentrators (MPCs) on the Packet Forwarding Engine for improved performance and decreases join and leave latency.

To enable distributed IGMP on static interfaces, include the **distributed** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.

To enable it on dynamic interfaces, include the **distributed** statement at the **[edit dynamic-profiles *profile-name* protocols igmp interface \$junos-interface-name]** hierarchy level.

You must also enable enhanced IP networking services at the **[edit chassis network-services enhanced-ip]** 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 Distributed IGMP](#) .]

- **Support for expanded traffic rate adjustment for DSL access lines (MX Series)**—Starting in Junos OS Release 17.4R1, the traffic rate adjustment feature is expanded to support PPPoE intermediate agent (PPPoE-IA) tags by processing the Vendor-Specific-Tags TLV in PADI and PADO packets received from the access node. Now both PPPoE subscriber connections (terminated and tunneled) and ANCP-triggered Layer 2 wholesale service connections are subject to the same class and quality-of-service management transformations.

Configuration for traffic rate adjustment and reporting for both AAA and CoS is moved to the new **[edit system access-line]** hierarchy level. In earlier releases, DSL line traffic rate adjustment is available only for the ANCP agent and uses statements at the **[edit protocols ancp]** and **[edit protocols ancp qos-adjust]** hierarchy levels.

[See [Traffic Rate Reporting and Adjustment by the ANCP Agent](#) and [Setting a Global Adjustment Factor per DSL Subscriber Line for ANCP Agent-Reported Traffic Rates](#).]

- **Displaying accurate subscriber accounting statistics (MX Series)**—Starting in Junos OS Release 17.4R1, you can enable the router to display accurate subscriber accounting statistics for dynamic interfaces by including the **actual-transit-statistics** statement in the dynamic profile that creates the interface. The aggregate statistics counters show the subscriber traffic bytes and packets arriving on and leaving from the interface; these are the same traffic values reported to RADIUS. The counters exclude overhead byte adjustments, dropped or discarded packets, and control packets. When enabled, use the **show subscribers id accounting-statistics** command to display counts for the specified subscriber session and the **show subscribers interface accounting-statistics** command to display counts for all subscriber sessions on the specified interface.

[See [Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI](#).]

- **Automatic 64-bit mode and maximum configuration database size (MX Series)**—Starting in Junos OS Release 17.4R1, when enhanced IP network services and enhanced subscriber management are enabled and a Routing Engine in the system has at least 32 GB of RAM, subscriber management daemons on

that Routing Engine run in 64-bit mode. For consistent operation, all Routing Engines in the system must have the same amount of memory.

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

- **DSL line attributes support for L2TP LNS (MX Series)**—Starting in Junos OS Release 17.4R1, an MX Series router configured as an LNS can process subscriber access line information that it receives from the LAC. This information includes access line attributes conveyed in ICRQ messages, initial Tx/Rx connect speeds (AVP 24/38) in ICCN messages, and connect speed updates in CSUN messages. The rate information enables CoS shaping on the subscriber session to be more accurate, but updates are subject to CoS adjustment control profiles. You can configure processing for information received from all LACs, or for only LACs you specify by address.

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

- **Enhancement to Gx-Plus Application (MX Series)**—Starting in Junos OS Release 17.4R1, the following enhancements to the Gx-Plus client application on the BNG are available:
 - When a monitored service is deactivated separate from a subscriber logout, the CCR-U indicates that the service is no longer active and includes the service's usage data.
 - The router updates the monitoring key and threshold values when they are received in a RAR message from the PCRF.
 - A CCR-U is sent to the PCRF after the router sends an RAA message in response to an RAR message that requests service activations or deactivations.
 - When the PCRF returns threshold values that are lower than the current values, the new threshold becomes the sum of the current value and the returned value.
 - The PCEF has default minimum threshold values. If the change between the current value and the value returned by the PCRF is less than the minimum value, then the new value is adjusted to the minimum.
 - The CCR-I message includes the Diameter AVP Subscription-Id attribute (443) with the Subscription-Id-Type Diameter AVP sub-attribute (450) set to 4 (END_USER_PRIVATE) and the Subscription-Id-Data Diameter AVP sub-attribute (444) set to **reserved**.

[See [Understanding Gx-Plus Interactions Between the Router and the PCRF](#) and [Messages Used by Diameter Applications](#).]

- **RADIUS attributes added to LNS messages (MX Series)**—Starting in Junos OS Release 17.4R1, the LNS includes the following RADIUS attributes when it sends an Access-Request message to the RADIUS server:
 - Tunnel-Type (64)
 - Tunnel-Medium-Type (65)
 - Tunnel-Client-Endpoint (66)
 - Tunnel-Server-Endpoint (67)

- Acct-Tunnel-Connection (68)
- Tunnel-Assignment-Id (82)
- Tunnel-Client-Auth-Id (90)
- Tunnel-Server-Auth-Id (91)

System Logging

- **Debugging firewall ukern-trace log toggle persisting across FPC reboot (MX Series)**—Starting in Junos OS Release 17.4R1, you can enable or disable ukern-trace logging for the debugging firewall (DFW) on a specific FPC slot by using the **set chassis fpc slot ukern-trace log app-type dfw logging (off | on)** command. The new logging value of each DFW log takes effect immediately and persists if the FPC slot reboots.

[See [ukern-trace](#)]

User interface and Configuration

- **Monitoring, detecting, and taking action on degraded physical 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet links to minimize packet loss (MX Series routers with MPC5E, MPC6E, and 2x10GE MIC on MPC3E)**—Starting with Junos OS Release 17.4R1, 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} .

Layer 2 and Layer 3 protocols support the monitoring of physical link degradation. An Ethernet link also supports monitoring of physical link degradation through the Link Fault Signaling (LFS) protocol. However, for both of these monitoring mechanisms, the BER threshold value range of 10^{-13} to 10^{-5} is very low. Because of the low BER threshold value, the physical link degradation goes undetected, causing disruption and packet loss on an Ethernet link.

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.

The supported exponent range is 1 through 16, and the default value is 7 for the **set** configuration and 12 for the **clear** configuration.

- 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. The recovery mechanism triggers the recovery of a degraded link.

You can view the link recovery status and the BER threshold values by using the **show interfaces interface-name** command.

VPNs

- **Support of BGP signaling for next-hop-based dynamic tunnels (MX Series)**—Starting in Junos OS Release 17.4R1, the next-hop-based dynamic GRE and UDP tunnels are signaled using BGP encapsulation extended community. BGP export policy is used to specify the tunnel types, advertise the sender side tunnel information, and parse and convey the receiver side tunnel information. A tunnel is created according to the received type tunnel community.

Multiple tunnel encapsulations are supported by BGP. On receiving multiple capability, the next-hop-based dynamic tunnel is created based on the configured BGP policy and tunnel preference. The tunnel preference should be consistent across both the tunnel ends for the tunnel to be set up, and by default, MPLS-over-UDP (MPLSoUDP) tunnel is preferred over GRE tunnels.

[See [Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels](#) and [Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels](#).]

SEE ALSO

[Changes in Behavior and Syntax | 126](#)

[Known Behavior | 136](#)

[Known Issues | 142](#)

[Resolved Issues | 157](#)

[Documentation Updates | 199](#)

[Migration, Upgrade, and Downgrade Instructions | 200](#)

[Product Compatibility | 207](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [EVPNs | 126](#)
- [High Availability \(HA\) and Resiliency | 127](#)
- [Interfaces and Chassis | 127](#)
- [Management | 128](#)
- [MPLS | 128](#)
- [Multicast | 130](#)
- [Network Management and Monitoring | 130](#)
- [Routing Protocols | 131](#)
- [Security | 131](#)
- [Services Applications | 132](#)
- [Software Defined Networking | 133](#)
- [Software Installation and Upgrade | 133](#)
- [Software Licensing | 133](#)
- [Subscriber Management and Services | 133](#)
- [User Interface and Configuration | 136](#)

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 17.4R2 for MX Series routers.

EVPNs

- **Changes in the output of `show route table` command**—Starting in Junos OS Release 17.4R2, the output for `show route table` no longer displays the loopback address as the route distinguisher for MAC address virtual routing and forwarding (MAC-VRF) routing instances route entries. Instead, the output now displays the route distinguisher for the evpn and virtual switch instance type.
- **Support for LSP on EVPN-MPLS**—Starting in Junos OS Release 17.4R2, Junos supports the mapping of EVPN traffic to specific label-switched paths (LSPs). Prior to this release, the traffic policies mapping extended community to specific LSPs did not work properly.
- **Changes in the `show route extensive` output**—Starting in Junos OS Release 17.4R2, the output for `show route extensive` displays unknown evpn, opaque, and experimental extended communities as follows:
 - EVPN: unknown iana evpn 0xtype:0xsubtype:0xvalue

- OPAQUE: unknown iana opaque Oxtype:Oxsubtype:Oxvalue
- EXP: unknown Oxtype:Oxsub-type:Oxvalue

where type, sub-type, and value are defined in RFC 4360 *BGP Extended Communities Attribute*, RFC7153 *IANA Registries for BGP Extended Communities*. Internet Assigned Numbers Authority (IANA) maintains a registry with information on the type and subtype field values at <https://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xhtml>

High Availability (HA) and Resiliency

- **Command 'show chassis in-service-upgrade' not available (MX10003)**—In this release, the command "show chassis in-service-upgrade" is not available for MX10003 routers. If you enter this command, the following output is shown: "error: command is not valid on the JNP10003 [MX10003]". Earlier, the output shown for this command was "error: Unrecognized command (chassis-control)".

Interfaces and Chassis

- **Deprecated maximum transmission unit configuration option for virtual tunnel interfaces**—In Junos OS Release 17.4R1, you cannot configure the maximum transmission unit (MTU) size for virtual tunnel (vt) interfaces, because the **mtu bytes** option is deprecated for vt interfaces. Junos OS sets the MTU size for vt interfaces by default to *unlimited*.
- **Modified output of the request vmhost zeroize command**—Starting with Junos OS Release 17.2, the command **request vmhost zeroize**, upon execution, prompts the user for confirmation to proceed. The following line is displayed:

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

- **Modified output of the show chassis ethernet-switch command**—The ports 24 and 26 on the MX240, MX480, and MX960 routers with the RE-S-X6-64G Routing Engines are dedicated for external Ethernet connectivity. The **show chassis ethernet-switch** command on these routers displays the link status for these ports as **External Ethernet**.
- **Recovery of PICs that are stuck because of prolonged flow controls (MS-MIC, MS-MPC, MS-DPC, MS-PIC 100, MS-PIC 400, and MS-PIC 500)**—Starting in Junos OS Release 16.1R7, if interfaces on an MS-PIC, MS-MIC, MS-MPC, or MS-DPC are in stuck state because of prolonged flow control, Junos OS restarts the service PICs to recover them from this state. However, if you want the PICs to remain in stuck state until you manually restart the PICs, configure the new option **up-on-flow-control** for the **flow-control-options** statement at the **[edit interfaces mo-fpc/pic/port multiservice-options]** hierarchy level. In releases before Release 16.1R7, there is no action taken to recover service PICs from this state

unless one of the options for the **flow-control-options** statement is configured, or service PIC is manually restarted.

Management

- **Changes to Junos OS YANG module naming conventions (MX Series)**—Starting in Junos OS Release 17.4R1, the native Junos OS YANG modules use a new naming convention for the module's name, filename, and namespace. The module name and filename include the device family and the area of the configuration or command hierarchy to which the schema in the module belongs. In addition, the module filename includes a revision date. The module namespace is simplified to include the device family, the module type, and an identifier that is unique to each module and that differentiates the namespace of the module from that of other modules.

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

MPLS

- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (MX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP ressignaling events. This is because the LSP is responsive to even minor increases or decreases in the utilization when current reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of around 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP ressignaling events for small LSPs of around 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP ressignaling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless they are huge, for example, 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value-based threshold along with the percentage-based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at the **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.
- **Support for label history for MPLS protocol (MX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at the **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entries per label . By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range. The **clear mpls label history** command clears the label history details.

- **Support for default time out duration for self-ping on an LSP instance (MX Series)**—Starting in Junos OS 17.4R1, the default time out duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also configure the self-ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is enabled. The LSP types like CCC, P2MP, VLAN-based, and non-default instances do not support self-ping. You can configure **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.
- **Support for Flap and MBB counter for LSP (MX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on the master routing engine (RE) only:
 - Flap counter-- Counts the number of times a LSP flaps down or up.
 - MBB counter— Counts the number of times a LSP incurs MBB.

The **clear mpls lsp counters** command resets the flap and the MBB counter to zero.

- **Support for inet.0 and inet.3 labeled unicast BGP route for protocol LDP (MX Series)**--- Starting in Junos OS Release 17.4R2, LDP egress policy is supported on both inet.0 and inet.3 routing Information bases (RIBs) also known as routing table for labeled unicast BGP routes. If a routing policy is configured with a specific (inet.0 and inet.3) RIB, the egress policy is applied on the specified RIB. If no RIB is specified and a prefix is present on both inet.0 and inet.3 RIBs for labeled unicast BGP routes, then inet.3 RIB is preferred. However, prior to Junos OS Release 12.3R1 and starting with Junos OS Release 16.1R1, LDP egress policy is always preferred on inet.0 RIB and support for inet.3 RIB egress policy for labeled unicast BGP routes was disabled. In Junos OS Release 12.3R1 and later releases up to Junos Release 16.1R1, LDP egress policy was supported in inet.3 RIBs, in addition to inet.0 RIBs, for labeled-unicast BGP routes.
- **New output fields to monitor LSP resigaling count**—Starting in Junos OS Release 17.4R1, the **show mpls lsp** command output displays the **Flap Count** and **MBB Count** output fields, that capture the historical count of the number of times a specific LSP has been resigaled because of autobandwidth-triggered reservation change, or other changes along the path. The flap count displays the number of times an LSP flaps down and up, and the MBB count displays the number of times an LSP incurred a make before break.
- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (MX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
- Starting in Junos OS Release 17.4R1, a new configuration statement - **adjust-threshold-absolute** - is introduced at the **[edit protocols mpls]** hierarchy level to specify the changes in the average label-switched path (LSP) utilization to trigger automatic bandwidth adjustment in bits per second (bps).

Currently, this change is specified as a percentage using the **adjust-threshold statement**. The **adjust-threshold-absolute** statement (bps) can be used in conjunction with the existing **adjust-threshold statement** (percent).

- Starting in Junos OS Release 17.4R1, the **spring-traffic-engineering** statement at the **[edit protocols]** hierarchy level is replaced with the **source-packet-routing** statement, although the support for the **spring-traffic-engineering** statement is provided as an alias. This replacement does not introduce any functionality change, and is intended for maintaining consistency across the terms used in Source Packet Routing in Networking (SPRING) or segment routing features.
- **Loss of traffic over bypass MPLS LSPs**—If RSVP link or node protection is enabled along with global RSVP authentication, there is loss of traffic over bypass MPLS LSPs at the time of local repair, when the point of local repair (PLR) and the merge point devices have different versions of the Junos OS software installed on them. That is, one device is running a release prior to Junos OS Release 16.1, and the other device is running a release starting with Junos OS Release 16.1R4-S12.

Multicast

- **Support for rpf-selection statement for PIM protocol at global instance level (MX Series)**—Starting in Junos OS 17.4R1, the **rpf-selection** statement for the PIM protocol is available at global instance level. You can configure **group** and **source** statements at the **[edit protocols pim rpf-selection]** hierarchy level.

Network Management and Monitoring

- **Customer-visible SNMP trap name changes (MX Series)**—In Junos OS Release 17.4R1, on Enhanced Switch Control Board (SCBE), name changes include the CB slot when `jnxTimingFaultLOSSet` and `jnxTimingFaultLOSClear` traps are generated in the case of BITS interfaces (T1 or E1). SNMP traps for the backup Routing Engine clock failure event have been added and the control board name is included in the SNMP trap interface name (`jnxClkSyncIntfName`), for example, value: "external(cb-0)".

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

- **SNMP syslog messages changed (MX Series)**—In Junos OS Release 17.4R1, 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](#).]

- **Change in default log level setting (MX Series)**—In Junos OS Release, 17.4R1, the following changes were made in default logging levels:

Before this change:

- `SNMP_TRAP_LINK_UP` was `LOG_INFO` for both the physical (IFD) and logical (IFL) interfaces.
- `SNMP_TRAP_LINK_DOWN` was `LOG_WARNING` for both the physical (IFD) and logical (IFL) interfaces.

After this change:

- IFD LinkUp -> LOG_NOTICE (because this is an important message but less frequent)
- IFL LinkUp -> LOG_INFO (no change)
- IFD and IFL LinkDown -> LOG_WARNING (no change)

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

- **New context-oid option for trap-options configuration statement to distinguish the traps which come from a non-default routing instance and non-default logical system (MX Series)**—In Junos OS Release 17.4R2, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

Routing Protocols

- **Option to configure SPRING bandwidth utilization change threshold in percentage (MX Series)**—Starting in Junos OS Release 17.4R1, you can specify a change threshold in percentage beyond which RSVP triggers IGP updates. To configure the change threshold percentage, configure **percent percent** at the **[edit protocols rsvp interface update-threshold-max-reservable]** hierarchy level.
- **BGP enterprise trap jnxBgpM2BackwardTransition notification for IPv4 neighbors (MX Series)**—Starting in Junos OS Release 17.4R2, when an IPv4 BGP neighbor transitions from a higher state to a lower state, an enterprise trap **jnxBgpM2BackwardTransition** is sent in addition to an existing standard trap notification **bgpM2BackwardTransition**. In earlier Junos OS releases only **bgpBackwardTransition** trap notification was generated when a BGP IPv4 neighbor's state transitioned to a lower state.

Security

- **Support to log the SSH key changes**—Starting with Junos OS 17.4R1, the configuration statement **log-key-changes** is introduced at the **[edit system services ssh]** hierarchy level. When the **log-key-changes** configuration statement is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time the **log-key-changes** configuration statement was enabled. If the **log-key-changes** configuration statement was never enabled, then Junos OS logs all the authorized SSH keys.
- **Support for SSH protocol version 2**—Starting in Junos OS Release 17.4R1, SSH protocol version 1 (SSHv1) is not supported. SSH protocol version 2 (SSHv2) is the default protocol-version option available under the **[edit system services ssh]** hierarchy level.

[See [protocol-version](#)]

Services Applications

- **Accurate value in exported inline flow monitoring records for MPLS-over-GRE tunnels**—Starting in Junos OS Release 17.4R1, the exported flow records for inline flow monitoring of traffic entering MPLS-over-GRE tunnels (also known as next-hop-based dynamic GRE tunnels) contain the correct values in the gateway address and outgoing interface fields. Prior to Junos OS Release 17.4R1, these fields contained a value of 0.
- **New syslog message displayed during NAT port allocation error (MX Series Routers with MS MPC)**—With address pooling paired (APP) enabled, an internal host is mapped to a particular NAT pool address. In case, all the ports under a NAT pool address are exhausted, further port allocation requests from the internal host results in a port allocation failure. The following new syslog message is displayed during such conditions:

JSERVICES_NAT_OUTOF_PORTS_APP

This syslog message is generated only once per NAT pool address.

Software Defined Networking

- **The 32-bit libstdc++ package no longer required for Junos Node Slicing setup**—Starting in Junos OS Release 17.4R2, you need not install the additional 32-bit **libstdc++** package for Red Hat Enterprise Linux (RHEL) or Ubuntu to set up Junos Node Slicing.

Software Installation and Upgrade

- **ZTP is supported on MX PPC platforms (MX Series)**—As of Junos OS Release 17.4R2, zero touch provisioning (ZTP) is supported on MX PPC platforms (which are MX5, MX10, MX40, MX80, and MX104 routers). Before the fix, the ZTP process did not start to load image and configuration for MX PPC routers.

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

Software Licensing

- **Key generator adds one day to make the duration of license show as 365 days (MX Series)**—Starting in Junos OS Release 17.4R1, the duration of subscription licenses as generated by the **show system license** command and shown in the output is correct to the numbers of days. Before this fix, for example, for a 1-year subscription license, the duration was generated as 364 days. After the fix, the duration of the 1-year subscription now shows as 365 days.

[See [show system license](#).]

Subscriber Management and Services

- **Correct SNMP index value in exported inline flow monitoring records for BNG subscribers**—Starting in Junos OS Release 17.4R1, the exported flow records for inline flow monitoring report the SNMP index of the broadband network gateway (BNG) subscriber's interface. Prior to Junos OS Release 17.4R1, the flow records reported the SNMP index of the underlying interface (PPPoE encapsulated interface), which caused incorrect values in the derived fields (mask, outgoing interface, gateway address).

Configure **nexthop-learning enable** at the **[edit services flow-monitoring (version-ipfix | version9) template *template-name*]** hierarchy level to get the correct outgoing interface and gateway address values for subscriber traffic in the following situations:

- Ingress and egress VRF are not the same.
- Traffic is load balanced.
- Traffic is forwarded through a composite next hop (for example, an MPLS over GRE tunnel).

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

- **Memory mapping statement removed for Enhanced Subscriber Management (MX Series)**— Starting in Junos OS Release 17.4R1, use the following command when configuring database memory for Enhanced Subscriber Management:

set system configuration-database max-db-size

CLI support for the **set configuration-database virtual-memory-mapping process-set subscriber-management** command has been removed to avoid confusion. Using the command for subscriber management now results in the following error message:

WARNING: system configuration-database virtual-memory-mapping not supported. error: configuration check-out failed.

[See [Interface Configuring Junos OS Enhanced Subscriber Management](#) for an example of how to use the **max-db-size** command.]

- **Support for IPv6 all-routers address in nondefault routing instance (MX Series)**—Starting in Junos OS Release 17.4R2, the well-known IPv6 all-routers multicast address, FF02::2, is supported in nondefault routing instances. In earlier releases it is supported only for the default routing instance; consequently IPv6 router solicitation packets are dropped in nondefault routing instances.
- **Correction to CLI for L2TP tunnel keepalives (MX Series)**—Starting in Junos OS Release 17.4R2, the CLI correctly limits to 3600 seconds the maximum duration that you can enter for the hello interval of an L2TP tunnel group. In earlier releases, the CLI allows you to enter a value up to 65,535, even though only 3600 is supported.

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

- **Wildcard supported for show subscribers agent-circuit-identifier command (MX Series)**—Starting in Junos OS Release 17.4R2, you can specify either the complete ACI string or a substring when you issue the **show subscribers agent-circuit-identifier** command. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

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

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

- **Changed behavior for framed routes without a subnet mask (MX Series)**—Starting in Junos OS Release 17.4R2, the router connects the session but ignores a framed route when it is received from RADIUS in the Framed-Route attribute (22) without a subnet mask.

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

- **DHCPv6 lease renewal for separate IA renew requests (MX Series)**—Starting in Junos OS Release 17.4R2, the `jdhcpcd` process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

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

In earlier releases:

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

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

- **Bandwidth options match for inline services and tunnel services (MX Series)**—Starting in Junos OS Release 17.4R2, you can configure the same bandwidth options for inline services with the **bandwidth** statement at the `[edit chassis fpc slot-number pic number inline-services]` hierarchy level as you can configure for tunnel services with the **bandwidth** statement at the `[edit chassis fpc slot-number pic number tunnel-services]` hierarchy level.

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

- **Change to ICRQ message inclusion of the ANCP Access Line Type AVP (MX Series)**—Starting in Junos OS Release 17.4R2, the ICRQ message includes the ANCP Access Line Type AVP (145) when the received ANCP Port Up message includes a DSL-type of 0 (OTHER). In earlier releases, the AVP is not sent when the value is 0.

User Interface and Configuration

- Junos OS prohibits configuring ephemeral configuration database instances that use the name **default** (MX Series)—Starting in Junos OS Release 17.4R2, user-defined instances of the ephemeral configuration database, which are configured using the **instance *instance-name*** statement at the **[edit system configuration-database ephemeral]** hierarchy level, do not support configuring the name **default**.

SEE ALSO

[New and Changed Features | 93](#)

[Known Behavior | 136](#)

[Known Issues | 142](#)

[Resolved Issues | 157](#)

[Documentation Updates | 199](#)

[Migration, Upgrade, and Downgrade Instructions | 200](#)

[Product Compatibility | 207](#)

Known Behavior

IN THIS SECTION

- [General Routing | 137](#)
- [EVPN | 139](#)
- [Interfaces and Chassis | 139](#)
- [Layer 2 Ethernet Services | 140](#)
- [MPLS | 140](#)
- [Routing Protocols | 140](#)
- [Services Applications | 140](#)
- [Software Installation and Upgrade | 141](#)
- [Subscriber Management and Services | 141](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for MX Series routers.

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

General Routing

- On MX Series routers with MS-MPC/MS-MIC, memory leaks will be seen with `jnx_msp_jbuf_small_oc` object, upon sending millions of Point-to-Point Tunneling Protocol control connections (3-5M) alone at higher cells per second (cps) (greater than 150K cps). This issue is not seen with up to 50,000 control connections at 10,000-30,000 cps. [PR1087561](#)
- Source-prefix filtering and protocol filtering of the CGNAT sessions are incorrect. For example, the **show services sessions extensive protocol udp source-prefix <0:7000::2>** command displays incorrect filtering of the sessions. [PR1179922](#)
- Chef for Junos OS supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev-resources. The netdev-resource developed for interface configuration has a limitation to configure the XE interface. The netdev-interface resource determines that speed is a configurable parameter that is supported on a GE interface but not on an XE interface. Hence, the netdev-interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages `chef-11.10.4_1.1.*.tgz` `chef-11.10.4_2.0.*.tgz` in all platforms {i386/x86-32/powerpc}. [PR1181475](#)
- In certain interface scaling scenarios, during configuration commit/rollback, you might see an `fpcx` error message. You can safely ignore this message because of the FPGA monitor mechanism on DPC cards for logical interface mapping (`ifl_map`). Between the deletion of a physical interface and the monitoring event, this mechanism checks through the stored logical interfaces. While the mechanism tries to find the family of a recently deleted logical interface that was not cleaned from the `ifl_map`, harmless messages might populate the log file. [PR1210877](#)
- There is no unified ISSU from Junos OS Release 15.1 and earlier releases to Junos OS Release 16.2R1. [PR1222540](#)
- The Routing Engine shows it is using Spring LSP, but forwarding actually uses L-ISIS label. The problem is, when some route or next hop has been created by the app, it is assumed that it can propagate to the rest of the system. The KRT asynchronously picks up this state for propagation. There is no reverse indication to the app, if there was an error in propagating the state. The system is supposed to eventually reconcile. So, if SPRING-TE produces a pair that looks legal from the app standpoint, but KRT is not able to download it to the kernel, because kernel rejected the NH, the sort of gets stuck in RPD. In the meantime, the previous version of the route (L-ISIS in this case) that was downloaded still lingers in the kernel and the Packet Forwarding Engine. [PR1253778](#)
- On MX104 routers, `JTASK_SCHED_SLIP` is seen on commit, randomly. [PR1281016](#)
- At reboot the RHEL 7.3 servers report `libvirtd[6282]: segfault at 10 ip 00007f87eab09bd0`. There is no core file generated and no operational impact is known. [PR1287808](#)

- When LLDP is configured on multihomed extended ports, the peer might have duplicate entries for a duration of the hold timer (default: 120 seconds) during catastrophic configuration events such as redundancy group ID change and redundancy group name change. The duplicate entry would be deleted after the LLDP hold timer expires on the peer. [PR1291519](#)
- A race condition is observed on Ubuntu based external servers, G-ARP might not be sent from the jmgmt0 interface, resulting in loss of connectivity to management IP of JDM. [PR1291836](#)
- This is a limitation/expected behavior for smart SFPs. When you insert a smart-sfp, it is observed that the link remains up for some time; for example, during smart SFP firmware initialization, the green LED on the transceiver glows green. [PR1293522](#)
- The af interface bandwidth that is shown is based on the peer GNF's Packet Forwarding Engine type. The local FPC on the GNF could have a higher capacity for throughput than af interface's statically configured bandwidth. Also, the fabric capacity of the Packet Forwarding Engine is slightly higher than that of WAN interface of same bandwidth. Since the fabric can accept more traffic, the af interface shows higher throughput rate than what the Packet Forwarding Engine is capable of. This is the expected behavior until the CoS shaping is supported on the interface. [PR1295050](#)
- Rpd sends a KStat request to the kernel, every time the **show dynamic-tunnels database** command is processed. Because Kstat is an asynchronous call and the CLI is not blocked until rpd receives a response from the kernel, there might be a mismatch in statistics between the Packet Forwarding Engine and kernel for some time. Eventually the statistics will be updated in rpd, whenever the response for the last statistics request is received. These statistics will be reflected in the next **show dynamic-tunnels database** command. [PR1297913](#)
- For CFP2-DCO-T-WDM-1 pluggable, Rx payload type shown incorrectly (shown 0 vs 7). [PR1300423](#)
- The UDP setup rate for DetNat64 is approximately 10 percent lesser than the setup rate of stateful-nat64 for 15M sessions on a single NPU. DetNat64 needs extra processing while creating sessions and hence, it's setup rate is 10 percent less than setup rate of stateful-nat64. [PR1307451](#)
- Support for enterprise profile support is with only 10G interfaces. 40G & 100G may result in phase alignment issue. [PR1310048](#)
- Parametrized (converged) HTTP redirect/rewrite services (CPCD) are not supported on MX104 routers with MS-MIC. Note that other flavors of CPCD continue to work fine with this combination, MX104 router with MS-MIC. [PR1330340](#)

EVPN

- Routing instances of type EVPN configured with a VLAN ID will advertise MAC (type 2) routes with the VLAN value in the Ethernet tag field of the MAC route. Advertising MAC routes with a nonzero VLAN is incompatible with the EVPN VLAN-based service type. To enable interoperability between a Junos OS routing instance of type EVPN and a remote EVPN device operating in VLAN-based mode, the Junos routing instance should be configured with **vlan-id none** so that the Ethernet tag in advertised MAC routes is set to zero. [PR945247](#)
- A PE device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE device. The IGP instance running in the VRF on the PE might be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- In scaled up EVPN-VPWS configurations (approximately 8000 EVPN VPWS), during a Routing Engine switchover, the rpd scheduler slip messages might be seen. [PR1225153](#)
- In an EVPN network with VXLAN encapsulation configured for direct-nexthop mode ("pure type 5" mode without overlay gateway addresses), at least one type 5 route per VRF from a remote endpoint must be received and installed in the local routing table of a device, to enable the local device to forward inbound type 5 traffic received from the remote endpoint. If the local device has not installed at least one route with a next hop pointing toward a specific remote endpoint, type 5 VXLAN-encapsulated IP traffic sent by the remote endpoint toward the local device will not be forwarded correctly. [PR1305068](#)
- When changing encapsulation from VXLAN to MPLS or vice versa, you must deactivate and reactivate the instance. [PR1326430](#)

Interfaces and Chassis

- In a node slicing context, issuing the command **set chassis fpc slot-number power off** on the base system (BSYS) powers off even those FPCs that are assigned to guest network functions (GNFs) in which unified in-service software upgrade (ISSU) is in progress.

Learn more about [Junos Node Slicing](#).

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

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

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

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

Layer 2 Ethernet Services

- Junos Fusion device supports Aggregate Interface with 16 member links. [PR1300504](#)

MPLS

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

Routing Protocols

- The BGP NSR replication starts after some delay in certain cases. [PR1256965](#)

Services Applications

- Session counters for **cleartext traffic** are not updated after decryption. Decrypted packet count can, however, be obtained by running the **show security group-vpn member ipsec statistics** command. [PR1068094](#)
- Broadband-edge platforms do not support service-set integration with dynamic profiles when the service set is representing a carrier-grade NAT configuration. As a workaround, you can use next-hop service set configurations and routing options to steer traffic to a multiservices (ms) interface where NAT functionality can be exercised. The following configuration snippet shows the basics of statically configuring the multiservices interface next hop and a next-hop service set. Traffic on which the service is applied is forced to the interface inside the network by configuring that interface as the next hop. This configuration does not show other routing-options or NAT configurations relevant to your network.

```
routing-options {
  static {
    route 0.0.0.0/0 {
      next-hop ms-3/0/0.1;
      preference 0;
    }
  }
  ...
}
```

```

services {
  service-set CGN {
    nat-rules CGN_SAMPLE;
    next-hop-service {
      inside-service-interface ms-3/0/0.1;
      outside-service-interface ms-3/0/0.2;
    }
  }
  nat {
    ...
  }
}

```

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

Software Installation and Upgrade

- **Unified ISSU with active BBE subscribers using advanced services supported only to 17.4R2 and later 17.4 releases**—If you have active broadband edge subscribers that are using advanced services, you cannot perform a successful unified in-service software upgrade (ISSU) to a Junos OS 17.4 release earlier than 17.4R2. If you perform an ISSU to a 17.4 release earlier than 17.4R2, the advanced services PCC rules are not attached to subscribers.

Subscriber Management and Services

- The **all** option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the **all** option with the **clear services l2tp destination**, **clear services l2tp session**, or **clear services l2tp tunnel** statements in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.
- Before you make any changes to the underlying interface for a demux0 interface, you must ensure that no subscribers are currently present on that underlying interface. If any subscribers are present, you must remove them before you make changes.
- For dual-stacked clients over the same PPP over L2TP LNS session, enhanced subscriber management does not support configurations where both of the following are true:
 - The CPE sends separate DHCPv6 solicit messages for the IA_NA and the IA_PD.
 - The solicit messages specify a type 2 or type 3 DUID (link-layer address).

As a workaround, you must configure the CPE to send a single solicit message for both IA_NA and IA_PD when the other configuration elements are present.

SEE ALSO

New and Changed Features		93
Changes in Behavior and Syntax		126
Known Issues		142
Resolved Issues		157
Documentation Updates		199
Migration, Upgrade, and Downgrade Instructions		200
Product Compatibility		207

Known Issues

IN THIS SECTION

- [General Routing](#) | [143](#)
- [Class of Service \(CoS\)](#) | [148](#)
- [EVPN](#) | [148](#)
- [Forwarding and Sampling](#) | [149](#)
- [High Availability \(HA\) and Resiliency](#) | [150](#)
- [Infrastructure](#) | [150](#)
- [Interfaces and Chassis](#) | [150](#)
- [Layer 2 Features](#) | [151](#)
- [Layer 2 Ethernet Services](#) | [151](#)
- [Multiprotocol Label Switching \(MPLS\)](#) | [151](#)
- [Platform and Infrastructure](#) | [152](#)
- [Routing Protocols](#) | [155](#)
- [Services Applications](#) | [157](#)
- [VPNs](#) | [157](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R2 for MX Series routers.

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

General Routing

- An intermittent issue occurs when an aggregated Ethernet interface is configured with the **bypass-queuing-chip** configuration statement. The follow-up configuration changes are such that, removing a child link from an aggregated Ethernet bundle and configuring per-unit-scheduler on the removed child link in a single commit causes intermittent issues with the per-unit-scheduler configuration updates to cosd and the Packet Forwarding Engine. Hence, dedicated scheduler nodes might not be created for all units or logical interfaces. [PR1162006](#)
- After loading a CoS-related configuration on MPC5E, MPC6E, MPC2E-NG, or MPC3E-NG line cards, the following error messages might be seen: **jnh_ifl_topo_handler_pfe(11591): ifl=495 err=1 updating channel table nexthop and _insert_ifl_channel:6449 ifl 495 chan_index 495 NOENT**. [PR1186645](#)
- The source-address based filter forwarding is used under forwarding-options to steer the packets towards the AMS bundle in the Vodafone configuration. When you remove the source-address condition from the filter, the reverse traffic gets looped back into the AMS bundle. Under this condition, prolonged flow control core files are seen. The source-address configured in the SFW rule should have dropped the packets, which are getting looped back into the AMS bundle, but this is not happening, even though SFW functionality works as expected for other packets. [PR1192184](#)
- With MPC8/9 MRATE MIC and plug-in optics module(QSFP28-100GBASE-LR4), bit errors might be seen. [PR1200010](#)
- Upgrading using unified ISSU might trigger a flap in the interfaces on MX Series routers and the following message might be seen: **SFP: pointer Null, sfp_set_present**. [PR1200045](#)
- After system boot up or after PSM reset, you might see the PSM INP1 or INP0 circuit Failure error message. [PR1203005](#)
- The SMID process has stopped responding to the management requests after a jl2tpd (L2TP process) crashes on an MX960 broadband network gateway. [PR1205546](#)
- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, **fpc7 KERNEL/PFE APP=NH OUT OF SYNC: error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none**. There is no service impact found in MPC2 and MPC3 type cards. [PR1205593](#)
- The following error messages occur during GRES and unified ISSU: **syslog errors @ agentd_rts_async_rtbm_msg : FLM : Failed to create private**. [PR1232636](#)
- When virtual switch type is changed from IRB type to regular bridge, interfaces under the OpenFlow protocol are removed. The openflow process fails to program any flows. [PR1234141](#)
- After configuring PCEP following log seen - **pccd: [89798] Could not decode message from rpd**. This might impact in growth of memory of pccd process over time, which can be cleared by restarting the process. [PR1235692](#)

- Sometimes, when PPPoE subscribers log in and log out from Junos OS Release 16.1 and later, the following messages are generated: `user@devcie> show log messages | match authd authd[5208]: sdb_app_access_line_entry_read_by_uifl: uifl key 'demux0.xxxxxxxx': snapshot failed (-7) authd[5208]: sdb_app_access_line_entry_read: uifl key 'demux0.xxxxxxxx': read failed` These messages indicate that **authd daemon for subscriber authentication is attempting to read private data for an underlying interface which no longer exists (-7 = SDB_DATA_NOT_FOUND)**. These messages have no impact and can be safely ignored, where the authd process is asking sdb for a record that no longer exists. [PR1236211](#)
- When gRPC subscription for telemetry data has a 2-second frequency, the jsd process might crash. [PR1247254](#)
- On MX Series routers with an XM chipset (such as, MPC3E/MPC4E/MPC5E/MPC6E/MPC2E-NG/MPC3E-NG), the MPC might reboot after a unified ISSU completion. [PR1256145](#)
- Error messages such as `mshpmand[190]: msvcs_session_send: Plugin id 3 not present in the svc chain for session ..`are seen. They are usually cosmetic. [PR1258970](#)
- When both the OAM protocol and the MACsec protocol are configured on an interface the interface does not come online. This issue occurs when an interface comes online and both OAM protocol and MACsec Key Agreement (MKA) protocol try to establish their respective sessions. Because of contention between these two protocols, OAM takes down the interface and MKA fails to establish connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
- On an MX Series Virtual Chassis system in a scaled subscriber management scenario, if a unified ISSU is performed while the BGP protocol sessions are active and such BGP sessions are clients of BFD, then these BGP sessions might go down and come back up again, causing traffic loss. [PR1265407](#)
- During a unified ISSU, only the Packet Forwarding Engine gets wedged. This very specific issue occurs when the Packet Forwarding Engine is oversubscribed with unknown unicast flooding with no MAC learning, which is not a common configuration. However, this issue is not seen when the Packet Forwarding Engine is oversubscribed with Layer 3 traffic or with Layer 2 traffic with MAC learning. [PR1265898](#)
- Guest network functions (GNFs) in a node-slicing setup currently do not support Junos Snapshot Administrator or recovery mechanisms. [PR1268943](#)
- Dynamic endpoint (DEP) does not support dh group group19, encryption algorithm aes-256-cbc, and hash sha-384 in its list of default proposals. These proposals must be configured explicitly in the configuration. [PR1269160](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- There are incorrect counters for output packets on child links of the ae0 interface when configured with the new feature **revertive**. [PR1273983](#)
- For inline jflow, when both packets and seconds interval are configured for the **template-refresh-rate** and **option-refresh-rate** configuration options, the packets interval configuration is not working. [PR1274206](#)

- A routine within an internal Junos OS sockets library is vulnerable to a buffer overflow. Malicious exploitation of this issue might lead to a denial-of-service (kernel panic) or be leveraged as a privilege escalation through local code execution. The routines are only accessible through programs running on the device itself, and verixec restricts arbitrary programs from running on Junos OS. There are no known exploit vectors utilizing signed binaries shipped with Junos OS itself. See, <https://kb.juniper.net/JSA10792>. [PR1282562](#)
- On an MX Series Virtual Chassis, when using a channelized configuration on MPC7/8/9 MRATE PIC QSFP interfaces for VCP connections between members, a VCP interface needs to be configured on channel 0 of each QSFP to activate the port. [PR1283283](#)
- Due to a code limitation, an ungraceful removal of summit MACsec TIC from the chassis might cause a crash or an unpredictable result. [PR1284040](#)
- On MX10003, the **chassisd hard restart** command is not supported due to an infrastructure limitation. The FPC power off does not happen cleanly as the old chassisd process initiates the **fpc power off** command and exits. Restarting chassisd hard with GRES on an MX10003 causes a new chassisd process to open, reconnect a window, and wait for the connection. The Routing Engine and FPC go out of synchronization and FPC reconnect is not handled, which causes the FPC to be restarted multiple times. Finally, FPC comes online. [PR1293314](#)
- Fixes committed in Junos OS Releases 15.1R5-S4, 16.1R4-S3, 16.1R5, and 17.3R1 with XM-based linecards (such as, MPC3E/4E/5E/6E/2E-NG/3E-NG) might report the chassisd error log message **DDR3 TEMP ALARM**. [PR1293543](#)
- If OC package upgrade is triggered when telemetry is occurring, the xmlproxyd process might crash. It will recover automatically and xmlproxyd related streaming will restart as the process comes up again. We recommend that you stop the streaming and then upgrade the OC package. [PR1295831](#)
- In some Junos OS for MX Series deployments, random syslog messages are observed for FPC cards **fpcx ppe_img_ucode_redistribute Failed to evict needed instr to GUMEM - xxx left**. These messages are not an issue and might not have any service impact. These messages are addressed as INFO level messages. On the Packet Forwarding Engine, there are dedicated UMEM and shared GUMEM memory blocks. This INFO message indicates some evicting events between UMEM and GUMEM and can be safely ignored. [PR1298161](#)
- When a GRES or NSR is performed on a BSYS, the master Routing Engine on the GNFs (virtual nodes/network slices) will detect the BSYS chassisd restart and enter an NSR hold-down delay. During this time, CLI commands evoke a switchover on the master Routing Engine indicating that the system is not NSR ready. This situation is similar to that of a standalone MX Series router in which chassisd is restarted on the master Routing Engine. Note that a CLI command on the BU Routing Engine will succeed. This too, is similar to the behavior on a standalone MX Series router. [PR1298571](#)
- The iLatency (calculated by differing producer timestamp and gRPC server timestamp) value can sometimes be negative for Packet Forwarding Engine related telemetry packets because of a drift in the Routing Engine and the Packet Forwarding Engine NTP servers. [PR1303376](#)

- The mgd might crash when an Ephemeral database is used. This type of crash indicates simultaneous operation on an ephemeral instance. When a process wants to open an ephemeral configuration in merge view, some other activity (such as purging, deletion, or re-creation) is being carried out on this ephemeral instance. The occurrence of this crash is rare. [PR1305424](#)
- The message **LIBJNX_REPLICATE_RCP_ERROR** is repeated multiple times in the syslog log files in the master Routing Engine, when the backup is not reachable. Although the message is marked as an error in the syslog, you can ignore this error because it will not have any impact on the system. [PR1305660](#)
- Telemetry thread on the FPC might overuse the CPU thread in case of certain telemetry features like per service identifier in SR Statistics. This is a generic issue in the FPC telemetry code but gets exposed easily when **per prefix stats** is enabled through telemetry. This occurs because you walk a lot of prefix (a lot of which might not have any data to export) but do not yield until the buffer is attempted to be filled up. [PR1308513](#)
- Support for enterprise profile is only provided for 10-Gigabit Ethernet interfaces. Use of 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. [PR1310048](#)
- For sensors belonging to the same producer (for example, BGP and MPLS coming from rpd), if you use the same reporting intervals, then they are not streamed in parallel but are streamed sequentially. As a workaround, use a different reporting rate for sensors that belong to the same producer. [PR1315517](#)
- An alarm is raised if mixed AC PEMs are present. The criteria has been changed to check whether mixed AC is present. If the PEM is AC (high), then the first bit of pem_voltage is set, and if it is AC (low), then the second bit of pem_voltage is set. So if both first and second bit are set, then the mixed AC is present. [PR1315577](#)
- Making changes in **services traffic-load-balance** for one instance, might lead to a refresh of existing instances. [PR1318184](#)
- When an xmlproxy YANG file is configured through the **request system yang add package <package-name> proxy-xml module <module-name>** CLI command, then a notification related to new rendering schema is sent to all the Junos OS processes instead of being sent only to a limited set of processes (xmlproxyd and agentd). This might cause some processes, such as chassisd and jsd to restart, resulting in a telemetry session bounce as well. [PR1320211](#)
- In JDM (running on a secondary server), a jdmd process might generate a core file if GNF add-image is aborted by pressing Ctrl+C. [PR1321803](#)
- With **commit full**, the na-grpd process might restart causing a disconnection of the streaming telemetry. [PR1326366](#)
- Under some race conditions with fail-over and multiple core interface flapping on Ethernet virtual private network (EVPN) / Virtual Extensible LAN (VXLAN) network, the rpd process can be with high CPU causing some issues in intercommunication with the l2ald process, then causing the l2ald process to coredump and restart. [PR1333823](#)

