

Release Notes

Published
2021-09-17

Junos[®] OS 17.4R1 Release Notes

SUPPORTED ON

- ACX Series, EX Series, Junos Fusion Enterprise, Junos Fusion Data Center, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series

HARDWARE HIGHLIGHTS

- Support for MX204 Universal Routing Platform
- Support for PTX10016 Packet Transport Router

SOFTWARE HIGHLIGHTS

- Support for RADIUS reauthentication of DHCPv4 and DHCPv6 clients (MX Series)
- Support for enhancements to composite next-hops (MX Series)
- Support for SPRING for EVPN (MX series)
- Support for PTP over Ethernet, Hybrid mode, and G.8275.1 profile (MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E)
- Support for broadband edge (BBE) telemetry sensors (MX Series)
- Support for Junos Telemetry Interface for virtual MX Series routers (vMX)
- Support for static adjacency segment identifier for IS-IS (MX Series)
- Support for segment routing policy for traffic engineering (MX Series)
- Support for Large Scale Packet-Forwarding Features (PTX10000)
- Support for topology Independent Loop-Free Alternate using SPRING for IS-IS (PTX Series)

- Support for EVPN pure type-5 route (QFX5110)
- Support for resilient hashing for LAGs (QFX10000)
- Support for enterprise profile for Precision Time Protocol (PTP) (QFX10002)
- Support for Precision Time Protocol (PTP) transparent clock (QFX5200)
- Support for BGP MPLS-based Ethernet VPN (QFX10000)
- Support for PCEP (QFX5100, QFX5110, QFX5200)
- Support for MLD snooping versions 1 and 2 (QFX5100 and Virtual Chassis)
- Support for multicast-only fast reroute (MoFRR) (QFX5100, QFX5110, and QFX5200)
- Support for IPFIX templates for flow aggregation (QFX10008 and QFX10016)
- Support for PPPoE (SRX Series and vSRX)
- Support for IPv6 address for PPP AutoVPN networks using traffic selectors (SRX Series and vSRX)
- Support for UDP flood screen whitelist (SRX Series)
- Support for increased number of IKE security associations (SRX Series)

Release Notes: Junos[®] OS Release 17.4R1 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

17 September 2021

Contents	Introduction 13
	Junos OS Release Notes for ACX Series 13
	New and Changed Features 14
	Management 14
	Timing and Synchronization 14
	Changes in Behavior and Syntax 15
	Management 16
	Security 16
	Known Behavior 16
	Known Issues 17
	Interfaces and Chassis 17
	Resolved Issues 18
	Documentation Updates 18
	Migration, Upgrade, and Downgrade Instructions 19
	Upgrade and Downgrade Support Policy for Junos OS Releases 19

Product Compatibility | 20

Hardware Compatibility | 20

Junos OS Release Notes for EX Series Switches | 21

New and Changed Features | 22

Hardware | 23

Authentication, Authorization and Accounting (AAA) | 23

EVPNs | 23

Junos OS XML API and Scripting | 25

Layer 2 Features | 25

Management | 25

Multicast | 27

Routing Protocols | 27

Software Installation and Upgrade | 28

Changes in Behavior and Syntax | 29

Management | 29

Network Management and Monitoring | 29

Security | 30

Software Licensing | 30

Known Behavior | 31

Authentication, Authorization, and Accounting (AAA) | 31

High Availability (HA) and Resiliency | 31

Interfaces and Chassis | 32

Junos Fusion Enterprise | 32

Management | 32

Platform and Infrastructure | 32

Known Issues | 33

Authentication, Authorization, and Accounting (AAA) | 34

Hardware | 34

High Availability (HA) and Resiliency | 34

Infrastructure | 34

Interfaces and Chassis | 35

Junos Fusion Enterprise | 35

Layer 2 Features | 35

Management | 35

MPLS	35
Network Management and Monitoring	35
Operation, Administration, and Maintenance	36
Platform and Infrastructure	36
Spanning-Tree Protocols	36
Virtual Chassis	36
Resolved Issues	37
Resolved Issues: 17.4R1	37
Documentation Updates	39
Migration, Upgrade, and Downgrade Instructions	40
Upgrade and Downgrade Support Policy for Junos OS Releases	40
Product Compatibility	41
Hardware Compatibility	41
Junos OS Release Notes for Junos Fusion Data Center	42
New and Changed Features	42
Changes in Behavior and Syntax	43
Known Behavior	43
Known Issues	44
Resolved Issues	44
Documentation Updates	45
Migration, Upgrade, and Downgrade Instructions	45
Basic Procedure for Upgrading an Aggregation Device	46
Preparing the Switch for Satellite Device Conversion	47
Autoconverting a Switch into a Satellite Device	50
Manually Converting a Switch into a Satellite Device	53
Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology	55
Configuring Satellite Device Upgrade Groups	56
Converting a Satellite Device to a Standalone Device	58
Upgrade and Downgrade Support Policy for Junos OS Releases	60
Downgrading from Release 17.4R1	60
Product Compatibility	61
Hardware Compatibility	61

Junos OS Release Notes for Junos Fusion Enterprise | 63

New and Changed Features | 63

Junos Fusion