- On MX204, MX10003, MPC7E, MPC8E, or MPC9E, the 100-Gigabit interface might keep flapping or stay down due to an interoperation issue between the Juniper Networks device and the remote transport device it is connected to. [PR1337327](#)
- In an MPLS-EVPN environment, when label-switched path (LSP) flapping causes RSVP LSP reroute, LSP might stick in down state with **Record route: <self> ...incomplete**. [PR1343289](#)
- On MX Series routers with 100M SFP used on MIC-3D-20GE-SFP-E/MIC-3D-20GE-SFP-EH, SFP might not work if it is not from Fiberxon or Avago. [PR1344208](#)
- There is a possibility of MACsec sessions not establishing if FPCs go through a continuous cycle of offline or online many times (greater than 10 times), followed by restarting the dot1xd process. [PR1344358](#)
- The Junos OS hidden hierarchies are not published in the Junos OS YANG schema and hence Junos OS should not emit these hidden hierarchies as part of the configuration. But in case of hidden choices, Junos OS is publishing a list without a key value because the key is hidden in the Junos OS schema. Hence, the ODL controller is not able to parse get-config response from Junos. As a workaround, you can remove such problematic hidden configurations from the device. The impact is limited only to the OpenDaylight controller. [PR1348503](#)
- On a single Routing Engine system, after the graceful Routing Engine switchover (GRES) configuration is removed, the Routing Engine mastership keepalive timer is not resumed to the default value with GRES enabled. [PR1349049](#)
- In some cases, OIR (removal followed by reinsertion) of a MIC on a FPC can lead to black holing of traffic destined to the FPC. The only way to recover from this is to restart the FPC. The issue will not be seen if you use the corresponding CLI commands to offline and then online the MIC. [PR1350103](#)
- On MX platform with the subscriber-management feature enabled, if the combination of an Ascend-Data-Filter (ADF) and a redirect filter is applied to the subscribers, it may cause a leak in the Broadband Edge (BBE) filter index. The index is not released when the subscriber logs out. Due to this issue, new subscribers are not able to connect when all the indexes are used up. [PR1353672](#)
- The system might take a longer period of time to reboot or the kernel might panic if rebooted during a broadcast storm on the mgmt port. [PR1351977](#)
- On an MX10003, a vmcore is observed **Kernel panic - not syncing: NMI: Not continuing**. [PR1353158](#)
- On MX Series routers with the subscriber management feature enabled, if the combination of an Ascend Data Filter (ADF) and a redirect filter is applied to the subscribers, it might cause a leak in the Broadband Edge (BBE) filter index. The index is not released when the subscriber logs out. Because of this issue, new subscribers are not able to connect when all the indexes are used up. [PR1353672](#)
- The "ipv4-flow-table-size" is used to configure the size of the IPv4 flow table in units of 256K entries. However, in "inline-jflow" scenario, if the knob "ipv6-extended-attrb" is configured, changing flow table configuration or clear the flow entries might lead to the condition that even the "ipv4-flow-table-size" has been changed to a number larger than 149, the maximum number of IPv4 flows still remains at 37372900. [PR1355095](#)
- DHCP subscriber unable to reach gateway as arp request dropped under pfe as dv discard. [PR1356101](#)

- When you use "show agent sensors verbose" FPC VTY command on MPC7E, the FPC might crash. [PR1366249](#)
- On ISSU to this release, there could be some impact to forwarding of packets of some destinations. [PR1366811](#)
- In some configurations, ISSU prepare time on MPC5E takes longer than usual. As a result, the chassisd triggers restart/crash of the MPC . The ISSU completes after the crash. [PR1369635](#)

Class of Service (CoS)

- A CoS scheduler update can fail when all of the following conditions are met: (1) Dynamic subscribers exist on an aggregated Ethernet bundle. (2) CoS traffic-control-profile or scheduler-map (or both) applied to these dynamic subscribers is from a static configuration. (3) The relevant static CoS is modified in the same configuration commit as a modification to the aggregated Ethernet bundle (either a leg add or leg remove) containing the subscribers. (4) The leg add or leg remove in the commit is the first or last leg to be added or removed from a line card. To avoid this issue, do not commit a bundle change in the same commit as a static CoS change. In this event, one of the following logs will be displayed in the message system log: **subscriber cos update not applied to interface <interface-name> status <id>** or **subscriber cos update not applied to interface-set <interface-set-name> status <id>**. This message indicates that the last update to the subscriber or interface set was not applied. If this event occurs, the workaround to fix the state is to: (1) Remove the last class-of-service update. (2) Commit the configuration. (3) Re-apply the class-of-service update. (4) Commit the configuration. [PR1276459](#)

EVPN

- The Layer 2 learning process (l2ald) might generate core files in a scaled Layer 2 setup, including bridge-domain, VPLS, EVPN, and so on. The l2ald process generation of core files usually follows a kernel page fault. In most cases, the issue is recovered on its own after l2ald generates the core file. In some cases, you can manually restart the process to recover. Logs: **/kernel: %KERN-3-BAD_PAGE_FAULT: pid 69719 (l2ald), uid 0: pc 0x88beb5ce got a read fault at 0x6ca, x86 fault flags = 0x4 /kernel: %KERN-6: pid 69719 (l2ald), uid 0: exited on signal 11 (core dumped) init: %AUTH-3: l2-learning (PID 69719) terminated by signal number 11. Core dumped!** [PR1142719](#)
- In an EVPN scenario with static MAC configured in the EVPN instance, the remote EVPN instance can see the MAC route information. However, after deactivating and activating the static MAC in the EVPN instance, and then checking the MAC route information in the remote EVPN instance, no such MAC route is found in the EVPN route table. [PR1193754](#)
- In an EVPN network with VXLAN encapsulation configured for **direct-nexthop** mode ("pure type 5" mode without overlay gateway addresses), at least one type 5 route per VRF from a remote endpoint must be received and installed in the local routing table of a device. This enables the local device to forward inbound type 5 traffic received from the remote endpoint. If the local device has not installed

at least one route with a next hop pointing toward a specific remote endpoint, type 5 VXLAN-encapsulated IP traffic sent by the remote endpoint toward the local device will not be forwarded correctly. [PR1305068](#)

- The issue is applicable to MAC-in-MAC PNN EVPN and does not affect any other scenario. When the provider backbone bridging (PBB) EVPN configuration is reloaded on MX Series routers, error logs are seen while deleting interfaces related to the backbone bridge component. These errors do not result in any functional issues. [PR1323275](#)
- The PBB EVPN will not be able to flood traffic towards the core. Traffic recovers by performing the **restart l2-learning** command. In addition to this, there is a limitation in PBB EVPN active/active (A/A) unicast traffic forwarding. If entropy in the traffic is not sufficient, then uneven load balancing causes a problem on the MH peer A/A routers. This will cause a drop for return traffic. These issues are applicable to PNN EVPN and do not affect any other scenario. [PR1323503](#)
- In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) deployment, the rpd process might crash on the new master after performing a Graceful Routing Engine Switchover (GRES). [PR1333754](#)
- On the Junos OS platform, the l2ald process might crash during the MAC address processing. The MAC learning process will be impacted during the period of l2ald crash. The l2ald will recover by itself. [PR1347606](#)
- The bidirectional Layer 2 traffic floods for around 5 second for streams from SH to MH, when the **clear mac table** command is executed on MX Series routers because MACs getting populated in the system are taking time. The **clear mac table** command is disruptive, which deletes all dynamic MACs in the system. [PR1360348](#)

Forwarding and Sampling

- When a policing filter is applied to an active LSP carrying traffic, the LSP resignals and drops traffic for approximately 2 seconds. It can take up to 30 seconds for the LSP to come up under the following conditions: (1) Creation of the policing filter and application of the same to the LSP through the configuration occurs in the same commit sequence. (2) Load override of a configuration file that has a policing filter and policing filter application to the LSP is followed by a commit. [PR1160669](#)
- In some stress test conditions, the sampled process crashes and generates a core file when connecting to L2BSA and EVPN subscribers aggressively. [PR1293237](#)
- A heap memory leak occurs on DPC when the flow specification route is changed. [PR1305977](#)
- Firewall Filter not applied as input filter to Extended Port when used for Layer 2 VPN. [PR1311013](#)

High Availability (HA) and Resiliency

- To avoid such kind of error make sure that space available in /var is twice the size of target image. This is the basic requirement for ISSU to proceed. [PR1354069](#)

Infrastructure

- The configuration statement **set system ports console log-out-on-disconnect** logs the user out from the console and closes the console connection. If the configuration statement **set system syslog console any warning** is used with the earlier configuration and when there is no active telnet connection to the console, the process tries to open the console and hangs as it waits for a "serial connect" that is received only by telnetting to the console. As a workaround, remove the latter configuration by using the **set system syslog console any warning** command, which solves the issue. [PR1230657](#)

Interfaces and Chassis

- Junos OS now checks ifl information under the aggregated Ethernet interface and prints only if it is part of it. [PR1114110](#)
- A Junos OS upgrade involving a CFM configuration can cause a cfmd crash after upgrade. This issue occurs because of the presence of the old version of the `/var/db/cfm.db` file. [PR1281073](#)
- The LAG member links running LACP in slow mode might get disassociated from the LAG bundle with a combination of restart interface-control and FPC offline or online trigger. The issue is seen with scale configuration on the device under test. The scale details are: 2800 CFM sessions, 2800 BFD sessions, 2043 BGP peers, and 3400 VRF instances. [PR1298985](#)
- The Y.1731 delay measurement is not supported on MPC6. [PR1303672](#)
- In a subscriber management scenario with Dynamic demultiplexer (demux) Interfaces configured, some subscribers belonging to one aggregated Ethernet interface might be migrated to a newly configured aggregated Ethernet interface. Subscribers might fail to access the device after deleting the old aggregated Ethernet configuration. [PR1322678](#)

Layer 2 Features

- This issue affects routers equipped with following line cards: T4000-FPC5-3D, MX-MPC3E-3D, MPC5E-40G10G, MPC5EQ-40G10G, MPC6E, and MX2K-MPC6E. If the router is working as a VPLS PE, due to MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)

Layer 2 Ethernet Services

- After changing an outer vlan-tags, the ifl is getting programmed with incorrect stp state (discarding), so the traffic is getting dropped. [PR1121564](#)

Multiprotocol Label Switching (MPLS)

- When using **mpls traffic-engineering bgp-igp-both-ribs** with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit area border router (ABR) when there are two or more such ABRs. This causes the interdomain RSVP LSPs to break. The RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on the other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. [PR1048560](#)
- This issue occurs when graceful Routing Engine switchover (GRES) is done between the master and backup Routing Engines of different memory capabilities. For example, one Routing Engine has only enough memory to run a routing protocol process (rpd) in 32-bit mode while the other is capable of 64-bit mode. The situation could be caused by using Junos OS Release 13.3 or later with the configuration statement **auto-64-bit** configured, or by using Junos OS Release 15.1 or later even without the configuration statement. Under these conditions, the rpd might crash on the new master Routing Engine. As a workaround, this issue can be avoided by using the CLI command **set system processes routing force-32-bit**. [PR1141728](#)
- When Flow-Label (FL) is enabled for PW, the OAM packets were not sent with Flow-Label because RPD is not aware of the Flow-Label values assigned by PFE software. Hence the packets were getting dropped by PFE at the tail-end PE. The remote PE was expecting the packet with FL and PW label. [PR1217566](#)
- In a CE-CE setup, traffic loss might be observed over a secondary LSP on a primary failover. [PR1240892](#)
- A minimal discrepancy between MPLS statistics and adjusted bandwidth is reported because of the current way of calculating bandwidth. [PR1259500](#)
- It takes longer to set up Layer 3 VPN egress protection starting in Junos OS Release 16.1R1. [PR1278535](#)
- In case of CSPF disabled LSPs, if the Primary path ERO is changed to unreachable strict hop, sometimes the Primary Path stays UP with the old ERO. The LSP does not switch to Standby Secondary. [PR1284138](#)
- Swapping the binding SID between colored and non-colored static SR LSPs might cause rpd to generate a core file. [PR1310018](#)

- There are some LSPs for which a router has link protection available, and the primary link failure is caused by an FPC restart. [PR1317536](#)
- In an LDP over RSVP setup, when the RSVP label-switched paths (LSPs) have protection and a route can be reached through both LDP direct neighbor (IP next hop) and LDP remote neighbor over RSVP LSPs (RSVP next hop), the LDP route next hop is transitioned between the IP next hop and the RSVP LSP next hop. Then RSVP LSP make-before-break (MBB) can happen, and the LDP route might use stale RSVP LSP next hop because of a timing issue. This might cause the rpd process to crash. [PR1318480](#)
- Executing a **restart chassisd** in a MXVC router with the following elements configured might result in a core. 1) IGP OSPF/OSPF3 (area 0, LFA) ISIS (level 2, LFA) LDP synchronization ipv4 and ipv6 2) IBGP dual, redundant route reflection ipv4 and ipv6 3) MPLS LDP (IGP synchronization, track IGP metric) RSVP (node link protection, adaptive, auto bandwidth, refresh reduction) 4) L3VPN OSPF OSPF3 BGPv4 BGPv6 RIPv2 static MBGP NGEN-MVPN l3vpn cnh with ext space any to any hub and spoke MPLS access Ethernet access multicast extranet per vpn and per prefix labels SRX based network address translation SRX based firewall 5) Direct Internet Access EBGp 6) CoS BA/MF classification policing/shaping queuing/scheduling hierarchical queuing/shaping/scheduling 8 traffic classes 7) BFD/OAM/CFM liveness detection 8) Load Balancing L2 aggregate ethernet IP equal cost multi path MPLS equal cost multi path 9) High Availability GRES/NSR ISSU fabric redundancy tail end protection BGP prefix independent convergence edge 10) Security loopback filter arp policers control plane traffic policers urpf check with all feasible paths ttl filtering jflow/ipfix export only SRX based DDOS [PR1352227](#)
- On optimize timer expiry, when the ted version number match indicates a CSPF has already run for the path, if an optimization has not yet been done with that version, it will be run despite the version number match. (Having a per path optimize-seq-no that is updated with ted seq no only on optimization.) When path-cc-updated is false and CSPF fails for optimization, disable the path just like we do for the ones on avoid colors/invalid ERO, making sure this does not interfere with global repair/local reversion [PR1365653](#)
- With static label-switched path (LSP) for MPLS configured with next-hop, the next-hop might get stuck in dead state when only changing the network mask but keeping the IP address unchanged for the outgoing interface via which the LSP next-hop is reachable. [PR1372630](#)

Platform and Infrastructure

- Starting in Junos OS Release 13.1R1 and later, if **no-fast-sync** is used with **configure-private mode**, the commit operation might throw errors after the configuration statements under choice (such as **protocol [ospf pim tcp]**) are added or deleted. Also, after those configuration statements under choice are deleted or added, the whole hierarchy is shown as changed when the **show configuration | compare** command is used. This is a day one issue. [PR1042512](#)
- The **login_getclass: unknown class 'j-idle-timeout'** error is displayed when the user has not configured a timeout value for the root user. If the user has not a configured timeout value, j-idle-timeout entry is not present in the **login.conf** file and an error message is displayed because j-idle-timeout class is not found. To Reproduce: (1) Log in to router as a root user. (2) Clear log messages. (3) Exit and go to CLI mode and perform the **show log messages** command. The login error is logged in the messages.


```
User@MX-re0> start shell user root Password: root@MX-re0:/var/home/lab # cli User@MX-re0> clear
log messages all User@MX-re0> exit User@MX-re0:/var/home/lab # cli User@MX-re0> show log
messages Jan 5 14:55:06.132 MX-re0 mgd[96513]: %INTERACT-6-UI_CHILD_STATUS: Cleanup child
'/usr/libexec/ui/clear-log', PID 96517, status 0 Jan 5 14:55:06.132 MX-re0 mgd[96513]:
%INTERACT-6-UI_FILE_CLEARED: 'messages' logfile cleared by user 'lab' Jan 5 14:55:08.047 MX-re0
mgd[96513]: %INTERACT-6-UI_CMDLINE_READ_LINE: User 'lab', command 'exit' Jan 5 14:55:08.048
MX-re0 mgd[96513]: %INTERACT-6-UI_LOGOUT_EVENT: User 'lab' logout Jan 5 14:55:10.310 MX-re0
cli: %USER-3: login_getclass: unknown class 'j-idle-timeout' <<<<<<<<< Login error Jan 5 14:55:10.318
MX-re0 mgd[96527]: %DAEMON-7: check_regex_add: 1059 regex_add = 0 Jan 5 14:55:10.319 MX-re0
mgd[96527]: %INTERACT-6-UI_AUTH_EVENT: Authenticated user 'root' at permission level 'super-user'
Jan 5 14:55:10.320 MX-re0 mgd[96527]: %INTERACT-6-UI_LOGIN_EVENT: User 'lab' login, class
'super-user' [96527], ssh-connection ", client-mode 'cli' Jan 5 14:55:15.496 MX-re0 mgd[96527]:
%INTERACT-6-UI_CMDLINE_READ_LINE: User 'lab', command 'show log messages ' User@MX-re0>
exit root@MX-re0:/var/home/lab # cat /var/etc/csh.login.inc | grep autologout
root@MX-re0:/var/home/lab # cat /var/etc/login.conf | grep j-idle No idle timeout values are seen in
"/var/etc/csh.login.inc and /var/etc/login.conf" files. PR1097799
```

- On MX2000 routers, the **show chassis hardware detail** might show MICs are installed even after MICs are removed. [PR1216413](#)
- The error message **LUCHIP(5) GUMEM1[77a0] mismatch** might be seen after an MX MPC card with an LU chipset goes offline or online [PR1221195](#)
- When certain hardware transient failures occur on an MQ-chip based MPC, traffic might be dropped on the MPC, and syslog errors **Link sanity checks** and **Cell underflow** are reported. There is no major alarm or self-healing mechanism for this condition. [PR1265548](#)
- MAC addresses are not learned on bridge-domains after an XE/GE interface flap. This issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have XE/GE links toward the downstream switch and LAG bundles as uplinks toward the upstream routers. The XE/GE link is part of the physical loop in the topology. Spanning tree protocols such as VSTP, RSTP, or MSTP are used for loop avoidance. Some MAC addresses are not learned on a device under test when LAG bundles that are part of such bridge domains are flapped and other events such as spanning tree root bridge change occur. [PR1275544](#)
- With a unified ISSU, momentary traffic loss is expected. In EVPN E-Tree, in addition to traffic loss, the known unicast frames can be flooded for around 30 seconds during unified ISSU before all forwarding states are restored. This issue does not affect BUM traffic. As a workaround, nonstop bridging (NSB) can be configured at **set protocols layer2-control nonstop-bridging**. This reduces traffic flooding to around 10 seconds in a moderate setup. [PR1275621](#)
- Due to a transient hardware error condition, the **CPQ Sram parity error** and **CPQ RLDRAM double bit ECC error** syslog errors on an MQCHIP raise a major CM alarm. [PR1276132](#)
- There is an accuracy issue with three-color policers of both types single rate and two rate, where for certain policer rate and burst-size combinations the policer accuracy varies. This issue is present since Junos OS Release 11.4 on all platforms that use a trio ASIC. [PR1307882](#)
- Traffic statistics might not match on PS after clearing the interface statistics. [PR1328252](#)

- On all JunOS platforms, execution of Python scripts through enhanced automation does not work on veriexec images. [PR1334425](#)
- You can configure host syslog from Junos OS guest. Host side: The facility is one of the following keywords: **auth**, **authpriv**, **cron**, **daemon**, **kern**, **lpr**, **mail**, **mark**, **news**, **security** (same as **auth**), **syslog**, **user**, **uucp** and **local0** through **local7**. The keyword **security** should not be used anymore and the mark is only for internal use and therefore should not be used in applications. However, you might want to specify and redirect these messages. The facility specifies the subsystem that produced the message, that is, all mail programs log with the mail facility (LOG_MAIL), if they log using syslog. The priority is one of the following keywords, in ascending order: **debug**, **info**, **notice**, **warning**, **warn** (same as **warning**), **err**, **error** (same as **err**), **crit**, **alert**, **emerg**, **panic** (same as **emerg**). The keywords **error**, **warn**, and **panic** are deprecated and should not be used anymore. The priority defines the severity of the message. Guest side: https://www.juniper.net/documentation/en_US/junos/topics/reference/general/syslog-facilities-severity-levels.html remote : sync the syslog server configuration from Junos OS to Linux and modify rsyslog.conf set vmhost/app-engine syslog host and **set vmhost/app-engine syslog host match xxx**. [PR1341549](#)
- For MPC5 , the inline-ka PPP echo requests are not transmitted when anchor-point is lt-x/2/x or lt-x/3/x in a pseudowire deployment. [PR1345727](#)
- When ephemeral DB instance is configured, if committing changes which are unrelated to IGMP/MLD (such as "set interfaces ge-0/0/1.0 description"), and the number of ephemeral commits reaches to ephemeral DB maximum size, the ephemeral DB purge might happen. Then it would purge all the commits and rollover. On this purge the mgd gives all the applications a FULL COMMIT view. And on this FULL COMMIT view IGMP/MLD deletes all configurations and adds it back again. This might cause PIM to prune the groups on those interfaces and send join messages again. Finally, the multicast traffic flapping and drop might be seen. [PR1352499](#)
- In a Layer 3 VPN topology, when you trace route to a remote PE device for a CE-facing network, you see that the ICMP TTL is expired and receive reply with a source address of only one of the many CE-facing networks. In Junos OS Releases 15.1R5, 16.1R3, and 16.2R1 and onwards there is a kernel sysctl value, **icmp.traceroute_l3vpn**. Setting this to 1 will change the behavior to select an address based on the destination specified in the **traceroute** command. This PR adds the option to the configuration. [PR1358376](#)
- If a tunnel interface is anchored on Trio-based FPC and the 'class-of-service host-outbound-traffic ieee-802.1 rewrite-rules' knob is configured, the host outbound traffic might get dropped when the traffic goes through this tunnel interface. [PR1371304](#)

- Two multicast tunnel (mt) interfaces are seen for each of the PIM neighbors after VPN-Tunnel-Source activation or deactivation. However, ideally, the same tunnel source should be used for both IPv4 and IPv6 address families, if both are using the same PIM tunnel. [PR1281481](#)
- When eBGP multihop sessions exchanging EVPN routes are configured, a core can result due to an internal error. [PR1304639](#)
- In rare cases, RIP replication might fail as a result of performing NSR Routing Engine switchovers when the system is not NSR ready. [PR1310149](#)
- The rpd process generates core files at 0x094680ac in task_reconfigure_complete (ctx=0x9dfe940 <task_args>, seqnum=570) at `../src/junos/lib/libjtask/mgmtlib/./module/task_reconfigure.c:172`. As a workaround, avoid doing additions and deletions in a single commit. Instead, first do the fwdclass deletion, wait for a while, and then do the fwdclass addition. [PR1319930](#)
- In a resource public key infrastructure (RPKI) scenario, the validation replication database might have much more entries than the validation database after restarting the RPKI cache server and the validation session is reestablished. [PR1325037](#)
- When route target filtering (RTF) is configured for VPN routes and multiple BGP session flaps, there is a possibility that some of the peers might not receive the VPN routes after the flapped sessions come up. [PR1325481](#)
- When the **clear validation database** command is issued back-to-back multiple times, it ends up with partial validation database. This eventually recovers after up to 30 minutes (half of the record lifetime), when you do periodical full updates. [PR1326256](#)
- When configuring any cast and prefix segments in SPRING for IS-IS, prefix-segment index 0 is not supported, even though you are allowed to configure 0 as an index. [PR1340091](#)
- Starting in Junos OS Release 16.1 and later, the **show bgp neighbor** command does not show the correct **Last traffic (seconds)** correctly. [PR1361899](#)
- On Junos platform, when openconfig is running with sensor for `/network-instances/network-instance/protocols/protocol/bgp`, changing BGP import or export policy may cause rpd core. [PR1366696](#)
- If IS-IS shortcut is enabled and ISIS "topologies ipv6-unicast" is configured, when any link with no IPv6 address configured in the MPLS LSP path is flapping (or bring down and then up), the route entry go through this flapping link might be missing for about 10 minutes, which might lead to traffic loss. The issue is because when the flapping link is down and then up, the flash route update checks both IPv4 and IPv6 address family, since IPv6 is not configured for this link, the flash route update is not triggered, hence the route entry is missing. [PR1372937](#)

Services Applications

- We do not recommend configuring the ms-interface when AMS bundle in one-to-one mode has the same member interface. [PR1209660](#)

VPNs

- A VLAN-CCC logical interface for l2ckt remains in CCC-Down when switching from l2ckt to EVPN-VPWS, unless it is deactivated and re-activated manually. [PR1312043](#)

SEE ALSO

[New and Changed Features | 93](#)

[Changes in Behavior and Syntax | 126](#)

[Known Behavior | 136](#)

[Resolved Issues | 157](#)

[Documentation Updates | 199](#)

[Migration, Upgrade, and Downgrade Instructions | 200](#)

[Product Compatibility | 207](#)

Resolved Issues

IN THIS SECTION

● [Resolved Issues: 17.4R2 | 158](#)

● [Resolved Issues: 17.4R1 | 184](#)

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: 17.4R2

Application Layer Gateways (ALGs)

- IKEv2 negotiation might fail with IKE ESP ALG enabled in an IKEv2 redirection scenario. [PR1329611](#)

Authentication and Access Control

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

Class of Service (CoS)

- CoS wildcard configuration is applied incorrectly after a router restart. [PR1325708](#)
- Remove CoS IDL from the jet IDL package. [PR1347175](#)
- The Routing Engine might get into amnesiac mode after restarting if **excess-bandwidth-share** is configured. [PR1348698](#)
- The aggregated Ethernet link-protection feature is not supported. [PR1355498](#)

EVPN

- EVPN traffic mapping to specific LSPs is not working. [PR1281415](#)
- The rpd might crash on platform using junos with evpn and nsr enabled after restarting the rpd process in EVPN environment [PR1320408](#)
- An EVPN discard route is installed on the local provider edge (PE) device after connection flaps on a remote PE device in a multihome EVPN topology. [PR1321125](#)
- If host is multihomed then all PEs should install the /32 host IP address pointing to its local IRB interface as long as its local multihomed ES interface is up. [PR1321187](#)
- The rpd crash might happen during EVPN/VXLAN configuration changes. [PR1321839](#)
- RPD crash on backup Routing Engine if NSR and IS-IS SR enabled. [PR1323980](#)
- The FPC might crash after deleting the VPLS configuration. [PR1324830](#)
- A core link flap might result in an inconsistent global MAC count. [PR1328956](#)
- On a deactivated end system identifier (ESI) for PS at a physical interface level, the rpd process generates core files for EVPN VPWS PWHT. [PR1332652](#)
- On doing **restart routing**, the rpd process might generate core files on a PE router that has a EVPN-VXLAN configuration. [PR1333331](#)
- MPLS label leak leads to label exhaustion and the rpd process crash [PR1333944](#)
- In an EVPN scenario with nonstop active routing (NSR) enabled, the rpd crashes and generates core files on the backup Routing Engine while any configuration changes on the master Routing Engine. [PR1336881](#)
- The rpd process might crash when executing CLI command "show route evpn-ethernet-tag-id" [PR1337506](#)

- In an EVPN-VXLAN environment, the BFD flap causes the VTEP to flap, causing the Packet Forwarding Engine to crash. [PR1339084](#)
- Traffic loss might be observed in an EVPN-VPWS scenario if the remote PE's interface comes down. [PR1339217](#)
- On EVPN-VXLAN scenarios, the traffic might get black-holed to interfaces that are down, but LACP is up. [PR1343515](#)
- The rpd might crash if the IRB interface and routing instance are deleted together in the same commit. [PR1345519](#)
- Traffic might be lost on a Layer 2 and Layer 3 spine node in a multihome EVPN scenario. [PR1355165](#)
- EVPN IRB configured with **no-gratuitous-arp-request** is still sending gratuitous ARP. [PR1356360](#)
- The rpd might crash if the EVPN instance refers to a vrf-export policy which doesn't have 'then community'. [PR1360437](#)
- Proxy ARP may not work as expected in an EVPN environment. [PR1368911](#)

Forwarding and Sampling

- The pfd process generates a core file in `pfed_process_session_state_notification_msg, pfed_timer_manager_c::remove_serv_id, pfed_delete_timer_id_by_serv_sid (serv_sid=0, serv_info=0x0)` at `../../../../src/junos/usr.sbin/pfed/pfed_timer.cc:16`. [PR1296969](#)
- Remote CE1 MAC address might take more time to clear after clearing MAC. [PR1304866](#)
- The dfwd process might crash during execution of **show firewall templates-in-use** command. [PR1305284](#)
- The second archive site in the accounting-file configuration is not used when the first one uses SFTP and is not reachable. [PR1311749](#)
- Accounting files with no records might be unexpectedly uploaded to the archive site. [PR1313895](#)
- The FPC CPU might reach 100 percent constantly if shared bandwidth policer is configured. [PR1320349](#)
- The error messages about `dfw_gencfg_handler` might be seen during a unified ISSU. [PR1323795](#)
- Ukernel leaks 6x40 bytes heap nodes upon each IPC path when handshaking or establishment occurs between I2alm and I2ald. [PR1326921](#)
- DHCP service crashes after the device is set to factory default by zeroize. [PR1329682](#)
- Some firewall filter counters might not be created in SNMP. [PR1335828](#)
- The error logical interface under VPLS might be blocked after MAC moving if the logical interfaces are on the same physical interface. [PR1335880](#)
- In EVPN-VXLAN **clear ethernet-switching table** might not work correctly. [PR1341328](#)
- Junos allows firewall filters with the same name under **edit firewall** and **edit firewall family inet** hierarchy levels [PR1344506](#)

- Commit failed when attempting to delete any demux0 unit numbers that are greater or equal to 1000000000. [PR1348587](#)
- The remote MAC might not be added in the forwarding table, which will cause a traffic drop in an EVPN scenario with RSVP and CBF configured. [PR1353555](#)
- The backup Routing Engine is writing dummy interface accounting records. [PR1361403](#)

General Routing

- In timing hybrid mode, MX Series MPC2 cards are not working with ACX with VLAN (native-vlan-id). [PR1076666](#)
- An rpd memory leak is caused by repeated RSVP reservation state block (RSB) deletes. [PR1115686](#)
- No warning is raised when the bridge family is configured with interface-mode trunk but without vlan-tagging or flexible-vlan-tagging. [PR1154024](#)
- An unexpected **MobileNext Gateway Activation license** alarm is observed when TDF gateway is configured. [PR1162518](#)
- The replacement PIC might bounce when PIC PB-4OC3-4OC12-SON-SFP (4x OC-12-3 SFP) is replaced with PB-4OC3-1OC12-SON2-SFP (4x OC-3 1x OC-12 SFP) and a CLI commit is made. [PR1190569](#)
- Agentd process crashes with core-dump [PR1197608](#)
- The **Unable to deregister sub error (131072) for error(0x1b0001) for module MIC** error messages are seen on the MPC5E card. [PR1221337](#)
- The error log **cc_mic_irq_status: CC_MIC(5/2) irq_status(0x1d) does not match irq_mask(0x20), enable(0x20), latch(0x1d)** is seen continuously for MIC-3D-4OC3OC12-1OC48. [PR1231084](#)
- The **chassisd[9132]: LIBJSNMP_NS_LOG_NOTICE: NOTICE: netsnmp_ipc_client_connection: unix connection error: socket(-1) main_session(0x9812f80)** error messages are seen after a chassis-control restart. [PR1243364](#)
- The GNF sometimes resets its MPC type 9 at NSR at a high scale. [PR1259910](#)
- On a vMX FPC, the software FPC might restart unexpectedly with the following message: **panic (format_string=format_string@entry=0x9e509c4 "Thread %s attempted to %s with irq priority at %d\n")**. [PR1263117](#)
- The **show chassis FPC** command does not show temperature. [PR1263315](#)
- The load-based throttling functionality is not enabled by default. [PR1271739](#)
- Flexible PIC concentrator (FPC) crash/reboot is observed when bringing up about 12K Layer 2 Bit Stream Access(L2BSA) subscribers simultaneously. [PR1273353](#)
- Error messages observed on vty session while running script for IGMP Snooping over EVPN-VXLAN. [PR1276947](#)
- On an MX104 platform with GRES enabled, the chassis network-services might not get set as "Enhanced-IP". [PR1279339](#)

- BSYS logs messages are reporting that GNF owned PICs do not support power off configuration at commit when no such configuration is present. [PR1281604](#)
- The kernel might crash when an NSR enabled device has BGP peer flapping. [PR1282573](#)
- The enhancement of reporting total SBE errors when the corrected single-bit errors threshold of 32 is exceeded for MPC7E/MPC8E/MPC9E. [PR1285315](#)
- The LC, PFH, and Packet Forwarding Engine interfaces do not come up on Routing Engine 1. [PR1285606](#)
- The missing statement **Shared bandwidth policer not supported for interface ge-x/x/x** is seen during a failed commit in Junos OS Release 16.1R3. [PR1286330](#)
- The oneset or leaf-list configuration might not get deleted with the delete operation through JSON. [PR1287342](#)
- PPPoE cannot dial in due to all padi dropped as "unknown iif" when deactivated/activated AE configuration. [PR1291515](#)
- During PPPoE subscriber login errors like **vbf_flow_src_lookup_enabled** and **Failed to find iff structure, ifl** were seen on FPC. [PR1294710](#)
- The KRT queue might be stuck with the **RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0** error. [PR1295756](#)
- Some random number of ports on a 10-Gigabit MPC7E card might not come up after the remote system or line card restarts or interface flaps. [PR1298115](#)
- The log message about the shutdown time is incorrect when the system exceeds chassis over the temperature limit. [PR1298414](#)
- When the subscriber limit feature is configured, any new login request after the maximum number of subscribers is denied. [PR1298924](#)
- The error messages about PEM might be seen in the MX Series platform with AC PEM. [PR1299284](#)
- A chassisd core file is seen after the insertion of REMX2K-X8-64 in MX2000 line routers with the older RE-S-1800x4. [PR1300083](#)
- The ICMP/ICMPv6 error messages might be discarded while forwarding through an AMS interface. [PR1301188](#)
- Reported same IFD KV by two different sensors. [PR1301858](#)
- The rpd might crash when NSR is enabled and routing-instance specific configurations are committed. [PR1301986](#)
- Continuous interface flapping might lead to an unwanted MIC reset. [PR1302246](#)
- The multicast resolve-rate value might go back to default after system upgrade or reboot. [PR1303134](#)
- Internal latency is high during the initial subscription of sensors. [PR1303393](#)
- Fan speed changes frequently on MX Series after an upgrade to Junos OS software. [PR1303459](#)

- The fabric planes might go into "check" state after restarting the line cards with SFB2 used on the MX2010 or the MX2020. [PR1304095](#)
- The **start shell pfe network fpc** command is not working on the MX960. [PR1306236](#)
- /Frame: messages might be seen with Telemetry enabled. [PR1308513](#)
- FPC syslog errors with **pfeman_inline_ka_steering_gencfg_handler: nh not found** could mean that steering rules are not installed correctly. [PR1308884](#)
- After a smooth upgrade from SFB to SFB2, if one plane/SFB is restarted, link training fails between those planes and MPC6 cards. [PR1309309](#)
- First access-request is failing for L2BSA subscribers when changing the MTU of LACP aggregate Ethernet A10NSP interface. [PR1309599](#)
- Subscribers might not be able to access the device if dynamic VLAN is used. [PR1309770](#)
- Ninety percent of subscribers might go down after a unified ISSU from Junos OS Release 16.1 to Junos OS Release 17.3. [PR1309983](#)
- Local IPv6 interface address from the NDRA prefix is not removed from the service interface when the subscriber dual-stack session is removed. [PR1310752](#)
- The utilization of "commit check" just after setting the master-password can trigger an improper decoding of configuration secrets. [PR1310764](#)
- After guest network functions (GNFs) Routing Engine switches mastership as expected, the rpd might be unresponsive. [PR1310765](#)
- The incorrect error number might be reported for syslog messages with a prefix of %DAEMON-3-RPD_KRT_Q_RETRIES. [PR1310812](#)
- Fragmented UDP packet might be incorrectly parsed as a uBFD packet and dropped. [PR1311134](#)
- Suppress chassis alarm for switched off PEMs. [PR1311574](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- The rpd process generates a core file after multiple session flaps on a scale setup. [PR1312169](#)
- The PEM alarms and I2C failures are observed on MX240, MX480, MX960, EX92, and SRX5K. [PR1312336](#)
- A false over temperature SNMP trap could be seen when using MPC5/6/7/8/9 on an MX2020. [PR1313391](#)
- The IPv6 router-solicit (RS) packets are dropped in nondefault RI, but for default RI it is working. [PR1313722](#)
- The **show version detail** command gives severity **error log traffic-dird[20126]: main: swversion pkg: 'traffic-dird' name: 'traffic-dird' ret: 0**. [PR1313866](#)
- The jdmd subsystem is not responding after an upgrade. [PR1313964](#)

- The mspmand process generates a core file because of a flow-control seen while clearing CGNAT+SFW sessions. [PR1314070](#)
- When ccc is configured on a umic interface, ARP is not resolving and observing traffic loss. [PR1314149](#)
- The JDM link is incorrectly shown to be up when the underlying physical link is down. [PR1314180](#)
- The **show version detail | no-more** command hangs for more than 120 seconds in the master Routing Engine and more than 60 seconds in the backup Routing Engine. [PR1314242](#)
- The smgd process generates a core file with reference to bbe_cos_ifl_publish() bbe_cos_if.c:6543. [PR1314651](#)
- The rpd process might crash in a MoFRR scenario. [PR1314711](#)
- For MPC7E, there is an **IR-mode** commit failure. [PR1314755](#)
- The L2TP LAC might drop packets that have an incorrect payload length while sending packets to the LNS. [PR1315009](#)
- Continuous logs from vhlclient are seen for all the commands executed. [PR1315128](#)
- FPC crash is observed when a route has unilist next-hops in a RSVP scenario. [PR1315228](#)
- The **show version detail** command gives severity **error log mobiled: main Neither BNG LIC nor JMOBILE package is present,exit mobiled**. [PR1315430](#)
- The **show version detail** command might generate severity error log **main: name: SRD ret: 0**. [PR1315436](#)
- Sensors belong to the same producer with identical reporting interval are not streamed in parallel [PR1315517](#)
- The rpd process generates a core file when a **show route inetcolor.0** command is executed from the CLI. [PR1316078](#)
- The fan speed might frequently keep changing between normal and full for the MX Series platform. [PR1316192](#)
- The demux interface sends a neighbor solicitation with source link-address of all zeros 00:00:00:00:00:00 MAC. [PR1316767](#)
- The **show configuration <> | display json** command might not be properly enclosed in double quotes. [PR1317223](#)
- Linux-based microkernel might panic due to a concurrent update on mutable objects. [PR1317961](#)
- CoA shaping rate is not applied successfully after a unified ISSU from Junos OS Release 15.1R6.7 to Release 16.1R6.2. [PR1318319](#)
- The rpd process might crash when the link flaps on an adjacent router. [PR1318476](#)
- The bbe-smgd process might crash after performing GRES. [PR1318528](#)
- The FPC crashes on a configuration change for the Packet Forwarding Engine sensors. [PR1318677](#)

- Changed text reported in the **show chassis hardware** output for CFP2-DCO optical transceivers. [PR1318901](#)
- MS-MPC and MS-MIC might crash after a new IPsec tunnel is added. [PR1318932](#)
- The MPC with specific failure hardware might impact other MPCs in the same chassis. [PR1319560](#)
- The kernel might generate a core file if the number of routing instances created are more than 256. [PR1319781](#)
- The task replication might not be complete to certain network protocols after multiple GRES. [PR1319784](#)
- The error log message of **MIB2D_COUNTER_DECREASING: pfes_stats_delta: counter** might be seen on VMX. [PR1319996](#)
- Loading xmlproxy YANG files cause telemetry session and some daemons to restart. [PR1320211](#)
- The chassis MIB SNMP OIDs for VC-B member chassis are not available after an MX Series Virtual Chassis unified ISSU. [PR1320370](#)
- The **show subscriber summary** command displays an incorrect terminated subscriber count. [PR1320717](#)
- The PPP inline keepalive does not work as expected when CPE aborts the subscriber session. [PR1320880](#)
- The rpd process crashes during the BGP configuration change and telemetry streaming with OpenConfig. [PR1320900](#)
- MX Series routers send the IPv6 router advertisements and the DHCPv6 advertisements before sending IPCPv6 ACK from CPE. [PR1321064](#)
- CoS is not applied to the Packet Forwarding Engine when the VCP link is added. [PR1321184](#)
- The bbe-smgd process generates core files after massive clients log out and log in, in a PPPoE dual stack subscriber scenario. [PR1321468](#)
- A CoA-NAK with "Error-Cause = Invalid-Request" is sent back to the RADIUS server when a drop policy is applied under radius-flow-tap in an L2TP subscriber scenario. [PR1321492](#)
- The **show system schema module hierarchy** command is broken in the CLI. [PR1321682](#)
- In commit fast-synchronize mode, the commit operation might get stuck after the commit check is performed. [PR1322431](#)
- The rpd process might crash when two next hops are installed with the same next-hop index. [PR1322535](#)
- The rpd process might crash when the OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)
- MS-MIC interface IFLs remain down after many iterations of offline/online. [PR1322854](#)
- An incorrect output is observed while verifying the command **show subscribers client-type vlan subscriber-state active logical-system default routing-instance default**. [PR1322907](#)
- NCP Conf-Ack/Conf-Req packets might be dropped constantly from the MLPPP client. [PR1323265](#)

- CLI commands in **show system subscriber-management route routing-instance <XXX>** hierarchy show unexpected outputs. [PR1323279](#)
- JDM Management is unreachable after flapping physical JDM and GNF/VNF management interfaces. [PR1323519](#)
- The **request vmhost halt routing-engine other** command does not halt the backup Routing Engine. [PR1323546](#)
- Memory leaks in the MGD-API process during Get API Requests and Error Handling during Set API Request. [PR1324321](#)
- Subscribers might fail to log in after the interface is deactivated or activated. [PR1324446](#)
- A memory leakage is seen in the mosquito-nossl process. [PR1324531](#)
- The SNMP interface filter does not work when "interface-mib" is part of the dynamic-profile. [PR1324573](#)
- KRTQ entries are waiting in an async queue. [PR1324669](#)
- The VLAN rewrite function might put the wrong VLAN ID when an Ethernet OAM is configured on DPCE cards. [PR1325070](#)
- The SNMP values might not be increased monolithically. [PR1325128](#)
- The MPC cards might drop traffic under a high temperature. [PR1325271](#)
- Non-MACsec interfaces are impacted when first time MACsec is configured on one of the interfaces or respective FPC is rebooted. [PR1325282](#)
- IS-IS adjacency fails to establish because packets drop on Packet Forwarding Engine. [PR1325311](#)
- MACsec session might fail to establish on MX10003. [PR1325331](#)
- The VLAN demux interface does not respond to the ARP request in a subscriber scenario with an MX Series router after Junos OS Release 15.1 with subscriber-management enabled. [PR1326450](#)
- MACsec MKA transmit Interval is changed to the upper limit. [PR1326526](#)
- In an MX Series BNG, the CoS service object is not deleted properly for TCP and scheduler. [PR1326853](#)
- Some **show** commands were issued twice when a **request support information** is executed. [PR1327165](#)
- With auto-installation USB configured, interface related commits might not take effect due to a dcd error. [PR1327384](#)
- Minor alarm **LCM Peer Connection un-stable** is observed on an MX150 after the chassisd process startup or restart. [PR1328119](#)
- Only 5.5M TCP sessions can be established for a NAPT44_SFW_APP_EIM/EIF configuration on an MS-MIC. [PR1328510](#)
- The following message is constantly logged: **fm_feacap_sys_feature_get:Attribute DB init not yet done, reading from pvid (id: 18)**. [PR1328868](#)

- For the **show class-of-service interface demux0 <demux interface>** command, the Adjustment overhead-accounting mode does not provide the expected output. [PR1329212](#)
- When an AMS bundle has a single MAMs added to it, the subinterfaces do not recover after the subinterface has been disabled. [PR1329498](#)
- Host-outbound traffic is not rewriting IEEE-801.pbits for a dynamic subscriber IFL over a PS interface. [PR1329555](#)
- SNMP walks of Interfaces related MIB objects are slower than expected in a scaled configuration. [PR1329931](#)
- The **show services nat mappings address-pooling-paired** command times out and fails. [PR1330207](#)
- The **Too many supplies missing in Lower/Upper zone** alarm flaps (set/clear) every 20 seconds if a zone does not have the minimum required PSMs. [PR1330720](#)
- The packets might be dropped if one route is advertised by BGP, where the session is established through the subscriber interface. [PR1330737](#)
- The rpd process generates core files on the new backup Routing Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler** after disabling NSR+GRES [PR1330750](#)
- The FPC might be wedged when the LSQ interface receives fragmented packets. [PR1330998](#)
- Under very high scale, replication is not started for BGP and is stuck in progress for RIP and LDP after a NSR. [PR1331145](#)
- Chassis FPC temperature with non-NEBS optics is higher after a software upgrade. [PR1331186](#)
- The bbe-smgd process might crash after executing the **clear ancp access-loop circuit-id <circuit id of interface set>** command. [PR1332096](#)
- Inaccurate Jflow records might be seen for an output interface and a next hop. [PR1332666](#)
- On an MX150 platform, the **set chassis alarm management-ethernet link-down ignore** command is not ignoring the alarm for the FPC Mgt 0 interface. [PR1332799](#)
- The subinfo process might crash and it might cause the PPPoE subscribers to get disconnected. [PR1333265](#)
- JDID thrashes continuously and continuous log messages are observed in syslog. [PR1333632](#)
- Active/active (A/A) Multihoming EVPN VXLAN in some race conditions can trigger constant high CPU usage on the backup Routing Engine. [PR1334235](#)
- Two subscribers cannot reach the online state at the same time if they have an identical Frame-Route attribute value. [PR1334311](#)
- MPC8E or MPC9E reports high temperature alarms and fan speed changing continuously through full and normal speed iterations. [PR1334750](#)
- The rpd process crashes when performing the BGP configuration change. [PR1334846](#)

- The UID limit is reached in a large-scale subscriber scenario. [PR1334886](#)
- When using the **show subscribers** command and when the FPC number has two digits, the interface and IPv6 address get connected together for DHCPv6 PD. [PR1334904](#)
- The IPsec rule might not work if both IPv4 ANY-ANY term and IPv6 ANY-ANY term are configured for it. [PR1334966](#)
- Traffic drops on the MX Series LNS because of software error/unknown family exception when traffic goes to or comes from an MLPPP subscriber if 'routing-services' is presented in the dynamic-profile used by this subscriber. [PR1335276](#)
- The master LED glows on the master and the backup RCB, while performing the image upgrade on the master with GRES/NSR enabled. [PR1335514](#)
- There are hitless key chain rollover feature limitations on MIC-MACSEC-MRATE. [PR1335644](#)
- The RIP route updates might be partially dropped when NSR is enabled. [PR1335646](#)
- The MAC_STUCK might be seen on the MS-MPC or the MS-MIC. [PR1335956](#)
- JET application might not respawn after a normal exit. [PR1336107](#)
- Subscriber might experience SDB DOWN event and drop the clients' connections when issuing the **show subscribers** commands. [PR1336388](#)
- On an MX2000 with an SFB card installed, high traffic volume on an MPC7E, MPC8E or MPC9E might cause traffic drops with cell underflow messages. [PR1336446](#)
- The bbe-smgd might crash when doing a CoS configure of the interface set. [PR1336852](#)
- The **set protocols lldp neighbour-port-info-display port-id** command might not take effect. [PR1336946](#)
- The error log message **sdb_db_interface_remove: del ifl:si-<index> with licnese cnt non zero on** can be seen on LTS during a subscriber logout. [PR1337000](#)
- AI-script does not get an auto reinstall upon a Junos OS upgrade on a next-generation Routing Engine. [PR1337028](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- The MX10003 MPC offline button is not effective. [PR1340264](#)
- The CLI shows CB states online after pressing RCB offline button for 4 seconds or more. [PR1340431](#)
- Upon a reboot from a cold state (or after a Junos OS software upgrade), MX150 might not forward multicast traffic, including VRRP packets, from the Packet Forwarding Engine to the Routing Engine. [PR1341044](#)
- There might be traffic loss on some subscriber sessions when more than 32,000 L2TP subscriber sessions are anchored in the ASI interface. [PR1341659](#)
- The reboot of the Routing Engine might occur if the PPPoE interface is configured over an aggregated Ethernet or RETH interface. [PR1341968](#)

- With discard Interfaces (configured with IGMPv3), the KRT queue gets stuck while deleting a multicast next hop (MCNH) with the error **EPERM -- Jtree walk in progress**. [PR1342032](#)
- An SNMP walk might fail for LLDP-related OIDs. [PR1342741](#)
- The vFPC might get absent resulting in the total loss of traffic. [PR1343170](#)
- Support required for the **show system resource-monitor subscribers-limit chassis extensive** command on Summit. [PR1343853](#)
- An MX Series router is sending IPv6 RA and the DHCPv6 advertisements before IPCPv6 ACK from CPE. [PR1344472](#)
- Unable to route over an RLT interface after upgrading from Junos OS Release 15.1 to Release 17.3. [PR1344503](#)
- The ancpd process might generate a core file when clearing ancp subscribers in a scaled scenario when enhanced-ip is configured. [PR1344805](#)
- The Framed-Route "0.0.0.0/0" will not be installed on an MX Series platform with Junos OS enhanced subscriber management releases. [PR1344988](#)
- The ARP packet uses the VRRP/virtual-gateway MAC address in the Ethernet header instead of the IRB MAC address. [PR1344990](#)
- A dot1x re-authentication issue. [PR1345365](#)
- The rpd process crash might be seen if **no-propagate-ttl** is set in a routing instance that has a specific route. [PR1345477](#)
- The MAC address of multiple interfaces are found to be duplicates. [PR1345882](#)
- The Routing Engine model changed from JNP10003-RE1 to RE-S-1600x8. [PR1346054](#)
- New PPPoE users might fail to log in. [PR1346226](#)
- The AC system error counter in the **show pppoe statistics** command is not working. [PR1346231](#)
- The VCCP-ADJDOWN detection is delayed on the Virtual Chassis backup router (VC-Bm) when deleting one VCP link on Virtual Chassis master router (VC-Mm). [PR1346328](#)
- Statistics daemon PFED might generate a core file on an upgrade between certain releases. [PR1346925](#)
- The twice-napt-44 sessions are not syncing to the backup SDG with stateful sync configured. [PR1347086](#)
- IPv6 MAC resolve will fail if the DHCPv6 client uses a non-EUI64 link-local address. [PR1347173](#)
- Remove libstdc++ dependency on the hypervisor to install the JDM rpm/deb package. [PR1347921](#)
- There is an issue with handling the community_action ("add") in a RPC call. [PR1348082](#)
- The FPC might crash due to a MIC error interrupt hogging. [PR1348107](#)
- Packet loop is detected when virtual routing and forwarding (VRF) multipath is enabled with **equal-external-internal** under an Layer 3 VPN instance and **install-nexthop** is enabled in a forwarding-table export policy regarding that Layer 3 VPN route. [PR1348175](#)

- A chassisd memory leak is observed on an MX10003 and an MX204 platform and it would eventually cause a Routing Engine switchover and crash. [PR1348753](#)
- The DHCPv6 solicit packet might be dropped on an MX Series Virtual Chassis with L2TP LNS when the packet is received over a VCP port and the anchor si- interfaces exist on the same Packet Forwarding Engine as the VCP port. [PR1348846](#)
- The **Major PEM 0 Input Failure** major alarm might be observed for a DC PEM. [PR1349179](#)
- The mspmand process might crash when executing the **show services nat deterministic-nat nat-port-block** command. [PR1349228](#)
- The mgd process generates a core file because of an issue in the nsindb infra. [PR1349288](#)
- The pccd might crash after a delegated LSP is removed in PCEP scenario. [PR1350240](#)
- The MTU value for subscriber's interface might be programmed incorrectly if **routing-services** or **protocol pim** is configured in dynamic-profile. [PR1350535](#)
- The subinfo process might crash when executing the **show subscribers address <> extensive** command for a DHCP IPv6 address. [PR1350883](#)
- The VCP port might not come back up after removing and adding it again. [PR1350845](#)
- The PPE Errors async xtxn error is observed when FPC is restarted or removed. [PR1350909](#)
- The pfed process might consume high CPU if subscriber or interface statistics are used at large scale. [PR1351203](#)
- A high CPU usage for the bbe-smgd process might be seen when L2BSA subscribers get stuck. [PR1351696](#)
- After GRES, the BGP neighbors at the master Routing Engine might reset and the BGP neighbors at the backup Routing Engine might take a long time to establish. [PR1351705](#)
- The bbe-smgd process might restart in a subscriber environment. [PR1352546](#)
- The DHCP relay-reply packets are dropped in the DHCPv6 relay scenario. [PR1352613](#)
- The offlining of MIC6-100G-CFP2 MIC through the CLI command might trigger the FPC card to crash. [PR1352921](#)
- The rpd process is permanently overusing CPU due to a logical system configuration commit. [PR1353548](#)
- Traffic interruption is observed after multiple Routing Engine switchover. [PR1354002](#)
- The dfw_bbe_filter_bind:1125 BBE filter bind type 0x84 index 167806251 returned 1. [PR1354435](#)
- The rpd might generate core files when adding an inter-region template in routing-instances. [PR1354629](#)
- Aggregated Ethernet operational state goes up even though some of the member interfaces configured under the Aggregated Ethernet are down. [PR1354686](#)
- The ifinfo process might crash in an MX BNG running an L2BSA service. [PR1354712](#)
- JSSCD static-subscribers do not properly update firewall information on the Packet Forwarding Engine when dynamic configuration changes are made to active subscribers. [PR1354774](#)

- A memory leak is found in agentd while running valgrind. [PR1354922](#)
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100Gbps). [PR1355168](#)
- Packets destined to Routing Engine might be dropped in the kernel when LACP is configured. [PR1355299](#)
- The fabric chip failure alarms are observed in a GRES scenario. [PR1355463](#)
- Syslog messages : **ui_client_connect_to_kmd_instance: KMD-SHOW connect to kmd-instance failed kmd-instance RE, fpc slot 0, pic slot 0.** [PR1355547](#)
- The flex-flow-sizing is not working on an MX204. [PR1356072](#)
- The rpd process will crash when issuing the **show dynamic-tunnels database terse** command for RSVP automatic mesh tunnels. [PR1356254](#)
- The L2C messages from PEM/PSM are reported if SNMP is enabled. [PR1356259](#)
- The **show pppoe underlying-interfaces** command in a scaled environment might cause a bbe-smgd memory leak. [PR1356428](#)
- The bbe-smgd generates core files in recursive loop between functions bbe_autoconf_if_l2_input and bbe_if_l3_input. [PR1356474](#)
- DHCP subscribers fail after a reconfiguration of the port from tagged to un-tagged mode. [PR1356980](#)
- Upgrading from Junos OS Release 15.1F2-S20 to Junos OS Release 15.1X12 using **validate** throws a Fabric Mixed Mode error. [PR1357423](#)
- A Routing Engine switchover during backup Routing Engine being not GRES ready might cause linecard restart, which causes the Routing Engine kernel to crash and multiple chassisd crashes. [PR1357427](#)
- Traffic might be sent to a wrong RLT member interface after RLT switchover. [PR1358320](#)
- An incorrect traffic load balance might be seen even if **locality-bias** is configured on MX Series Virtual Chassis. [PR1358635](#)
- FPC was offline with the **Disconnected after ISSU and before switchover** message during a unified ISSU from Junos OS Release 17.4 to Junos OS Release 18.2. [PR1359282](#)
- The **FRU-model-number** is not displayed for a few FRUs in the component sensor for an MX10008 and an MX10003. [PR1359300](#)
- The IPv6 subscriber might fail to access network. [PR1359520](#)
- The rpd cores at **Assertion failed rpd[10169]: file**
`"../../../../../../../../src/junos/usr.sbin/rpd/lib/rt/rt_attrib.c, line 3329: rt_template_get_rtn_ngw(nhp)`
`<= 1` on doing Routing Engine switchover with SRTE routes. [PR1360354](#)
- The rpd scheduler slip might be seen when frequently deleting, modifying, and adding groups which are applied on top level. [PR1361304](#)
- Spontaneous bbe-smgd core file might be seen on the backup Routing Engine. [PR1362188](#)

- The route stuck might be seen after BGP neighbor and route flapping. [PR1362560](#)
- Unexpected DCD_PARSE_ERROR_SCHEDULER messages are logged when MS-MPC/MS-MIC is brought offline or online. [PR1362734](#)
- A quick memory leak in bbe-smgd is observed if the dynamic profile variable name and the default associated value are configured to be the same. [PR1362810](#)
- The non-default routing-instance is not supported correctly for NTP packet in subscriber scenario. [PR1363034](#)
- Traffic destined to the MAC or IP address of VRRP VIP gets dropped on the platforms which have common TFEB terminals such as MX5/10/40/80/104. [PR1363492](#)
- A **pmbus_read_volt: sfb-07 - MAX20751-PF1-0.9v: pmbus** read failed for cmd 0x8b. [PR1363587](#)
- The xmlproxyd for internal interfaces is reporting uint32 instead of uint64. [PR1363766](#)
- The l2circuit on MPC7E/8E/9E with asynchronous-notification and ccc configured might keep flapping when the circuit is going up. [PR1363773](#)
- A traffic loop might occur even though that port is blocked by RSTP in a ring topology. [PR1364406](#)
- The traffic is still forwarded through the member link of an Aggregated Ethernet bundle interface even with **Link-Layer-Down** flag set. [PR1365263](#)
- Midplane attributes are not getting exported. [PR1365303](#)
- The next-hop of MPLS path might be stuck in hold state which might cause traffic loss. [PR1366562](#)
- Snmp mib walk for udp flood gives different output statistics than CLI. [PR1366768](#)
- The **show system resource-monitor fpc** might show non-existing Packet Forwarding Engine. [PR1367534](#)
- The **commit** or **commit check** might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)
- Subscriber filter not removed from the Packet Forwarding Engine when routing-services are enabled in the dynamic profile on an L2TP LNS. [PR1369968](#)
- Kernel crash might be seen after committing DEMUX related configuration. [PR1370015](#)
- The packet which size exceeds 8000 might be dropped by MS-MPC in ALG scenario. [PR1370582](#)
- FPC high CPU utilization or crash during hot-banking condition. [PR1372193](#)
- PCE initiated LSPs remain **Control status became local** after removing PCE configuration. [PR1374596](#)

High Availability (HA) and Resiliency

- After server links flap, the GNFs associated with the ports on the Control Board show the status message: **Switchover Status: Not Ready** message. [PR1306395](#)
- The ksyncd process might crash continuously on the new backup Routing Engine after performing GRES. [PR1329276](#)

- There is insufficient available space on the hard disk lead by the crashinfo files that are generated by the ksyncd process when GRES is configured in a large-scale configuration scenario. [PR1332791](#)
- VC-Bm cannot sync with VC-Mm when the the Virtual Chassis splits then reforms. [PR1361617](#)

Infrastructure

- The syscalltrace.sh might create a huge output file, which might cause the router to run out of storage space. [PR1306986](#)
- A cleanup at the thread exit is causing memory leaks. [PR1328273](#)
- On all Junos OS platforms, on a port configured with both dot1x static mac by-pass and normal authentication, the hosts configured for static mac by-pass may not be able to send traffic. [PR1335125](#)
- The kernel might crash and the system might reboot in an SNMP query reply scenario. [PR1351568](#)
- Junos OS is no longer going to database prompt at ~ +Ctrl+b. [PR1352217](#)

Interfaces and Chassis

- RL-dropped packets are not displayed by **show interfaces <ifl> detail/extensive** commands. [PR1249164](#)
- Out of sequence packets seen with LSQ interface. [PR1258258](#)
- L2TP subscribers might not be cleared if the access-internal routes fail to install. [PR1298160](#)
- Some CFM sessions do not come up after a DUT with MPC9 line cards is rebooted with scale configuration. [PR1300515](#)
- The MPC CPU might reach 100 percent when optical transport network (OTP) ultra forward error correction (UFEC) is configured. [PR1311154](#)
- Observing jpppd core **telemetry_start_timer,mosquitto_handle_connack,telemetry_mqtt_publisher** [PR1311396](#)
- The jpppd process generates a core file at **telemetry_start_timer,mosquitto_handle_connack,telemetry_mqtt_publisher**. [PR1311396](#)
- The ifinfo process might crash and generate core files when executing the **show interfaces name** command with a name greater than 128 characters. [PR1313827](#)
- The MX Series Virtual Chassis unified ISSU emits a benign error message if unsupported FRUs are present. [PR1316374](#)
- There is no route to an IP address from the directly connected route. [PR1318282](#)
- The **show interfaces interface-set** command is displaying wrong logical interface. [PR1319682](#)
- The IPv6 framed Interface ID field (from the **show subscribers extensive** command output) is not properly matching the negotiated one. [PR1321392](#)
- IPCP negotiation might fail for dual stack PPPoE subscribers. [PR1321513](#)

- Unexpected log messages might be seen if a BGP session flaps in a dynamic-tunnels GRE scenario. [PR1326983](#)
- Unexpected log messages might be seen on a router for a subscriber management scenario. [PR1328251](#)
- Traffic loss might be seen after deleting aggregated Ethernet bundle unit 1. [PR1329294](#)
- The cfmd process generates core files. [PR1329779](#)
- The interface might not work properly after the FPC restarts. [PR1329896](#)
- The dcd process might crash due to a memory leak and cause a commit failure. [PR1331185](#)
- The last IFL digit is sometimes truncated in jpppd trace logs. [PR1332483](#)
- The transportd process might crash when you run an snmp query on the jnxoptIfOChSinkCurrentExtTable with an unsupported interface index. [PR1335438](#)
- The MX Series router might occasionally drop the first LCP configure request packet when operating in PPPoE subscriber management configuration. [PR1338516](#)
- The 100G DWDM interface might be going down for 15 seconds after a loss of signal event. [PR1343535](#)
- When eth-oam is deactivated with a scale PM configuration (under hardware-assited-pm-mode), the FPC might become unstable and generate core files. [PR1347250](#)
- Suppressing cfmd logs : `jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0`. [PR1347650](#)
- The jpppd process generates core files spontaneously on the backup Routing Engine in a longevity test at `../../../../../../../../src/junos/usr.sbin/jpppd/pppMain.cc:400`. [PR1350563](#)
- The VRRP VIP becomes unreachable after deleting one of the logical interfaces. [PR1352741](#)
- The FPC might be stuck at 100 percent for a long time when MC-AE with enhanced-convergence is configured with large-scale logical interfaces. [PR1353397](#)
- The FPC generates a core file related to cfmmman. [PR1358192](#)
- Clients might not get an IPv4 address in a PPPoE dual-stack scenario. [PR1360846](#)
- Approximately 50 percent of PPPoE subscribers (PTA and L2TP) and all ESSM sub lost after post unified ISSU during DT CST stress test. [PR1360870](#)
- On all Junos OS products, the CLI allows to configure more than 2048 sub-interfaces on LAG interface from 17.2R1. [PR1361689](#)
- The EOAM LTM messages might not get forwarded after system reboot in CFM scenario configured with CCC interface. [PR1369085](#)
- Subscriber cannot negotiate MLPPP session with MX LNS when dynamic-profile name contains more than 30 characters. [PR1370610](#)

Layer 2 Features

- The rpd process memory leak is observed upon any changes in a VPLS configuration such as deleting or re-adding VPLS interfaces. [PR1335914](#)

- The VPLS instance stays in NP state after the LDP session flaps. [PR1354784](#)
- The Routing Engine kernel might crash when OSPFv3 is configured with an IPsec key authentication over an IRB interface. [PR1357430](#)

Layer 2 Ethernet Services

- The MAC address might not be learnt due to spanning-tree state discarding in kernel table after a Routing Engine switchover. [PR1205373](#)
- The MX Series platforms might display a false positive CB alarm **PMBus Device Fail**. [PR1298612](#)
- DHCP IPv6 traffic might be dropped in a subscriber scenario. [PR1316274](#)
- The jdhcpd process generates core files after making DHCP configuration changes. [PR1324800](#)
- The on-demand-address-allocation under dual-stack-group does not work for IPv6. [PR1327681](#)
- The snmpget for OID: dot3adInterfaceName might not work. [PR1329725](#)
- A memory leak might happen in l2cpd if the l2-learning process is disabled. [PR1336720](#)
- The DHCPv6 second Solicit message might not be processed when IA_NA and IA_PD are sent in a separate Solicit message. [PR1340614](#)
- DHCP client is not able to connect if VLAN is modified on the aggregate Ethernet interface associated with the IRB. [PR1347115](#)
- ZTP infra scripts are not included for MX PPC routers. [PR1349249](#)
- When DHCP subscribers are in an bound (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state if dhcp-service is restarted then the subscribers in this state are logged out. [PR1350710](#)
- The DHCP relay agent will discard a DHCP request message silently if the requested IP address has been allocated to the other client. [PR1353471](#)
- Restarting an FPC that hosts the micro-BFD link might cause LACP to generate a core file. [PR1353597](#)
- DHCPv6 relay ignores replies from server when renewing. [PR1354212](#)
- The DHCP lease query message is replied with incorrect source address. [PR1367485](#)
- DHCP Relay Binding state - rebinding state counter added to dhcpv4 and dhcpv6 binding sensors. [PR1368392](#)

MPLS

- When minimum-bandwidth and bandwidth commands are present in the configuration, the bandwidth selection of the lsp is inconsistent. [PR1142443](#)
- Ingress RSVP LSP fails to come up after issuing the **clear rsvp lsp all** command on the egress router. [PR1275563](#)
- The rpd might crash in an LDP Layer 2 circuit scenario. [PR1275766](#)
- LDP egress policy not advertising label for inet.3 BGP labeled-unicast route. [PR1289860](#)

- Traffic drop is observed during an NSR switchover for RSVP P2MP provider tunnels used by MVPN. [PR1293014](#)
- The traffic in P2MP tunnel might be lost when NG-MVPN uses RSVP-TE. [PR1299580](#)
- The rpd process might crash in rare conditions where **traffic-engineering** is configured. [PR1303239](#)
- The RSVP node-hello packet might not work correctly after the next hop for a remote destination is changed. [PR1306930](#)
- The kysncd process might crash after removing and inserting backup RE in analytics and "mpls sensor" scenario. [PR1303491](#)
- The RSVP node-hello packet might not work correctly after the next-hop for remote destination is changed. [PR1306930](#)
- The rpd process might crash if LDP updates the label for a BGP route. [PR1312117](#)
- The output of the **show mpls container-lsp** command is delayed. [PR1314960](#)
- An RSVP node-neighbor is found even when node-hello has been disabled. [PR1317241](#)
- The IPv4/IPv6 multicast traffic might get dropped in an MX Series Virtual Chassis scenario when the traffic comes in through an Layer 2 circuit and goes out through an aggregated Ethernet member interface across Virtual Chassis members. [PR1320742](#)
- The rpd might crash when LDP P2MP recursive is configured. [PR1321626](#)
- The rpd might crash due to a memory leak in an RSVP scenario. [PR1321952](#)
- Receipt of specially crafted UDP packets over MPLS may bypass stateless IP firewall rules. [PR1326402](#)
- SNMP OID counters for mplsLspInfoAggrOctets show constant value for some LSPs even though traffic is constantly increasing in **show mpls lsp statistics**. [PR1327350](#)
- In Junos OS Release 17.2X75-D40, a new feature related to "per AE member OAM" introduced additional processing on pfeman thread during link flaps. [PR1327988](#)
- Packet loss might be observed when **auto-bandwidth** is enabled for CCC connections. [PR1328129](#)
- The rpd might crash on the backup Routing Engine due to memory exhaustion. [PR1328974](#)
- Fate-sharing group cost does not re-set to the default value after a CLI change, removing explicit cost configuration. [PR1330161](#)
- After a MPLS LSP link flap and local repair, a new LSP instance is tried to be signaled but it may get stuck. [PR1338559](#)
- Whenever there is a decrease in the stats value across an LSP, the mplsLspInfoAggrOctets value takes two intervals to get updated. [PR1342486](#)
- An LDP label is generated for a serial interface subnet route unexpectedly. [PR1346541](#)
- The MPLS LSP does not come up after changing admin-group mapping. [PR1348208](#)
- The rpd crash might happen in an RSVP setup-protection scenario. [PR1349036](#)

- In a very rare scenario, the rpd might crash when LDP failed to allocate a self-ID for the P2MP FEC. [PR1349224](#)
- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. [PR1359956](#)
- Layer 2 Circuit might flap after an interface goes down even if the LDP session stays up when l2-smart-policy is configured. [PR1360255](#)
- The process rpd might crash during P2MP LSPs churn. [PR1363408](#)
- The rpd process might crash after RSVP is deactivated and then re-activated globally for multi times. [PR1366243](#)
- The rpd might crash in BGP LU and LDP scenario. [PR1366920](#)

Multicast

- DHCP6 Relay is not working unless DHCP is restarted. [PR1316210](#)
- Multicast traffic is not forwarded on the newly added P2MP branch or receiver. [PR1317542](#)
- Some IGMP groups might have wrong upstream interface due to discard route is installed in PIM. [PR1337591](#)

Network Management and Monitoring

- The syslog might generate duplicate entries of hostname and timestamp. [PR1304160](#)
- The mib2d might crash when SNMP polling occurs on interface mibs and while the FPC restarts or the interface flaps. [PR1318302](#)
- SNMP stops or becomes very slow after a very long period of time. [PR1328455](#)
- With interface-mib, the MX Series router is responding with **type : NoSuchInstance** for OIDs when multiple OIDs are polled in one SNMPGET request. [PR1329749](#)
- The eventd process fails to start up with the syslog configuration. [PR1353364](#)
- The jnxDcuStatsEntry and jnxScuStatsEntry OIDs are missing in a post interface configuration change. [PR1354060](#)
- The SNMP process crashes during polling the CFM stats. [PR1364001](#)

Platform and Infrastructure

- On MX Series routers, if a large number of routes are processed, then the Packet Forwarding Engine of the MS-MPC might crash. [PR1277264](#)
- Executing the **show services inline ip-reassembly statistics** command might cause a ukern sheaf memory leak. [PR1285833](#)
- The **apply-path** prefix is not inherited under policy after modifying the interface address. [PR1286987](#)
- The output values of command **show system resource-monitor** are not accurate. [PR1287592](#)

- The **interface-mac-limit** might fail for an aggregated Ethernet interface. [PR1303293](#)
- The source MACs might leak (or not learn) between different VPLS instances at the receiving end of VPLS PE devices. [PR1306293](#)
- An rpm probe with a probe interval of 1 second fails on MX Series routers. [PR1308952](#)
- Error messages are not observed during telnet with a username longer than an acceptable limit. [PR1312265](#)
- The mgd process might crash and a session gets terminated after the load override from netconf. [PR1313158](#)
- The issue addresses the ICMP error messages in the Packet Forwarding Engine and is forwarded to the correct pic in the AMS bundle. [PR1313668](#)
- VPLS instance fails to learn MAC addresses upon pseudowire switchover. [PR1316459](#)
- Rate-limit configured with a small temporal buffer size might cause packet loss. [PR1317385](#)
- Multicast traffic might get duplicated when MoFRR is configured. [PR1318129](#)
- The GNF FPC hangs at reboot during a unified ISSU. [PR1318394](#)
- The default severity of the correctable ECC errors on MX Series routers with MPC2E NG Q, MPC3E NG Q, or MPC5E has been changed from fatal to major. [PR1320585](#)
- Errors might be observed when the **fabric-header-crc-enable** feature is enabled. [PR1320874](#)
- The traffic with more than 2 VLAN tags might be incorrectly rewritten and sent out. [PR1321122](#)
- The RPM probes delegated to MS-MIC get stuck when any change is made to the BGP group statement. [PR1322097](#)
- The **no-propagate-ttl** option might not take effect if **chained-composite-next-hop ingress l3vpn extended-space** is configured. [PR1323160](#)
- The MAC might not be learned on MX Series routers with MPCs or MIC-based line cards due to the negative value of the bridge MAC table limit counter. [PR1327723](#)
- The packet might get dropped in an LSR if MPLS pseudowire payload does not have a control word and its destination MAC starts with '4'. [PR1327724](#)
- Traffic loss might be observed on the LT interface. [PR1328371](#)
- Directories and files under **/var/db/scripts** lose execution permission or directory 'jet' is missing under **/var/db/scripts** causing an **error: Invalid directory: No such file or directory** error during commit. [PR1328570](#)
- The tcpdump filter might not work in the egress direction on PS and LT logical interfaces. [PR1329665](#)
- The router hits the database prompt at **netisr_process_workstream_proto**. [PR1332153](#)
- RPM MIB's pingResultsMinRtt, pingResultsMaxRtt, and pingResultsAverageRtt response is "1" while target address is unreachable, it should be "0". [PR1333320](#)

- Traffic loss might be seen for some flows due to network churn. [PR1335302](#)
- Commit might fail with error reading from commit script handler **error: commit script failure**. [PR1335349](#)
- The MPC might crash after setting **max-queues** to a very large number. [PR1338845](#)
- Route corruption occurs in the Packet Forwarding Engine with CFM enabled on the aggregated Ethernet interface. [PR1338854](#)
- Configuring the same DHCP server in different routing instances is not supported in a DHCP relay scenario. [PR1342019](#)
- Commit error is observed when configuring the same VLAN ID on different logical interfaces of the same LT physical interface and the **ethernet-bridge** encapsulation is configured. [PR1342229](#)
- Route corruption in the Packet Forwarding Engine with connectivity-fault-management is enabled for I2ckt. [PR1342881](#)
- ZTP is not supported for vmhost images on next-generation Routing Engines on the MX Series platforms. [PR1343338](#)
- The IPv4 GPRS traffic over the aggregated Ethernet interface might be dropped if gtp-tunnel-endpoint-identifier is configured. [PR1347435](#)
- Output policing action does not work on IRB interfaces for VNIs. [PR1348089](#)
- FPC CPU utilization with LT interfaces is pegged continuously at 100 percent. [PR1348840](#)
- Running RSI through the console port might cause a system crash and reboot. [PR1349332](#)
- The ICMP error messages are not generated if 'don't fragment' packets exceed the MTU of the multiservice interface. [PR1349503](#)
- When viewing IPv6 addresses, **display rfc5952** does not work when combined with **display set**. [PR1349949](#)
- The chassisd process memory leak is observed. [PR1353111](#)
- The kernel crashes because the initialization of the logical Interface MAC filter function is missing for Packet Forwarding Engine extended port devices. [PR1353498](#)
- The FPC might crash due to the memory leak caused by the VTEP traffic. [PR1356279](#)
- Traffic is discarded silently along with **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH** messages. [PR1357707](#)
- When forwarding-class-accounting knob is enabled, on an interface, inside of a routing-instance of instance-type vrf, aggregate input forwarding-class statistics do not increment (egress statistics work fine). [PR1357965](#)
- Select CLI functions are not triggering properly (set security ssh-known-hosts load-key-file, set system master-password). [PR1363475](#)
- Same vlan-id not allowed on multiple IFLs of the same GR interface. [PR1365640](#)

- Subscribers over AE interface might have tail drops which will affect the fragmented packets due to QXCHIP buffer getting filled up. [PR1368414](#)
- The logical tunnel interface might be unable to send out control packets generated by RE. [PR1372738](#)

Routing Policy and Firewall Filters

- Condition based policy fails to take action even though condition is matched [PR1300989](#)
- The policy configuration might not be evaluated if the policy expression is changed. [PR1317132](#)
- Access-internal route might fail to be leaked between routing instances when **from instance** is configured in the policy. [PR1339689](#)
- The policy might not clean up after deleting configuration and cause the rpd to generate a core file. [PR1357724](#)

Routing Protocols

- The **show bgp summary** results are incorrect while assisting GR. [PR1045151](#)
- BGP extended communities with sub-type 4 erroneously displayed at LINK_BANDWIDTH. [PR1216696](#)
- The rpd generates core files in the ASBR when BGP is deactivated in the ASBR before all stale labels have been cleaned up. [PR1233893](#)
- The rpd might crash after deactivating or activating BGP. [PR1272202](#)
- After a bfdd restart, the issue is seen with a next-generation MVPN and Layer 2 VPN route exchange causing MVPN and VPLS traffic drop. [PR1278153](#)
- Routing loops might be seen after configuring BGP Prefix Independent Convergence (BGP PIC). [PR1282520](#)
- Few adj-sid details are not updated in an IS-IS database with a LAN + adjset scenario. [PR1288331](#)
- Multihop BFD sessions flap continuously. [PR1291340](#)
- The Impd crashes repeatedly when a logical system is configured on the same device. [PR1294166](#)
- The rpd process might crash because of the AS PATH check error that occurs when RIB groups are added first and later the routing instances are added. [PR1298262](#)
- MSDP sessions might flap when NSR or GRES is enabled. [PR1298609](#)
- While the device is booting up with the Junos OS Release 17.4R1 image, **error: channel 0: chan_shutdown_read: shutdown() failed for fd 10 [i0 o3]: Socket is not connected** messages might show up. [PR1300409](#)
- IBGP route damping is not taking effect on an IBGP inet-vpn address family. [PR1301519](#)
- Observed mcsnoopd core file at
`__raise,abort,__task_quit,__task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal
(enable_slip_detector=true,no_exit=true) at ../../../../src/junos/lib/libjtask/base/task_scheduler.c:275`
[. PR1305239](#)

- BGP traceoption logs are still written when it is deactivated. [PR1307690](#)
- The rpd might generate a core file in `bgp_rt_send_message` at `../../../../src/junos/usr/sbin/rpd/bgp/bgp_io.c:1460`. [PR1310751](#)
- The BGP session might flap when the connection between the master Routing Engine and the backup Routing Engine keeps flapping with NSR configured. [PR1311224](#)
- The rpd might crash when the neighbor IS-ISv6 router is restarted, causing a route churn. [PR1312325](#)
- Unexpected route age refresh might be observed if BGP PIC is configured. [PR1312538](#)
- The IS-IS SPF might be triggered by LSP updates containing changes only in reservable bandwidth in a TE extension. [PR1313147](#)
- The rpd might crash and generate a core file with distributed IGMP. [PR1314679](#)
- The rpd might constantly consume a high percentage of CPU in a BGP setup. [PR1315066](#)
- On a chassis with BMP configured, the rpd might crash when the rpd process is gracefully terminated. [PR1315798](#)
- The primary path of an MPLS LSP might switch to another address. [PR1316861](#)
- If a loop free alternative is configured, an Isdb entry cleanup might cause the rpd to crash. [PR1317023](#)
- The inactive route cannot be installed in a multipath next hop after disabling and enabling the next hop interface in an Layer 3 VPN scenario. [PR1317623](#)
- A BGP-LU update oscillates with a BGP-PIC. [PR1318093](#)
- IS-IS might choose a suboptimal path after the metric change in ECMP links. [PR1319338](#)
- Traffic might get discarded temporarily when BGP GR is triggered and the direct interface flaps. [PR1319631](#)
- There is an issue with tracing of the BGP Layer 2 VPN DF election community. [PR1323596](#)
- The rpd crash is seen when deactivating the static route if the next-hop interface is type P2P. [PR1323601](#)
- When the prefix limit is reached, increasing maximum-prefixes does not take effect. [PR1323765](#)
- The rpd process might crash continuously on both Routing Engines when **backup-spf-options remote-backup-calculation** is configured in the IS-IS protocol. [PR1326899](#)
- Multiple next hops might not be installed for an IBGP multipath route after an IGP route update. [PR1327904](#)
- With BGP/LDP/IS-IS configurations, deleted IS-IS routes might still be visible in the RIB. [PR1329013](#)
- The rpd might crash on the backup Routing Engine after BGP peer is deleted. [PR1329932](#)
- Manual GRES with an MX Series Virtual Chassis results in some packet loss on core facing interfaces. [PR1329986](#)
- The conditional route policy cannot withdraw all routes in a BGP add-path scenario. [PR1331615](#)

- LDP route in inet.3 is missing when both OSPF rLFA and LFA protections are available and rejected by the backup selection policy. [PR1333198](#)
- Discard next hop being installed when the primary LSP interface drops. When primary interface returns, discard next hop remains until BGP LU neighbor is cleared. This only impacts the cloned route (S=0). [PR1333570](#)
- For Junos OS Release 15.1 and later, IGMP joins are not processed with the **passive allow-receive** command configured on the IGMP interface. [PR1334913](#)
- BGP sessions get stuck in an active state after the remote end restarts the device. [PR1335319](#)
- The rpd crash might occur when receiving BGP updates. [PR1341336](#)
- Changes to the displayed value of AIGP in the **show route ... extensive** command. [PR1342139](#)
- Traffic black hole might be seen if a local device is receiving BFD-down. [PR1342328](#)
- The rpd might crash when BGP flaps. [PR1342481](#)
- The rpd generates a core file while running streaming telemetry test. [PR1347431](#)
- The rpd might crash if a route for RPF uses a qualified-next-hop. [PR1348550](#)
- The rpd might crash while restart routing or deactivate IS-IS. [PR1348607](#)
- The rpd might crash when the BGP route damping and the BGP multipath feature are configured. [PR1350941](#)
- Source-as community is not appended to the rendezvous point. The display issue is in the **show route** detail output. [PR1353210](#)
- Static Route flaps on commit when configured with resolve statement. [PR1366940](#)

Services Applications

- PCP mappings cannot be manually cleared when a NAT pool is shared between PCP and standard NAT. [PR1284261](#)
- The L2TP subscribers might get stuck in a terminating state during login. [PR1298175](#)
- LTS clients experience packet drop for large packets due to fragmentation in LTS. [PR1312691](#)
- AVP 145 is not present in IRQ when *ANCP DSL-type = 0*. [PR1313093](#)
- L2TP tunnel Tx and Rx byte count sometimes decrease when subscriber sessions are reduced within the tunnel. [PR1318133](#)
- SNMP MIBs are not yielding data related to sp-interfaces. [PR1318339](#)
- The MRU might be changed to 1492 instead of the default 1500 in an L2TP scenario. [PR1319252](#)
- IPCP active mode is not getting enabled for MLPPP on LNS. [PR1319580](#)
- Long route remains in forwarding table after subscriber session goes down. [PR1322197](#)

- The L2TP LTS might drop the first CHAP success packet from LNS due to delayed programming of /136 route on the Packet Forwarding Engine. [PR1325528](#)
- The jl2tpd might crash if the RADIUS server returns 32 tunnel-server-endpoints. [PR1328792](#)
- A few CSURQ messages might not respond when the number of sessions addressed in CSURQ is more than 107. [PR1330150](#)
- The l2tpd might crash when multiple l2tp related commands are executed together. [PR1337406](#)
- The **show services stateful-firewall flows count** command shows an incorrect flow count after a services configuration change. [PR1338704](#)
- Output of **show interfaces si-x/y/z.xxxxx extensive** CLI command shows an incorrect inet/inet6 MTU value for an MLPPP subscriber on MX Series L2TP LNSs. [PR1346049](#)
- The bbe-smgd process might crash if there are 65,535 L2TP sessions in a single L2TP tunnel. [PR1346715](#)
- Session limit per tunnel on LAC does not work as expected. [PR1348589](#)
- After performing an SNMP walk on the IKE SA that is deleted, IPsec tunnels might go down and an infinite loop scenario might be seen. [PR1348797](#)
- The UDP checksum inserted by an MS-DPC after a NAT64 is not valid when an incoming IPv4 packet has UDP checksum set to 0. [PR1350375](#)
- The **show services stateful-firewall flows counter** command shows high numbers. [PR1351295](#)
- The JI2tpd process might crash shortly after one of the L2TP destinations becomes unavailable. [PR1352716](#)
- L2TP tunnel-switch clients in subscriber session database reference the wrong routing instance. [PR1355396](#)
- In some corner cases, a few tunneled PPPoE subscribers might get stuck in a terminating state. [PR1363194](#)
- The L2TP subscribers might not be able to log in successfully due to the jl2tpd memory leak. [PR1364774](#)
- Actual Data Rate Downstream value not included in the L2TP ICRQ message from the LAC. [PR1370699](#)

Software Installation and Upgrade

- New versions of Junos OS do not have the tool for accessing an aux port - /usr/libexec/interposer. [PR1329843](#)
- Commit might fail in single-user mode [PR1368986](#)

Subscriber Access Management

- A memory leak might happen after clearing a subscriber either with a script or manually. [PR1312517](#)
- Service interim is missing for random users in a JSRC scenario. [PR1315207](#)
- The PPPoE subscribers might encounter a connection failure during login. [PR1317019](#)
- The unified ISSU is allowed to proceed when the account is suspended. [PR1320038](#)

- IP addresses are assigned discontinuously from the linked IP pools. [PR1323829](#)
- Authd considers RADIUS attribute *Framed-IPv6-Prefix* = *::/64* or *Delegated-IPv6-Prefix* = *::/56* as valid parameters. [PR1325576](#)
- An MX204 does not send a **RADIUS Accounting-Off** message. [PR1327822](#)
- Multiple RADIUS servers having a different dynamic-request-port is not supported. [PR1330802](#)
- Subscriber might get stuck in a terminated state when JSRC synchronization state is stuck in a FULL-SYNC in progress state. [PR1337729](#)
- In dual stack subscribers scenario with NDRA pool configured, the linked pools are not used when the first NDRA pool is exhausted. [PR1351765](#)
- When attempting to scale clients saw `sdbsts_lock_holder.bbe-smgd.pid10686.core` core files. [PR1358339](#)

User Interface and Configuration

- There is an increase in commit times. [PR1029477](#)
- The CLI session might die while issuing the **show configuration | compare rollback 1** command. [PR1331716](#)
- The **max-db-size** configuration might not work on some MX platforms. [PR1363048](#)

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the Layer 2 circuit flaps repeatedly. [PR1282875](#)
- Un-hide **set protocols pim mvpn family inet6 disable configuration** to allow users to disable inet6 on MVPN. [PR1317767](#)
- The rpd might crash after a unified ISSU in a large scale scenario with a PIM configuration. [PR1322530](#)
- Moving MC-LAG from LDP based pseudowire to BGP based pseudowire might cause the rpd to crash. [PR1325867](#)
- The multicast might be rejected when Junos OS PE devices received a C-Mcast route from other vendor PE devices. [PR1327439](#)
- MVPN sender-site configuration is not allowed with S-PMSI. [PR1328052](#)
- The rpd generates a core file on the backup Routing Engine with an next-generation MPVPN and NSR configuration. [PR1328246](#)
- The rpd might crash after committing interface related parameters (for example, MTU change, VRF RD or RT, QOS) on the PS interface with vlan-ccc encapsulation and no vlan-id. [PR1329880](#)

- The rpd might continuously crash on the backup Routing Engine and some protocols might flap on the master Routing Engine if hot-standby is configured for Layer 2 circuit or VPLS backup-neighbor. [PR1340474](#)
- The rpd might crash on the backup Routing Engine when changing the Layer 2 circuit **virtual-circuit-id** in an NSR scenario. [PR1345949](#)

Resolved Issues: 17.4R1

Class of Service (CoS)

- The Routing Engine level **scheduler-hierarchy** command misses a forwarding class when the "per-unit-scheduler" mode is configured. [PR1281523](#)

Forwarding and Sampling

- The Sampled process stops collecting data on Routing Engine based sampling supported platforms. [PR1270723](#)
- Firewall filter might not be matched when wildcard (*.*) is specified as the matching condition. [PR1274507](#)
- The sampled route reflector process (srrd) might crash in a large routes churn situation. [PR1284918](#)
- The mib2d process generated a core file @fw_counter_key2components. [PR1286448](#)
- The sampled process might crash and generate a core file if traceoptions are enabled. [PR1289530](#)
- Some accounting files might be missed if the remote archive site is unreachable. [PR1300764](#)
- There is memory leak on mib2d when polling firewall MIBs. [PR1302553](#)
- ACCT_FORK_LIMIT_EXCEEDED log level is ERROR even when backup-on-failure feature is enabled for accounting files. [PR1306846](#)
- The commit might fail if enabling nexthop-learning knob for J-Flow v9. [PR1316349](#)

General Routing

- Enhanced IP/enhanced Ethernet and MS-DPC compatibility. [PR1035484](#)
- Ksyncd might crash due to transient replication errors between Routing Engines. [PR1161487](#)
- On MX240/480/960 platforms, due to a I2C bus hardware issue, error messages might appear. [PR1174001](#)
- SNMP trap sent for **PEM Input failure** alarm. [PR1189641](#)
- Stale VBF states occur without SDB sessions. [PR1204369](#)
- The rpd might crash on the backup Routing Engine after a Routing Engine switchover in MX Series subscriber environment. [PR1206804](#)
- The rpd might crash on platforms with 64-bit X86 RE if IPv6 is configured. [PR1224376](#)

- MPC2E-NG/MPC3E-NG generates a core file with specific MIC due to tight loop of PCI Express critical exceptions. [PR1231167](#)
- The MS-MPC card might crash when OSPFv3 IPv6 traffic goes through it. [PR1233459](#)
- FPCs on MX960 platform might be stuck in offline state with **FPC Incompatible with SCB** due to delayed PEM startup. [PR1235132](#)
- With vLNS (vBNG), a commit generates the message **warning: requires 'l2tp-inline-lns' license** even if a valid license is installed. [PR1235697](#)
- The "multicast-replication" setting cannot be reflected in the redundancy environment after rebooting both Routing Engines. [PR1240524](#)
- In a BGP/MPLS scenario, if the next-hop type of label route is indirect, disabling and enabling the "family mpls" of the next-hop interface might cause the route to go into a dead state. [PR1242589](#)
- XM chip-based line card might drop traffic under high temperature. [PR1244375](#)
- On MX2000 with MPC6E, EOAM LFM adjacency flaps when an unrelated MIC accommodated in the same MPC6E slot is brought online by configuring OAM pdu-interval 100 ms and pdu-threshold 3. [PR1253102](#)
- The "validation-state:unverified" routing entry might not be shown with proper location in show route output. [PR1254675](#)
- The rpd might crash during the next-hop change, if unicast reverse-path- forwarding (uRPF) is used. [PR1258472](#)
- Status LED for the ge-0/0/0 interface does not glow. [PR1259112](#)
- MPC might report a parity error with the **fast-lookup-filter** command configured. [PR1266879](#)
- When ISSU is performed under scaled scenarios where the Packet Forwarding Engine next-hop memory uses more than 4 Million Dwords, PPE traps and traffic loss might be observed during software-sync phase until the end of hardware-sync. [PR1267680](#)
- On MX Series routers, the **show chassis led** command should not be displayed in possible completions of the **show chassis** command. [PR1268848](#)
- A low memory condition putting the Service PIC into the red zone on the MS-MIC or MS-MPC card might cause the SIP ALG to generate a core file. [PR1268891](#)
- The FPC might go offline and the ABB fan might crash after enabling MACsec. [PR1270121](#)
- The mspmand log incorrectly generates messages about memory zone level. This occurs every 49.7 days and will recover by itself. This is a display issue and will not affect traffic. [PR1273901](#)
- CLI commands fail to execute for **show subscribers detail**, **show subscribers extensive**, **show subscribers count client-type <>** and other commands because the subscriber management database is unavailable. [PR1274464](#)
- Link stays down after a flap on MPC next-generation cards with QSFP+-40G direct attach copper (DAC) cable. [PR1275446](#)

- The Packet Forwarding Engine of service DPC might crash with large scale of routes for MX Virtual Chassis. [PR1277264](#)
- Layer 2 control BUS stuck causes SFP+ thread hogging and restarting of MPC. [PR1277467](#)
- Multicast traffic when using iflsets in universal call admission control policy mode does not flow as expected in certain use cases, and bbe-smgd might generate a core file. [PR1278543](#)
- VLAN out-of-band subscriber session fails in autoconfigured mode. The physical interface goes down even if it is physically up. [PR1279612](#)
- After a MS-MPC-PIC is turned offline or online or bounced(because of an AMS configuration change), sometimes the PIC can take approximately 400 seconds to come up. [PR1280336](#)
- **MIC Error code: 0x1b0001** alarm might not be cleared for MIC on MPC7/8/9 when the voltage has returned to normal. [PR1280558](#)
- Authenticated subscriber dynamic VLAN interface might get disconnected immediately after a successful connection. [PR1280990](#)
- jfirmware upgrade support is not available for Routing Engine BIOS. [PR1281050](#)
- The **ingress service-accounting-deferred** command is not providing the correct IP traffic statistics for for L2BSA subscribers. [PR1281201](#)
- Establishment of IPsec SAs for link-type tunnels might fail under certain conditions. [PR1281223](#)
- Subscribers might not be able to connect to MX BNG in certain scenarios. [PR1281896](#)
- DHCP/PPPoE subscribers fail to bind after FPC restart and smgd restart with BBE_RTsock_GET_RTsock_IFL_FAIL_TERMINATED counter going up. [PR1281930](#)
- Inline J-Flow unrelated configuration changes related to a routing instance result in invalid or incomplete J-Flow data packets. The **commit full** command resumes proper functionality. [PR1282580](#)
- In a specific CE device environment in which **asynchronous-notification** is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. [PR1282875](#)
- Error messages related to "IFRT: 'IFL'", "IFRT: 'Aggregate interface'" and "IFRT: 'IFD'" are seen on configuration change. [PR1282938](#)
- VBF flows are not programmed correctly on aggregated Ethernet interfaces. [PR1282999](#)
- The MX: **show interfaces** command should display the cause for Intf down when the Packet Forwarding Engine disabled. [PR1283323](#)
- GRE OAM fails to come up when GRE tunnel source and family inet address are the same. [PR1283646](#)
- PPTP session could not be established on MS-MPC when both stateful firewall and NAT were enabled. Also, the address could not be translated. [PR1285207](#)
- The J-Flow data template sequence number is zero for MPLS flows. [PR1285975](#)
- With CoS-based forwarding, when the primary path of one of the next-hop LSPs flaps, traffic carried by the other next-hop LSP could get load-balanced across the primary and secondary path. [PR1285979](#)

- Internal latency increases the overtime for Packet Forwarding Engine sensors with streaming telemetry. [PR1286286](#)
- Unified ISSU is not supported from Junos OS Release 15.1 or later, because the source release includes one or more BBE features such as logical interface (IFL) options, CoS fragmentation map, MLPPP, advisory options, advanced services, and multicast distribution. [PR1286507](#)
- DDS culprit flows are not reported by CLI or logs during login to a MX Series router with a single Packet Forwarding Engine. [PR1286521](#)
- The routing protocol process (rpd) crashes during subscriber login or logout with multicast service enabled while performing GRES switchover. [PR1286653](#)
- Framed routes might get stuck in KRT queue. [PR1286849](#)
- A10NSP interface is not getting attached to the L2 routing instance after the routing instance name is renamed. [PR1287070](#)
- The rpd might generate a core file when the routing-options dynamic-tunnels configuration is changed. [PR1287109](#)
- **Host 0 RTC Battery failure** error messages are seen on PTX1000 and QFX10000-line after upgrading to Junos OS Release 16.1. [PR1287128](#)
- LTS functionality is not working on Junos OS 16.1R4-S2 if the **rewrite-rule** statement is applied to the dynamic profile. [PR1287788](#)
- SNMP query for IF-MIB::ifOutQLen reports **Wrong Type should be Gauge32 or Unsigned32** for a dynamic VLAN DEMUX0 interface. [PR1287852](#)
- The **services-oids-ev-policy.slax** and **services-oids.slax** files built in the Junos OS image are not the latest versions. [PR1287894](#)
- After offlining and onlining back fabric planes, a few planes are stuck in offline state in MX480. [PR1287973](#)
- The bbe-smgd process might crash and generate a core file on the standby Routing Engine during a reboot upgrade with active locally terminated PPPoE subscribers. [PR1288121](#)
- During unified ISSU upgrade micro BFD flap is observed. [PR1288433](#)
- The smg-service process (daemon) might generate core files in the backup Routing Engine with a distributed IGMP configuration. [PR1288465](#)
- Performance issues can be seen when nontranslated traffic is introduced to a service-set using a large number of NAT terms. [PR1288510](#)
- After GRES **smid** was thrashing and was not restarted after a fatal SDB error. [PR1288871](#)
- Kernel "rtdata" memory leak is found on an MX Series Virtual Chassis with the **heartbeat** command enabled. [PR1289363](#)
- FPC memory leak might happen in a BBE subscriber environment. [PR1289365](#)

- The interfaces might got to a down state after performing GRES. [PR1289493](#)
- The **request system zeroize** command deletes the **/var/db/scripts** directory, which does not get re-created until the next USB/Netboot recovery. [PR1289692](#)
- The jnxContainersType MIB is not displayed for PIC and MIC as correctly as it is displayed on other Juniper platforms. [PR1289778](#)
- If the vmhost application is not running, then the alarm string will have "Application" name embedded in it. [PR1290150](#)
- NAT-T and DPD functionality do not work for aggressive mode. [PR1290689](#)
- Incorrect temperature is displayed for MPCP5/MPC7 in **show chassis fpc** output. [PR1290771](#)
- When IGMP protocol is enabled, there can be a leak of 56 bytes in the bbe-smgd process (daemon) during logout for every subscriber who had joined any multicast group during the session. [PR1290918](#)
- Rpd core file might be generated when restarting the process via CLI. [PR1291110](#)
- JDI-RCT-RPD: Device going to the DB prompt "db@jsr_jsm_send_ka_after_merge,send_proto_keepalive" was observed on master Routing Engine. [PR1291247](#)
- l2tp iccn fast retransmission occurs after tunnels go down. [PR1291557](#)
- The bbe-smgd process might crash and subscribers might get stuck when a large group of different types of subscribers login/logout. [PR1291969](#)
- The local preference cannot work correctly for EVPN type 5 route in multipath scenario. [PR1292234](#)
- An error in **vbf_filter_add_orphan_check** might be seen when the subscribers using filters log out or log in. [PR1292582](#)
- Error message might be seen while bringing up the subscriber in a subscriber management environment. [PR1293057](#)
- CPCDD might generate core files while using Routing Engine based http-redirect. [PR1293553](#)
- The **show extensible-subscriber-services sessions** command is displaying incorrect timestamp after a unified ISSU. [PR1293800](#)
- Loss of DHCP/PPPoE subscribers is observed during unified ISSU from Junos OS Release 16.1-20170718_161_r4_s5.0 to Release 16.1-20170718_161_r4_s5.0. [PR1294709](#)
- The krt queue might be stuck with the error of "RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0". [PR1295756](#)
- Unable to edit dynamic profiles after scaling up to 400 dynamic profiles. [PR1295446](#)
- The bbe-smgd process might generate a core file at bbe_mcast_ifl_vbf_encoder on service activation or deactivation along with smg-service process (daemon) restart. [PR1295938](#)
- The service-profile's CoS might be overrode by the client-profile's CoS when second family DHCP session added in dual-stack subscriber scenario. [PR1296002](#)

- TACACS remote user is unable to run JET applications because of a bad stored heap. [PR1296237](#)
- The mspmand process might crash if you use SCG services on MS-MPC/MS-MIC. [PR1296422](#)
- The continuous kernel might crash when a lot of terms are configured for firewall filters. [PR1296884](#)
- In ECMP fast reroute scenario, traffic might get silently dropped or discarded because of a next hop in "hold" state. [PR1297251](#)
- A memory leak is seen when **set protocols mld XXX** is changed and committed. [PR1297454](#)
- Multiple bbe-smgd core files are seen during a subscriber binding configuration with DT CST with as little as 200-300 subscribers and continual core files while scaling. Maximum scale cannot be achieved with multicast- enabled subscribers (related to IPTV profile). [PR1297612](#)
- During InFlight Daemon Kill test, rpd core files are seen with PPPoE and L2BSA flapping. [PR1298587](#)
- Commit error is thrown when trying to commit a configuration with apply groups. [PR1298649](#)
- The bbe-smgd process might crash when traceoption is enabled due to an invalid username character. [PR1298667](#)
- The bbe-smgd process constantly generates core files while ESSM+PPPoE stress test with concurrent GRES is running. [PR1298742](#)
- MX Series BNG does not respond to PADI after GRES on some ports/VLANs. [PR1298890](#)
- Junos Telemetry Interface: DREND errors are seen for components "mpcs-software-rev", "rom-software-rev", "software-rev", and "firmware-rev". [PR1299470](#)
- The "asynchronous-notification" feature cannot be implemented properly in a circuit that has MIC-3D-20GE-SFP-E/Tri Rate Copper SFP(740-013111). [PR1299574](#)
- Flat accounting files are not generated according to the configured timers. [PR1299597](#)
- Subscriber database is stuck in not-ready state after GRES. [PR1299940](#)
- After IS-IS-TE routes and BGP routes attribute change, traffic loss might be seen because BGP routes point to some stale labels. [PR1300425](#)
- Junos Telemetry Interface: The error **error: the SDN-Telemetry subsystem is not responding to management requests** is seen on issuing the CLI command **show agent sensors** if traceoptions is enabled for services analytics. [PR1300829](#)
- Configured logical interface might not be created correctly after commit. [PR1301823](#)
- The rpd might crash when toggling the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statement. [PR1302504](#)
- The log message **jam_cache_get.636 ERR:entity 0x997 not found, get cache failed** is continuously seen in jam_chassisd log file. [PR1302975](#)
- chassisd.core-tarball.0.tgz found during ISSU is aborted in FRU upgrade phase. [PR1303086](#)

- Incorrect MTU might be seen on PPP interfaces when PPP MTU is not defined in the dynamic profile. [PR1303175](#)
- The list of available routing instances is no longer provided for output of **show subscribers routing-instance ?command**. [PR1303199](#)
- Blocking PPPoE/DHCP to initiate VLAN auto-sensing if VLAN-OOB connected is in pending state. [PR1303338](#)
- MX Series MIB polling returns a value that has "sdg". Polling result should include "svc" generic value. [PR1303848](#)
- Truncated output appears for the **show pppoe lockout** CLI command. [PR1304016](#)
- Effective rate of E3 in framed mode is limited to 30 Mbps on certain channelized MICs. [PR1304344](#)
- RPF check strict mode is causing traffic drop in next-generation subscriber management release. [PR1304696](#)
- On MX2000 platform with MPC9E and SFB2 installed, certain high amount traffic volume might cause traffic drops with cell underflow messages. [PR1304801](#)
- Commit fails with error: **ffp_intf_ifd_hier_tagging_config_verify: Modified IFD "si-1/1/0" is in use by BBE subscriber, active L2TP LNS client**. [PR1304951](#)
- Inline J-Flow VMX: OIF field of VPLS data records sometimes reports the SNMP index value of the LSI interface instead of the egress physical interface. [PR1305411](#)
- MX Series router is sending immediate-interim for the services pushed by SRC. [PR1305425](#)
- Customers running 32-bit Junos OS might generate rpd core file when traceoptions are enabled. [PR1305440](#)
- Going forward, JET daemonize applications will not get respawned on a normal exit, which should be the ideal behavior of any App. [PR1305615](#)
- L2BSA subscriber connection attempts failed with vlan profile-request-error. [PR1305962](#)
- L2BSA subscribers came up, but no new ANCP session got established during the RADIUS disaster backup procedure. [PR1306872](#)
- Smihelperd generates core files when SNMP is polling for JUNIPER-SUBSCRIBER-MIB::jnxSubscriberGeneral.7.0. [PR1306966](#)
- Split horizon label is not allocated after switching a configuration of ESI from single-active to all-active. [PR1307056](#)
- The kmd process error UI_DBASE_OPEN_FAILED is seen because of too many open files. [PR1308380](#)
- License lost during Routing Engine switchover in scale-subscriber scenario. [PR1308620](#)
- CoS applied to a subscriber demux logical interface (IFL) is not working. [PR1308671](#)
- All the MICs on FPC, with ps interfaces configured, went offline during the restart of FPC in another slot. [PR1308995](#)

- Error message: %PFE-3: fpc0 vbf_var_iflset_add:633: vbf container 11 not found in the msg for ifl .demux.6514 is often seen after MPC restart. [PR1309013](#)
- Incorrect values are found in the event-timestamp of RADIUS Accounting-Stop packets for L2BSA subscribers. [PR1309212](#)
- RPT BBE REGRESSIONS: DHCP client is stuck in selecting state while verifying untagged DHCP subscribers after modifying router configuration. [PR1309730](#)
- In next-generation subscriber-management release, bbe-smgd process memory leak is seen after deleting or adding the address pool. [PR1310038](#)
- The MS-MIC/MS-MPC memory utilization might stay at high level in the subscriber management scenario. [PR1310064](#)
- SPD_CONN_OPEN_FAILURE and SPC_CONN_FAILURE log messages are seen in the log for SI interfaces when running SNMP walk on Service PIC NAT OIDs. [PR1310081](#)
- The krt_junos_sanity_check_ctrl_resp: rtsock request finally succeeded after error 16' syslog message in the Junos OS Release 17.1R1.8. [PR1310678](#)
- After bsys reboot sometimes rpd is unresponsive on one or more GNFs. [PR1310765](#)
- In streaming telemetry, when a user logs in and logs out quickly from TACACS, the following message is displayed: bad stored heap: heap-ptr=0x0 data-ptr=0x1481cbf8. [PR1311482](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Counter at PPPoE session logical interface (IFL) incremented wrongly cause accounting packet contains wrong Acct-input-packets value and wrong Acct-input-octets value. [PR1312998](#)
- Rpd core is seen when any **show route inetcolor.0** command is executed from CLI. [PR1316078](#)
- **show auto-configuration out-of-band** CLI command with different configuration statements shows the same output. [PR1316661](#)
- After NSR to re1, switch back to RE0 has replication stuck for BGP and LDP. [PR1319784](#)
- Rpd core seen during configuration changes with BGP neighbors. [PR1320900](#)
- Commit operation gets stuck when commit check is performed with fast-synchronize option is enabled. [PR1322431](#)
- JDM Management is unreachable after flapping physical JDM and GNF/VNF management interfaces. [PR1323519](#)

High Availability (HA) and Resiliency

- Line Card reboots after GRES. [PR1286393](#)
- After flapping server CB ports GNFs shows "Switchover Status: Not Ready". [PR1306395](#)

Infrastructure

- "Last flapped " time stamp is not getting updated for fxp0 interface as it should be. [PR1244502](#)
- The **show system users** CLI command output displays users that are not using the router. [PR1247546](#)
- When **set system ports console log-out-on-disconnect** is enabled, system reboot or switchover can result in processes remaining in the wait state and failure of the syslog feature. [PR1253544](#)
- The device might fail to upgrade. [PR1298749](#)
- The syscalltrace.sh might create huge output file which could cause the router to run out of storage space. [PR1306986](#)

Interfaces and Chassis

- The output value is incorrect when querying the optical power of OTN interfaces in the router. [PR1216153](#)
- EX Series Packet Forwarding Engine and MX Series MPC7E/8E/9E PFE crash when fetching interface statistics with extended-statistics enabled (CVE-2017-10611). [PR1247026](#)
- At a high logical interface scale, an ifinfo process (daemon) generates a core file on executing the command **show interfaces extensive | no-more**. [PR1254189](#)
- The MRU of ae interface might reset to default value. [PR1261423](#)
- The MTU configuration option for vt interfaces should be removed because the MTU on this interface is already set to unlimited. [PR1277600](#)
- Monitor interface on aggregated Ethernet logical interfaces displays incorrect bps value compared to **show interface** output. [PR1283831](#)
- Interface flap while executing Routing Engine switchover if the member links of an ae interface are configured with framing settings. [PR1287547](#)
- No L2TP sessions come up on some si interfaces after an MPC restart followed by a Routing Engine switchover. [PR1290562](#)
- PPPoE/PPP subscriber might not be brought up with **reject-unauthorized-ipv6cp** configured. [PR1291181](#)
- Change in history records supported per EOAM performance-monitoring session. [PR1294123](#)
- Family inet shows as not-configured after adding or deleting the loopback address. [PR1294267](#)
- A VRRP track interface down does not trigger a mastership election immediately. [PR1294417](#)
- IRB interface shows incorrect bandwidth value. [PR1302202](#)
- AFEB might not come up if LFM is deactivated. [PR1306707](#)
- After executing the **request system reboot both** CLI command, the Juniper PPP daemon might become unresponsive. [PR1310909](#)

- The PPPoE subscriber might not login correctly after authentication failure in subscriber scenario. [PR1311113](#)
- MX Series Virtual Chassis unified ISSU emits benign error message if unsupported FRUs are present. [PR1316374](#)

Layer 2 Ethernet Services

- DHCP is not using the configured IRB MAC as the source MAC in DHCP offer unicast replies. [PR1272618](#)
- DHCPV6 client bound to IA_PD prefix on reception of DHCV6 Request for IA_NA, MX deletes the existing binding. [PR1286359](#)
- ARP requests not generated for IRB configured in VPLS over GRE tunnel. [PR1295519](#)
- PPPoE/DHCP clients cannot login to PPPoE/DHCP dual-stack subscriber scenario. [PR1298976](#)
- Multiple jdncpd core files are observed in jdncpd_update_groups at `../../../../../../../../src/junos/usr/sbin/jdncpd/jdncpd_config.c:2290`. [PR1311569](#)

Layer 2 Features

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

MPLS

- RSVP p2mp sub-LSPs having more than one sub-LSP in down state might not get re optimized after transit path goes down. [PR1174679](#)
- The rpd might crash when moving static LSP from one routing instance to another [PR1238698](#)
- Created time value in **show mpls lsp extensive** drifts by a second when the show command is issued multiple times. [PR1274612](#)
- Next generation MVPN mLDLP at the receivers' PE device does not join to P2MP LSP on changing the root PE device route from IGP/LDP to LBGP. [PR1277911](#)
- MPLS I2ckt ping packet incorrectly parsed by the output loopback filter. [PR1288829](#)
- The routing protocol process (rpd) crashes due to LDP defect during NSR-enabled Routing Engine switchover. [PR1290789](#)
- Received MTU might not get updated in RSVP MTU signaling. [PR1291533](#)
- Stale RSVP LSP entry after NSR switchover and session is not refreshed. [PR1292526](#)
- The rpd might crash if the MPLS LSP path change occurs. [PR1295817](#)
- The rpd process might crash when performing MPLS traceroute. [PR1299026](#)

- When using IS-IS traffic engineering database, if an LSP's state changes, the routing protocol process might loose track of memory. [PR1303239](#)
- BGP multipath might not work if interface flaps. [PR1305228](#)
- Feature explicit-null might block host-bound traffic incoming from LSP. [PR1305523](#)
- The rpd process might crash during interface-down when UHP-based LSPs are configured. [PR1309397](#)

Network Management and Monitoring

- Command Esc-Q does not work when the syslog is disabled. The syslog message is still seen even if it is disabled by Esc-Q. [PR1269274](#)
- MIB2D-related syslog message **MIB2D_RTSLIB_READ_FAILURE: rtslib_iflm_snmp_pointchange** is seen when configurations are removed or restored. [PR1279488](#)
- MIB2D logs **RLIMIT curr 1048576000 max 1048576000** every time a commit is done. [PR1286025](#)
- The mib2d process might crash when polling the OID ifStackStatus.0 after a logical interface (IFL) of lo0 is deleted. [PR1286351](#)
- An alarm-mgmd core file is seen after upgrade due to an old version of the alarm.db file. [PR1296597](#)
- Implement prefix compression for subinterfaces from mib2d. [PR1297447](#)
- The **show arp no-resolve interface X** output for inexistent interface X is showing all unrelated static ARP entries. [PR1299619](#)
- After SNMP configuration activation the snmpd process started to consume a lot of CPU time. [PR1300016](#)

Platform and Infrastructure

- Traffic drop might occur under a large-scale firewall filter configuration. [PR1093275](#)
- The traffic might not be transmitted correctly from MPC/FPC in rare condition. [PR1170527](#)
- FPC crashes with the MAC accounting feature enabled. [PR1173530](#)
- The "forwarding-class-accounting enhanced" feature is not supported in combination with "forwarding-options hyper-mode". Using both features together results in traffic being silently discarded or dropped. [PR1198021](#)
- Packet Process Engine UCODE rebalancing getting enabled by default. [PR1207532](#)
- With a commit script configured, the mgd process might crash when configure anything in private configuration mode. [PR1244015](#)
- The RPM loss percentage values for "over all tests" via SNMP might be incorrect. [PR1272566](#)
- EVPN-VXLAN traffic gets dropped as **Incorrect vxlan fw path executed** due to a sampling configuration on the core interface. [PR1280539](#)
- The **request routing-engine login other-routing-engine** command might require password. [PR1283430](#)

- The traffic might be classified into the wrong queue when aggregated Ethernet interfaces with child legs are anchored on an MQ-based MPC without a queuing chip. [PR1284264](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- Administratively disabling an interface might cause high FPC CPU usage. [PR1285673](#)
- Transit traffic that has the second LSB set in the first octet of destination MAC will be punted to the Routing Engine when **mac-learn-enable** is configured. [PR1285874](#)
- Generate-event time-interval usage now triggers the event only on the actual expiry of the time interval. [PR1286803](#)
- Incorrect load-balancing on the aggregated Ethernet interface might occur if traffic goes from MS-DPC to MPC in enhanced-ip mode. [PR1287086](#)
- Packet Forwarding Engine heap memory leak is found in three routers with PPPoE subscribers. [PR1287870](#)
- mgd: error: **Couldn't open library: /usr/lib/render/libvccpd-render.tlv**. [PR1289158](#)
- Syslog error appears: not a proper library: **/usr/lib/render/libdcd-render.so: Cannot open "/usr/lib/render/libdcd-render.so"**. [PR1289974](#)
- The source MAC learned from Packet Forwarding Engines across ae interface might bounce between ae member Packet Forwarding Engines for a long time and might cause MLP-ADD storm. [PR1290516](#)
- Dynamic MAC learning might fail on GRE tunnel interface. [PR1291015](#)
- RMOPD might get stuck at sbwait upon receiving a specific response from the HTTP agent. [PR1292151](#)
- Transient flow control asserted by XLP MAC after upgrading the MX Series router to Junos OS Release 16.1. [PR1293232](#)
- The scale-subscriber license might leak on the backup Routing Engine during bulk subscriber logout. [PR1294104](#)
- The mgd process generates a core file after GRES in a subscriber environment. [PR1298205](#)
- **RMOPD_HW_TIMESTAMP_INVALID** is reported two to four times a day which raises an alarm when polled via jnxRpmResSumPercentLost MIB. [PR1300049](#)
- MPC might reset in firewall filter scenario during loading configuration on MX Series platform. [PR1300990](#)
- All traffic can be Tail/RED-dropped on some interfaces when **chassis fpc max-queues** is configured. [PR1301717](#)
- Classifier does not get applied on the aggregated Ethernet member links on DPC (I-chip) based platforms with CoS configured. [PR1301723](#)
- MX Series FPC wedges when creating more than 4000 logical tunnel interfaces per Packet Forwarding Engine. [PR1302075](#)
- When you execute the **mk destroy-all** command, it gives the error **Could not find jnx.wrlsb.mk**. [PR1302974](#)

- The interface-mac-limit might fail for aggregated Ethernet interface. [PR1303293](#)
- The Two-Way Active Measurement Protocol (TWAMP) Request-TW-Session message's Type-P Descriptor format is not RFC-compliant. [PR1305752](#)
- On MX Series routers with MPCs or MICs, the resource monitor (RSMON) thread might be stuck in a loop consuming 100 percent of FPC CPU. [PR1305994](#)

Routing Protocols

- No multicast forwarding in ASM mode occurs after unified ISSU. [PR1146621](#)
- RLFA computation might still consider a PQ-node not reachable via LDP, when LDP is deactivated. [PR1202392](#)
- The routing protocol process (rpd) on the backup Routing Engine might restart unexpectedly upon the addition of a new L2VPN routing instance. [PR1233514](#)
- When the **advertise-from-main-vpn-tables** configuration statement is used under BGP and the route reflector functionality is added, a refresh message is not sent, resulting in some missing routes. [PR1254066](#)
- MPLS over UDP tunnel creation fails in the absence of a VRF table. [PR1270955](#)
- A few BFD sessions are flapping while coming up after FPC restart/reboot. [PR1274941](#)
- Error messages might be seen when receiving BGP update messages with UNREACH NLRI. [PR1276758](#)
- After Routing Engine switchover (GRES+GR), default mdt failed to come up and core-facing interface flap was seen. [PR1279459](#)
- BGP updates might not be advertised to peers completely in certain condition. [PR1282531](#)
- The rpd process might crash due to a certain chain of events in a BGP-LU protection scenario. [PR1282672](#)
- The second multicast packet might be discarded on the rendezvous point router. [PR1282848](#)
- The rpd process might crash while deactivating the routing instance of pim static. [PR1284760](#)
- Some BGP-related traceoptions flag settings will not be effective immediately after the configuration commit, until the BGP sessions are flapped. [PR1285890](#)
- The rpd will run into a loop if bootstrap messages exceed the interface MTU size. [PR1287467](#)
- The rpd might crash if the dynamic rendezvous point goes down in ECMP topology and also PIM **join-load-balance automatic** is configured. [PR1288316](#)
- The rpd might crash after loading merge and rollback configuration with BGP traceoption. [PR1288558](#)
- Multicast flow reset might occur on OIF for RPT joined branch when PIM prune comes on another interface. [PR1293900](#)
- The rpd might crash if BGP flap happens. [PR1295062](#)
- ISSU might take more time to complete and the MPC card might go offline during ISSU reboot. [PR1298259](#)

- Inline BFD on IRB will be broken after GRES/NSR switchover, and the anchor FPC subsequent goes offline. [PR1298369](#)
- BGP might send an incorrect AS path when the alias is enabled and multiple peers are under the BGP group. [PR1300333](#)
- The rpd process might crash with a core file while deleting a multipath route. [PR1302395](#)
- Junos OS Release 16.2 and later releases might give the following error: **Request failed: OID not increasing: ospflflpAddress.0.0.0.0.0.** [PR1307753](#)
- Qualified next-hop resolution fails in some scenarios when there is a next-hop interface specified. [PR1308800](#)
- BGP labeled-unicast protection might break multicast Reverse Path Forwarding (RPF). [PR1310036](#)
- An rpd core file is observed while importing IS-IS routes. [PR1312325](#)
- BGP prefixes with three levels of recursion for resolution will get stuck with a stale next-hop at the first level after a link-down event. [PR1314882](#)

Services Applications

- Business service fails to get deactivated after Routing Engine switchover. [PR1280074](#)
- Backup Routing Engine goes to the database prompt with a vmcore if the configuration for the ASI interface that has gone down is deleted. [PR1281882](#)
- TLVs in ICRQ for actual-rate-downstream/actual-data-rate-upstream do not reflect PPPoE-IA value. [PR1286583](#)
- mspmand cored "@_arena_mALLOc" seen in Backup SDG's MS70. [PR1291664](#)
- L2TP subscribers are down after a GRES while verifying framed IPv6 route support for L2TP network server (LNS) at a higher scale with a maximum number of framed IPv6 routes. [PR1293783](#)
- Each subscriber session gets its own L2TP tunnel without "Tunnel-Client-Endpoint" from RADIUS. [PR1293927](#)
- The jl2tpd process might crash shortly after a GRES switchover. [PR1295248](#)
- [OC/ST] Continuous generation of *jl2tpd_era_Ins* log files occurs even though l2tp is not configured. [PR1302270](#)

Software Installation and Upgrade

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

Subscriber Access Management

- The DHCP subscriber might not get an IP address if the address pool utilization is tight. [PR1274870](#)
- Some RADIUS attributes might not be filtered out of the accounting-on/accounting-off message on an MX Series. platform. [PR1279533](#)

- IP assigned by RADIUS is incorrectly counted by the local pool after a Virtual Chassis switchover. [PR1286609](#)
- The authd process generates a core file at DynamicRequestEntry::addHistory authd_aaa_dyn_req. [PR1289215](#)
- Service interim for DHCP subscriber is not working in JSRC scenario. [PR1303553](#)
- The **show network-access aaa accounting** command might display additional entries. [PR1304594](#)
- Incorrect **Acct-Delay-Time in Radius Accounting-On** message is seen after rebooting the MX Series router acting as a BNG. [PR1308966](#)
- The delegated prefix from RADIUS is incorrectly parsed when the prefix is fewer than 20 bytes long. [PR1315557](#)

User Interface and Configuration

- Increasing commit times are seen. [PR1029477](#)
- The commitd process might generate a core file when removal of certain configuration is followed by a commit operation. [PR1267433](#)
- The commit might fail with the error of "Could not open configuration database" and "foreign file propagation (ffp) failed". [PR1287539](#)

VPNs

- Next generation MVPN SG entry and MVPN route persist after data stop. [PR1236733](#)
- Rpd memory leak is observed in a next generation MVPN environment. [PR1259579](#)
- Next generation MVPN IPv6 RP bootstrap type 3 S-PMSI AD route prefix ff02::d persist after BSR data stop. [PR1269234](#)
- L2circuits stitched via It peer interfaces might be stuck in "LD" (local site signaled down) status. [PR1305873](#)

SEE ALSO

[New and Changed Features | 93](#)

[Changes in Behavior and Syntax | 126](#)

[Known Behavior | 136](#)

[Known Issues | 142](#)

[Documentation Updates | 199](#)

[Migration, Upgrade, and Downgrade Instructions | 200](#)

[Product Compatibility | 207](#)

Documentation Updates

IN THIS SECTION

- [Subscriber Management Provisioning guide | 199](#)

This section lists the errata and changes in Junos OS Release 17.4R2 documentation for MX Series.

Subscriber Management Provisioning guide

- The *Broadband Subscriber Sessions User Guide* did not report that you can suspend AAA accounting, establish a baseline of accounting statistics, and resume accounting. This feature was introduced in Junos OS Release 15.1R4.
- Starting in Junos OS Release 15.1, the *Broadband Subscriber Sessions User Guide* and the [CLI Explorer](#) incorrectly included information about the **show extensible-subscriber-services accounting** command. This command is not present in the CLI. Instead, you can use accounting profiles to collect statistics from the Packet Forwarding Engine for Extensible Subscriber Services Manager (ESSM) subscribers. See [Flat-File Accounting Overview](#) for information about accounting for ESSM subscribers.

SEE ALSO

New and Changed Features 93
Changes in Behavior and Syntax 126
Known Behavior 136
Known Issues 142
Resolved Issues 157
Migration, Upgrade, and Downgrade Instructions 200
Product Compatibility 207

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.4 | 201](#)
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS | 201](#)
- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | 203](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 205](#)
- [Upgrading a Router with Redundant Routing Engines | 206](#)
- [Downgrading from Release 17.4 | 206](#)

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

Starting in Junos OS 17.4R1 release, FreeBSD 11.x is the underlying OS for all Junos OS platforms that were previously running on FreeBSD 10.x-based Junos OS. FreeBSD 11.x does not introduce any new features or modifications but is the latest version of FreeBSD.

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.x-based Junos OS	FreeBSD 11.x-based Junos OS
MX5,MX10, MX40,MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 17.4

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

Issue the following command:

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

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

For more information about the installation process, see [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).

Procedure to Upgrade to FreeBSD 11.x-Based Junos OS

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install FreeBSD 11.x-based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.

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

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

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.4R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.4R1.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.4R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.4R1.9-limited.tgz
```

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

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

- `http://hostname/pathname`
- `scp://hostname/pathname`

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.x) to Junos OS (FreeBSD 11.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 11.x, and Junos OS (FreeBSD 6.x) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 17.4 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. See the VM Host software administrative commands in the [Installation and Upgrade Guide](#).

Procedure to Upgrade to FreeBSD 6.x-Based Junos OS

Products impacted: MX5, MX10, MX40, MX80, MX104.

To download and install FreeBSD 6.x-based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.

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

- All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/jinstall-ppc-17.4R1.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot  
source/jinstall-ppc-17.4R1.9-limited-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.

- For software packages that are downloaded and installed from a remote location:

- `ftp://hostname/pathname`
- `http://hostname/pathname`
- `scp://hostname/pathname`

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.4 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. 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 17.4

To downgrade from Release 17.4 to another supported release, follow the procedure for upgrading, but replace the 17.4 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 93
Changes in Behavior and Syntax 126
Known Behavior 136
Known Issues 142
Resolved Issues 157
Documentation Updates 199
Product Compatibility 207

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 207](#)

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 93
Changes in Behavior and Syntax 126
Known Behavior 136
Known Issues 142
Resolved Issues 157
Documentation Updates 199
Migration, Upgrade, and Downgrade Instructions 200

Junos OS Release Notes for NFX Series

IN THIS SECTION

- New and Changed Features | 208
- Changes in Behavior and Syntax | 209
- Known Behavior | 210
- Known Issues | 210
- Resolved Issues | 211
- Documentation Updates | 212
- Migration, Upgrade, and Downgrade Instructions | 212
- Product Compatibility | 214

These release notes accompany Junos OS Release 17.4R2 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R2 New and Changed Features | 209
- Release 17.4R1 New and Changed Features | 209

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

Release 17.4R2 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 17.4R2.

Release 17.4R1 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 17.4R1.

NOTE: vSRX version 15.1X49-D100 is compatible with the Junos OS Release 17.4R1 for NFX Series devices.

SEE ALSO

- [Changes in Behavior and Syntax | 209](#)
- [Known Behavior | 210](#)
- [Known Issues | 210](#)
- [Resolved Issues | 211](#)
- [Documentation Updates | 212](#)
- [Migration, Upgrade, and Downgrade Instructions | 212](#)
- [Product Compatibility | 214](#)

Changes in Behavior and Syntax

There are no changes in behavior and syntax for NFX Series in Junos OS Release 17.4R2.

SEE ALSO

- [New and Changed Features | 208](#)
- [Known Behavior | 210](#)
- [Known Issues | 210](#)
- [Resolved Issues | 211](#)

[Documentation Updates | 212](#)

[Migration, Upgrade, and Downgrade Instructions | 212](#)

[Product Compatibility | 214](#)

Known Behavior

There are no known limitations in Junos OS Release 17.4R2 for NFX Series.

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

SEE ALSO

[New and Changed Features | 208](#)

[Changes in Behavior and Syntax | 209](#)

[Known Issues | 210](#)

[Resolved Issues | 211](#)

[Documentation Updates | 212](#)

[Migration, Upgrade, and Downgrade Instructions | 212](#)

[Product Compatibility | 214](#)

Known Issues

IN THIS SECTION

- [Known Issues: 17.4R2 | 211](#)

- [Known Issues: 17.4R1 | 211](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R2 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.

Known Issues: 17.4R2

- When dscp and dscp-ipv6 classifiers are configured together on L2 interface, the functionality may not work as expected

[PR1169529](#)

Known Issues: 17.4R1

There are no known issues for NFX Series in Junos OS Release 17.4R1.

SEE ALSO

New and Changed Features 208
Changes in Behavior and Syntax 209
Known Behavior 210
Resolved Issues 211
Documentation Updates 212
Migration, Upgrade, and Downgrade Instructions 212
Product Compatibility 214

Resolved Issues

There are no fixed issues in Junos OS Release 17.4R2 for NFX Series.

SEE ALSO

New and Changed Features 208
Changes in Behavior and Syntax 209
Known Behavior 210
Known Issues 210
Documentation Updates 212
Migration, Upgrade, and Downgrade Instructions 212
Product Compatibility 214

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 documentation for NFX Series.

SEE ALSO

[New and Changed Features | 208](#)

[Changes in Behavior and Syntax | 209](#)

[Known Behavior | 210](#)

[Known Issues | 210](#)

[Resolved Issues | 211](#)

[Migration, Upgrade, and Downgrade Instructions | 212](#)

[Product Compatibility | 214](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 212](#)
- [Basic Procedure for Upgrading to Release 17.4 | 213](#)

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

Basic Procedure for Upgrading to Release 17.4

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [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 17.4R2:

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 device or to your internal software distribution site.
10. Install the new jinstall package on the device.

SEE ALSO

[New and Changed Features | 208](#)

[Changes in Behavior and Syntax | 209](#)

[Known Behavior | 210](#)

[Known Issues | 210](#)

[Resolved Issues | 211](#)

[Documentation Updates | 212](#)

[Product Compatibility | 214](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 215](#)

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/>.

NFX250 Software Version Compatibility

This section lists the vSRX and Cloud CPE Solution software releases that are compatible with the Junos OS releases on the NFX250 platform:

Table 1: Software Compatibility Details with vSRX and Cloud CPE Solution

NFX250 Junos OS Release	vSRX	Cloud CPE Solution
15.1X53-D40.3	15.1X49-D40.6	Cloud CPE Solution 2.0
15.1X53-D41.6	15.1X49-D61	Cloud CPE Solution 2.1
15.1X53-D102.2	15.1X49-D61	Cloud CPE Solution 3.0
15.1X53-D47.4	15.1X49-D100.6	Cloud CPE Solution 3.0.1

Table 2: Software Compatibility Details with Only vSRX Installed

NFX250 Junos OS Release	vSRX
15.1X53-D40.3	15.1X49-D40.6
15.1X53-D41.6	15.1X49-D40.6
15.1X53-D45.3	15.1X49-D61
15.1X53-D47.4	15.1X49-D78.3
17.2R1	15.1X49-D75
17.3R1	15.1X49-D100
17.4R1	15.1X49-D100

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 208
Changes in Behavior and Syntax 209
Known Behavior 210
Known Issues 210
Resolved Issues 211
Documentation Updates 212
Migration, Upgrade, and Downgrade Instructions 212

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- [New and Changed Features | 217](#)
- [Changes in Behavior and Syntax | 229](#)
- [Known Behavior | 236](#)
- [Known Issues | 237](#)
- [Resolved Issues | 240](#)
- [Documentation Updates | 247](#)
- [Migration, Upgrade, and Downgrade Instructions | 247](#)
- [Product Compatibility | 252](#)

These release notes accompany Junos OS Release 17.4R2 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R2 New and Changed Features | 217](#)
- [Release 17.4R1 New and Changed Features | 217](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R2 for the PTX Series.

Release 17.4R2 New and Changed Features

There are no new features or enhancements to existing features for PTX Series in Junos OS Release 17.4R2.

Release 17.4R1 New and Changed Features

Hardware

- **PTX10016 Packet Transport Router**—Starting in Junos OS Release 17.4R1, the PTX10016 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, telco, 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, 21 rack unit (21RU) modular chassis provides 48 Tbps of throughput and 32 Bpps of forwarding capacity. The PTX10016 router has 16 slots for the line cards that can support a maximum of 2304 10-Gigabit Ethernet ports, 576 40-Gigabit Ethernet ports, or 480 100-Gigabit Ethernet ports.

You can deploy the PTX10016 router in the core of the network for the following functions:

- Label switching routing
- IP core routing
- Internet peering

PTX10016 Packet Transport Router supports two PTX10K 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. 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 breakout cable using standard structured cabling techniques.

For more information, see [PTX10016 Packet Transport Router Hardware Guide](#) .

- **Support for the CFP2-DCO-T-WDM-1 transceiver on the P2-100GE-OTN PIC (PTX)**—Starting in Junos OS Release 17.4R1, you can install the CFP2-DCO-T-WDM-1 transceiver on the P2-100GE-OTN PIC. The CFP2-DCO-T-WDM-1 transceiver is a 100-Gigabit digital pluggable CFP2 digital coherent optical module.

The CFP2-DCO-T-WDM-1 transceiver supports the following:

- International Telecommunication Standardization (ITU-T) OTN performance monitoring and alarm management
- 100-Gigabit Ethernet quadrature phase shift keying (QPSK) with differential encoding mode and soft-decision forward error correction (SD-FEC)
- proNX Service Manager (PSM)
- Junos OS YANG extensions
- Firmware upgrade

[See [100-Gigabit Ethernet OTN PIC with CFP2 \(PTX Series\)](#) .]

High Availability (HA) and Resiliency

- **Resiliency Support for PTX10K-LC1101 and PTX10K-LC1102 (PTX10016)**—Starting with Junos OS Release 17.4R1, resiliency support is enabled for the following components:
 - PTX10K-LC1101 and PTX10K-LC1102
 - Routing and Control Boards
 - Switch Interface Boards

Interfaces and Chassis

- **Fabric Management Support (PTX100016)**—Starting in Junos OS Release 17.4R1, you can set up and manage the fabric connections between the Packet Forwarding Engines in the PTX100016 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**

[See [Fabric Management Overview](#).]

- **Support for large-scale packet-forwarding features (PTX10000)**—Starting with Junos OS Release 17.4R1, PTX10000 router supports large scaling IPv4 and IPv6 forwarding information base (FIB). A maximum of 4 million routes are supported.
- **Support for pre-FEC BER monitoring when using the CFP2-DCO-T-WDM-1 transceiver (PTX Series)**—Starting in Junos OS Release 17.4R1, you can monitor the condition of an OTN link by using the pre-forward error correction (pre-FEC) bit error rate (BER) when using the CFP2-DCO-T-WDM-1 transceiver.

[See [Understanding Pre-FEC BER Monitoring and BER Thresholds](#).]

- **Support for a 16 Slot Chassis (PTX10016)**—Starting with Junos OS Release 17.4R1, the PTX10016 has 16 slots and supports core and edge profiles.

IPv6

- **Support for IPv6 statistics on PTX Series routers**—Starting in Junos OS Release 17.4R1, you can obtain the transit IPv6 statistics at both the physical interface and logical interface levels on third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008 by using both a CLI command and SNMP MIB counters. Use the **show interfaces statistics** command to display both physical interface and logical interface statistics. You can view only logical interface statistics if you use SNMP MIB counters. However, for aggregated Ethernet interfaces, the accounting is not done at the level of the child links and, thus, IPv6 statistics for child links are not displayed.

To start getting IPv6 statistics on third-generation FPCs, use the **route-accounting** statement at the **[edit forwarding-options family inet6]** hierarchy level. PTX Series routers with first-generation and second-generation FPCs do not display IPv6 statistics for physical interfaces or logical interfaces, and transit statistics on child links in aggregated Ethernet interfaces are also not taken into account.

NOTE: Egress accounting for IPV6 traffic is not performed for cases where MPLS packets arrives on TCC interface and egress out of the router as IPV6 packets.

[See [route-accounting](#) and [show interfaces extensive](#).]

Junos OS XML API and Scripting

- **Automation script library additions and upgrades (PTX Series)**—Starting in Junos OS Release 17.4R1, devices running Junos OS include new and upgraded Python modules as well as upgraded versions of Junos PyEZ and libslax. On-box Python automation scripts can use features supported in Junos PyEZ Release 2.1.4 and earlier releases to perform operational and configuration tasks on devices running Junos OS. Python automation scripts can also leverage new on-box Python modules including **ipaddress**, **jxmlease**, **pyang**, **serial**, and **six**, as well as upgraded versions of existing modules. In addition, SLAX automation scripts can include features supported in libslax release 0.22.0 and earlier releases.

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

Layer 2 Features

- **Support for Layer 2 protocols (PTX 10016)**—Starting in Junos OS Release 17.4R1, Layer 2 protocols are supported on PTX10016 routers that have third-generation FPCs installed. 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

- **Support for Layer 3 protocols (PTX 10016)**—Starting in Junos OS Release 17.4R1, Layer 3 protocols are supported on PTX10016 routers that have third-generation FPCs installed. 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 multiple, smaller configuration YANG modules (PTX Series)**—Starting in Junos OS Release 17.4R1, the YANG module for the Junos OS configuration schema is split into a root configuration module that is augmented by multiple, smaller modules. The root configuration module comprises the top-level configuration node and any nodes that are not emitted as separate modules. Separate, smaller modules augment the root configuration module for the different configuration statement hierarchies. Smaller configuration modules enable YANG tools and utilities to more quickly and efficiently compile and work with the modules, because they only need to import the modules required for the current operation.

[See [Understanding the YANG Modules That Define the Junos OS Configuration](#).]

- **Support for IS-IS sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export data for the IS-IS routing protocol through the Junos Telemetry Interface. Only gRPC streaming is supported. To export statistics for IS-IS, include the `/network-instances/network-instance[name_'instance-name']/protocols/protocol/isis/levels/level/` and `/network-instances/network-instance[name_'instance-name']/protocols/protocol/isis/interfaces/interface/levels/level/` set of paths. Use the `telemetrySubscribe` RPC to specify telemetry parameters and provision the sensor. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

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

- **Support for Packet Forwarding Engine traffic sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export Packet Forwarding Engine traffic statistics through the Junos Telemetry Interface. Both UDP and gRPC are supported. This sensor tracks reporting of Packet Forwarding Engine statistics counters and provides visibility into Packet Forwarding Engine error and drop statistics. The resource name for the sensor is `/junos/system/linecard/packet/usage/`. The OpenConfig path is `/components/component/subcomponents/subcomponent[name='FPC<id>:NPU<id>']/properties/property/`, where NPU refers to the Packet Forwarding Engine. To provision the sensor to export data through gRPC, use the `telemetrySubscribe` RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the `[edit services analytics]` hierarchy level.

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

- **Enhancements to LSP events sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, telemetry data streamed through gRPC for LSP events and properties is reported separately for each routing instance. To export data for LSP events and properties, you must now include

`/network-instances/network-instance[name_'instance-name']` in front of all supported paths. For example, to export LSP events for RSVP Signaling protocol attributes, use the following path: `/network-instances/network-instance[name_'instance-name']/mpls/signaling-protocols/rsvp-te/`. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Enhancement to BGP sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can specify to export the number of BGP peers in a BGP group for telemetry data exported through gRPC. To export the number of BGP peers for a group, use the following OpenConfig path: `/network-instances/network-instance[name_'instance-name']/protocols/protocol/bgp/peer-groups/peer-group[name_'peer-group-name']/state/peer-count/`. The BGP peer count value exported reflects the number of peering sessions in a group. For example, for a BGP group with two devices, the peer count reported is 1 (one) because each group member has one peer. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

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

- **Support for bypass LSP statistics for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export statistics for bypass label-switched paths (LSPs). Previously, only statistics for the primary LSP path were exported. The ability to export bypass LSP statistics helps to monitor the efficiency of global convergence when the bypass LSP is used to carry traffic during a link or node failure. Statistics are exported for the following:

- Bypass LSP originating at the ingress router of the protected LSP
- Bypass LSP originating at the transit router of the protected LSP
- Bypass LSP protecting the transit LSP as well as the locally originated LSP

When the bypass LSP is active, traffic is exported both on the bypass LSP and the ingress (protected) LSP. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module. You must also include the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level.

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

- **Support for BGP routing table sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can provision Junos Telemetry Interface sensors to export data for BGP routing tables (RIBs) for IPv4 and IPv6 routes. Each address family supports exporting data for five different tables. Only gRPC streaming is supported.

The tables are:

- **local-rib**—Main BGP routing table for the main routing instance.
- **adj-rib-in-pre**—NLRI updates received from the neighbor before any local input policy filters have been applied.
- **adj-rib-in-post**—Routes received from the neighbor eligible for best-path selection after local input policy filters have been applied.
- **adj-rib-out-pre**—Routes eligible for advertising to the neighbor before output policy filters have been applied.
- **adj-rib-out-post**—Routes eligible for advertising to the neighbor after output policy filters have been applied.

To stream data for the main BGP routing table for IPv4 routes, include the **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/loc-rib/** set of paths. To stream data for the main BGP routing table for IPv6 routes, include the **/bgp-rib/afi-safis/afi-safi/ipv6-unicast/loc-rib/** set of paths.

For the neighbor BGP routing tables for IPv4 routes, include the following sets of paths:

- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-in-pre/**
- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-in-post/**
- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-out-pre/**
- **/bgp-rib/afi-safis/afi-safi/ipv4-unicast/neighbors/neighbor/adj-rib-out-post/**

To stream data for IPv6 routes change **ipv4-unicast** **ipv6-unicast** in any of the paths.

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

- **Support for bidirectional authentication for gRPC for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can configure gRPC to require client authentication as well as server authentication. Previously, only the client initiating an RPC request was able to authenticate the server, that is, Juniper device, using SSL certificates. To enable bidirectional authentication, include the **mutual-authentication** statement at the **[edit system-services extension-service request-response grpc ssl]** hierarchy level. You must also configure and reference a certificate-authority profile. Include the **certificate-authority profile name** statement at the **[edit system services extension-service request-response grpc ssl]** hierarchy level. For **profile-name**, include the name of **certificate-authority** profile configured at the **[edit security pki ca-profile]** hierarchy level. This profile is used to validate the certificate provided by the client.

[See [gRPC Services for Junos Telemetry Interface](#).]

- **Enhancements to MPLS sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export statistics for MPLS through the Junos Telemetry Interface in the following categories:
 - Shared Risk Link Groups (SRLGs)

- Traffic engineering global attributes
- Traffic engineering interface attributes

Additional RSVP Signaling Protocol attributes, such as counters and interfaces, that were not previously available are also supported. Only gRPC streaming is supported.

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

- **FPC1 and FPC2 support for CPU and NPU sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export data for CPU memory and NPU memory and utilization for FPC1 and FPC2 on PTX Series routers through the Junos Telemetry Interface. Previously, only FPC3 was supported on these sensors. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

[See [sensor \(Junos Telemetry Interface\)](#) and [Guidelines for gRPC sensors.](#)]

MPLS

- **Support for static adjacency segment identifier for aggregate Ethernet member links using single-hop static LSP (PTX Series)**—Starting with Junos OS Release 17.4R1, you can configure a transit single-hop static label switched path (LSP) for a specific member link of an aggregate Ethernet (AE) interface. A static labeled route is added with next-hop pointing to the AE member link of an aggregate interface. Label for these routes is picked from the segment routing local block (SRLB) pool of the configured static label range. This feature is supported for AE interfaces only.

A new **member-interface** CLI command is added under the **next-hop** configuration at the **[edit protocols mpls static-label-switched-path lsp-name transit]** hierarchy to configure the AE member interface name. The static LSP label is configured from a defined static label range.

[See [Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-hop Static LSP.](#)]

- **Support for static adjacency segment identifier for IS-IS (PTX Series)**—Starting with Junos OS Release 17.4R1, you can configure static adjacency segment ID (SID) labels for an interface. You can configure two IPv4 adjacency SIDs (protected and unprotected), IPv6 adjacency SIDs (protected and unprotected) per level per interface. You can use the same adjacent SID for multiple interfaces by grouping a set of interfaces under an interface-group and configuring the adjacency-segment for that interface-group. For static adjacency SIDs, the labels are picked from either a static reserved label pool or from segment routing global block (SRGB).

[See [Static Adjacency Segment Identifier for ISIS.](#)]

- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (MX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage, which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP resignalling events. This is because the LSP is responsive to even minor increase or decrease in the utilization when current

reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of say 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP ressignalling events for small LSPs of say 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP ressignaling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless they are huge, for example, 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value based threshold along with the percentage based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at the **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.

- **Support for default time-out duration for self-ping on an LSP instance (PTX Series)**—Starting in Junos OS 17.4R1, the default time out duration for which the self-ping runs on an LSP instance is reduced from 65535 (runs until success) to 1800 seconds. You can also configure the self ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is enabled. The LSP types like CCC, P2MP, VLAN-based, and non-default instances do not support self-ping. You can configure **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.
- **Support for flap and MBB counter for LSP (PTX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on master routing engine only:
 - Flap counter-- Counts the number of times an LSP flaps down or up.
 - MBB counter— Counts the number of times an LSP incurs MBB.

The **clear mpls lsp counters** command resets the flap and the MBB counter to zero.

- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (PTX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
 - **Support for label history for MPLS protocol (PTX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entry per label. By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range.
- The **clear mpls label history** command clears the label history details.

Routing Protocols

- **Support for importing IGP topology information into BGP-LS (PTX Series)**—Starting in Junos OS Release 17.4R1, you can import interior gateway protocol (IGP) topology information into BGP-Link State (BGP-LS) in addition to RSVP-traffic engineering (RSVP-TE) topology information through the `Isdist.0` routing table. This allows you to monitor both IGP and traffic engineering topology information.

To install IGP topology information into the traffic engineering database, use the **set igp-topology** configuration statement at the **[edit protocols isis traffic-engineering]** and **[edit protocols ospf traffic-engineering]** hierarchy levels. To import IGP topology information into BGP-LS from `Isdist.0`, use the **set bgp-ls** configuration statement at the **[edit protocols mpls traffic-engineering database import igp-topology]** hierarchy level.

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

- **BGP supports segment routing policy for traffic engineering (PTX Series)**—Starting in Junos OS Release 17.4R1, a BGP speaker supports traffic steering based on a segment routing policy. The controller can specify a segment routing policy consisting of multiple paths to steer labeled or IP traffic. This feature enables BGP to support a segment routing policy for traffic engineering at ingress routers. The segment routing policy adds an ordered list of segments to the header of a packet for traffic steering. Static policies can be configured at ingress routers to allow routing of traffic even when the link to the controller fails.

To enable BGP IPv4 segment routing traffic engineering capability for an address-family, include the **segment-routing-te** statement at the **[edit protocols bgp family inet]** hierarchy level.

[See [Understanding Ingress Peer Traffic Engineering for BGP SPRING.](#)]

- **Topology-independent loop-free alternate for IS-IS (PTX Series)**—Starting in Junos OS Release 17.4R1, topology-independent loop-free alternate (TI-LFA) with segment routing provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. You can enable TI-LFA for IS-IS by configuring the **use-post-convergence-lfa** statement at the **[edit protocols isis backup-spf-options]** hierarchy level. TI-LFA provides protection against link failure, node failure, and failures of fate-sharing groups.

You can enable the creation of post-convergence backup paths for a given interface by configuring the **post-convergence-lfa** statement at the **[edit protocols isis interface *interface-name* level *level*]** hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

You can enable **node-protection** and/or **fate-sharing-protection** mode for a given interface at the **[edit protocols isis interface *interface-name* level *level* post-convergence-lfa]** hierarchy level. To use a particular fate-sharing group as a constraint for the fate-sharing-aware post-convergence path, you need to configure the **use-for-post-convergence-lfa** statement at the **[edit routing-options fate-sharing group *group-name*]** hierarchy level.

[See [Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS.](#)]

- **Support for network instance-based BGP configuration (PTX Series)**—Starting in Junos OS Release 17.4R1, you can configure BGP in a specific network instance. After the network instance is configured, you will be prompted with options for BGP configuration such as global bgp, neighbor bgp, and so on.

[See [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

- **DDoS protection support (PTX3000, PTX-5000, PTX1000, PTX10000)**—Starting with Junos OS Release 17.4R1, protection from DDoS attack is provided on PTX3000, PTX 5000, PTX1000, and PTX10000 routers only if they have PE-based FPCs installed.

If the total amount of traffic that a Routing Engine can handle exceeds its limit, the Routing Engine becomes overloaded and is unable to handle the routing protocol messages and other important control plane packets. This results in an inconsistent control plane protocol state and that is termed as DDoS attack.

With the support for DDoS protection, the firewall filters and policers available in Junos OS are used to discard or rate-limit control plane traffic so that such malicious traffic does not overwhelm and bring down the Routing Engine. The Packet Forwarding Engine does not support rate-based policers; therefore, DDoS protection works based on bandwidth.

DDoS protection is supported with the following protocols:

- L3 protocols— IGMP v4/v6, OSPF-Hello, OSPF, LDP-Hello, LDP, PIM-Ctrl, PIM-Data, RSVP, RIP, BFD, MHOP BFD, MSDP, BGP, TELNET, FTP, SSH, SNMP, NTP, TACACS, DNS, GRE, ICMP, MLD, NDP, and EGPv6
- L2 protocols— STP, LACP, LLDP, OAM-CFM, OAM-LFM, ISIS, ISO-TCC, ETH-TCC, and PVST

Exceptions to DDoS protection support include the following:

- L3 protocols are per protocol level and not at packet type level.
- Unsupported L3 protocols— DHCP v4/v6, PTP, VRRP, DTCP, RADIUS-SERVER, RADIUS-ACCT, RADIUS-AUTH, DIAMETER, DIAMETER-TCP, DIAMETER-SCTP, L2TP, LMP, BFDv6, Martian-address, and PIM-REGISTER
- Unsupported L2 protocols— STP, DOT1X, GARP, FC, Bridge control, and PVST
- FPC1 and FPC2 on PTX5000 router are not supported.

For more information, see [Distributed Denial-of-Service \(DDoS\) Protection Overview](#).

- **Support for EBGp route server (PTX Series)**—Starting in Junos OS Release 17.4R1, BGP feature is enhanced to support EBGp route server functionality. A BGP route server is the external BGP (EBGP) equivalent of an internal IBGP (IBGP) route reflector that simplifies the number of direct point-to-point EBGp sessions required in a network. EBGp route server propagates unmodified BGP routing information between external BGP peers to facilitate high scale exchange of routes in peering points such as Internet Exchange Points (IXPs). When BGP is configured as a route server, EBGp routes are propagated between peers unmodified, with full attribute transparency (NEXT_HOP, AS_PATH, MULTI_EXIT_DISC, AIGP, and Communities).

The BGP JET `bgp_route_service.proto` API has been enhanced to support route server functionality as follows:

- Program the EBGp route server.

- Inject routes to the specific route server RIB for selectively advertising it to the client groups in client-specific RIBs.

The BGP JET `bgp_route_service.proto` API includes a peer-type object that identifies individual routes as either EBGP or IBGP (default).

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

- **Support for BGP advertising aggregate bandwidth across external BGP links for load balancing (MX Series)**—Starting in Junos OS Release 17.4R1, BGP uses a new link bandwidth extended community, **aggregate-bandwidth**, to advertise aggregated bandwidth of multipath routes across external links. BGP calculates the aggregate of multipaths that have unequal bandwidth allocation and advertises the aggregated bandwidth to external BGP peers. A threshold to the aggregate bandwidth can be configured to restrict the bandwidth usage of a BGP group. In earlier Junos OS releases, a BGP speaker receiving multipaths from its internal peers advertised the link bandwidth associated with the active route. To advertise aggregated bandwidth of multipath routes and to set a maximum threshold, configure a policy with **aggregate-bandwidth** and **limit bandwidth** actions at the [edit policy-options policy-statement *name* then] hierarchy level.

[See [Advertising Aggregate Bandwidth Across External BGP Links for Load Balancing Overview](#).]

Security

- **Support for Layer 2 circuit pass-through (PTX Series)**—Starting in Junos OS Release 17.4R1, you can configure PTX Series routers to allow LACP, LLDP, OAM LFM, and OAM CFM packets to cross the Layer 2 circuit. To configure Layer 2 circuit pass-through, include the **l2circuit-control-passthrough** statement at the [set forwarding-options] hierarchy level.

NOTE: LACP can be configured only when the aggregated interface is configured with the ethernet-ccc encapsulation.

[See [l2circuit-control-passthrough](#).]

Services Applications

- **Reporting of true outgoing interface packets for inline flow monitoring (PTX Series)**—Starting in Junos OS Release 17.4R1, you can configure inline flow monitoring to report true packets for the outgoing interface. For ECMP, the actual outgoing interface used for a given flow is the true outgoing interface. To enable a true outgoing interface, include the **nexthop-learning enable** statement at the [set services flow-monitoring (version9 | version-ipfix) template *template-name*] hierarchy level.

[See [template \(Flow Monitoring IPFIX Version\)](#) or [version9 \(Flow Monitoring\)](#).]

- **Reporting of the true incoming interface for the sampled packets for inline flow monitoring (PTX Series)**—Starting in Junos OS Release 17.4R1, inline flow monitoring reports the true incoming interface

for the GRE-encapsulated packets entering the router for the configured inline flow monitoring filter criteria.

[See [Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers.](#)]

- **Support for inline JFlow version 9 flow templates (PTX 10016)**—Starting in Junos OS Release 17.4R1, you can use inline-J-Flow export capabilities with version 9 flow templates to define a flow record template suitable for IPv4 or IPv6 traffic.

Software Installation and Upgrade

- **Device serial number added to DHCP option 60 (PTX1000)**—Starting in Junos OS Release 17.4R1, DHCP option 60 (Vendor Class Identifier) includes the serial number of the device when you use zero touch provisioning to automate provisioning of the device configuration and software image. The serial number can uniquely identify the device in a broadcast network. The serial number appears in the format *Juniper-model-number*. For example, a PTX1000 router numbered DA000 appears as *Juniper-ptx1000-DA000*.

SEE ALSO

Changes in Behavior and Syntax	 229
Known Behavior	 236
Known Issues	 237
Resolved Issues	 240
Documentation Updates	 247
Migration, Upgrade, and Downgrade Instructions	 247
Product Compatibility	 252

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\)](#) | [230](#)
- [Interfaces and Chassis](#) | [230](#)
- [Management](#) | [232](#)
- [MPLS](#) | [232](#)
- [Multicast](#) | [233](#)
- [Network Management and Monitoring](#) | [233](#)

- Security | 235
- Software Licensing | 235
- Subscriber Management and Services | 235

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 17.4R2 for the PTX Series.

Class of Service (CoS)

- **Changes in configuration of hardware-based queue priority (PTX Series)**—Starting in Junos OS Release 17.4R1, the mapping of output queue priority values in the Junos OS to the output queue priorities supported by physical interfaces on PTX Series routers has changed. For shared scheduling, when **strict-high** is not configured, setting the priority to high maps to the hardware priority high. And for strict-priority scheduling, setting the priority to **high** maps to the hardware priority high. For the full mapping of output queue priorities, see [Understanding Scheduling on PTX Series Routers](#).

Interfaces and Chassis

- **Secondary interface (em2) raises an alarm when the link is down (PTX1000)**—Starting in Junos OS Release 17.4R1, secondary interface (em2) raises alarm when the link goes down. Earlier, no alarm was raised when an em2 (secondary interface) went down. Currently, the behavior is changed and an alarm will be raised when the interface link goes down as shown below:

```
user@host# run show chassis alarms
3 alarms currently active
Alarm time           Class  Description
2017-09-12 23:41:20 PDT  Major  FPC Management2 Ethernet Link Down
2017-09-12 23:38:45 PDT  Major  FPC0: PEM 2 Not Powered
2017-09-12 23:38:45 PDT  Major  FPC0: PEM 0 Not Powered
```

- **Modified output of the request vmhost zeroize command**—Starting with Junos OS Release 17.2, the command **request vmhost zeroize**, upon execution, prompts the user for confirmation to proceed. The following line is displayed:

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

- **Power supply alarm is not raised when the input switch status is OFF or power not connected (PTX10008, PTX10016)**—Starting in Junos OS Release 17.4R2, the power supply alarm **A power supply input has failed** will not be raised if the INP1/INP2 switch status is off and the power is not connected. Earlier, an alarm was raised for the power entry module (PEM) that it was not powered on as **Not Powered** irrespective of the switch state. For the power supply status, execute the **show chassis power** or **show chassis power detail** CLI command. The **DC input** is the new output parameter that provides information about the status of the input feed.

Previous behavior:

user@host> show chassis power

```

PEM 0:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC output:  864 W (zone 0, 72 A at 12 V, 34% of capacity)

PEM 1:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC output:  864 W (zone 0, 72 A at 12 V, 34% of capacity)

System:
  Zone 0:
    Capacity:      7500 W (maximum 7500 W)
    Allocated power: 6525 W (975 W remaining)
    Actual usage:   2616 W
    Total system capacity: 7500 W (maximum 7500 W)
    Total remaining power: 975 W

...

```

Current behavior:

user@host> show chassis power

```

PEM 0:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC input:   OK (No feed expected, Both feed connected)
  DC output:  576 W (zone 0, 48 A at 12 V, 23% of capacity)

PEM 1:
  State:      Online

```

```
Capacity: 2500 W (maximum 2500 W)
DC input: OK (No feed expected, Both feed connected)
DC output: 576 W (zone 0, 48 A at 12 V, 23% of capacity)
```

```
...
```

[See [show chassis power](#).]

Management

- **Changes to Junos OS YANG module naming conventions (PTX Series)**—Starting in Junos OS Release 17.4R1, the native Junos OS YANG modules use a new naming convention for the module's name, filename, and namespace. The module name and filename include the device family and the area of the configuration or command hierarchy to which the schema in the module belongs. In addition, the module filename includes a revision date. The module namespace is simplified to include the device family, the module type, and an identifier that is unique to each module and that differentiates the namespace of the module from that of other modules.

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

MPLS

- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (PTX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP resignaling events. This is because the LSP is responsive to even minor increases or decreases in the utilization when current reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of around 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP resignalling events for small LSPs of around 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP resignaling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless they are huge, for example, 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value-based threshold along with the percentage-based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at the **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.
- **Support for label history for MPLS protocol (PTX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at the **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entries per label . By default, label history is off and there is no maximum limit

for the number of entries for each label. The **show mpls label history *label-value*** command displays the label history for a given label value and the **show mpls label history label-range *start-label end-label*** command displays the history of labels between the given label range.

The **clear mpls label history** command clears the label history details.

- **Support for default time out duration for self-ping on an LSP instance (PTX Series)**—Starting in Junos OS 17.4R1, the default time out duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also configure the self ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration *value*** command at the **[edit protocols mpls label-switched-path *label-switched-path*]** hierarchy level. By default, self-ping is enabled. The LSP types like CCC, P2MP, VLAN-based, and non-default instances do not support self-ping. You can configure **no-self-ping** command at the **[edit protocols mpls label-switched-path *label-switched-path*]** hierarchy level to override the behavior of self-ping running by default.
- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (PTX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
- **Support for flap and MBB counter for LSP (PTX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on the master routing engine only:
 - Flap counter-- Counts the number of times an LSP flaps down or up.
 - MBB counter— Counts the number of times an LSP incurs MBB.

The **clear mpls lsp counters** command resets the flap and the MBB counter to zero.

Multicast

- **Support for rpf-selection statement for PIM protocol at global instance level (PTX Series)**—Starting in Junos OS 17.4R1, the **rpf-selection** statement for the PIM protocol is available at global instance level. You can configure **group** and **source** statements at the **[edit protocols pim rpf-selection]** hierarchy level.

Network Management and Monitoring

- **Change in default log level setting (PTX Series)**—In Junos OS Release, 17.4R1, the following changes were made in default logging levels:

Before this change:

- SNMP_TRAP_LINK_UP was LOG_INFO for both the physical (IFD) and logical (IFL) interfaces.
- SNMP_TRAP_LINK_DOWN was LOG_WARNING for both the physical (IFD) and logical (IFL) interfaces.

After this change:

- IFD LinkUp -> LOG_NOTICE (because this is an important message but less frequent)
- IFL LinkUp -> LOG_INFO (no change)
- IFD and IFL LinkDown -> LOG_WARNING (no change)

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

- **SNMP syslog messages changed (PTX Series)**—In Junos OS Release 17.4R1, 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](#).]

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a non-default routing instance with a non-default logical system (PTX Series)**—Starting in Junos OS Release 17.4R2, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

Security

- **Support for logging SSH key changes**—Starting with Junos OS Release 17.4R1, the configuration statement **log-key-changes** is introduced at the `[edit system services ssh]` hierarchy level. When the **log-key-changes** is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time **log-key-changes** was enabled. If **log-key-changes** was never enabled, then Junos OS logs all the authorized SSH keys.

Software Licensing

- **Key generator adds one day to make the duration of license show as 365 days (PTX Series)**—Starting in Junos OS Release 17.4R1, the duration of subscription licenses as generated by the **show system license** command and shown in the output duration is correct to the numbers of days. Before this fix, for example, for a 1-year subscription license, the duration was generated as 364 days. After the fix, the duration of the 1-year subscription now shows as 365 days.

[See [show system license](#).]

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (PTX Series)**—Starting in Junos OS Release 17.4R2, the `jdhcpd` process handles the second renew request differently if the DHCPv6 client CPE device does both of the following:
 - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
 - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

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

In earlier releases:

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

binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

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

SEE ALSO

New and Changed Features 217
Known Behavior 236
Known Issues 237
Resolved Issues 240
Documentation Updates 247
Migration, Upgrade, and Downgrade Instructions 247
Product Compatibility 252

Known Behavior

IN THIS SECTION

- [General Routing | 236](#)
- [Interfaces and Chassis | 237](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for PTX Series.

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

General Routing

- For CFP2-DCO-T-WDM-1 pluggable, Rx payload type is shown incorrectly. [PR1300423](#)
- On PTX (GingerAle PIC on Gladiator FPC), when backward frr is enabled on far end the convergence time is higher as extra delay (average 500 msec) incurred in triggering FRR, due to software based polling. [PR1303820](#)

- Forwarding-option filter (FTF) with DSCP action is not supported on PTX1000 and other PE chipset platforms. [PR1310747](#)
- In the specific case of semi-graceful RCB reboot initiated by the internal shell command: 'vhclient init 0', GRES takes longer to complete i.r 3 minutes as opposed to 21 seconds. The regular cli command: 'request vmhost reboot' (graceful) and a jack-out-jack-in of the RE (ungraceful) do not exhibit this delay. [PR1312065](#)

Interfaces and Chassis

- On PTX10008 and PTX10016 routers, if you remove the redundant Switch Interface Board (SIB) after upgrading Junos OS from Release 17.4R1 or Release 17.2X75-D90 to a later release, then an alarm is not generated. This is a known behavior and has no impact on the performance of the router.

SEE ALSO

[New and Changed Features | 217](#)

[Changes in Behavior and Syntax | 229](#)

[Known Issues | 237](#)

[Resolved Issues | 240](#)

[Documentation Updates | 247](#)

[Migration, Upgrade, and Downgrade Instructions | 247](#)

[Product Compatibility | 252](#)

Known Issues

IN THIS SECTION

- [General Routing | 238](#)
- [Interfaces and Chassis | 240](#)
- [MPLS | 240](#)
- [Platform and Infrastructure | 240](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R2 for the PTX Series.

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

General Routing

- While upgrading from Junos OS Release 15.1F based images to Junos OS 16.x and later releases or downgrading from Junos OS Release 16.x to Junos OS Release 15.1F images, if the "validate" option is enabled, chassisd might crash and upgrade or downgrade will fail. This issue should not be seen if both base and target images are from Junos OS Release 15.1F or Junos OS Release 16.x and later. [PR1171652](#)
- PTX Series FPC3 might receive noise on the FPC console port and interpret it as valid signals. This might cause a login failure on the console port and generate core files, or even reloads. [PR1224820](#)
- On PTX platforms with FPC3, PTX1000 with build-in chassis and QFX10000 platforms, an FPC major alarm might be seen if the system detects parity error, and the error messages **DLU: ilp memory cache error** and **DLU: ilp prot1 detected_imem_even error** might appear. The alarm might be cleared without intervention. This error might also be accompanied by traffic loss. [PR1251154](#)
- When an FPC goes offline or restarts, FPC 'x' sends traffic to FPC 'y'. The following error messages are seen on the destination FPC. A corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline.

```
Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error.
```

[PR1268678](#)
- On PTX 5000 with FPC type 3 in rare condition FPC might crash during lo0.0 inet6 input filter. [PR1268875](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- When msata installation fails in Linux-based linecards msata is removed from the boot list thinking its bad hardware. Linecard continues to pxie boot from network. There will be alarm on Routing Engine side indicating this failure. If msata installation is successful, then it can be manually added back. [PR1279344](#)
- The FPC state is displayed incorrectly. [PR1300795](#)
- When CFP2-DCO-T-WDM-1 plugged in PTX PIC, after repeated configuration rollback, link sometime might take longtime to come-up. [PR1301462](#)
- When CFP2-DCO-T-WDM-1 plugged in PTX PIC, after FPC restart sometimes carrier frequency offset tca is raised even when tca not enabled. [PR1301471](#)

- iLatency (calculated by differing producer timestamp and gRPC server timestamp) might sometimes be negative for Packet Forwarding Engine related telemetry packets due to drift in Routing Engine and Packet Forwarding Engine NTP servers. [PR1303376](#)
- The crash indicates simultaneous operation on an ephemeral instance. When a process wants to open ephemeral configuration in merge view, other activities (like purging, deletion/recreation) are being carried out on this ephemeral instance. The occurrence of this core is rare. [PR1305424](#)
- When the command **set forwarding-options l2circuit-control-passthrough** is configured on a working LACP bundle, the interface goes down and no traffic will pass. [PR1320407](#)
- When DCO hot-plug is done *before* link is UP, the FPC might crash. [PR1322260](#)
- On PTX Series platform with FPC type 3, the error message could be observed when FPC card goes online or off-line. [PR1322491](#)
- On PTX Series platform with broadband cards (for example, FPC1, FPC2) and class of service (CoS) used, a high priority queue might not get the entire configured bandwidth. [PR1324853](#)
- In streaming telemetry scenario, when **commit full** command is executed, na-grpd daemon might disconnect streaming telemetry. [PR1326366](#)
- In the event of Routing Engine switchover, if there are existing sensor subscriptions, they will continue to show in **show agent sensors** output after the switchover. These stale sensors will be cleared when the device becomes master again. [PR1347779](#)
- This issue applies to PTX3000 FPC-SFF-PTX-P1-A/FPC-SFF-PTX-T. When BFD is configured for BGP, any L3 packet injects from linecard might lead to the result that BFD sessions does not come up on PTX3000. [PR1352112](#)
- If output firewall filter is configured with "syslog" option, the host interface might be wedged on a PTX1000, or a PTX Series platform with FPC type 3. [PR1354580](#)
- In case of link flap with LDP adjacencies, there is a possibility that the traffic drop is up to 1-2 seconds. [PR1357925](#)
- When a Routing Engine reboots and comes up again, it sends gratuitous ARP packets to the internal interfaces in order to advertise its MAC address. These packets get in to the UKERN running on the FPC, which drops these packets. The messages seen here are printed just before dropping these packets. These error messages are harmless and do not disrupt working of any feature. [PR1374372](#)

Interfaces and Chassis

- Junos OS upgrade involving Junos OS Release 14.2R5 and later maintenance releases, and Junos OS Release 16.1 later mainline releases with CFM configuration might cause cfmd crash after upgrade. This is due the old version of /var/db/cfm.db. [PR1281073](#)

MPLS

- LDP to BGP stitching with eBGP indirect next hop having an implicit null label does not work. It does work when BGP indirect next hop has a real label. As a workaround, ensure the peer advertises a real label by adding another router between the egress and ingress PE devices. Use IBGP that gets resolved over LDP or RSVP-TE LSPs. This ensures that the BGP indirect next hop has a real label. [PR1254702](#)

Platform and Infrastructure

- Execution of Python scripts through enhanced automation is only supported on non-veriexec images. [PR1334425](#)

SEE ALSO

New and Changed Features 217
Changes in Behavior and Syntax 229
Known Behavior 236
Resolved Issues 240
Documentation Updates 247
Migration, Upgrade, and Downgrade Instructions 247
Product Compatibility 252

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R2 | 241](#)
- [Resolved Issues: 17.4R1 | 244](#)

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

Resolved Issues: 17.4R2

General Routing

- The commands **restart na-grpc-server** and **restart na-mqtt** for MTRE does not work. [PR1284121](#)
- PTX1000s routers are periodically exporting IPFIX flow packets with high octet values. [PR1286427](#)
- On a PTX1000, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3-20170721 fails frequently on enabling sampling. [PR1296533](#)
- Interfaces might go down when Packet Forwarding Engine encounters **TOE::FATAL ERROR**. [PR1300716](#)
- On a PTX3000 platform, powering on a FPC (PTX-IPLC-B-32) card might cause the other FPC cards to reboot. [PR1302304](#)
- Internal latency is high during the initial subscription of sensors. [PR1303393](#)
- Repeated log message **%PFE-3 fpcX expr_nh_index_tree_ifl_get and expr_nh_index_tree_ipaddr_get** is observed when the sampling packet is discarded with a log(or syslog) configuration statement under the firewall filter. [PR1304022](#)
- The "interface hold-time down" timer does not take effect on a PTX5000 with an optical interface. [PR1307302](#)
- Packet Forwarding Engine error messages are flooding as **expr_sensor_update_cntr_to_sid_tree** after a deletion and rollback as **protocols isis source-packet-routing node-segment** . [PR1309288](#)
- On a PTX10000, a suppress chassis alarm for switched off PEM is seen. [PR1311574](#)
- The SIB LED on the FPD comes to GREEN-STEADY state even before an SIB comes online. [PR1311632](#)
- PTX10000 router does not bounce FPC without warning or alarm for different port speed settings. [PR1311875](#)
- The rpd process generates a core file after multiple session flaps on a scale setup. [PR1312169](#)
- Many logs are generated after executing many Vhclient related commands. [PR1315128](#)
- Memory leak in chassisd daemon is noticed while streaming active telemetry subscriptions. [PR1315672](#)
- The Packet Forwarding Engine on a PTX FPC3 or PTX10000 line card might be disabled if the interface connecting to the Packet Forwarding Engine goes down. [PR1315823](#)
- The physical interfaces might generate framing errors when ports are connected to an odd interface. [PR1317827](#)

- On MX0016 router, after Jack-out/Jack-in, FPCs shows up as "No-Power" for some time; FPC, however, comes up. [PR1319156](#)
- There is no traffic flowing with IPv6 payload prefixes on PTX Series platform. [PR1319273](#)
- PTX10000 routers for 100G LR4 optics with part number 740-061409 changes the display of the **show chassis hardware** command to QSFP-100G-LR4-T2. [PR1322082](#)
- The rpd process might crash when the OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)
- On Junos OS MPC7, MPC8, and MPC9, PTX-FPC3 (FPC-P1, FPC-P2), PTX3000-FPC3, and PTX1000 line cards might crash upon receipt of specific MPLS packet (CVE-2018-0030). [PR1323069](#)
- On a PTX1000, the local time on an FPC might be different from the local time on a Junos-VM or VM-host. [PR1325048](#)
- The GRE traffic is not decapsulated by the firewall filter. [PR1325104](#)
- PTX Series routers MKA sessions are not coming up after changing CA parameters such as transmit-interval, and key-server-priority. [PR1325392](#)
- MPLS traceroute fails across the PTX Series platform. [PR1327609](#)
- On a PTX5000 with FPC3 line cards, on a PTX10000, and on PTX1000 platforms, output firewall filters that are configured with "syslog" and "discard" actions do not perform the "syslog" action. [PR1328426](#)
- PTX10000 line card might reboot continuously after upgrading to Junos OS Release 17.2R1 or later if HMC BIST fails. [PR1330618](#)
- There is a link instability after a link-down event on PTX Series routers. [PR1330708](#)
- A PTX5000 FPC might reboot in certain rare scenarios when **interface-specific policer** is configured. [PR1335161](#)
- The directory where the init and configuration files are stored are removed when **request system zeroize** is executed. Hence, the telemetry process does not start after the zeroize is done. [PR1336004](#)
- A member of an IPv4 unicast next hop might get stuck in "Replaced" state after an interface flap. [PR1336201](#)
- Disabling a breakout 10G port on interface et-0/0/5 will unexpectedly disable another breakout 10G port on interface et-0/0/5. [PR1337975](#)
- FPC/FPC2/FPC E on a PTX Series does not forward traffic. [PR1339524](#)
- The link goes down on a PTX3000. The PTX5000 with an FPC3 is inserted after the router reboots or the link flaps. [PR1340612](#)
- The interface might flap continuously after reboot on PTX5000 and PTX3000 platform with P3-24-U-QSFP28 PIC. [PR1342681](#)
- On a PTX1008, the 30-port coherent line card (DWDM-IC) does not come up. [PR1344732](#)
- The FPC was rebooted a few minutes after loading the configuration. [PR1346467](#)

- Sensors are not getting cleared up after doing Routing Engine switchover. [PR1347779](#)
- MPLS traceroute for P2MP LSPs configured with link-protection causes FPC crash. [PR1348314](#)
- The interface of 15 100G ports PIC might delay 60 seconds to come up. [PR1357410](#)
- P2MP LSP replication traffic loss on an aggregated Ethernet bundle after a member link is down on PTX Series routers. [PR1359974](#)
- The route stuck might be seen after BGP neighbor and route flapping. [PR1362560](#)
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with **Link-Layer-Down** flag set. [PR1365263](#)
- On a PTX Series IPLC (OPT3-SFF-PTX FPC), a first J-UKERN crash triggers multiple secondary J-UKERN crashes. [PR1365791](#)
- The **commit** or **commit check** might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)

Infrastructure

- The ixlv interface statistics are not accounting properly. [PR1313364](#)
- PTX Series routers might get to abnormal state because of the malfunction of the protection mechanism for the F-Label. [PR1336207](#)

MPLS

- Traffic drops during NSR switchover for RSVP P2MP provider tunnels used by MVPN. [PR1293014](#)
- Traffic loss for static LSP configured with the **stitch** configuration statement. [PR1307938](#)
- The rpd process might crash on the backup Routing Engine due to memory exhaustion. [PR1328974](#)
- The rpd might crash with MPLS traceoption configured. [PR1329459](#)
- MPLS LSP statistics are not shown in the CLI command. **show mpls lsp ingress statistics**. [PR1344039](#)
- On PTX1000, PTX10000, or QFX10000 platform, when hybrid memory cube (HMC) error occurs, label switched paths (LSPs) might go down because of the incorrect bandwidth requested for auto-bandwidth adjustment. [PR1374102](#)

Platform and Infrastructure

- Continuous log messages occur. For example, you see **tftpd[23724]: Timeout #35593 on DATA block 85**. [PR1315682](#)
- Traffic black hole seen along with **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH messages**. [PR1349332](#)
- Traffic black hole seen along with **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH messages**. [PR1357707](#)

Routing Protocols

- With BGP LU FRR in an inter-AS scenario, a very high FRR time is visible once the link is up. [PR1307258](#)
- The rpd process might constantly consume high CPU in a BGP setup. [PR1315066](#)
- The primary path of an MPLS LSP might switch to another address. [PR1316861](#)
- The rpd process might crash after deactivating the passive interface under IS-IS. [PR1318180](#)
- The rpd process might crash continuously on both Routing Engines when **backup-spf-options remote-backup-calculation** is configured in an IS-IS protocol. [PR1326899](#)
- The rpd process generates a core file at `ispfc_incrementally_mend_one_postf_sp_tree (postf_spf_res=<optimized-out>, topo=<optimized-out>)` at `../../../../../../../../src/junos/usr.sbin/rpd/lib/igp-spf-compute/igp_spf_compute_ti_lfa.c:3364` PTX1K. [PR1339296](#)
- A protocol churn might lead the rpd process to crash. [PR1341466](#)
- The rpd process generates a core file while running a streaming telemetry test. [PR1347431](#)

VPNs

- In a specific CE device environment in which asynchronous notification is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. [PR1282875](#)

Resolved Issues: 17.4R1

General Routing

- PTX1000 : `ch_get_product_attribute.324: Cannot find chassisd` error message appears while loading image. [PR1217505](#)
- The rpd might crash on platforms with 64-bit X86 RE if IPv6 is configured. [PR1224376](#)
- On PTX Series platforms, chassisd thread is not getting CPU resources for 200 seconds and multiple chassisd core files are continuously generated. [PR1226992](#)
- The "validation-state:unverified" routing entry might not be displayed with proper location when using the show route command. [PR1254675](#)
- The routing protocol process (rpd) might crash after flapping BGP sessions and routes. [PR1269327](#)
- 100Base-ER4 (740-045420) is displayed as "UNKNOWN" in **show chassis hardware** in Junos OS Release 15.1R5.5. [PR1280089](#)
- FPC cards might go offline due to fabric healing in PTX3000 with SIB-SFF-PTX-240-S platform. [PR1282983](#)
- "Host 0 RTC Battery failure" error messages are seen on PTX1000, QFX10000-series after upgrade to Junos version 16.1. [PR1287128](#)

- The MPLS TTL might be reset to 255 on third-generation PTX Series FPCs if **mpls no-propagate-ttl** protocols configuration statement is configured. [PR1287473](#)
- LSP traffic gets silently dropped or discarded after link goes down in bypass path. [PR1291036](#)
- The rpd core file might be generated when restarting the process through CLI. [PR1291110](#)
- Incorrect SNMP OID values are sent in SNMP traps for removal or insertion of front panel display on PTX Series routers. [PR1294741](#)
- LINK LED is “RED” when the port is disabled on PTX Series routers. [PR1294871](#)
- The rpd core might be generated after interface or BGP flapping. [PR1294957](#)
- The chassisd process might run out of memory and restart on PTX1000 platform. [PR1295691](#)
- PTX5K/SyncE (ESMC): clock is not getting locked if the source interface is a member link of an ae bundle. [PR1296015](#)
- CoS escalation: Alarms and syslog errors are seen with priority strict-high on AF4 queue, on the oversubscription cases (1X100G egress to 1X10G egress setup). [PR1297343](#)
- **PE Chip: FATAL ERROR!! from pe0[0]: HMCIF:** might trigger FPC crash or slow route/next-hop installation processing. [PR1300180](#)
- PTX Series FPC3 will drop MPLS packets if its oif has inet MTU that is less than the MPLS packet size. [PR1302256](#)
- Heap memory leak might be observed on PTX Series FPCs during a multicast route installation into the Packet Forwarding Engine. [PR1302303](#)
- The third-generation FPC (FPC3-SFF-PTX) is not booting up on Control Board/Routing Engine systems. [PR1303295](#)
- On PTX3000 and PTX5000 platforms, the 100G interfaces might not come up. [PR1303324](#)
- If MPLS LSP self-ping is enabled (self-ping is enabled by default), kernel might panic with the error message **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)
- PTX3000 with RCB-PTX Routing Engine might not be online or recognize IPLCs. [PR1304124](#)
- The 10g interface might flap if it is set to 100g speed. [PR1315079](#)
- The physical interfaces might generate framing errors when ports are connecting odd interfaces. [PR1317827](#)
- The physical interfaces might generate framing errors when ports are connecting to odd interface. [PR1317827](#)
- No traffic is flowing with IPV6 payload prefixes on PTX platform. [PR1319273](#)
- PFT : RCB restarts continuously after executing **request system reboot**. [PR1320977](#)

Infrastructure

- The **show system users** CLI command output displays a larger number of users than that are actually using the router. [PR1247546](#)

Interfaces and Chassis

- The interface might flap when performing Routing Engine switchover if the member link of an ae interface is configured with framing settings. [PR1287547](#)
- 100-Gigabit Ethernet interfaces might not come up if **otn-options laser-enable** is configured on PTX Series platforms. [PR1297164](#)
- LFM discovery state appears as Fault for aggregate Ethernet interface after GRES. [PR1299534](#)

Multiprotocol Label Switching (MPLS)

- Stale RSVP LSP entry after NSR switchover and session is not refreshed. [PR1292526](#)
- The rpd might crash if the MPLS LSP path change occurs. [PR1295817](#)

Platform and Infrastructure

- Mgd generates core file when downgrading from Junos OS Release 17.3-20170721 to 16.1X65D40.2. The mgd core is also overwritten if attempted multiple times. [PR1296504](#)

Routing Protocols

- A few BFD sessions are flapping while coming up after FPC restart/reboot. [PR1274941](#)
- The rpd generated core files multiple times when it received an “OPEN” message from an existing BGP peer. [PR1299054](#)
- With BGP LU FRR in Inter-As scenario, a very high FRR time is seen once link is up. [PR1307258](#)
- Assignment of SUB-TLV values for Segment routing TE policy SUB-TLVs. [PR1315486](#)

SEE ALSO

[New and Changed Features | 217](#)

[Changes in Behavior and Syntax | 229](#)

[Known Behavior | 236](#)

[Known Issues | 237](#)

[Documentation Updates | 247](#)

[Migration, Upgrade, and Downgrade Instructions | 247](#)

[Product Compatibility | 252](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 documentation for PTX Series.

SEE ALSO

[New and Changed Features | 217](#)

[Changes in Behavior and Syntax | 229](#)

[Known Behavior | 236](#)

[Known Issues | 237](#)

[Resolved Issues | 240](#)

[Migration, Upgrade, and Downgrade Instructions | 247](#)

[Product Compatibility | 252](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 247](#)
- [Upgrading a Router with Redundant Routing Engines | 248](#)
- [Basic Procedure for Upgrading to Release 17.4 | 248](#)

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. 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 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 17.4

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 17.4R2:

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 17.4R2 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-17.4R2.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-17.4R2.SPIN-export-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the `request vmhost software add` command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 17.4 `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	217
Changes in Behavior and Syntax	229
Known Behavior	236
Known Issues	237
Resolved Issues	240
Documentation Updates	247
Product Compatibility	252

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 252](#)

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 217
Changes in Behavior and Syntax 229
Known Behavior 236
Known Issues 237
Resolved Issues 240
Documentation Updates 247
Migration, Upgrade, and Downgrade Instructions 247

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- New and Changed Features | 253
- Changes in Behavior and Syntax | 266
- Known Behavior | 272
- Known Issues | 276
- Resolved Issues | 280
- Documentation Updates | 291
- Migration, Upgrade, and Downgrade Instructions | 292
- Product Compatibility | 305

These release notes accompany Junos OS Release 17.4R2 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.


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

New and Changed Features

IN THIS SECTION

- Release 17.4R2 New and Changed Features | 255
- Release 17.4R1 New and Changed Features | 255

This section describes the new features for the QFX Series switches in Junos OS Release 17.4R2.



NOTE: The following QFX Series platforms are supported in Release 17.4R2: QFX5100, QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016.

Release 17.4R2 New and Changed Features

Restoration Procedures and Failure Handling

- **Device recovery mode support in Junos OS with upgraded FreeBSD (QFX Series)**—Starting in Junos OS Release 17.4R2, devices running Junos OS with an upgraded FreeBSD and a saved rescue configuration have an automatic device recovery mode should the system go into amnesiac mode. The new process has the system automatically reboot with the saved rescue configuration. Then, the system displays "Device is in recovery mode" in the CLI (in both operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

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

Release 17.4R1 New and Changed Features

Hardware

- **QFX10000-30C-M line card (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.4R1-S2, the QFX10000-30C-M line cards provides 30 ports of either 100-gigabit or 40-gigabit QSFP28 with MACsec features.

Class of Service (CoS)

- **Priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.4R1, to support lossless traffic across Layer 3 connections to Layer 2 subnetworks on QFX5110 and QFX5200 switches, you can configure priority-based flow control (PFC) to operate using 6-bit DSCP values from Layer 3 headers of untagged VLAN traffic, rather than IEEE 802.1P priority values in Layer 2 VLAN-tagged packet headers. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).

To enable DSCP-based PFC, map a forwarding class to a PFC priority using the **pfc-priority** statement, define a congestion notification profile to enable PFC on traffic specified by a 6-bit DSCP value, and set up a classifier for the DSCP value and the PFC-mapped forwarding class.

[See [Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.](#)]

EVPNs

- **Support for LACP in EVPN active-active multihoming (QFX5100, QFX5100 Virtual Chassis, QFX5110, and QFX5200 switches)**—Starting with Junos OS Release 17.4R1, an extra level of redundancy can be achieved in an Ethernet VPN (EVPN) active-active multihoming network by configuring the Link Aggregation Control Protocol (LACP) on both the endpoints of the link between the multihomed customer edge (CE) and provider edge (PE) devices. The link aggregation group (LAG) interface of the multihomed CE-PE link can either be in the active or in the standby state. The interface state is monitored and operated by LACP to ensure fast convergence on isolation of a multihomed PE device from the core. When there is a core failure, a traffic black hole can occur at the isolated PE device. With the support

for LACP on the CE-PE link, at the time of core isolation, the CE-facing interface of the multihomed PE device is set to the standby state, thereby blocking data traffic transmission from and toward the multihomed CE device. After the core recovers from the failure, the interface state is switched back from standby to active.

To configure LACP in EVPN active-active multihoming network:

- On the multihomed CE device include the `lacp active` statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy.
- On the multihomed PE device include the `lacp active` statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy, and include the `service-id` number statement at the `[edit switch-options]` hierarchy.

[See [Understanding LACP for EVPN Active-Active Multihoming](#).]

- **EVPN pure type-5 route support (QFX5110 switches)**—Starting with Junos OS Release 17.4R1, 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. 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. To enable two-level equal-cost multipath (ECMP) next hops in an EVPN-VXLAN overlay network, you must also include the `overlay-ecmp` statement at the `[edit forwarding-options vxlan-routing]` hierarchy level.

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

- **SPRING support for EVPN (QFX10000 switches)**—Starting in Junos OS Release 17.4R1, Junos OS supports using Source Packet Routing in Networking (SPRING) as the underlay transport in EVPN. SPRING tunnels enable routers to steer a packet through a specific set of nodes and links in the network.

To configure SPRING, use the `source-packet-routing` statement at the `[edit protocols isis]` hierarchy level.

[See [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

- **Support for duplicate MAC address detection and suppression (QFX10000 switches)**— When a MAC address relocates, PE devices can converge on the latest location by using sequence numbers in the extended community field. Misconfigurations in the network can lead to duplicate MAC addresses. Starting in Junos OS Release 17.4R1, Juniper supports duplicate MAC address detection and suppression. You can modify the duplicate MAC address detection settings on the switch by configuring the detection window for identifying duplicate MAC address and the number of MAC address moves detected within the detection window before duplicate MAC detection is triggered and the MAC address is suppressed. In addition, you can also configure an optional recovery time that the switch waits before the duplicate MAC address is automatically unsuppressed.

To configure duplicate MAC detection parameters, use the **detection-window**, **detection-threshold**, and **auto-recovery-time** statements at the **[edit routing instance *routing-instance-name* protocols evpn duplicate-mac-detection]** hierarchy level.

To clear duplicate MAC suppression manually, use the **clear evpn duplicate-mac-suppression** command.

[See [Overview of MAC Mobility](#).]

General Routing

- **Enhancement to show chassis forwarding-options command (QFX5200 Virtual Chassis)**—Starting in Junos OS Release 17.4R1, the **show chassis forwarding-options** command displays information about memory banks for QFX5200 Virtual Chassis only for the master. This information is not displayed for all the other members. Memory banks can be partitioned among different types of forwarding table entries through the Unified Forwarding Table feature. Values remain the same across all members. All configuration changes for the Unified Forwarding Table are made through the Master.

[See [show chassis forwarding-options](#).]

Interfaces and Chassis

- **Support for resilient hashing for LAGs and ECMP (QFX10000)**—Starting with Junos OS Release 17.4R1 on QFX10000 switches, you can prevent the reordering of flows to active paths in link aggregation groups (LAGs) or ECMP when one or more paths fail. Only flows that are on inactive paths are redirected. It overrides the default behavior of disrupting all existing, including active, TCP connections when an active path fails. You can optionally set a specific value for the resilient-hash seed that differs from the hash-seed value that will be used by the other hash functions on the switch. A resilient hashing configuration on ECMP is applied through use of a route policy.

[See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping](#).]

- **Enterprise profile for Precision Time Protocol (PTP) (QFX10002 switches)**—Starting with Junos OS Release 17.4R1, the enterprise profile, which is based on PTPv2, provides the ability for enterprise and financial markets to timestamp on different systems and to handle a range of latency and delays.

The enterprise profile supports the following options:

- IPv4 multicast transport
- Ordinary and boundary clocks
- 1-Gigabit SFP grandmaster port
- 512 downstream slave clocks

You can configure the enterprise profile at the **[edit protocols ptp *profile-type*]** hierarchy.

[See [Understanding Transparent Clocks in Precision Time Protocol](#).]

- **Support for Precision Time Protocol (PTP) transparent clock (QFX5200 switches)**—Starting with Junos OS Release 17.4R1, 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, UDP over IPv4 and IPv6 and unicast and multicast transparent clock are supported.

[See [Understanding Transparent Clocks in Precision Time Protocol](#).]

Junos OS XML API and Scripting

- **Automation script library additions and upgrades (QFX Series)**—Starting in Junos OS Release 17.4R1, devices running Junos OS include new and upgraded Python modules as well as upgraded versions of Junos PyEZ and libslax. On-box Python automation scripts can use features supported in Junos PyEZ Release 2.1.4 and earlier releases to perform operational and configuration tasks on devices running Junos OS. Python automation scripts can also leverage new on-box Python modules including **ipaddress**, **jxmlease**, **pyang**, **serial**, and **six**, as well as upgraded versions of existing modules. In addition, SLAX automation scripts can include features supported in libslax release 0.22.0 and earlier releases.

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

Management

- **Enhancements to LSP events sensor for Junos Telemetry Interface (QFX5110, QFX5200, and QFX10000 switches)** —Starting with Junos OS Release 17.4R1, telemetry data streamed through gRPC for LSP events and properties is reported separately for each routing instance. To export data for LSP events and properties, you must now include `/network-instances/network-instance[name_ 'instance-name']/` in front of all supported paths. For example, to export LSP events for RSVP Signaling protocol attributes, use the following path:

`/network-instances/network-instance[name_ 'instance-name']/mpls/signaling-protocols/rsvp-te/`. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

- **Enhancement to BGP sensor for Junos Telemetry Interface (QFX5110, QFX5200, and QFX10000 switches)**—Starting with Junos OS Release 17.4R1, you can specify to export the number of BGP peers in a BGP group for telemetry data exported through gRPC. To export the number of BGP peers for a group, use the following OpenConfig path:

`/network-instances/network-instance[name_ 'instance-name']/protocols/protocol/bgp/peer-groups/peer-group[name_ 'peer-group-name']/state/peer-count/`. The BGP peer count value exported reflects the number of peering sessions in a group. For example, for a BGP group with two devices, the peer count reported is 1 (one) because each group member has one peer. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

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

- **Support for multiple, smaller configuration YANG modules (QFX Series)**—Starting in Junos OS Release 17.4R1, the YANG module for the Junos OS configuration schema is split into a root configuration module that is augmented by multiple, smaller modules. The root configuration module comprises the top-level configuration node and any nodes that are not emitted as separate modules. Separate, smaller modules augment the root configuration module for the different configuration statement hierarchies. Smaller configuration modules enable YANG tools and utilities to more quickly and efficiently compile and work with the modules, because they only need to import the modules required for the current operation.

[See [Understanding the YANG Modules That Define the Junos OS Configuration](#).]

Multicast

- **Support for static multicast route leaking for VRF and virtual-router instances (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 17.4R1, you can configure your switch to share IPv4 multicast routes among different virtual routing and forwarding (VRF) instances or different virtual-router instances. Only multicast static routes with a destination-prefix length of /32 are supported for multicast route leaking. Only Internet Group Management Protocol version 3 is supported. To configure multicast route leaking for VRF or virtual-router instances, include the **next-table routing-instance-name.inet.0** statement at the [edit routing-instances routing-instance-name routing-options static route destination-prefix/32] hierarchy level. For **routing-instance-name**, include the name of a VRF or virtual-router instance.

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

- **MLD snooping versions 1 and 2 (QFX5100 switches and Virtual Chassis)**—Starting with Junos OS Release 17.4R1, QFX5100 switches and QFX5100 Virtual Chassis support Multicast Listener Discovery (MLD) snooping version 1 (MLDv1) and version 2 (MLDv2). MLD snooping constrains the flooding of IPv6 multicast traffic on VLANs. When MLD snooping is enabled on a VLAN, the switch examines MLD messages encapsulated within ICMPv6 packets transferred between hosts and multicast routers. The switch learns which hosts are interested in receiving traffic for a multicast group, and forwards multicast traffic only to those interfaces in the VLAN that are connected to interested receivers instead of flooding the traffic to all interfaces. You configure MLD snooping parameters and enable MLD snooping using configuration statements at the [edit protocols] mld-snooping vlan **vlan-name** hierarchy.

[See [Understanding MLD Snooping on Switches](#).]

- **Multicast-only fast reroute (MoFRR) (QFX5100, QFX5110, and QFX5200 switches)**—Starting in Junos OS Release 17.4R1, QFX5100, QFX5110, and QFX5200 switches support MoFRR, which minimizes multicast packet loss in PIM domains when there are link failures. With MoFRR enabled, the switch maintains both a primary and a backup multicast packet stream toward the multicast source, accepting traffic received on the primary path and dropping traffic received on the backup path. Upon primary path failure, the backup path becomes the primary path and quickly takes over forwarding the multicast traffic. If alternative paths are available, a new backup path is created. When enabling MoFRR, you can optionally configure a policy for the (S,G) entries to which MoFRR should apply; otherwise MoFRR applies to all multicast (S,G) streams.

[See [Understanding Multicast-Only Fast Reroute on Switches.](#)]

- **Support for rpf-selection statement for PIM protocol at global instance level (QFX Series)**—Starting in Junos OS 17.4R1, the **rpf-selection** statement for the PIM protocol is available at global instance level. You can configure **group** and **source** statements at the **[edit protocols pim rpf-selection]** hierarchy level.

MPLS

- **Support for BGP MPLS-based Ethernet VPN (QFX10000 Series switches)**—Starting with Junos OS Release 17.4R1, you can use MPLS-based Ethernet VPN (EVPN) to route MAC addresses using BGP over an MPLS core network. An EVPN enables you to connect dispersed customer sites by using a Layer 2 virtual bridge. As with other types of VPNs, an EVPN consists of a customer edge (CE) device (host, router, or switch) connected to a provider edge (PE) switch. The QFX10000 acts as a PE switch at the edge of the MPLS infrastructure. The switch can be connected by an MPLS Label Switched Path (LSP) which provides the benefits of MPLS technology, such as fast reroute and resiliency. You can deploy multiple EVPNs within a service provider network, each providing network connectivity to a customer while ensuring that the traffic sharing on that network remains private.

[See [EVPN Overview.](#)]

- **Support for static adjacency segment identifier for ISIS (QFX Series)**—Starting with Junos OS Release 17.4R1, you can configure static adjacency segment ID (SID) labels for an interface. You can configure two IPv4 adjacency SIDs (protected and unprotected), IPv6 adjacency SIDs (protected and unprotected) per level per interface. You can use the same adjacent SID for multiple interfaces by grouping a set of interfaces under an interface-group and configuring the adjacency-segment for that interface-group. For static adjacency SIDs, the labels are picked from either a static reserved label pool or from segment routing global block (SRGB).

[See [Static Adjacency Segment Identifier for ISIS.](#)]

- **Support for static adjacency segment identifier for aggregate Ethernet member links (QFX Series)**—Starting with Junos OS Release 17.4R1, you can configure a transit single-hop static label switched path (LSP) for a specific member link of an aggregate Ethernet (AE) interface. A static labeled route is added with next-hop pointing to the AE member link of an aggregate interface. Label for these routes is picked from the segment routing local block (SRLB) pool of the configured static label range. This feature is supported for AE interfaces only.

A new **member-interface** CLI command is added under **[edit protocols mpls static-label-switched-path lsp-name transit]** hierarchy to configure the AE member interface name. The static LSP label is configured from a defined static label range.

[See [Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP.](#)]

- **Support for PCEP (QFX5100, QFX5110, QFX5200 switches)**—Starting with Junos OS Release 17.4R1, MPLS RSVP-TE functionality was extended to provide a partial client-side implementation of the stateful Path Computation Element (PCE) architecture (draft-ietf-pce-stateful-pce). The PCE computes path for the traffic engineered LSPs (TE LSPs) of ingress routers that are configured for external control. The

ingress router that connects to a PCE is called a Path Computation Client (PCC). The PCC is configured with the Path Computation Client Protocol (PCEP) (defined in RFC 5440, but limited to the functionality supported on a stateful PCE only) to facilitate external path computing by a PCE. In this new functionality, the active stateful PCE sets parameters for the PCC's TE LSPs, such as bandwidth, path (ERO), and priority.

[See [PCEP Overview](#).]

- **Support for Flap and MBB counter for LSP (QFX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on master routing engine (RE) only:

- Flap counter-- Counts the number of times a LSP flaps down or up.
- MBB counter— Counts the number of times a LSP incurs MBB.

The **clear mpls lsp counters** command resets the flap and the MBB counter to zero.

- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (QFX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
- **Support for default timeout duration for self-ping on an LSP instance (QFX Series)**—Starting in Junos OS 17.4R1, the default timeout duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also manually configure the self-ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is enabled. The LSP types such as CCC, P2MP, VLAN-based , and non-default instances do not support self-ping . You can configure the **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.
- **Support for label history for MPLS protocol (QFX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entry per label . By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range.

The **clear mpls label history** command clears the label history details.

- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (QFX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage that is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP resigalling events. This is because the LSP is responsive to even minor increase or decrease in the utilization when current reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of say 1G

to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP resigalling events for small LSPs of say 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP resigalling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless it is huge, for example 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value based threshold along with the percentage based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.

Network Management and Monitoring

- **Real-time performance monitoring (RPM) (QFX5100 switches)**—Starting in Junos OS Release 17.4R1-S1, real-time performance monitoring (RPM) on QFX5100 switches enables you to configure active probes to track and monitor traffic across the network and to investigate network problems.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.
- Set thresholds to trigger SNMP traps when values are exceeded.

You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test.

- Determine automatically whether a path exists between a host router or switch and its configured BGP neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

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

Port Security

- **Media Access Control Security (MACsec) support (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 17.4R1-S2, MACsec is supported on all 30 interfaces of the QFX10000-30C-M 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 \(MACsec\)](#).]

Routing Protocols

- **Topology-independent loop-free alternate for IS-IS (QFX Series)**—Starting in Junos OS Release 17.4R1, topology-independent loop-free alternate (TI-LFA) with segment routing provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. You can enable

TI-LFA for IS-IS by configuring the **use-post-convergence-lfa** statement at the [edit protocols isis backup-spf-options] hierarchy level. TI-LFA provides protection against link failure, node failure, and failures of fate-sharing groups.

You can enable the creation of post-convergence backup paths for a given interface by configuring the **post-convergence-lfa** statement at the [edit protocols isis interface *interface-name* level *level*] hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

You can enable **node-protection** and/or **fate-sharing-protection** mode for a given interface at the [edit protocols isis interface *interface-name* level *level* post-convergence-lfa] hierarchy level. To use a particular fate-sharing group as a constraint for the fate-sharing-aware post-convergence path, you need to configure the **use-for-post-convergence-lfa** statement at the [edit routing-options fate-sharing group *group-name*] hierarchy level.

[See [Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS](#).]

- **Support for EBGp route server (QFX Series)**—Starting in Junos OS Release 17.4R1, BGP feature is enhanced to support EBGp route server functionality. A BGP route server is the external BGP (EBGP) equivalent of an internal IBGP (IBGP) route reflector that simplifies the number of direct point-to-point EBGp sessions required in a network. EBGp route server propagates unmodified BGP routing information between external BGP peers to facilitate high scale exchange of routes in peering points such as Internet Exchange Points (IXPs). When BGP is configured as a route server, EBGp routes are propagated between peers unmodified, with full attribute transparency (NEXT_HOP, AS_PATH, MULTI_EXIT_DISC, AIGP, and Communities).

The BGP JET **bgp_route_service.proto** API has been enhanced to support route server functionality as follows:

- Program the EBGp route server.
- Inject routes to the specific route server RIB for selectively advertising it to the client groups in client-specific RIBs.

The BGP JET **bgp_route_service.proto** API includes a peer-type object that identifies individual routes as either EBGp or IBGP (default).

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

- **Support for BGP advertising aggregate bandwidth across external BGP links for load balancing (QFX Series)**—Starting in Junos OS Release 17.4R1, BGP uses a new link bandwidth extended community, **aggregate-bandwidth**, to advertise aggregated bandwidth of multipath routes across external links. BGP calculates the aggregate of multipaths that have unequal bandwidth allocation and advertises the aggregated bandwidth to external BGP peers. A threshold to the aggregate bandwidth can be configured to restrict the bandwidth usage of a BGP group. In earlier Junos OS releases, a BGP speaker receiving multipaths from its internal peers advertised the link bandwidth associated with the active route. To advertise aggregated bandwidth of multipath routes and to set a maximum threshold, configure a policy with **aggregate-bandwidth** and **limit bandwidth** actions at the [edit policy-options policy-statement *name* then] hierarchy level.

See [\[Advertising Aggregate Bandwidth Across External BGP Links for Load Balancing Overview\]](#).

Services Applications

- **Support for IPFIX templates for flow aggregation (QFX10008 and QFX10016)**—Starting with Junos OS Release 17.4R1, 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.

This feature was previously introduced on QFX10002 switches in Junos OS Release 17.2R1.

[See [Configuring Flow Aggregation to Use IPFIX Flow Templates.](#)]

Software Installation and Upgrade

- **Support for personality files (QFX5100 switches)**—Starting in Junos OS Release 17.4R1, when a switch in a data center network goes down because of a hardware failure, replacing that switch can be time-consuming and error-prone, because you have to ensure that the crucial elements that you had running on the downed switch are exactly replicated on the new switch. To save time and to avoid errors in configuration and state when you replace a switch, create a “personality” file for your current switch while the switch is still up and save that personality file on a remote server. The “personality” of a switch could include (but is not limited to) its running configuration, SNMP indices, and installed scripts and packages. If the current switch goes down, retrieve the personality file from the server, install it on a new switch, and then bring that new switch online in place of the downed switch.

[See [Personality File for Easy Switch Replacement.](#)]

Virtual Chassis

- **Virtual Chassis support (QFX5200 switches)**—Starting in Junos OS Release 17.4R1, QFX5200 switches can be interconnected into a Virtual Chassis as one logical device managed as a single chassis. A QFX5200 Virtual Chassis can contain up to 3 members that must be QFX5200-32C switches (no mixed mode support). Any non-channelized 100-Gbps QSFP28 ports or 40-Gbps QSFP+ ports can be configured as Virtual Chassis ports (VCPs) to interconnect member switches. Configuration and operation are the same as for other QFX Series Virtual Chassis.

[See [Understanding QFX Series Virtual Chassis.](#)]

SEE ALSO

[Changes in Behavior and Syntax | 266](#)

[Known Behavior | 272](#)

[Known Issues | 276](#)

[Resolved Issues | 280](#)

[Documentation Updates | 291](#)

[Migration, Upgrade, and Downgrade Instructions | 292](#)

[Product Compatibility | 305](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 267](#)
- [EVPNs | 267](#)
- [General Routing | 267](#)
- [Management | 267](#)
- [MPLS | 267](#)
- [Network Management and Monitoring | 269](#)
- [Routing Policy and Firewall Filters | 270](#)
- [Security | 270](#)
- [Software Licensing | 270](#)
- [Virtual Chassis | 270](#)

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 17.4R2 for the QFX Series.

Class of Service (CoS)

- When you configure a **transmit-rate**, you must also configure a **guaranteed-rate** under **traffic-control-profiles**. If you commit a configuration of a **transmit-rate** without a **guaranteed-rate**, a warning message is displayed and the default scheduler map is applied.

EVPNs

- **Change to the show vlans evpn command (QFX5100 switches)**—Starting with Junos OS Release 17.4R2, the **show vlans evpn** command is replaced by the **show ethernet-switching evpn** command.

General Routing

- **Change in default value for port ID TLV for QFX5200 switches**—In Junos OS Release 17.4R1, for QFX5200 switches, the default value used for port ID TLV in LLDP messages is interface name, not SNMP index.

Management

- **Changes to Junos OS YANG module naming conventions (QFX Series)**—Starting in Junos OS Release 17.4R1, the native Junos OS YANG modules use a new naming convention for the module's name, filename, and namespace. The module name and filename include the device family and the area of the configuration or command hierarchy to which the schema in the module belongs. In addition, the module filename includes a revision date. The module namespace is simplified to include the device family, the module type, and an identifier that is unique to each module and that differentiates the namespace of the module from that of other modules.

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

MPLS

- **Support for Flap and MBB counter for LSP (QFX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on the master routing engine (RE) only:
 - Flap counter-- Counts the number of times a LSP flaps down or up.
 - MBB counter— Counts the number of times a LSP incurs MBB.

The **clear mpls lsp counters** command resets the flap and the MBB counter to zero.

- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (QFX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO)

for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.

- **Support for default timeout duration for self-ping on an LSP instance (QFX Series)**—Starting in Junos OS 17.4R1, the default timeout duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also manually configure the self-ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is enabled. The LSP types such as CCC, P2MP, VLAN-based , and non-default instances do not support self-ping . You can configure the **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.
- **Support for label history for MPLS protocol (QFX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at the **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entries per label . By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range. The **clear mpls label history** command clears the label history details.
- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (QFX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP ressignaling events. This is because the LSP is responsive to even minor increases or decreases in the utilization when current reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of around 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP ressignalling events for small LSPs of around 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP ressignaling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless they are huge, for example, 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value-based threshold along with the percentage-based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at the **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.
- When the **no-propagate-ttl** statement is configured on a QFX5200 switch in an MPLS network, the TTL value is not is not copied and decremented on the transit devices during a swap operation. When the switch acts as an ingress device for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the switch acts as the penultimate provider switch, it pops the MPLS header without writing the MPLS TTL into the IP packet. PR1368417

Network Management and Monitoring

- **Change in default log level setting (QFX Series)**—In Junos OS Release, 17.4R1, the following changes were made in default logging levels:

Before this change:

- SNMP_TRAP_LINK_UP was LOG_INFO for both the physical (IFD) and logical (IFL) interfaces.
- SNMP_TRAP_LINK_DOWN was LOG_WARNING for both the physical (IFD) and logical (IFL) interfaces.

After this change:

- IFD LinkUp -> LOG_NOTICE (because this is an important message but less frequent)
- IFL LinkUp -> LOG_INFO (no change)
- IFD and IFL LinkDown -> LOG_WARNING (no change)

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

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a non-default routing instance with a non-default logical system (QFX Series)**—Starting in Junos OS Release 17.4R2, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

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

- **SNMP syslog messages changed (QFX Series)**—In Junos OS Release 17.4R1, 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 Policy and Firewall Filters

- **Support for configuring the GTP-TEID field for GTP traffic (QFX5000 line of switches)**—Starting in Junos OS Release 17.3R3 and 17.4R2, the **gtp-tunnel-endpoint-identifier** statement is supported to configure the hash calculation of IPv4 or IPv6 packets that are included in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations. The **gtp-tunnel-endpoint-identifier** configuration statement is configured at the **[edit forwarding-options enhanced-hash-key family inet]** hierarchy level.

In most of the cases, configuring **gtp-tunnel-endpoint-identifier** statement is sufficient for enabling GTP hashing. After enabling, if GTP hashing does not work, it is recommended to capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In this cases, use **gtp-header-offset** statement to set a proper offset value. Once the header offset value is resolved, run **gtp-tunnel-endpoint-identifier** command for enabling GTP hashing successfully.

[See [gtp-tunnel-endpoint-identifier](#) and [gtp-header-offset](#).]

Security

- **Support to log the SSH key changes**—Starting with Junos OS 17.4R1, the configuration statement **log-key-changes** is introduced at the **[edit system services ssh]** hierarchy level. When the **log-key-changes** configuration statement is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time the **log-key-changes** configuration statement was enabled. If the **log-key-changes** configuration statement was never enabled, then Junos OS logs all the authorized SSH keys.

Software Licensing

- **Key generator adds one day to make the duration of license show as 365 days (QFX Series)**—Starting in Junos OS Release 17.4R1, the duration of subscription licenses as generated by the **show system license** command and shown in the output is correct to the numbers of days. Before this fix, for example, for a 1-year subscription license, the duration was generated as 364 days. After the fix, the duration of the 1-year subscription now shows as 365 days.

[See [show system license](#).]

Virtual Chassis

- **Adaptive load balancing (ALB) feature (Virtual Chassis Fabric)**—Starting in Junos OS Release 17.4R1, the adaptive load balancing (ALB) feature for Virtual Chassis Fabric (VCF) is being deprecated to avoid potential VCF instability. The **fabric-load-balance** configuration statement in the **[edit forwarding-options**

enhanced-hash-key] hierarchy is no longer available to enable and configure ALB in a VCF. When upgrading a VCF to a Junos OS release where ALB is deprecated, if the configuration has ALB enabled, you should delete the **fabric-load-balance** configuration item before initiating the upgrade.

[See [Understanding Traffic Flow Through a Virtual Chassis Fabric](#) and [fabric-load-balance](#).]

- **New configuration option to disable automatic Virtual Chassis port conversion (QFX5100 Virtual Chassis)**—Starting in Junos OS Release 17.4R2, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in a QFX5100 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:
 - LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
 - The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
 - The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

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

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

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

SEE ALSO

[New and Changed Features | 253](#)

[Known Behavior | 272](#)

[Known Issues | 276](#)

[Resolved Issues | 280](#)

[Documentation Updates | 291](#)

[Migration, Upgrade, and Downgrade Instructions | 292](#)

[Product Compatibility | 305](#)

Known Behavior

IN THIS SECTION

- [Class of Service \(CoS\) | 272](#)
- [EVPN | 272](#)
- [Interfaces and Chassis | 274](#)
- [Layer 2 Features | 274](#)
- [MPLS | 274](#)
- [Routing Protocols | 274](#)
- [Platform and Infrastructure | 275](#)
- [Virtual Chassis | 275](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for the QFX Series.

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

Class of Service (CoS)

- With pechip version 1.1, if dot1p rewrites are configured on an interface, then packets that are not matching to a rewrite rule will not retain their previous value. Set the rewrite rule value to 0. This functionality is fixed in pechip version 2.0 [PR1294471](#)

EVPN

- A provider edge (PE) device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE device. The IGP instance running in the VRF on the PE might be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- A QFX10000 switch running Junos OS Release 17.4R1 or later might experience a small and continuous traffic loss under the following conditions:
 - The switch is configured as a Layer 2, Layer 3 or both VXLAN gateway in an EVPN-VXLAN topology with either a two-layer or collapsed IP fabric.

- The switch has default ARP and MAC aging timer values.

Under these conditions, the following types of traffic flows might be impacted:

- Bidirectional Layer 3 traffic in a multihomed topology.
- Unidirectional Layer 3 traffic in a single-homed topology.

Note that this issue does not impact bidirectional Layer 3 traffic in a single-homed topology.

To prevent loss in these traffic flows, you must set the **aging-timer** configuration statement in the **[edit system arp]** hierarchy level so that the value is less than the value of the **global-mac-table-aging-time** configuration statement in the **[edit protocols l2-learning]** hierarchy level. [PR1309444](#)

- Even though an ARP route is learned locally, the **show arp** command output on the provider edge (PE) device on which the route was learned might display the route as **permanent remote**. In Junos OS releases earlier than Junos OS Release 17.4R1, *permanent remote* means that the ARP route was learned from a remote PE device such as an EVPN Type 2 route (MAC+IP route).

This issue might occur under the following conditions:

- A customer edge (CE) device is multihomed to QFX10000 switches in an EVPN-VXLAN topology with a two-layer IP fabric or collapsed IP fabric.
- The QFX switches function as Layer 3 only, or Layer 2 and Layer 3 PE devices.
- The QFX switches run Junos OS Release 17.4R1 or later.

To work around this issue, you can view locally learned ARP routes by entering the **show evpn database origin local** command on the PE devices. [PR1324824](#)

Interfaces and Chassis

- Configuring link aggregation group (LAG) hashing with the **edit forwarding-options enhanced-hash-key inet vlan-id** statement uses the VLAN ID in the hashing algorithm calculation. On some switching platforms, when this option is configured for a LAG that spans FPCs, such as in a Virtual Chassis or Virtual Chassis Fabric (VCF), packets are dropped due to an issue with using an incorrect VLAN ID in the hashing algorithm. As a result, the **vlan-id** hashing option is not supported in a Virtual Chassis or VCF containing any of the following switches as members: EX4300, EX4600, QFX5100, or QFX5110 switches. Under these conditions, use any of the other supported **enhanced-hash-key** hashing configuration options instead. [PR1293920](#)

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging has been enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. [PR1075230](#)

MPLS

- Layer 2 circuits on aggregated Ethernet interfaces are not supported on QFX5100, QFX5110, and QFX5200 switches. [PR1333730](#)
- On QFX5100, QFX5110, QFX5200 switches with Layer 2 circuit configured on the PE switches, enabling VLAN bridge encapsulation on a CE interface drops packets if flexible Ethernet services and VLAN CCC encapsulation are configured on the same logical interface. You can configure only one encapsulation type, either **set interfaces xe-0/0/18 encapsulation flexible-ethernet-services** or **set interfaces xe-0/0/18 encapsulation vlan-ccc**. [PR1329451](#)

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](#)
- For the QFX10002 and QFX10008 switches, you might observe an increase in the convergence time of OSPF routes when compared to Junos OS Release 17.3. An average increase of 1.5 seconds is seen for 100,000 OSPFv3 routes. [PR1297541](#)
- A QFX10000 switch running Junos OS Release 17.3Rx or 17.4Rx software might experience a small and continuous traffic loss under the following conditions: 1) The switch is configured as a Layer 2, Layer 3 or both VXLAN gateway in an EVPN-VXLAN topology with either a two-layer or collapsed IP fabric. 2) The switch has default ARP and MAC aging timer values. Under these conditions, the following types of traffic flows might be impacted: 1) Bidirectional Layer 3 traffic in a multihomed topology, and 2)

Unidirectional Layer 3 traffic in a single-homed topology. Note that this issue does not impact bidirectional Layer 3 traffic in a single-homed topology. [PR1309444](#)

Platform and Infrastructure

- On a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On a QFX5110-32C switch, if a splitter cable is connected to a peer end device capable of 10G CV/MX card, ports will not come up due to varied pre-empt settings for the splitter and DAC cables. There is a hardware limitation where we have no way in EEPROM to differentiate between splitter and DAC cable to apply different settings. As a workaround, use manual channelisation on the QFX5110-32C side. [PR1280593](#)
- ERPS convergence takes time after a GRES switchover and hence traffic loss is observed for a brief period. [PR1290161](#)
- On QFX Series, the logical interface (IFD) and the physical interface (IFL) go down when traffic exceeds the ratelimit. Storm control is supported only on interfaces configured in family Ethernet-switching. Moreover, in this family, only one IFL is supported per IFD. Thus, bringing down the IFD is acceptable. Flexible VLAN tagging is not supported on the interfaces enabled for storm control. [PR1295523](#)
- Traffic drop occurs when sending Layer 3 traffic across an MPLS LSP. [PR1311977](#)
- Traffic drop occurs when sending traffic over "et" interfaces due to CRC errors. [PR1313977](#)
- On Junos OS Automation Enhancement images there is a way to use the Python interpreter in interactive mode. When Python interpreter is used in an interactive mode on a shell, the prompt does not seem to return immediately. This is an example of a session: -- % python Python 2.7.8 (default, Nov 10 2017, 01:45:13) [GCC 4.2.1 (for JUNOS)] on junos Type "help", "copyright", "credits" or "license" for more information. >>> >>> print "hello" >>> hello -----> waiting here, hit 'enter' here to return the python prompt >>> quit() >>> % -- The regular script is not impacted. [PR1324124](#)

Virtual Chassis

- Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on an QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur and its considered to be known behavior. Release note reference: https://www.juniper.net/documentation/en_US/junos/information-products/topic-collections/release-notes/17.2/topic-118735.html [PR1347902](#)

SEE ALSO

Changes in Behavior and Syntax	266
Known Issues	276
Resolved Issues	280
Documentation Updates	291
Migration, Upgrade, and Downgrade Instructions	292
Product Compatibility	305

Known Issues

IN THIS SECTION

- [EVPN | 276](#)
- [Layer 2 Features | 277](#)
- [MPLS | 277](#)
- [Platform and Infrastructure | 277](#)
- [Routing Protocols | 279](#)

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.4R2.

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

EVPN

- In a scaled setup, if MAC move is triggered more than 4 times, the MAC move detection might not be reliable. [PR1284315](#)
- CNH (chained-composite-next-hop) is must for EVPN pure type 5 with VXLAN encapsulation. Without this Packet Forwarding Engine wouldn't program the tunnel next hop. You have to explicit set it on QFX5110. set routing-options forwarding-table chained-composite-next-hop ingress evpn. QFX10000 it is applied as part of default configuration. **user@router> show configuration routing-options forwarding-table | display inheritance defaults.** [PR1303246](#)
- In an EVPN collapsed Layer 2 or Layer 3 multihomed gateway topology, when traffic is sent from an IP fabric towards EVPN, some traffic loss is seen. If the number of hosts behind the EVPN gateways is increased, the traffic loss becomes higher. This issue is seen with the QFX10000. [PR1311773](#)

- ARP gets deleted and relearned during the first ARP refresh with EVPN-VXLAN multihomed CE, so traffic drops and recovers for the first ARP refresh. [PR1327062](#)
- On QFX platforms (QFX5K/10K), VTEP's MAC address is not learned in the ethernet switching table though they are present in the EVPN database. [PR1371995](#)

Layer 2 Features

- When an FPC encounters a memory exhaustion condition, the FPC restarts unexpectedly with the **PPMAN: failed decoding IDL msg - retval -2 type 5 encode_len 208 length 208 data 0x344ff1b0** message. [PR1321117](#)

MPLS

- LDP to BGP stitching with an eBGP indirect next hop having an implicit null label does not work. It does work when BGP indirect next hop has a real label. As a workaround, perform the following: (1) Ensure the peer advertises a real label by adding another router between the egress and ingress PE devices. (2) Use IBGP, which gets resolved over LDP or RSVP-TE LSPs. This will ensure that the BGP indirect next hop has a real label. [PR1254702](#)
- On optimize timer expiry, when the ted version number match indicates a CSPF has already run for the path, if an optimization has not yet been done with that version, it will be run despite the version number match. (Having a per path optimize-seq-no that is updated with ted seq no only on optimization.) When path-cc-updated is false and CSPF fails for optimization, disable the path just like we do for the ones on avoid colors/invalid ERO, making sure this does not interfere with global repair/local reversion. [PR1365653](#)

Platform and Infrastructure

- While scaling more than 2000 VLAN or IRBs, Layer 3 multicast traffic does not converge to 100 percent and continuous drops are observed after bringing down or up the downstream interface or while an FPC comes online after an FPC restart. [PR1161485](#)
- When per-packet load balancing is removed or deleted, the next-hop index might change. [PR1198092](#)
- On PTX platforms with FPC3, PTX1000 with build-in chassis and QFX10000 platforms, a Flexible PIC Concentrator (FPC) major alarm might be seen if the system detects parity error, and the error messages **DLU: ilp memory cache error** and **DLU: ilp prot1 detected_imem_even error** might appear. The alarm might be cleared without intervention. This error may also be accompanied by traffic loss. [PR1251154](#)
- Single-bit and multiple-bit ECC errors are not logged on QFX5110 switches. [PR1251917](#)
- On the QFX10000-12C-DWDM coherent line card, it is possible that sometimes the link flaps when MACsec is enabled on Ethernet interfaces. [PR1253703](#)

- The management process (daemon) might crash if the Openconfig package is installed immediately or within minutes of Network Agent package installation. This is a transient issue and will not impact any functionality. There is no action needed from the user side in response to the crash. As a workaround, install Openconfig before installing Network Agent. [PR1265815](#)
- On QFX5100 switches, static LAG link protection switchover/revert is not working consistently. [PR1286471](#)
- When link protection with the backup port state "down" and LACP are configured, if backup state "down" is removed from the configuration, both ports should be up and the primary port should pass all egress traffic. In some instances, however, traffic might pass through the backup port instead of the primary port. [PR1297597](#)
- Traffic drop occurs on sending traffic over "et" interfaces due to CRC errors. [PR1313977](#)
- Family Ethernet-switching cannot be used when **flexible-vlan-tagging** is configured. It is unsupported. The behavior is non-deterministic with this configuration and there is a possibility of seeing a dcpfe core file. [PR1316236](#)
- Port 0 of Qfx5100-48t does not come up in a mixed VCF. As a workaround, use the **phy diag xe0 dsc** command as of now from the BCM shell upon reboot, which brings up the port and stays up continuously until the next reboot. [PR1323323](#)
- The management process (mgd) might panic after modifying aggregated Ethernet interface members under the **ethernet-switching vlan** stanza. After mgd panic, your remote session is terminated as a result. [PR1325736](#)
- In Streaming Telemetry scenario, if performing "commit full", na-grpd daemon might restart causing disconnection of streaming telemetry. [PR1326366](#)
- On QFX5100 Series platforms, in some cases, the CoS (class of Service) configuration is not properly applied in the Packet Forwarding Engine, leading to an unexpected egress traffic drop on some interfaces. [PR1329141](#)
- On QFX52xx standalone devices with Vxlan configured, user configured Ingress ACL scale limit is 256 terms. [PR1331730](#)
- On QFX5110, the FEC for 100g optics is not being displayed when expected behavior is for FEC to be shown as NONE. On QFX10002 Elit, the FEC for 40g optics is being displayed as NONE when expected behavior is for FEC not to be displayed. On QFX10008 Ultimat, the FEC for 40g optics is being displayed as NONE when expected behavior is for FEC not to be displayed. [PR1360948](#)
- When MCLAG is configured with Force-Up enabled on MCLAG Nodes, LACP admin key should not match with Access/CE device. [PR1362346](#)
- On QFX10000 platform with IRB enabled, traffic might not be forwarded on some of the child members when the member link of the aggregated Ethernet is added or deleted. [PR1362653](#)

Routing Protocols

- For single-hop eBGP session, upon interface down event, do not do GR helper logic. In problem state Peer: 8.3.0.2 AS 100 Local: 8.3.0.1 AS 101 Group: EBGP Routing-Instance: master Forwarding routing-instance: master Type: External State: Active Flags: <> Last State: Idle Last Event: Start Last Error: Cease Import: [reject] Options: <Preference PeerAS LocalAS Refresh> Holdtime: 90 Preference: 170 Local AS: 101 Local System AS: 0 Number of flaps: 2 Last flap event: Stop Error: 'Cease' Sent: 1 Recv: 0 NLRI we are holding stale routes for: inet-unicast Time until stale routes are deleted or become long-lived stale: 00:01:54 >>>>>>>> Time until end-of-rib is assumed for stale routes: 00:04:54 Table inet.0 RIB State: BGP restart is complete Send state: not advertising Active prefixes: 14 Received prefixes: 21 Accepted prefixes: 15 Suppressed due to damping: 0 Stale prefixes: 21 >>>>>>>>>>>>>>>> With the fix: Peer: 8.3.0.2 AS 100 Local: 8.3.0.1 AS 101 Group: EBGP Routing-Instance: master Forwarding routing-instance: master Type: External State: Active Flags: <> Last State: Idle Last Event: Start Last Error: Cease Import: [reject] Options: <Preference PeerAS LocalAS Refresh> Holdtime: 90 Preference: 170 Local AS: 101 Local System AS: 0 Number of flaps: 1 Last flap event: Stop Error: 'Cease' Sent: 1 Recv: 0 [PR1129271](#)
- On QFX10000 line platforms, during a route next-hop churn or an earliest deadline first (EDF) job priority changes, memory corruption might occur, leading to processing issues and constant packet drop. [PR1243724](#)
- We strongly recommend using BGP as the protocol for configuring the local-address for each multihop iBGP/eBGP peer configuration. We recommend that local-address be a routeable lo0 address. Using loopback address reduces dependency with interfaces. Note: Multihop is by default enabled for iBGP peers. [PR1323557](#)
- On a scaled setup, when the host table is full and the host entries are installed in an LPM table, OSPF sessions might take more time to come up. [PR1358289](#)

SEE ALSO

[New and Changed Features | 253](#)

[Changes in Behavior and Syntax | 266](#)

[Known Behavior | 272](#)

[Resolved Issues | 280](#)

[Documentation Updates | 291](#)

[Migration, Upgrade, and Downgrade Instructions | 292](#)

[Product Compatibility | 305](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R2 | 280](#)
- [Resolved Issues: 17.4R1 | 287](#)

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

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

Resolved Issues: 17.4R2

Class of Service (CoS)

- You cannot filter packets with DstIP as 224/4 and DST MAC = QFX_intf_mac on a loopback interface using a single match condition for source address 224.0.0.0/4. [PR1354377](#)

EVPN

- Next hop installation error messages are seen on QFX10000 line switches. [PR1258930](#)
- EVPN-VXLAN QFX10000: `jprds_dlu_alpha_add : 222 JPRDS_DLU_ALPHA KHT` addition failed. [PR1258933](#)
- VXLAN-EVPN: IPv6 packet loss after a normal traffic run rate. [PR1267830](#)
- Subinterfaces from the same physical port do not work if configured under the same VXLAN VLAN. [PR1278761](#)
- For a VLAN with an IRB interface as the routing interface, set the `vlan-id` parameter to "none" to ensure proper traffic routing. [PR1287557](#)
- QFX10000 VXLAN with MPLS underlay traffic loss is seen at the RSVP egress. [PR1289666](#)
- VXLAN traffic loss is observed after deleting and adding VLANs. [PR1318045](#)
- A core link flap might result in an inconsistent global MAC count. [PR1328956](#)
- The partial multicast traffic might be dropped in an EVPN-VXLAN multi homing scenario with non-default `virtual-switch/evpn routing-instance` configured. [PR1334408](#)
- The MAC movement between remote VTEP and local VTEP might cause traffic to be transmitted incorrectly in an EVPN-VXLAN scenario. [PR1335431](#)

- Configuring **encapsulate-inner-vlan** on the partial VXLANs might cause traffic impact. [PR1337953](#)
- In an EVPN-VXLAN environment, BFD flaps cause VTEP flaps and cause the Packet Forwarding Engine to crash. [PR1339084](#)
- Rpd has unreproducible cored with scaling EVPN-VXLAN configuration on QFX10K platform. [PR1339979](#)
- The rpd core might be seen if deleting the default switch in an EVPN-VXLAN environment. [PR1342351](#)
- In an EVPN-VXLAN scenario, the traffic might get dropped as the core-facing interfaces goes down. [PR1343515](#)
- Traffic might be lost on a Layer 2 and Layer 3 spine node in a multihome EVPN scenario. [PR1355165](#)
- The QFX10000 might drop transited traffic coming from MPLS network to VXLAN/EVPN. [PR1360159](#)
- Increased risk of a routing crash with temporary impact on traffic on QFX10000 or QFX5100 nodes with certain configuration changes or clearing L2 or L3 learning information in a high-scale EVPN-VXLAN configuration environment. [PR1365257](#)
- Proxy ARP may not work as expected in an EVPN environment. [PR1368911](#)
- QFX10k / Import default ipv6 route to VRF causes infinite entries to get created in 'evpn ip-prefix-database' and become unstable. [PR1369166](#)

High Availability (HA) and Resiliency

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

Infrastructure

- QFX5100: Enabling mac-move-limit stops ping on **flexible-vlan-tagging** enabled interface. [PR1357742](#)

Interfaces and Chassis

- Multicast data packets are looping in MC-LAG. [PR1281646](#)
- Upgrading might encounter a commit failure if **redundancy-group-id-list** is not configured under ICCP. [PR1311009](#)
- CVLANs range is 16, which might not pass traffic in a Q-in-Q scenario. [PR1345994](#)
- MC-LAG peer doesn't send ARP request to the host. [PR1360216](#)

Layer 2 Ethernet Services

- A jdhcpd core file is generated after making DHCP configuration changes. [PR1324800](#)

Layer 2 Features

- Device transmits packets that exceed the interface MTU. [PR1306724](#)
- NLB heartbeat packets might be dropped on a QFX10000. [PR1322183](#)
- ARP entry might be learned on STP blocking ports. [PR1324245](#)

- The DHCP discover packets might be looped in an MC-LAG and a DHCP-relay scenario. [PR1325425](#)
- QFX5100: With multiple logical units configured on an interface, **input-vlan-map POP** is not removing outer VLAN-tag when Q-in-Q and VXLAN are involved. [PR1331722](#)
- The operation of pushing a VLAN tag does not work for VXLAN local switching tunneled Q-in-Q traffic. [PR1332346](#)
- Interface with **flexible-vlan-tagging** and **family ethernet-switching** does not work on a QFX10000. [PR1337311](#)

MPLS

- QFX5100: ISSU is not supported with an MPLS configuration. [PR1264786](#)
- Traffic drop during a NSR switchover for RSVP P2MP provider tunnels used by MVPN . [PR1293014](#)
- MPLS forwarding might not happen properly for some LSPs. [PR1319379](#)
- The rpd process might crash on backup Routing Engine due to memory exhaustion. [PR1328974](#)
- The hot standby for the L2 circuit does not work on a QFX5000. [PR1329720](#)
- RSVP sessions go down for ingress LSPs with no-cspf enabled. [PR1339916](#)
- LSP is not received by QFX5110. [PR1351055](#)
- NO-propagate-TTL acts on MPLS Swap operation. [PR1366804](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition. [PR1374102](#)

Platform and Infrastructure

- After upgrading the QFX5100 to Junos OS Release 16.1 or later from Junos OS Release 15.1, the commit warning **/boot/ffp.cookie+** might be seen. [PR1283917](#)
- SFP management Ethernet port C0 might not come up. [PR1298876](#)
- Run-time pps statistics value might show zero for a subinterface of the aggregated Ethernet interface. [PR1309485](#)
- Traffic loss might be seen if traffic is sent through the 40G interface. [PR1309613](#)
- Some log messages are seen on the QFX5110 platform when plugging in an SFP-SX. [PR1311279](#)
- One aggregated Ethernet member cannot send out sFlow sample packets. [PR1311559](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- A memory leak is seen for dot1xd. [PR1313578](#)
- Some certain IGMP join packets cannot be processed correctly at a high rate. [PR1314382](#)
- Transit traffic over a GRE tunnel might hit the CPU and trigger a DDoS violation on the L3 next hop. [PR1315773](#)

- On an L2 next-generation switch platform (QFX5100/QFX10000), l2cpd might drop core files repeatedly if an interface is connected to a VoIP product with LLDP and LLDP-MED enabled. [PR1317114](#)
- Packets such as TDLS without an IP header are looped between virtual gateways. [PR1318382](#)
- The optic interface transmits power even after it has been administratively shutdown. [PR1318997](#)
- The packet might be dropped between 4-60 seconds when the master Routing Engine is rebooted in a virtual chassis. [PR1319146](#)
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX-VC ISSU. [PR1320370](#)
- The MAC address is stuck with "DR" flag on spine node even though packets are received on the interface from the source MAC. [PR1320724](#)
- FPCs go offline in some situations. [PR1321198](#)
- On the QFX10016 EVPN-VXLAN scaled testbed, it takes up to 3 minutes for traffic to converge when configured. [PR1323042](#)
- The openflow session cannot be established correctly with controller and interface options configured on QFX5100 switches. [PR1323273](#)
- Update new firmware versions for jfirmware package for 100G-PSM4 and 100G-AOC issues. [PR1323321](#)
- EVPN Type 5: Unicast traffic is getting dropped on the backup forwarder. [PR1323907](#)
- The next hop of _all_ces__ flood details might go missing. [PR1324739](#)
- The GRE traffic is not decapsulated by the firewall filter. [PR1325104](#)
- VLAN or VLAN bridge might not be added or deleted if there is an IFBD HW token limit exhaustion. [PR1325217](#)
- ARP request packets might not be flooded on a QFX5110. [PR1326022](#)
- The major alarm about 'Fan & PSU Airflow direction mismatch' might be seen by removing the management cable. [PR1327561](#)
- Deleting one VXLAN might cause a traffic loop on another VXLAN in a multi homing EVPN-VXLAN scenario with a service provider style interface. [PR1327978](#)
- QFX10002: Major alarm should be cleared once the chassis has more PEM units installed than the **minimum PEM** configuration. [PR1327999](#)
- Directories and files under **/var/db/scripts** lose execution permission or directory 'jet' is missing under **/var/db/scripts** causing **error: Invalid directory: No such file or directory** error during commit. [PR1328570](#)
- FAN tray removal or insertion trap is not generated for a backup FPC. [PR1329031](#)
- The **etherStatsCRCAlignErrors** counters might disappear in the SNMP tree. [PR1329713](#)
- After commit, members of Virtual Chassis or VCF are split and some members might get disconnected. [PR1330132](#)

- An rpd process core file generated on a new backup Routing Engine at `task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler` after disabling NSR+GRES. [PR1330750](#)
- The **out of HMC range** and **HMC READ failed** error messages are seen. [PR1332251](#)
- Traffic does not pass through VCP ports after rebooting the Virtual Chassis members. [PR1332515](#)
- EVPN-VXLAN: DF drops multicast traffic. [PR1333069](#)
- On QFX10K8/QFX10K16 platforms, SIB LEDs on the fan tray are off after the replacement of the Fan Tray Controllers (FTC). [PR1334006](#)
- The DHCPv6 SOLICIT message is dropped. [PR1334680](#)
- AI-script does not get auto re-install upon a JUNOS upgrade on Next Generation-Routing Engine. [PR1337028](#)
- The DF of an EVPN instance might flood all the ARP request back to the Ethernet Segment. [PR1337275](#)
- On QFX5100 platforms, LR4 QSFP can take up to 15 min to come up after Virtual Chassis reboot. [PR1337340](#)
- SNMP jnxBoxDescr OID returns different value when upgrading to Junos OS Release 17.2. [PR1337798](#)
- On the QFX10000 platforms, VRRP function does not work well when it is configured on sub-interfaces. [PR1338256](#)
- The traffic coming from the remote VTEP PE might be dropped. [PR1338532](#)
- The analyzer status might show as down when port mirroring is configured to mirror packets from an aggregated Ethernet member. [PR1338564](#)
- The VXLAN traffic might not be transmitted correctly with an IRB interface as the underlay interface of the VTEP tunnel. [PR1338586](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- Multicast traffic drop is seen if downstream IRB interfaces have snooping enabled. [PR1340003](#)
- On the QFX5200: there is an inconsistent result after using **deactivate xxx** command on 'pfc-priority' and 'no-loss' context. [PR1340012](#)
- L3 traffic is not getting converged properly upon disabling the ECMP link between spine and leaf with EVPN-VXLAN configurations. [PR1343172](#)
- BPDU packets might get dropped and bpdud-block-on-edge might not work. [PR1343330](#)
- Broadcast frames might be modified with the ethertype 0x8850. [PR1343575](#)
- EVPN-VXLAN: VLAN with **flexible-tag** mode , the xe statistics appears to not be updated for ingress. [PR1343746](#)
- LACP packets are getting dropped with native-vlan-id configured after reboot. [PR1361054](#)

- QFX5000 Virtual-Chassis acting as EVPN-VxLAN ARP Proxy might cause ARP resolution to fail. [PR1365699](#)
- Hashing does not work for the IPv6 packet encapsulated in VxLAN scenario. [PR1368258](#)
- When native-vlan-id is configured for aggregated Ethernet LACP session to multihomed server goes down. [PR1369424](#)
- A port might still work if it's deleted from an aggregated Ethernet interface. [PR1372577](#)
- Implement the **edit interfaces interface-name ether-options] configured-flow-control** option for the QFX Series. [PR1343917](#)
- For EVPN-VXLAN, the ARP packet uses VRRP/virtual-gateway MAC in an Ethernet header instead of an IRB MAC address. [PR1344990](#)
- In the QFX5100, fan RPM fluctuates when temperature sensor reaches its threshold. [PR1345181](#)
- FXPC process might generate a core file when removing VXLAN configuration. [PR1345231](#)
- Backup Routing Engine might experience a crash, causing vmcore to be generated on master Routing Engine, master Routing Engine performance will not be affected. [PR1346218](#)
- CPU and memory statistics not populating for the backup switch in a QFX5110 Virtual Chassis. [PR1346268](#)
- An incorrect inner VLAN tag is sent from the QFX10000 platform with Q-in-Q configured on the Layer 3 sub interface. [PR1346371](#)
- Statistics daemon pfed might generate core files on an upgrade between certain releases. [PR1346925](#)
- On QFX5110 switches, a DCPFE core file might be generated after removing Type-5 tunnel in an EVPN-VXLAN configuration. [PR1346980](#)
- A QFX5100-48T 10G interface might be auto negotiated at 100M speed instead of 10G. [PR1347144](#)
- On QFX5110-48S-4C platforms, part numbers and serial numbers are not displayed for any of the 10G optics/DAC connected. [PR1347634](#)
- The ARP might not update and packets might get dropped at the Routing Engine. [PR1348029](#)
- On a QFX5100, a BGP session flaps when changes are made on the extended-vni-list under the EVPN hierarchy and if the BGP neighborship is through an IRB. [PR1349600](#)
- QFX5100 40G port has an interoperability issue with some other vendors. [PR1349664](#)
- Blackholing traffic with destination MAC matching the virtual gateway MAC might be seen. [PR1348659](#)
- The pfed process might consume high CPU if subscriber or interface statistics are used at large scale. [PR1351203](#)
- A DCPFE process might crash on QFX10000 switches. [PR1351503](#)
- The GTP traffic might not be hashed correctly for an aggregated Ethernet interface. [PR1351518](#)

- Telemetry traffic does not leave the local box when telemetry server is reachable via a VR routing-instance. [PR1352593](#)
- QFX5100 arp fail after change interface MAC address. [PR1353241](#)
- RPC output not showing failure when running **request system software add** with software already staged. [PR1353466](#)
- SFP-LX10 on QFX5110 might fail to connect with another device. [PR1353677](#)
- The alarm errors might be seen during the bootup on a QFX10000. [PR1354582](#)
- Untagged packets might not be forwarded through the trunk port. [PR1355338](#)
- Commit error observed if box is downgraded from from 18.2/18.3 release to 17.3R3. [PR1355542](#)
- On QFX5110 platforms, LX10 SFP needs to be reinserted after autonegotiation is enabled or disabled. [PR1355746](#)
- EVPN-VXLAN: the VXLAN traffic might be lost in EVPN type 2 and type 5 scenario. [PR1355773](#)
- "Load averages" output under **show chassis routing-engine** shows "nan" periodically. [PR1356676](#)
- The IGMP membership report packets might not be forwarded over an interface on a QFX10000. [PR1360137](#)
- On QFX10k, virtual-gateway-address should be only configured on a irb interface associated with a vxlan VLAN. [PR1360646](#)
- Unable to create QFX5200 VC w/100G DACs. [PR1360721](#)
- The GTP traffic might not be hashed correctly on aggregated Ethernet interface. [PR1361379](#)
- The **clear services accounting statistics inline-jflow fpc-slot 0** command should be supported in QFX Series. [PR1362396](#)
- QFX5100VC: Unable to connect management address through vme interface. [PR1362437](#)
- On QFX10008, QFX10016, PTX1000, PTX5000, PTX10008, PTX10016 platforms, MPLS exp rewrite might not work for IPV6 and IPV4 traffic. [PR1364391](#)
- Root password recovery process doesn't work. [PR1365740](#)
- On QFX5100/QFX5110/QFX5200 platforms, ISIS adjacency goes down when mtu 9192 is configured. [PR1368913](#)
- On QFX10000 platforms, before the 17.3R3 code, the maximum number of ESI IFLs was 4000 in the Packet Forwarding Engine. [PR1371414](#)
- TPI-50840 BUM traffic received on 5110 is not flooded to all remote vteps. [PR1373093](#)

Routing Protocols

- Observed mcsnospd core file at `__raise,abort,__task_quit__,task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal(enable_slip_detector=true,no_exit=true)` at `../src/junos/lib/libtask/base/task_scheduler.c:275` [PR1305239](#)
- Packet drop is seen when programming for GRE traffic. [PR1308438](#)
- Diffserv bits/ToS bits are not getting copied from Inner IP header to GRE header. [PR1313311](#)
- Some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update. [PR1320723](#)
- On the QFX5100, consistent hashing is not getting programmed. [PR1322299](#)
- IS-IS Layer 2 hello packets are dropped when they come from another vendor's device. [PR1325436](#)
- The loopbacked IRB interface is not accessible to a remote network. [PR1333019](#)
- The dcpfe process crash is seen in a route leak scenario on the QFX10000. [PR1334714](#)
- The rpf-check-policy does not work as expected. [PR1336909](#)
- Ping fails if MTU is different on the interfaces. DF is not working as expected. [PR1345495](#)
- vrf-fallback on QFX5K is not supported in ALPM mode. [PR1345501](#)
- On QFX10000 platforms, Netconf SSH TCP port 830 traffic hitting host path/unclassified queue. [PR1345744](#)
- On QFX5100 platforms, parity errors in L3 IPv4 table in the Packet Forwarding Engine memory might cause traffic black holing. [PR1364657](#)

Software Installation and Upgrade

- Commit may fail in single-user mode. [PR1368986](#)

Virtual Chassis

- QFX-Virtual Chassis: Sometimes, the multicast packets are received 2x 3x times than expected. [PR1306239](#)

Resolved Issues: 17.4R1

Class of Service (CoS)

- On QFX5100 switches, traffic might be dropped when there is more than one forwarding class under `forwarding-class-sets`. [PR1255077](#)
- The transmit rate applied with `forwarding-class-set` does not work properly. [PR1277497](#)

EVPNs

- On QFX5100 switches with EVPN-VXLAN deployed, broadcast and multicast traffic might not be sent to other switches through VTEP interfaces. [PR1293163](#)
- On QFX10000 switches with EVPN deployed, packet corruption is seen with Packet Forward Engine trap code (129) egp.v4_chksum when sending L3 inter-VNI traffic with the underlay vlan-tagging inet interface. [PR1295491](#)
- The dynamic routing protocols might not work correctly over the IRB interface in an EVPN-VXLAN scenario with ECMP. [PR1301521](#)
- QFX5110-48S: L3 VPN traffic is dropped for some instances when EVPN-VXLAN configuration is removed and reapplied. [PR1307590](#)

Hardware

- FEC is disabled by default on 100G-LR optics for QFX5200 switches. [PR1286389](#)
- The 1G copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms. [PR1286709](#)
- ULC-60S-6Q LC on QFX10008: The port becomes unusable after inserting a third-party SFP-T optic. [PR1294394](#)
- Update new firmware versions for jfirmware package for 100G-PSM4 and 100G-AOC issues. [PR1323321](#)

High Availability (HA) and Resiliency

- Normal VRRP MAC is triggering a MAC move, and logical interfaces on the BD are getting shut down. [PR1285749](#)

Infrastructure

- Create new command: "enable-tcp-nodelay" and allow flash sub-jobs to run for max quantum. [PR1136167](#)
- Disabled 10-Gigabit Ethernet interfaces might stay up on QFX10000 line switches. [PR1300775](#)
- The 40-Gigabit Ethernet connection between two QFX5100-24Qs might not come up sometimes. [PR1178799](#)
- QFX10002 and QFX10008: BFD sessions over IRB interfaces with Junos OS Releases 17.1R1, 17.1R2, 17.2R1, and 17.3R1 are centralized. [PR1284743](#)

Interfaces and Chassis

- Random interfaces do not come up after a line card is rebooted. [PR1262839](#)
- Copper ports flap on QFX5100-48T when short-reach-mode is enabled. [PR1248611](#)
- The 40-Gigabit Ethernet interface might flap between QFX5100 and other products. [PR1273861](#)
- QFX10000-12C-DWDM: an ot- interface link flap is observed whenever an optics TCA alarm is raised; however, there is no LOS and no traffic loss is observed. [PR1279351](#)

- On QFX5100 switches, an AE interface might flap upon commit if an explicit speed is configured on an AE member interface [PR1284495](#)
- On QFX10000 line switches, the input and output rates for 10-Gigabit, 40-Gigabit, or 100-Gigabit Ethernet interfaces are not 0 if the interface is down. [PR1291412](#)
- Traffic might not be received on a 1-Gigabit Ethernet interface if autonegotiation is disabled and speed/duplex is configured on both the QFX Series switch and the peer host. [PR1292275](#)
- High heap memory utilization might be seen if multiple SFP-T optics are inserted or **set interface <> link-mode full-duplex** is enabled. [PR1294208](#)
- The 40-Gigabit Ethernet interface might not come up if a specific vendor's DAC cable is used. [PR1296011](#)
- QFX10008/10016: Commit error is seen when configured with mixed speed. [PR1301923](#)

Junos Fusion Satellite Software

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

Layer 2 Features

- To set up PTP BC forwarding on a QFX10002, configure routing on the interface or add a static ARP entry on the remote PTP device. [PR1275327](#)
- Feature swap-swap might not work as expected in a Q-in-Q scenario. [PR1297772](#)
- QFX5100 crashes and the fxcp process generates a core file. [PR1306768](#)

MPLS

- QFX10008 is dropping egress MPLS traffic, if the egress interface is an IRB with access L2 AE interface. [PR1279827](#)

Network Management and Monitoring

- UFT for non-local member is not shown in the CLI. [PR1243758](#)
- LAG interface input bytes counter continuously decreases when no packets come in. [PR1266062](#)
- SNMP process is not running on QFX Series switches with incorrect source addresses. [PR1285198](#)
- On QFX5100, an incorrect alarm type might be displayed. [PR1291622](#)
- Previous learned MAC address from remote ESI cannot be changed to local. [PR1303202](#)
- The sflow records are missing "extendedType ROUTER" fields as well as an outbound interface for traffic that is using BGP multipath. [PR1303236](#)
- QFX5110-48S: digital optical monitoring statistics cannot be received through the CLI in Junos OS Releases 15.1X53 through 17.x. [PR1305506](#)

Platform and Infrastructure

- A hostname synchronization issue occurs between the Junos OS VM instance and the Linux host in TVP platforms. [PR1283710](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- The dcpfe process might crash and restart on MC-LAG active and standby nodes when there is ARP/NDP next-hop change. [PR1299112](#)
- OSPFv3 authentication using IPsec SA does not work if you are using IPsec to authenticate OSPFv3 neighbors on some QFX Series platforms. [PR1301428](#)

Port Security

- On QFX10000 switches, MACsec sessions are not coming up on a Layer 3 logical interface. [PR1282995](#)
- Proxy-ARP and ARP suppression are not yet supported for the QFX10000 line. [PR1293707](#)

Routing Protocols

- When the static link protection mode configured backup state is down, the primary port goes to down state instead of the secondary port, and the secondary remains in up state. [PR1276156](#)
- Analytics JSON data format is reporting a incorrect value for 'rxbps' counter. [PR1285434](#)
- On QFX5100 switches, if a term with the policer action is configured, **dc-pfe: list_destroy()** messages might be displayed on commit. [PR1286209](#)
- GRE tunnel traffic does not switch over to the alternate path if the primary path to the tunnel destination changes. [PR1287249](#)
- UDP traffic with destination port 520 and 521 is discarded on QFX5110 switches after a Junos OS upgrade. [PR1287271](#)
- OVSDDB and Openflow have some limitations on QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016 switches running Junos OS Releases 17.1R1, 17.1R2, and 17.2R1. [PR1288227](#)
- Storm-control flags are not set after a Routing Engine switchover. [PR1290246](#)
- In a data center environment with EVPN-VXLAN and proxy MAC plus IP advertisement enabled on a Layer 3 gateway, the state for some MACs might be lost during MAC moves. [PR1291118](#)
- QFX5110-32C: Routable ICMP packets get flooded on one of the newly provisioned 100 VXLAN IRB interfaces on a non-collapsed VXLAN L3 gateway (same IP, same MAC profile). [PR1291406](#)
- The dcpfe process might crash after a period of idle time on QFX10000 switches. [PR1294055](#)

Software Licensing

- VXLAN license might display as invalid if QFX-ADV-FEATURE-LIC is installed. [PR1288916](#)

Virtual Chassis

- QFX5100 TVP: Not able to load TVP image on top of a non-TVP 5100 image while adding a QFX5100 switch to the Virtual Chassis. [PR1248145](#)
- QFX5100: The ovsdb-server daemon failed to start. [PR1288052](#)
- On QFX-5100, the fxpc process generates a core file. [PR1294033](#)
- QFX5200: New apply group not applying to the Virtual Chassis after a reboot. [PR1305520](#)

VLAN Infrastructure

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

SEE ALSO

New and Changed Features 253
Changes in Behavior and Syntax 266
Known Behavior 272
Known Issues 276
Documentation Updates 291
Migration, Upgrade, and Downgrade Instructions 292
Product Compatibility 305

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.4R2.

SEE ALSO

New and Changed Features 253
Changes in Behavior and Syntax 266
Known Behavior 272
Known Issues 276

[Resolved Issues | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 292](#)

[Product Compatibility | 305](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 292](#)
- [Installing the Software on QFX10002 Switches | 295](#)
- [Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches | 295](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 297](#)
- [Performing a Unified ISSU | 301](#)
- [Preparing the Switch for Software Installation | 302](#)
- [Upgrading the Software Using Unified ISSU | 302](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 305](#)

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.

3. Select **17.4** 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 17.4 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 sourcejinstall-host-qfx-10-f-x86-64-17.4
-R1.n-secure-signed.tgz reboot 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 17.4 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 17.4R1.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-f-x86-64-17.4
-R2.n-secure-signed.tgz reboot reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-17.4
-R1.n-secure-signed.tgz reboot reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```


After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-x86-64-17.4
-R1.n-secure-signed.tgz reboot
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the **request system reboot** command:

```
user@switch> request system reboot
```

NOTE: You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete <package-name>** command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

```
user@switch> show version
```

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
Routing Engine status:
Slot 1:
  Current state           Master
  Election priority       Backup (default)
```

14. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-17.4R1.n-secure-signed.tgz
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

15. Reboot the Routing Engine using the **request system reboot** command:

```
user@switch> request system reboot
```

NOTE: You must reboot to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your system. Instead, finish the installation and then issue the **request system software delete jinstall <package-name>** command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the **show version** command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
  Slot 0:
    Current state           Master
    Election priority       Master (default)
Routing Engine status:
  Slot 1:
    Current state           Backup
    Election priority       Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 302](#)
- [Upgrading the Software Using Unified ISSU on page 302](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:

NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication** is **Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.

Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, *jinstall-host-qfx-5-17.3R1-signed.tgz*.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-17.4
-R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-17.4
-R1.n-secure-signed.tgz completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
```

```

ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item              Status              Reason
  FPC 0             Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. 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>.

SEE ALSO

New and Changed Features	253
Changes in Behavior and Syntax	266
Known Behavior	272
Known Issues	276
Resolved Issues	280
Documentation Updates	291
Product Compatibility	305

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 306

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	253
Changes in Behavior and Syntax	266
Known Behavior	272
Known Issues	276
Resolved Issues	280
Documentation Updates	291
Migration, Upgrade, and Downgrade Instructions	292

Junos OS Release Notes for SRX Series

IN THIS SECTION

- New and Changed Features | 307
- Changes in Behavior and Syntax | 319
- Known Behavior | 321
- Known Issues | 325
- Resolved Issues | 328
- Documentation Updates | 342
- Migration, Upgrade, and Downgrade Instructions | 342
- Product Compatibility | 346

These release notes accompany Junos OS Release 17.4R2 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R2 New and Changed Features | 308
- Release 17.4R1-S1 New and Changed Features | 308
- Release 17.4R1 New and Changed Features | 310

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R2 for the SRX Series devices.

Release 17.4R2 New and Changed Features

There are no new features in Junos OS Release 17.4R2 for the SRX Series devices.

Release 17.4R1-S1 New and Changed Features

Junos OS Release 17.4R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/SRX4200, SRX5400, SRX5600, and SRX5800.

Junos OS Release 17.4R1-S1 supports SRX4600 device.

Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 15.1X49-D80 through 15.1X49-D100. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D100 are not available in 17.4 releases.

NOTE: Junos OS for SRX Series Software documentation includes information about SRX4600 Services Gateway.

New features for security platforms in Junos OS Release 17.4R1 and Junos OS Release 17.4R1-S1 include:

Chassis Cluster

- **Media Access Control Security (MACsec) (SRX4600)**—Starting in Junos OS Release 17.4R1-S1, Media Access Control Security (MACsec) is supported on HA control and fabric ports of SRX4600 devices in chassis cluster mode to secure point-to-point Ethernet links between two nodes in a cluster.

In the SRX chassis cluster implementation, the control and fabric link carry secure traffic between two nodes in clear text format. Because of this, it is important to encrypt the data between the two nodes. MACsec is an industry-standard security technology that provides secure communication and identifies and prevents most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks. MACsec can be used in combination with other security protocols to provide end-to-end network security.

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

Hardware

- **SRX4600 Services Gateway**—Starting with Junos OS Release 17.4R1-S1, SRX4600 Services Gateways are available as the next-generation, high-performance, and scalable security services devices. The services gateway supports 75-Gbps Internet mix (IMIX) throughput, is suited for large enterprises and small to medium data centers. The SRX4600 Services Gateway provides industry-leading next-generation firewall capabilities (AppID, UserFW, IPS, UTM, and so on) and advanced threat detection and mitigation capabilities features such as SecIntel and SkyATP. The Services Gateway features two high-performance Intel Xeon processors with 14 cores per processor.

Platforms and Infrastructure

- **Software support for SRX4600 devices**—Starting in Junos OS Release 17.4R1-S1, Junos OS supports the SRX4600 Services Gateway. The SRX4600 device is a high-end dynamic services gateway that consolidates security functionality, networking services, and uncompromised performance for medium to large enterprises. With advanced security and threat mitigation capabilities, SRX4600 device can be used for campus edge integrated firewall, data center edge firewall, data center core firewall, LTE security gateway, and Gi/SGi firewall.

SRX4600 device supports Juniper's Software-Defined Secure Network (SDSN) framework, including Sky Advanced Threat Prevention (Sky ATP), which is built around automated and actionable intelligence that can be shared quickly to recognize and mitigate threats.

The SRX4600 device supports the following software features:

- Stateful firewall
- Application security suite
- UTM (Sophos AV, Web filtering, content filtering, and antispam)
- IDP
- Advanced anti-malware
- High availability (Chassis cluster)

- Dual HA control ports (10G)
- MACsec support for HA ports
- Ethernet interfaces through QSFP28 (100G modes), QSFP+ (40G/4x10G modes) and SFP+ (10G mode)
- IPsec VPN, including AutoVPN and Group VPNv2
- QoS and network services
- J-Web
- Routing policies with multicast

The SRX4600 implements use of an individual thread for each session that is dedicated to management of that session and its flow. As a result, out-of-order packet problems that can occur with concurrent processing are eliminated.

Installation packages available for SRX4600 devices are, Preboot Execution Environment (PXE), USB install media package, and CLI upgrade.

You can use the **show chassis hardware** command to display the part number and the model number of the SRX4600 device.

You can use the **show security ipsec tunnel-distribution** command to display the number of VPN tunnels anchored in each thread ID.

[See [Understanding Flow Processing on the SRX4600 Device](#).]

Security

- **Secure Boot (SRX4600)**—Starting in Junos OS Release 17.4R1-S1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms.

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

Release 17.4R1 New and Changed Features

Junos OS Release 17.4R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/SRX4200, SRX5400, SRX5600, and SRX5800.

Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 15.1X49-D80 through 15.1X49-D100. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D100 are not available in 17.4 releases.

ALG

- **H.323 gateway-to-gateway support (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, the gateway-to-gateway call feature is supported on the H.323 ALG. This feature introduces

one-to-many mapping between an H.225 control session and H.323 calls as multiple H.323 calls go through a single control session.

[See [Understanding H.323 ALG.](#)]

- **NAT64 support for H.323 ALG (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, the H.323 ALG supports NAT64 rules in an IPv6 network.

[See [Understanding H.323 ALG.](#)]

Application Security

- **Advanced policy-based routing (APBR) with midstream support (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, SRX Series Services Gateways support advanced policy-based routing (APBR) with an additional enhancement to apply the APBR in the middle of a session (midstream support). With this enhancement, you can apply APBR for a non-cacheable application and also for the first session of the cacheable application.

You can fine-tune the outbound traffic with APBR configuration (for example, limiting route changes and terminating sessions) to avoid issues such as excessive transitions due to frequent route changes.

The enhancement provides more flexible traffic-handling capabilities that offer granular control for forwarding packets.

[See [Understanding Advanced Policy-Based Routing.](#)]

- **Application tracking enhancements to support category and subcategory (SRX Series, vSRX instances)**—Starting from Junos OS Release 17.4R1, AppTrack session create, session close, and volume update logs include new fields **category** and **subcategory**. AppTrack syslog message provide general information about the application type, and including category and subcategory of the application in the message, helps in categorizing the applications.

[[Understanding AppTrack.](#)]

Authentication and Access

- **User firewall support for IPv6 (SRX Series, vSRX instances)**—Starting in Junos OS Release 17.4R1, SRX Series devices support IPv6 addresses for user firewall (UserFW) authentication. This feature allows IPv6 traffic to match any security policy configured for source identity. Previously, if a security policy was configured for source identity and “any” was specified for its IP address, the UserFW module ignored the IPv6 traffic. IPv6 addresses are supported for the following authentication sources:
 - Active directory authentication table
 - Device identity with active directory authentication
 - Local authentication table
 - Firewall authentication table

[See [Overview of Integrated User Firewall](#).]

Chassis Cluster

- **Preemptive delay timer (SRX Series)**—Starting with Junos OS Release 17.4R1, a failover delay timer is introduced on SRX Series devices in a chassis cluster to limit the flapping of redundancy group state between the secondary and the primary nodes in a preemptive failover.

Back-to-back failovers of a redundancy group in a short interval can cause the cluster to exhibit unpredictable behavior because of flapping of the active and backup systems.

To prevent this, a delay timer can be configured to delay the immediate failover for a configured period of time—between 1 and 21,600 seconds. In addition, you can configure the preemptive limit to restrict the number of failovers (1 to 50) in a given time period (1 to 1440 seconds) when preemption is enabled for a redundancy group.

This enhancement enables the administrator to introduce a failover delay, which can reduce the number of failovers and result in a more stable network state due to the reduction in active / backup flapping within the redundancy group.

[[Understanding Chassis Cluster Redundancy Group Failover](#).]

Class of Service (CoS)

- **Support for CoS on dI0 Interface on SRX320, SRX340, SRX345, and SRX550M devices**— Starting with Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can configure the following class of service (CoS) features on the dI0 interface for 4G wireless modems: behavior aggregate classifiers, multifield classifiers, policers, shapers, schedulers, and rewrite rules. The dialer interface, dI0, is a logical interface for configuring properties for modem connections.

[See [LTE Mini-PIM Overview](#).]

- **Support CoS on Logical Tunnel Interface in a Chassis Cluster on SRX300, SRX320, SRX340, SRX345, and SRX550M devices**— Starting with Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, queuing is supported on logical tunnel (lt) interfaces to allow CoS configuration.

[See [CoS Queuing for Tunnels Overview](#).]

- **Support for port-based egress traffic shaping and policing on SRX Series devices**— Starting with Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can configure egress traffic shaping and policing at the physical port level, which limits the egress traffic rate of all logical interfaces on the port.

[See [shaping-rate \(CoS Interfaces\)](#).]

Flow-based and Packet-based Processing

- **Hash-based session distribution (SRX5400, SRX5600, SRX5800)**— Starting with Junos OS Release 17.4R1, traffic is hashed and distributed to different SPUs by the IOC, based on a hash-based session distribution algorithm. This enhancement provides an even hash distribution among all SPUs by using a larger fixed-length hash table. In earlier Junos OS releases, the traffic distribution was uneven among all SPUs in some cases due to a smaller fixed-length hash table.

[See [Understanding Load Distribution in SRX5800, SRX5600, and SRX5400 Devices and vSRX](#).]

GPRS

- **Support for GTP handover group (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 devices and vSRX instances)**—Starting with Junos OS Release 17.4R1, GTP handover group configuration is supported on GTP profiles. An administrator can configure a GTP profile and associate a GTP handover group to a GTP profile.

A GTP handover group is a set of SGSNs or serving gateway (SGW) with a common address-book library. When a GTP handover group name is referenced by a GTP profile, the device checks to see if the current SGSN/SGW address and the proposed SGSN/SGW address are contained within the same GTP handover group. If both the current and proposed SGSN/SGW addresses are contained within the same GTP handover group, then the handover is allowed. If both the current and proposed SGSN/SGW addresses are not within the same GTP handover group, then the profile for the default handover group is used.

This feature enables the administrator to define policies that determine whether handover can happen between individual SGSNs/SGW and/or groups of SGSNs/SGW for roaming.

[See [GTP Handover Group Overview](#).]

Hardware

- **SRX345 Services Gateway (DC power supply model)**—The SRX345 Services Gateway now includes a DC model. The DC model has a single internal power supply, which is not field-replaceable. The DC model supports the same features as those supported on the existing SRX345 Services Gateways. The minimum Junos OS release supported on the DC model is 17.4R1. The services gateway can be managed using the CLI, Junos Space, and J-Web.

[See [SRX345 Services Gateway Description](#).]

Interface and Chassis

- **MACsec support (SRX300, SRX320, SRX340 and SRX345)**—Starting in Junos OS Release 17.4R1, Media Access Control Security (MACsec) is supported on all MACsec-capable ports of SRX300, SRX320, SRX340 and SRX345 devices.

On SRX300 line devices MACsec is supported on the following ports:

- SRX300 and SRX320: 2 ports (on two fixed SFP interfaces.)
- SRX340 and SRX345: 16 ports (on eight fixed SFP interfaces + eight fixed Ethernet ports)

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

- **PPPoE support on SRX Series and vSRX devices**—Starting in Junos OS Release 17.4R1, SRX series devices and vSRX support Point-to-Point Protocol over Ethernet (PPPoE). You can connect multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. The hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

[See [Understanding PPPoE Interfaces.](#)]

- **RFC 4638 support for SRX300, SRX320, SRX340, SRX345, and SRX550M devices**— Starting in Junos OS Release 17.4R1, you can use the PPP-Max-Payload option to override the default behavior of the PPPoE client by providing a maximum size that the PPP payload can support in both sending and receiving directions. The PPPoE server might allow the negotiation of an MRU larger than 1492 and the use of an MTU larger than 1492.

[See [Understanding MTU and MRU Configuration for PPP Subscribers.](#)]

Installation and Upgrade

- **Upgraded FreeBSD support (SRX1500, SRX4100, SRX4200, and vSRX instances)**—Starting with Junos OS Release 17.4R1, the Junos Control Plane (JCP) virtual machine (VM) in the SRX Series devices is upgraded to support FreeBSD 11. Two virtual CPUs (VCPU) are allocated for JCP VM in the Linux host to improve Routing Engine performance for SRX4100 and SRX4200 devices and vSRX instances. For vSRX, additional vCPU will be allocated if you allocate more CPUs than the minimum required. For SRX1500 devices, no additional CPUs are available to allocate for JCP VM.

[See [Understanding Junos OS with Upgraded FreeBSD for SRX Series Devices.](#)]

Logical System

- **Logical system (LSYS) support (SRX1500)**—Starting in Junos OS Release 17.4R1, the logical system feature is supported on SRX1500 devices in addition to the existing support on SRX Series devices such as SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800. A logical system provides virtualization on a device that is partitioned into multiple logical administrative segments. Each segment can have its own security, routing, and bridging attributes.

[See [Understanding Logical Systems for SRX Series Services Gateways.](#)]

Management

- **Support for multiple, smaller configuration YANG modules (SRX Series)**—Starting in Junos OS Release 17.4R1, the YANG module for the Junos OS configuration schema is split into a root configuration module that is augmented by multiple, smaller modules. The root configuration module comprises the top-level configuration node and any nodes that are not emitted as separate modules. Separate, smaller modules augment the root configuration module for the different configuration statement hierarchies. Smaller configuration modules enable YANG tools and utilities to more quickly and efficiently compile and work with the modules, because they only need to import the modules required for the current operation.

[See [Understanding the YANG Modules That Define the Junos OS Configuration.](#)]

NAT

- **Source NAT resource allocation improved (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 17.4R1, source NAT resources handled by the central point architecture have been offloaded to the SPUs when the SPC number is more than four, resulting in more efficient resource allocation.

[See [Understanding Central Point Architecture Enhancements for NAT.](#)]

Routing Policy and Firewall Filters

- **Maximum number of addresses per security policy increased (SRX550M)**—Starting in Junos OS Release 17.4R1, the maximum number of addresses per policy has been increased from 1024 to 2048 for SRX550M. SRX300, SRX320, SRX340 and SRX345 devices already support 2048 source and 2048 destination addresses per policy.

Routing Protocols

- **Support for EBGp route server (SRX Series)**—Starting in Junos OS Release 17.4R1, BGP feature is enhanced to support EBGp route server functionality. A BGP route server is the external BGP (EBGP) equivalent of an internal IBGP (IBGP) route reflector that simplifies the number of direct point-to-point EBGp sessions required in a network. EBGp route server propagates unmodified BGP routing information between external BGP peers to facilitate high scale exchange of routes in peering points such as Internet Exchange Points (IXPs). When BGP is configured as a route server, EBGp routes are propagated between peers unmodified, with full attribute transparency (NEXT_HOP, AS_PATH, MULTI_EXIT_DISC, AIGP, and Communities).

The BGP JET **bgp_route_service.proto** API has been enhanced to support route server functionality as follows:

- Program the EBGp route server.
- Inject routes to the specific route server RIB for selectively advertising it to the client groups in client-specific RIBs.

The BGP JET **bgp_route_service.proto** API includes a peer-type object that identifies individual routes as either EBGp or IBGP (default).

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

System Logging

- **Support for log warning messages on throughput overuse (SRX4100)**—Starting with Junos OS Release 17.4R1, when Internet mix (IMIX) throughput exceeds the limitation for an SRX4100 device, new log warning messages are logged. These log warning messages remind you that there is throughput overuse.

[See [Log File Sample Content](#).]

- **On-box reporting enhancements (SRX Series, vSRX instances)**—Starting in Junos OS Release 17.4R1, SRX4600 devices support the on-box reporting feature, which is already supported on SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200 devices and vSRX instances. Also, the on-box reports are now enhanced to provide comprehensive and detailed reports.

The on-box reporting feature now provides the following enhancements:

- AppTrack API gets information on application category, subcategory, and risk level. An RTLOG module uses this API to get and send information to the local log management process (daemon).
- Reports for applications, categories, subcategories, risk levels, and botnet threats are now by count and volume.
- Application information is generated in UTM log reports.
- Logs can now be listed from latest to oldest. Previously, logs were sorted only from oldest to latest.
- SRX4600 devices now have a hard disk partition available to save traffic logs.

[See [Understanding On-Box Logging and Reporting](#).]

Screens

- **UDP flood screen whitelist (SRX300, SRX320, SRX340, SRX345, SRX1400, SRX4100, and SRX4200 devices, and vSRX instances)**—Starting with Junos OS Release 17.4, UDP flood whitelist mechanism is implemented on SRX300, SRX320, SRX340, SRX345, SRX1400, SRX4100, and SRX4200 devices, and vSRX instances.

When UDP is enabled in a zone, all the UDP traffic performs UDP flood attack detection. The UDP packets that are above the threshold level will be dropped. To avoid these packet drops and instead allow these packets to bypass UDP flood detection, the UDP flood screen whitelist is implemented. To support UDP flood whitelist, the traffic from addresses in the whitelist groups will bypass UDP flood check. Both IPv4 and IPv6 whitelists are supported and can be configured using a single address or a subnet address. UDP flood whitelist supports a maximum of 32 whitelist groups and each group has 32 or fewer IPv4 or IPv6 addresses.

See [Understanding Whitelists for UDP Flood Screens](#).

UTM

- **Custom URL category support for SSL forward proxy (SRX Series)**—Starting with Junos OS Release 17.4R1, the whitelisting feature is extended to include custom URL categories supported by UTM in the whitelist configuration of SSL forward proxy. In this implementation, the Server Name Indication (SNI) field is extracted by the UTM module from client hello messages to determine the URL category. SNI is an extension of the SSL/TLS protocol. Each URL category has a unique ID. The list of URL categories in the whitelist is parsed and the corresponding category IDs are pushed to the Packet Forwarding Engine for each SSL forward proxy profile. The SSL forward proxy then determines through APIs whether to accept the proxy or to ignore the session.

[See [SSL Proxy Overview](#)]

- **Enhanced Web Filtering (EWF) reputation and categorization behavior support for EWF category (SRX Series)**—Starting from Junos OS Release 17.4R1, predefined base filters, defined in a category file, are supported for individual EWF categories. Each EWF category has a default action in a base filter, which is attached to the user profile to act as a backup filter. If the categories are not configured in the user profile, then the base filter takes the action. Online upgradation of base filters is also supported. Further, users can apply global reputation values, provided by the Websense ThreatSeeker Cloud (TSC). For the non-category URLs, the global reputation value is used to perform filtering, and from this release onward, the reputation base scores are configurable.

[See [Understanding Enhanced Web Filtering Process](#).]

- **Local Web filtering enhancement to support custom category configuration (SRX Series)**—Starting from Junos OS Release 17.4R1, support for custom category configuration is available for EWF, local, and Websense redirect profiles. The **custom-message** option is also supported in a category for local Web filtering and Websense redirect profiles. You can create multiple URL lists (custom categories) and apply them to a UTM Web filtering profile with actions such as permit, permit and log, block, and quarantine.

To create a global whitelist or blacklist, apply a local Web filtering profile to a UTM policy and attach it to a global rule.

[See [Understanding Local Web Filtering](#).]

- **Support for new Websense EWF categories (SRX Series)**—Starting from Junos OS Release 17.4R1, you can download and dynamically load new Enhanced Web Filtering (EWF) categories. The downloading and dynamic loading of the new EWF categories do not require a software upgrade. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.

[See [Understanding Redirect Web Filtering](#).]

VPN

- **Increased number of IKE security associations supported (SRX5600, SRX5800)**—Starting from JunosOS Release 17.4R1, SRX5600 with 5 SPC2 cards, and SRX5800 with 10 SPC2 cards can support up to 50,000 IKE security associations (SAs) (each SPC2 card supports upto 20,000 IKE SAs (5,000 IKE SAs / SPU)) for AutoVPN networks in point-to-point secure tunnel mode with multiple traffic selectors. There are no changes in configuration.

[See [Understanding AutoVPN](#).]

- **IPv6 address support for point-to-point AutoVPN networks that use traffic selectors (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, AutoVPN networks that use secure tunnel interfaces in point-to-point mode support IPv6 addresses for traffic selectors and for IKE peers.

NOTE: IPv6 addresses are not supported for AutoVPN networks in point-to-multipoint secure tunnel mode.

[See [Understanding AutoVPN](#) and [Understanding AutoVPN with Traffic Selectors](#).]

- **IPsec VPN performance optimization (SRX5400, SRX5600, SRX5800)**—Starting with Junos OS Release 17.4R1, IPsec VPN performance is optimized when the VPN session affinity and performance acceleration features are enabled. Session affinity is enabled with the **set security flow load-distribution session-affinity ipsec** command, while performance acceleration is enabled with the **set security flow ipsec-performance-acceleration** command.

[See [Accelerating the IPsec VPN Traffic Performance](#) and [Understanding VPN Session Affinity](#).]

SEE ALSO

[Changes in Behavior and Syntax](#) | 319

[Known Behavior](#) | 321

[Known Issues](#) | 325

[Resolved Issues](#) | 328

[Documentation Updates | 342](#)[Migration, Upgrade, and Downgrade Instructions | 342](#)[Product Compatibility | 346](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Chassis Cluster | 319](#)
- [IDP | 320](#)
- [Forwarding and Sampling | 320](#)
- [System Logging | 321](#)
- [User Interface and Configuration | 321](#)

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 17.4R2.

Chassis Cluster

- **IP Monitoring**—Starting with Junos OS Release 17.4R2, on all SRX Series devices, if the reth interface is in bundled state, IP monitoring for redundant groups is not supported on the secondary node. This is because the secondary node sends reply using the lowest port in the bundle which is having a different physical MAC address. The reply is not received on the same physical port from which the request is sent. If the reply comes on the other interface of the bundle, then the internal switch drops it.
- **Power Entry Module**—Starting with Junos OS Release 17.4R2, when you use DC PEM on SRX Series devices operating in chassis cluster mode, the output of **show chassis power** command shows **DC input: 48.0 V input (57000 mV)**. The value **48.0 V input** is a fixed string and can be interpreted as a measured input voltage. The acceptable range of DC input voltage accepted by the DC PEM is 40 to 72 V. The **(57500 mV)** is a measured value, but is not related with the input. It is the actual output value of the PEM and the value is variable. The **DC input:** from **show chassis power** and **Voltage:** information from **show chassis environment pem** command output are removed for each PEM.
- SRX5400, SRX5600, and SRX5800 devices operating in a chassis cluster might encounter the em0 or em1 interface link failure on either of the nodes, which results in split-brain condition. That is, both devices are unable to detect each other. If the failure occurs on the secondary node, the secondary node is moved to the disabled state.

This solution does not cover the following cases:

- em0 or em1 failure on primary node
- HA process restart
- Preempt conditions
- Control link recovery

IDP

- Custom Attack (SRX Series)—Starting with Junos OS Release 17.4R2, the maximum number of characters allowed for a custom attack object name is 60. You can validate the statement using the CLI **set security idp custom-attack** command.

Forwarding and Sampling

- Support for Address Resolution Protocol (ARP) throttle and ARP detect [SRX5400, SRX5600, and SRX5800]—Starting in Junos OS Release 17.4R2, an ARP throttling mechanism is introduced for SRX Series devices.

Excessive ARP processing results in high utilization of Routing Engine CPU resources, resulting in deprivation of CPU resources to other Routing Engine processes. To provide protection against excessive ARP processing, you can now use the following configuration statements:

- **edit forwarding-options next-hop arp-throttle *seconds***
- **edit forwarding-options next-hop arp-detect *milliseconds***



CAUTION: We recommend that only advanced Junos OS users attempt to configure the ARP throttle and ARP detect feature. An improper configuration could result in high CPU utilization of the Routing Engine, which could affect other processes on your device.

[See [arp-throttle](#) and [arp-detect](#)].

System Logging

- **System log host support (SRX300, SRX320, SRX340, SRX345 Series devices)**— Starting in Junos OS Release 17.4R2, when the device is configured in stream mode, you can configure maximum of eight system log hosts.

In Junos OS Release 17.4R1 and earlier releases, you can configure only three system log hosts in the stream mode. If you configure more than three system log hosts, then the following error message is displayed **error: configuration check-out failed**.

User Interface and Configuration

- **Junos OS prohibits configuring ephemeral configuration database instances that use the name default (SRX Series)**—Starting in Junos OS Release 17.4R2, user-defined instances of the ephemeral configuration database, which are configured using the `instance instance-name` statement at the `[edit system configuration-database ephemeral]` hierarchy level, do not support configuring the name **default**.

SEE ALSO

New and Changed Features 307
Known Behavior 321
Known Issues 325
Resolved Issues 328
Documentation Updates 342
Migration, Upgrade, and Downgrade Instructions 342
Product Compatibility 346

Known Behavior

IN THIS SECTION

- [Authentication and Access | 322](#)
- [Chassis Clustering | 322](#)
- [Install and Upgrade | 323](#)

- Interfaces and Chassis | 323
- J-Web | 324
- Layer 2 Ethernet Services | 325
- User Interface and Configuration | 325
- VPNs | 325

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.4R2 for the SRX Series.

Authentication and Access

- On SRX Series devices with 256K user firewall authentication entries, in case of a failover or when PFE restart occurs, the **show services user-identification** command will generate response timeout. This timeout will last for at least 10 minutes. [PR1302269](#)
- On SRX Series devices, the traffic that is sourced-from or destined-to the SRX Series device itself is classified as UNKNOWN in AppTrack log messages. [PR1340338](#)

Chassis Clustering

- On SRX4600 devices, the dedicated Chassis Cluster fabric ports are not available. Instead, any 40G or 10G traffic ports can be used as chassis cluster fabric ports.
- IP monitoring for redundancy groups does not work on the secondary node if the reth interface has more than one physical interfaces configured. This is because the backup node sends traffic using the MAC address of the lowest port. If the reply does not come back on the same physical port, then the internal switch drops the traffic. [PR1344173](#)

Install and Upgrade

- On SRX Series devices, when you perform a downgrade from Junos OS Release 17.4R1-S2 or 17.4R2 to Junos OS Release 15.1X49-D125, using the **request system software add** command, downgrade fails. An error message mentioning that you need to force the downgrade process using the **force** CLI option is displayed. Use the **force** CLI option to force the downgrade. There is no need to use the force option when you downgrade from Junos OS Release 15.1 to any other release. [1350558](#)

Interfaces and Chassis

- On SRX4600 devices, the 10-Gigabit Ethernet and chassis cluster ports cannot be configured to operate as 1-Gigabit Ethernet ports.
- SRX4600 device interfaces only support the following two traffic port modes:
 - 4x40G (all four QSFP+ ports) + 8x10G (all eight SFP+ ports) by default.
 - 2x100G (first two QSFP+ ports) + 4x10G (first four SFP+ ports) by configuration as shown below:
 - **set chassis fpc 1 pic 0 pic-mode 100G**
 - **set chassis fpc 1 pic 0 number-of-ports 2**
 - **set chassis fpc 1 pic 1 number-of-ports 4**

NOTE: The system requires a reboot after committing the above configuration.

- On SRX4600 devices, the RAID-1 mirror feature is not available. The second SSD is not available for use, although it is present.
- On SRX4600 devices, precision Time Protocol (PTP) feature is not available.
- On SRX4600 devices, USB disk is not available for the Junos OS. However, the USB disk is available with full access for Host OS (Linux) and USB is still used in the booting process (install and recovery functions). [PR1283618](#)
- On SRX1500 devices, pp0.0 interface link status is not up. [PR1315416](#)
- USB stops working if the USB is removed while it is in initialization state. To avoid this issue, wait for few seconds before removing the USB. [PR1332360](#)

J-Web

- On SRX550M and SRX1500 devices, there is no option to configure Layer 2 firewall filters from J-Web, irrespective of the device mode. [PR1138333](#)
- On SRX Series devices in chassis cluster, if you want to use J-Web to configure and commit the configurations, you must ensure that all other user sessions are logged out including any CLI sessions. Otherwise, the configurations might fail. [PR1140019](#)
- On SRX1500 devices in J-Web, snapshot functionality under **Maintain->Snapshot->Target Media->Disk->Click Snap Shot** is not supported. [PR1204587](#)
- On SRX Series devices, DHCP relay configuration under **Configure > Services > DHCP > DHCP Relay** page is removed from J-Web in Junos OS Release 15.1X49-D60. The same DHCP relay can be configured using the CLI. [PR1205911](#)
- On SRX Series devices, DHCP client bindings under Monitor is removed. The same bindings can be seen in CLI using the **show dhcp client binding** command. [PR1205915](#)
- On SRX Series devices, if the configuration load is more than 5000 bytes then J-Web responds slowly and the navigation of pages might take more time. [PR1222010](#)
- On SRX Series devices, you cannot view the custom log files created for event logging in J-Web. [PR1280857](#)
- On SRX Series devices, generation of reports will work in IE and chrome browsers. To generate report in firefox, delete existing ff profile and relaunch firefox with new profile. [PR1303722](#)
- Uploading certificate using browse button, stores the certificate in device at **/jail/var/tmp/uploads/**, which is deleted when you execute the CLI **request system storage cleanup** command. [PR1312529](#)
- The values of address and address-range are not displayed in the inline address-set creation pop-up window of Juniper Identity Management Service (JIMS). [PR1312900](#)

Layer 2 Ethernet Services

- PPPoE + DHCPv6 cannot work in all SRX platforms with 15.1X49 and later versions. [PR1229836](#)

User Interface and Configuration

- On SRX1500 devices, committing a configuration with a huge number of logical systems will take more time. This issue occurs because taking backup of previous configurations might take a little longer to finish. [PR1339862](#)

VPNs

- On SRX5400, SRX5600, and SRX5800 devices, when CoS is enabled on the st0 interface and the incoming traffic rate destined for the st0 interface is higher than 300,000 packets per second (pps) per SPU, the device might drop some of the high-priority packets internally and shaping of outgoing traffic might be impacted. We recommended that you configure the appropriate policer on the ingress interface to limit the traffic below 300,000 pps per SPU. [PR1239021](#)

SEE ALSO

New and Changed Features 307
Changes in Behavior and Syntax 319
Known Issues 325
Resolved Issues 328
Documentation Updates 342
Migration, Upgrade, and Downgrade Instructions 342
Product Compatibility 346

Known Issues

IN THIS SECTION

- [Outstanding Issues | 326](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R2.

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

Outstanding Issues

Application Layer Gateways (ALGs)

- In a chassis cluster with logical systems are configured, any ALG (excepts DNS ALG) enabled, and NAT configured for the ALG sessions, the flowd process on the secondary node might not work. [PR1343552](#)

Chassis Cluster

- On SRX5600 and SRX5800 devices in chassis cluster mode, when the secondary Routing Engine is installed to enable dual control links, the **show chassis hardware** command might display the same serial number for both the routing engines on both the nodes. [PR1321502](#)
- On SRX Series devices, the forwarding plane might failover from node 0 to node 1 when an SPC stops unexpectedly. [PR1331809](#)
- On SRX4600 device with chassis cluster enabled, when a failover occurs the dedicated fabric link is down. [PR1365969](#)

Class of Service (CoS)

- On SRX Series devices, if the action of **forwarding-class** is configured in the output direction on a firewall filter, the host outbound traffic matching the same term of this firewall filter will be blocked. [PR1272286](#)

Flow-based and Packet-based Processing

- On SRX Series devices, sometimes the time range slider is not working for all events, as well individual events in Google Chrome or Firefox browser. [PR1283536](#)
- On SRX4600 device, when the next-hop is set to the st0 interface, the output of the **show route forwarding-table** command displays the next-hop IP address twice. [PR1290725](#)
- On all SRX Series devices, filter-based forwarding (FBF) does not work when applied on IPsec tunnel interface (st0.*). [PR1290834](#)
- On SRX Series devices with chassis cluster enabled, the ingress interface of the multicast session in the first logical system is reth2.0, which belongs to redundancy group 2. Redundancy group 2 is active on node 1. The ingress interface of multicast session in the second logical system will be the PLT interface, which belongs to redundancy group 1. Redundancy group 1 is active on node 0. So, the multicast session in the second logical system will be active on node 0. Due to this condition multicast session active/backup is not aligned with forwarding traffic. This issue occurs when multicast traffic goes across logical systems. As a workaround to make RG-1 and RG-2 active on the same node. [PR1295893](#)
- On SRX300, SRX320, SRX340, and SRX345 devices, if there is power outage many times in a short period of time, the device might end up getting stuck in the loader prompt. [PR1292962](#)

- On SRX Series devices, packet capture does not work after you change, delete, or add maximum capture size. [PR1304723](#)
- On SRX Series devices, when you run the command **clear nhdb statistics** on the SPU PIC, the SPC might reset. [PR1346320](#)
- The IPsec replay error for Z-mode traffic is observed. [PR1349724](#)
- The IPsec VPN traffic might be dropped on pass-through SRX Series device after an IKE rekey. [PR1353779](#)
- On the secondary control plane, a multicast session leak is observed for the PIM register. [PR1360373](#)

Intrusion Detection and Prevention (IDP)

- After an IDP signature automatic update is scheduled, the secondary node might not update the signatures. [PR1358489](#)

Interfaces and Routing

- Incorrect ingress packet per second is observed on the MPLS enabled interface. [PR1328161](#)

Interfaces and Chassis

- On SRX1500, if Junos OS Release 15.1X49-D70 or later is installed and you have a single PEM in slot 0, you will see an alarm saying PEM 1 is not present. [PR1265795](#)
- On SRX4600 device, the 1GE interface is not supported in Junos OS Release 17.4R2. [PR1315073](#)

Platform and Infrastructure

- The Secure Shell (SSH) to SRX fails if the **phone-home: kern.maxfiles** limit exceeds. [PR1357076](#)
- On SRX4100 and SRX4200 devices, the Network Time Protocol (NTP) server might not synchronize because device the clock often switched from NTP to local time. [PR1357843](#)

Routing Policy and Firewall Filters

- On SRX Series devices, DNS name entries in policies might not be resolved if the routing instance is configured under a system name server. [PR1347006](#)

Routing Protocols

- On SRX Series devices, RIP is supported in packet to packet DC mode on st0 interfaces. [PR1141817](#)
- A new CLI command **stickydr** is required to prevent traffic loss during the disaster recovery. [PR1352589](#)

VPNs

- IPsec uses ESP as the default protocol, if the user does not explicitly configure the protocol. [PR1061838](#)
- When an SRX Series device acts as an initiator behind the NAT, disabling NAT on the router in between causes an immediate new negotiation failure because of an attempt to disable NAT using the port 4,500. The next attempt succeeds by using the port 500. Disabling NAT and bringing down all the existing tunnels and re-establishing the tunnels with port 500 is the expected behavior. [PR1273213](#)

- On SRX Series devices, in case multiple traffic-selectors are configured for a peer with IKEv2 reauthentication, only one traffic-selector will rekey at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors will be cleared without immediate rekey. New negotiation of those traffic-selectors might trigger through other mechanisms such as traffic or by peer. [PR1287168](#)
- On SRX Series devices, when the VPN monitoring feature is enabled, the st interfaces go down immediately. [PR1295896](#)
- If a period . is present in the CA profile name then the PKID might face issues, if the PKID is restarted at any point. [PR1351727](#)
- On SRX5600 and SRX 5800 devices, during VPN to AutoVPN configuration migration, traffic loss is observed. [PR1362317](#)

SEE ALSO

[New and Changed Features | 307](#)

[Changes in Behavior and Syntax | 319](#)

[Known Behavior | 321](#)

[Resolved Issues | 328](#)

[Documentation Updates | 342](#)

[Migration, Upgrade, and Downgrade Instructions | 342](#)

[Product Compatibility | 346](#)

Resolved Issues

IN THIS SECTION

● [Resolved Issues: 17.4R2 | 329](#)

● [Resolved Issues: 17.4R1 | 338](#)

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

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

Resolved Issues: 17.4R2

Application Layer Gateways (ALGs)

- On SRX1400 device, the NFS traffic to port 2049 might drop. [PR1307763](#)
- The configure download URL displays warning message **requires appid-sig license**. [PR1324858](#)
- On SRX Series devices with SIP ALG enabled, the SIP ALG might drop SIP packets which have a **referred-by** or **referred-to header** field containing multiple header parameters. [PR1328266](#)
- SIP calls drop, when the limit per SPU crosses 10,000 calls. [PR1337549](#)

Authentication and Access Control

- On SRX Series devices, PFE might crash and huge number of core files might be generated within a short period of time. [PR1326677](#)
- On SRX Series devices, incomplete Request Support Information (RSI) might be seen. [PR1329967](#)
- On SRX Series devices, the sessions might close because of the idle **Timeout junos-fwauth-adapter** logs. [PR1330926](#)
- The uacd process is unstable after upgrading to Junos OS Release 12.3X48 and later releases. [PR1336356](#)
- On SRX Series devices, the **show version detail** command returns an error message: **Unrecognized command (user-ad-authentication)** while configuring the useridd settings. [PR1337740](#)
- A new configuration is available to configure the web-authentication timeout. [PR1339627](#)

Chassis Clustering

- The route information might not be synchronized between node0 and node1 when configuring the firewall filter or APBR to use the non-default routing-instance. [PR1292235](#)
- Flowd process core files are generated after adding 65536 VPN tunnels using traffic selector with the same remote IP. [PR1301928](#)
- On devices enabled with chassis cluster, the ISSU upgrade might fail and display an error message **ISSU aborted and exiting ISSU window**. [PR1306194](#)
- On SRX1500, SRX4100 and SRX4200 devices, ISSU might fail if LACP and interface monitoring are configured. [PR1305471](#)
- File Descriptor might leak on SRX Series chassis clusters with Sky ATP enabled. [PR1306218](#)
- When services offloading feature is enabled, the device changes TCP checksum value to 0x0000. [PR1317650](#)
- When ISSU is performed from a Junos OS Release prior to 15.1X49-D60 to a Junos OS Release 15.1X49-D60 or later, flowd process generates core files. [PR1320030](#)
- The device might stop forwarding traffic after RG1 failover from node0 to node1. [PR1323024](#)

- When RGO failover or primary node reboot happens, some of the logical interfaces might not be synchronized to the other node if the system has around 2,000 logical interfaces and 40,000 security policies. [PR1331070](#)
- After the primary node or the secondary node restarts, the FPC module goes offline on the secondary node. [PR1340116](#)
- In and active/active cluster, route change timeout does not work as expected. [PR1314162](#)

Class of Service (CoS)

- Packets go out of order on SPC2 cards with IOC1 or FIOC cards. [PR1339551](#)

Flow-Based and Packet-Based Processing

General Routing

- SRX1500 devices might power off unexpectedly because of incorrect device temperature readings which reported a too high temperature, leading to an immediate pro-active power-off of the device to protect the device from overheating. However in these cases the temperature was not actually too high and a power-off would not be required. When this occurs, the following log message is shown in file `/var/log/hostlogs/lcmd.log`: Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor. [PR1241061](#)
- On SRX4100 and SRX4200 devices, packet loss is observed when the value of packet per second (pps) through the device is very high. This occurs because of the update of the **application interval statistics** statement, which has a default timer value of 1 minute. You can avoid this issue by setting the interval to maximum using the **set services application-identification statistics interval 1440** command. [PR1290945](#)
- The **show host server name-server host** CLI command fails when the source address is specified under the name-server configuration. [PR1307128](#)
- A memory leak might occur in the appidd process while updating an application signature package. [PR1308863](#)
- On SRX4600 devices, when you run the **clear security flow session** command, time taken to clear the session depends on the total session number. For example, the clear session takes nine minutes to clear 57M session. [PR1308901](#)
- On SRX Series devices, if destination NAT and session affinity are configured with multiple traffic selectors in IPsec VPN, the traffic selector match might fail. [PR1309565](#)
- The flowd process might stop and generate a core file during failover between node 0 and node 1. [PR1311412](#)
- On SRX Series devices, the IPsec tunnel might fail to be established if datapath debug configuration include the options **preserve-trace-order**, **record-pic-history**, or both. [PR1311454](#)
- The SRX Series device drops packets citing the reason "Drop pak on auth policy, not authed". [PR1312676](#)
- The flowd process might stop if the SSL-FP profile is configured with whitelist. [PR1313451](#)

- If IDP and SSL forward proxy whitelist are configured together, the device might generate a core file. [PR1314282](#)
- On SRX550M devices, phone-home.core is generated after the zeroization procedure. [PR1315367](#)
- If the Sky ATP cloud feed updates, the Packet Forwarding Engine might stop causing intermittent traffic loss. [PR1315642](#)
- On SRX Series devices, the IPSec VPN tunnel with traffic-selector is configured and the packets TTL is set to 1, the flowd process stops and generates a core file on both the nodes. [PR1316134](#)
- Periodic PIM register loop is observed during switch failure. [PR1316428](#)
- On SRX Series devices, the **fin-invalidate-session** command does not work when the Express Path feature is enabled on the device. [PR1316833](#)
- Return traffic through the routing instance might drop intermittently after changing the zone and routing-instance configuration on the st0.x interface. [PR1316839](#)
- SRX300 devices DHCP client cannot obtain IP addresses. [PR1317197](#)
- Default route is lost after system zero. [PR1317630](#)
- SSL firewall proxy does not work if root-ca has fewer than four characters. [PR1319755](#)
- The OSPF peers are unable to establish neighbors between the LT interfaces of the logical systems. [PR1319859](#)
- On SRX Series devices, after logical system is configured, about 10 logical systems are not working. [PR1323839](#)
- The flowd process generates core files on both nodes causing an outage. [PR1324476](#)
- The MPC cards might drop traffic in the event of high temperatures. [PR1325271](#)
- Software next-hop table is full with log messages RT_PFE: NH IPC op 1 (ADD NEXTHOP) failed, err 6 (No Memory) peer_class 0, peer_index 0 peer_type 10. [PR1326475](#)
- If the serial number of the certificate for the SSL proxy has two consecutive zeros, the certificate authentication fails. [PR1328253](#)
- When you use CFLOW, the source address for flow packets is not displayed. [PR1328565](#)
- On SRX Series devices, the one-way jitter traps are not generated when the TWAMP is configured. [PR1328708](#)
- The FPC is dropped or hangs in the present state when the intermittent control link heartbeat is observed. [PR1329745](#)
- On SRX Series devices with stream logging configured, high CPU load is observed. [PR1331011](#)
- The IPv6 traffic does not work as expected on IOC3 with the services offloading (npcache) feature. [PR1331401](#)
- NTP synchronization fails and switches to a local clock. [PR1331444](#)

- Inaccurate Jflow records might be seen for output interface and next hop. [PR1332666](#)
- The whitelist function in syn-flood does not work. [PR1332902](#)
- The **show vlans detail no-forwarding** command in the RSI does not display any information, because the **no-forwarding** option is not supported. [PR1336267](#)
- Two-way active measurement protocol (TWAMP) client, when configured in a routing instance, does not work after a reboot. [PR1336647](#)
- On the front panel LED, the red alarm goes on after an RG0 failover is triggered when the flowd process stops. [PR1338396](#)
- The unfiltered traffic is captured after traceoptions are deactivated. [PR1339213](#)
- SSH to the loopback interface of SRX Series devices does not work properly when AppTrack is configured. [PR1343736](#)
- The flowd process might stop when SYN-proxy function is used. [PR1343920](#)
- SNMP MIB walk provides incorrect data counters for total current flow sessions. [PR1344352](#)
- SRX1500 devices might encounter a failure while accessing the SSD drive. [PR1345275](#)
- On SRX Series devices, when you upgrade to a Junos OS Release with "no-validate" option and if there are unsupported configurations with the new version, then configuration push fails and the ksyncd process stops. [PR1345397](#)
- The REST API is not working on the SRX320-POE device. [PR1347539](#)
- File download stops over a period of time when TCP proxy is activated through Antivirus or Sky ATP. [PR1349351](#)
- When a J-Flow related configuration is deleted, the forwarding plane begins to drop packets. [PR1351102](#)
- If the Trusted Platform Module (TPM) is enabled, the configuration integrity failure occurs when there is a power loss for few seconds after the commit. [PR1351256](#)
- On SRX1500 device, after the SSL forward proxy is configured, the system stops and generates a core file. [PR1352171](#)
- The flowd process generates a core file when the SIP ALG is enabled. [PR1352416](#)
- When the routing instance is configured, the **UTM Anti-Spam:DUT** process do not send the DNS query. [PR1352906](#)
- On SRX Series devices, if the memory buffer is accessed without checking the mbuf and the associated external storage, the flowd process might stop. [PR1353184](#)
- On SRX Series devices in a chassis cluster, if an IPv6 session is being closed and at the same time the related data-plane Redundancy Group (RG1+) failover occurs, this IPv6 session on the backup node might hang and cannot be cleared. [PR1354448](#)
- The PIM register might stop the message from the source First Hop Router (FHR). [PR1356241](#)

- On SRX300, SRX320, SRX340, and SRX345 devices, with LTE mini-PIM the DHCP relay packets are not forwarded. [PR1357137](#)
- On SRX5000 series devices, when the IPsec performance acceleration feature is enabled, packets going in or out of a VPN tunnel are dropped. [PR1357616](#)
- On SRX5400, SRX5600, and SRX5800 devices, the MIB walk tool is not working when screens are applied to the security zones. [PR1364210](#)

Interfaces and Chassis

- Unable to add IRB and aggregated Ethernet interfaces. [PR1310791](#)
- On SRX1500 devices, pp0.0 interface link status is not up. [PR1315416](#)
- An error is not seen at each commit or commit check if autonegotiation is disabled but the speed and duplex configurations are not configured on the interface. [PR1316965](#)
- If an interface is configured with the Ethernet switching family, we recommend that you do not configure **vlan-tagging**. [PR1317021](#)
- The interface might be brought down by IP monitoring at the time of committing a configuration because of incorrect interface status computing. [PR1328363](#)

Interfaces and Routing

- JIMS server stops responding to requests from SRX Series devices. [PR1311446](#)
- On SRX Series devices in a chassis cluster, the IRB interface does not send an ARP request after clearing the ARP entries. [PR1338445](#)
- Packet reorder occurs on the traffic received on the PPP interface. [PR1340417](#)
- On SRX Series devices, when the VPLS interface receives a broadcast frame, the device sends this frame back to the sender. [PR1350857](#)
- On the SRX1500, when the LACP is configured with interfaces ae0 and ae1, the mac address is displayed as 00:00:00:00:00:00 and 00:00:00:00:00:01 for interfaces ae0 and ae1 respectively. [PR1352908](#)
- The **set protocols rstp interface all** command does not enable RSTP on all interfaces. [PR1355586](#)

Intrusion Detection and Prevention (IDP)

- The control plane CPU usage is high when using IDP. [PR1283379](#)
- IDP signatures might not get pushed to the Packet Forwarding Engine if there is a policy in logical systems. [PR1298530](#)
- The IDP PCAP feature has been improved. [PR1297876](#)
- The output of **show security idp status** command does not accurately reflect the number of decrypted SSL or TLS sessions being inspected by the IDP. [PR1304666](#)
- The file descriptor might leak during a security package auto update. [PR1318727](#)

- On SRX4600 devices, the maximum SSLRP session count is observed to be approaching 100,000. In the CLI, configuring a maximum of 100,000 sessions are allowed, whereas in SSLFP, 600,000 sessions are allowed. Thus, the **set security idp sensor-configuration ssl-inspection sessions** command is now modified to allow a maximum of 600,000 sessions. However, for other devices the original session limit value of 100,000 is retained. [PR1329827](#)
- Loading IDP policy fails because of less available heap memory. [PR1347821](#)

J-Web

- J-Web system snapshot throws error. [PR1204587](#)
- In J-Web when you click the SKIP TO JWEB OPTIONS, the Google Chrome browser automatically redirects. [PR1284341](#)
- J-Web does not display all global address book entries. [PR1302307](#)
- On SRX300, SRX320, SRX340, and SRX345 devices, CPU usage is high when generating on-box reporting on the J-Web. [PR1310288](#)
- J-Web authentication fails when a password includes the backslash. [PR1316915](#)
- J-Web dashboard displays wrong last updated time. [PR1318006](#)
- J-Web display problems for security policies are observed. [PR1318118](#)
- J-Web displays the red alarm for temperature value within the threshold. [PR1318821](#)
- J-Web does not display wizards on the dashboard. [PR1330283](#)
- Unable to delete the dynamic VPN user configuration. [PR1348705](#)
- When the J-Web fails to get resource information, the Routing Engine CPU usage is displayed as 100 percent. [PR1351416](#)
- Security policies search button on the J-Web does not work with Internet Explorer version 11. [PR1352910](#)

Layer 2 Ethernet Services

- In DHCP relay configuration, the option **VPN** has been renamed to **source-ip-change**. [PR1318487](#)
- On SRX1500 devices, VLAN popping and pushing does not work over Layer 2 circuits. [PR1324893](#)
- DHCP rebind and renew packets is not calculated in BOOTREQUEST. [PR1325872](#)
- The default gateway route might be lost after the failover of RG0 in a chassis cluster. [PR1334016](#)
- The subnet mask address is not sent as a reply to the **DHCPINFORM** request. [PR1357291](#)

Network Address Translation (NAT)

- The default-gateway route received by DHCP when some interface in the chassis cluster has been configured as a DHCP client is lost in about 3 minutes after RG0 failover. [PR1321480](#)
- On SRX Series devices, the Sky ATP connection leak causes the service plane to be disconnected from the Sky ATP cloud. [PR1329238](#)

- Arena utilization on a FPC spikes and then resumes to a normal value. [PR1336228](#)

Network Management and Monitoring

- SRX300 device is unresponsive as a result of cf/var: filesystem full error. [PR1289489](#)
- CLI options are available to manage the packet forwarding engine handling the ARP throttling for NHDB resolutions. [PR1302384](#)

Platform and Infrastructure

- When you perform commits with apply-groups, VPN might flap. [PR1242757](#)
- The packet captured by datapath-debug on an IOC2 card might be truncated. [PR1300351](#)
- Inconsistent flow-control status on reth interface is observed. [PR1302293](#)
- On SRX5400, SRX5600, and SRX5800 devices, DC PEM is used on the box, the output of **show chassis environment pem** and **show chassis power** commands do not show DC input value correctly. [PR1323256](#)
- On SRX5400, SRX5600, and SRX5800 devices, SPC2 XLP stops processing packets in the ingress direction after repeated RSI collections. [PR1326584](#)
- When SecIntel is configured, IPFD CPU utilization might be higher than expected. [PR1326644](#)
- The log messages file contains **node*.fpc*.pic* Status:1000 from if_np for ifl_copnfig op:2 for ifl :104** message. [PR1333380](#)
- Log message **No Port is enabled for FPC# on node0** is generated every 5 seconds. [PR1335486](#)
- In RSI, a mandatory argument is missing for the **request pfe execute** and the **show usp policy counters** commands. [PR1341042](#)
- On SRX Series devices in a chassis cluster, configuration commit might succeed even though the external logical interface configuration (reth) associated with the Internet Key Exchange (IKE) VPN gateway configuration is deleted. This might lead to configuration load failure during the next device boot-up. [PR1352559](#)
- On SRX4100 devices, interfaces are shown as half-duplex, but there is no impact on the traffic. [PR1358066](#)

Routing Policy and Firewall Filters

- The firewall authentication does not list the correct polices when the NSD process is busy. [PR1312697](#)
- The number of address objects per policy for SRX5400, SRX5600, SRX5800 devices is increased from 4,096 to 16,000. [PR1315625](#)
- The flowd process stops when AppQoS is configured on the device. [PR1319051](#)
- Flowd process stops after configuring a huge number of custom applications. [PR1347822](#)

- On SRX Series devices, with a large number of firewall authentication entries, the flowd process might stop. [PR1349191](#)
- On SRX Series devices, a large scale commit, for example, 70,000 lines security policy might stop the NSD process on the Packet Forwarding Engine (PFE). [PR1354576](#)

Routing Protocols

- On SRX1500 devices, the IS-IS adjacency remains down when using an IRB interface. [PR1300743](#)
- Dedicated BFD does not work on SRX Series devices. [PR1312298](#)
- On a chassis with BMP configured, if the rpd termination timeout is happening while the BMP main task has failed to terminate and delete itself (seen when rpd is gracefully terminated), the rpd might stop. [PR1315798](#)
- When BGP traceoptions are configured and enabled, the traces specific to messages sent to the BGP peer (BGP SEND traces)are not logged The traces specific to received messages (BGP RECV traces) are logged correctly. [PR1318830](#)
- OpenSSL Security Advisory [07 Dec 2017]. Refer to <https://kb.juniper.net/JSA10851> for more information. [PR1328891](#)
- The ppmmd process might stop, after one node is upgraded and failover completes. [PR1347277](#)
- On SRX Series devices, dedicated BFD does not work. [PR1347662](#)

Software Installation and Upgrade

- The **request system reboot node in/at** command results in an immediate reboot instead of rebooting at the allotted time. [PR1303686](#)
- On SRX1500 devices, the fan speed often fluctuates. [PR1335523](#)

System Logs

- A warning syslog message is displayed when the number of security screens installed exceed the IOC capacity. [PR1209565](#)
- The following log messages are displayed on the device: **L2ALM Trying peer/master connection, status 26.** [PR1317011](#)

User Firewall and Authentication

- User firewall has a command to fetch the user-group mapping from the active directory server. [PR1327633](#)

Unified Threat Management (UTM)

- The ISSU upgrade might fail because of the Packet Forwarding Engine generating a core file. [PR1328665](#)

Upgrade and Downgrade

- The command **show system firmware** displays the old firmware image. [PR1345314](#)

VLAN Infrastructure

- On SRX Series devices in transparent mode, the flowd process might stop when matching the destination MAC. [PR1355381](#)

VPNs

- The IRB interface does not support VPN. [PR1166714](#)
- Next hop tunnel binding (NHTB) is not installed occasionally during rekey for VPN using IKEv1. [PR1281833](#)
- IPsec traffic statistic counters return 32-bit values. [PR1301688](#)
- Auto Discovery VPN (ADVPN) tunnels might flap with the spoke error no response ready yet, leading to IKEv2 timeout. [PR1305451](#)
- On SRX Series devices, core files are observed under certain conditions with VPN and when NAT-T is enabled. [PR1308072](#)
- PKID syslog for key-pair deletion is required for conformance. [PR1308364](#)
- On SRX Series devices, ESP packet drops in IPsec VPN tunnels with NULL encryption algorithm configuration are observed. [PR1329368](#)
- SNMP for jnxIpSecTunMonVpnName does not work. [PR1330365](#)
- The kmd process might generate a core file when all the VPNs are down. [PR1336368](#)
- On SRX5400, SRX5600, and SRX5800 devices, the chassis cluster control link encryption does not work. [PR1347380](#)
- The kmd process might stop if multiple IKE gateways uses the same IKE policy. [PR1337903](#)
- All IPsec tunnels are in both active and inactive state. [PR1348767](#)
- S2S tunnels are not redistributed after IKE or IPsec are reactivated in a configuration. [PR1354440](#)

Resolved Issues: 17.4R1

Application Layer Gateways (ALGs)

- On SRX Series devices SIP packet might drop when SIP traffic performs destination NAT. [PR1268767](#)
- The pfed process stops and generates core files. [PR1292992](#)
- H323 ALG decode Q931 packet error was observed even after disabling H323 ALG. [PR1305598](#)
- HTTP ALG is listed within **show security match-policies**, when the HTTP ALG does not exist. [PR1308717](#)

Chassis Cluster

- Node 0 is going into db prompt after applying Layer 2 switching configuration and rebooting. [PR1228473](#)
- HA configuration synchronization monitoring does not work if **encrypt-configuration-files** is enabled. [PR1235628](#)
- The ISSU or ICU operation might fail if upgrade is initiated from Junos Space on multiple SRX clusters. [PR1279916](#)
- ALG traffic and other traffic with tcp-proxy gets stuck after back-to-back RG1 failover when using PPPoE on the reth interface. [PR1286547](#)
- Warning messages are incorrectly tagged as errors in the RPC response from the SRX Series device when you configure a change through NETCONF. [PR1286903](#)
- After software upgrade, the cluster goes into a brief split-brain state when rebooting RG0 on the secondary node. [PR1288819](#)
- In an SRX1500 cluster, if control-link-recovery is configured, ISSU might not complete successfully and the cluster will end up with different software releases. [PR1303948](#)
- IP monitoring on the secondary node shows unknown status after rebooting. [PR1307749](#)
- On SRX Series devices, the traffic logging impact issue after ISSU is fixed. [PR1284783](#)

Class of Service (CoS)

- on SRX devices, self-generated TCP session from RE destined to an lt-0/0/0.x nexthop is not established. [PR1286866](#)

Flow-Based and Packet-Based Processing

- The software-NH value increases and and causes a traffic outage. [PR1190301](#)
- SRX1500 devices might power-off unexpectedly because of incorrect device temperature readings which reportedly is a too high temperature, leading to an immediate proactive power-off of the device to protect the device from overheating. When this condition occurs, the following log message is shown in file `/var/log/hostlogs/lcmd.log`: `Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor`. [PR1241061](#)
- Duplicate hops or a higher than expected hop count is seen in L2 traceroute. [PR1243213](#)

- Configuring dpd results in timeouts for TCP encapsulation sessions. [PR1254875](#)
- A down interface in the **mirror-filter** command might cause a core file in certain situations. [PR1270724](#)
- Core files are seen on SRX1500 when J-Flow is enabled. [PR1271466](#)
- SRX320 with MPIM: IPv6 static route on dl0.0 is not active, so it cannot work for dial-on-demand. [PR1273532](#)
- Multicast traffic sent to the downstream interface in the destination MAC address is set to all zeros. [PR1276043](#)
- Output hangs while checking pki ca-certificate ca-profile-group details. [PR1276619](#)
- SRX1500 randomly stops forwarding traffic. [PR1277435](#)
- When using integrated user firewall, the useridd process might consume high CPU. [PR1280783](#)
- When executing operational commands for creating rescue configuration, some errors will be reported but the rescue configuration will still be created. [PR1280976](#)
- User firewall users are not assigned their roles. [PR1282744](#)
- Certain SCTP packets are dropped. [PR1285089](#)
- The pfed process stop and core files are generated by committing traceoptions configure. [PR1289972](#)
- More CPU threshold warnings are seen than in the previous releases. [PR1291506](#)
- CoS scheduler and shaping does not work on IRB interface. [PR1292187](#)
- Cryptographic weakness is seen on SRX300 line devices TPM Firmware (CVE-2017-10606) [PR1293114](#)
- The APN profile password is displayed in cleartext. [PR1295274](#)
- On SRX Series devices running the user firewall feature, under some conditions, flowd or useridd might generate core files. The Packet Forwarding Engine might get restarted, and RG1+ failover occurs. [PR1299494](#)
- SRX Series device fail to upgrade the Junos image when you use the unlink and partition options at the same time. [PR1299859](#)
- When you run the **show interfaces queue rethx** command, the output displays ingress queue information. [PR1309226](#)
- On SRX Series devices, the Stream Control Transmission Protocol (SCTP) packet has an incorrect SCTP checksum after the payload is translated by the device. [PR1310141](#)

Interfaces and Chassis

- On SRX1500 devices with SFP+-10G-CU3M DAC, 10-Gigabit Ethernet interface does not work. [PR1246725](#)
- On SRX1500, 10-Gigabit Ethernet interface might not come up between the SRX Series device and another type of device when using SFP+-10G-CU3M DAC. [PR1279182](#)

- Ping to VRRP (VIP) address failed when VRRP on vlan-tagging. This only affected IOC2 and IOC3 cards in SRX5000 line devices. SRX1500, SRX4100, and SRX4200 devices are not impacted. [PR1293808](#)
- RPM packets do not go through the LT interface under certain configurations. [PR1303445](#)

J-Web

- SRX Series devices cannot be upgraded with Junos image using J-Web. [PR1297362](#)
- Configuration upload using J-Web does not work. [PR1300766](#)
- In J-Web, when logical system adds a custom application, the applications 'any' are not present in **Logical System Configure > Security > Security Policy > Add Policy**. [PR1303260](#)
- J-Web removes the backslash character on the source identity object when the commit changes. [PR1304608](#)

Layer 2 Ethernet Services

- ARP issues are seen when using Layer 2 switching with the IRB interface. [PR1266450](#)
- On SRX1500 devices in an Ethernet switching mode, an IRB interface located in a custom routing instance is not reachable. [PR1234000](#)
- The **change no-dns-propagation** command should be changed to **no-dns-install**. [PR1284852](#)
- DHCPv6 prefix delegation does not start with the first available subnet [PR1295178](#)

Network Address Translation (NAT)

- On SRX Series devices, the periodic execution of the **show security zones detail** command causes the NSD process to fail in releasing unused memory, causing memory leak. [PR1269525](#)
- The proxy-arp does not work intermittently after RGO failover. [PR1289614](#)
- Commit check might allow a Source NAT pool without addresses to be committed, leading to flowd core file generation when the misconfigured pool is utilized by traffic. [PR1300019](#)
- Active source NAT causes an NSD error and the session closes. [PR1313144](#)

Network Management and Monitoring

- On the SRX340 device, one Routing Engine does not reply for the SNMP request after power-on or RGO failover in a cluster. [PR1240178](#)
- On SRX Series devices, when J-Flow is enabled for multicast traffic **extern nexthop** is installed during the multicast composite next hop. However, when you uninstall the composite next hop, it does not free the **extern nexthop**, which results in the jtree memory leak. [PR1276133](#)
- The mib2d process might crash when polling the OID ifStackStatus.0 after a logical interface of lo0 is deleted. [PR1286351](#)
- The **show arp no-resolve interface X** command for nonexistent interface X is showing all unrelated static ARP entries. [PR1299619](#)

Platform and Infrastructure

- SRX300 line devices reboot when Juniper RE-USB-4G-S (yellow or orange) USB is inserted. [PR1214125](#)
- The flowd process might crash during route update. [PR1249254](#)
- Unexpected behavior with IP monitoring is seen. [PR1263078](#)
- The TTL (Time To Live) of some Z-mode packets is reduced to zero incorrectly, if IOC2 or IOC3 interface is configured as HA fabric port. [PR1270770](#)
- DNS cache does not get populated in multiple virtual router (VR) environments. [PR1275792](#)
- Memory leak occurs on SRX Series devices chassis cluster when em0 or em1 interface is down. [PR1277136](#)
- On SRX5000 line devices, under a heavy flood of IPv6 Neighbor Discovery Protocol (NDP) packets, some incoming IPv6 neighbor advertisements (NA) might be dropped because of a queue being full. This issue has been resolved by using a different queue for IPv6 NA packets. [PR1293673](#)
- XLP lost heartbeat (SPU hang) is not detected in a timely manner by hardware monitoring. [PR1300804](#)

Routing Policy and Firewall Filters

- Secured e-mail application is not available. [PR1273725](#)
- On SRX Series devices, the DNS configured in the address-book fails to resolve the IP address, if the case (uppercase or lowercase) in the DNS query and the DNS response do not match. [PR1304706](#)
- The NSD process might crash when replacing the name of a logical-system. [PR1307876](#)

System Logging

- The logs from syslog **RT_FLOW: FLOW_REASSEMBLE_SUCCEED: Packet merged** might cause high CPU usage on the Routing Engine. [PR1278333](#)

Unified Threat Management (UTM)

- The Packet Forwarding Engine CPU utilization is high when using the UTM antivirus feature. [PR1282719](#)

VPNs

- The st0 global counter statistics do not increment. [PR1171958](#)
- The second client is disconnected when the assigned IP address is changed in the access profile for the first client. [PR1246131](#)
- IPsec traffic through tunnel fails without configuring the authentication algorithm under the IPsec proposal on the SRX1500; however, it works on the SRX5600. [PR1285284](#)

SEE ALSO

New and Changed Features	 307
Changes in Behavior and Syntax	 319
Known Behavior	 321
Known Issues	 325
Documentation Updates	 342
Migration, Upgrade, and Downgrade Instructions	 342
Product Compatibility	 346

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R2 for the SRX Series documentation.

SEE ALSO

New and Changed Features	 307
Changes in Behavior and Syntax	 319
Known Behavior	 321
Known Issues	 325
Resolved Issues	 328
Migration, Upgrade, and Downgrade Instructions	 342
Product Compatibility	 346

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Scripts for Address Book Configuration](#) | 343

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

Upgrade and Downgrade Scripts for Address Book Configuration

IN THIS SECTION

- [About Upgrade and Downgrade Scripts | 343](#)
- [Running Upgrade and Downgrade Scripts | 344](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 345](#)

Beginning with Junos OS Release 12.1, you can configure address books under the **[security]** hierarchy and attach security zones to them (zone-attached configuration). In Junos OS Release 11.1 and earlier, address books were defined under the **[security zones]** hierarchy (zone-defined configuration).

You can either define all address books under the **[security]** hierarchy in a zone-attached configuration format or under the **[security zones]** hierarchy in a zone-defined configuration format; the CLI displays an error and fails to commit the configuration if you configure both configuration formats on one system.

Juniper Networks provides Junos operation scripts that allow you to work in either of the address book configuration formats (see [Figure 1 on page 344](#)).

About Upgrade and Downgrade Scripts

After downloading Junos OS Release 12.1, you have the following options for configuring the address book feature:

- **Use the default address book configuration**—You can configure address books using the zone-defined configuration format, which is available by default. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.
- **Use the upgrade script**—You can run the upgrade script available on the Juniper Networks support site to configure address books using the new zone-attached configuration format. When upgrading, the system uses the zone names to create address books. For example, addresses in the trust zone are created in an address book named **trust-address-book** and are attached to the trust zone. IP prefixes used in NAT rules remain unaffected.

After upgrading to the zone-attached address book configuration:


- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.

For information on how to configure zone-attached address books, see the Junos OS Release 12.1 documentation.

- **Use the downgrade script**—After upgrading to the zone-attached configuration, if you want to revert to the zone-defined configuration, use the downgrade script available on the Juniper Networks support site. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.

NOTE: Before running the downgrade script, make sure to revert any configuration that uses addresses from the global address book.

Figure 1: Upgrade and Downgrade Scripts for Address Books



Running Upgrade and Downgrade Scripts

The following restrictions apply to the address book upgrade and downgrade scripts:

- The scripts cannot run unless the configuration on your system has been committed. Thus, if the zone-defined address book and zone-attached address book configurations are present on your system at the same time, the scripts will not run.
- The scripts cannot run when the global address book exists on your system.
- If you upgrade your device to Junos OS Release 12.1 and configure logical systems, the master logical system retains any previously configured zone-defined address book configuration. The master

administrator can run the address book upgrade script to convert the existing zone-defined configuration to the zone-attached configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.

NOTE: You cannot run the downgrade script on logical systems.

For information about implementing and executing Junos operation scripts, see the *Junos OS Configuration and Operations Automation Guide*.

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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after.

For example, Junos OS Releases 12.3X48, 15.1X49, 17.3 and 17.4 are EEOL releases. You can upgrade from Junos OS Release 15.1X49 to Release 17.3 or from Junos OS Release 15.1X49 to Release 17.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

Upgrade from Junos OS Release 17.4 to successive Junos OS Release, is supported. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

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

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

SEE ALSO

[New and Changed Features | 307](#)

[Changes in Behavior and Syntax | 319](#)

[Known Behavior | 321](#)

[Known Issues | 325](#)

[Resolved Issues | 328](#)

[Documentation Updates | 342](#)

[Product Compatibility | 346](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 346](#)

Hardware Compatibility

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

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

SEE ALSO

[New and Changed Features | 307](#)

[Changes in Behavior and Syntax | 319](#)

[Known Behavior | 321](#)

[Known Issues | 325](#)

[Resolved Issues | 328](#)

[Documentation Updates | 342](#)

[Migration, Upgrade, and Downgrade Instructions | 342](#)

Upgrading Using ISSU

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

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

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

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

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information

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](#).


To access Software Release Notifications for Junos OS Service Releases, visit our Knowledge Center at <https://support.juniper.net/support/>. You'll need to log in to your Juniper Account. From the Knowledge Center, search by the specific release number, for example 17.4R1-S2. Use the Software Release Notifications to download software, and learn about known and resolved issues for specific service releases.

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://apps.juniper.net/feature-explorer/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

Revision History

17 September 2021—Revision 19, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 August 2021—Revision 18, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 17, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 February 2020—Revision 16, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

5 September 2019—Revision 15, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 July 2019—Revision 14, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 June 2019—Revision 13, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 May 2019—Revision 12, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 March 2019—Revision 11, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 March 2019—Revision 10, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 February 2019—Revision 9, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 February 2019—Revision 8, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 January 2019—Revision 7, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 December 2018—Revision 6, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 November 2018—Revision 5, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 October 2018—Revision 4, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 October 2018—Revision 3, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 September 2018—Revision 2, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2018—Revision 2, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 August 2018—Revision 1, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 February 2018—Revision 5, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 January 2018—Revision 4, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 January 2018—Revision 3, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 December 2017—Revision 2, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 December 2017—Revision 1, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

Copyright © 2019 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

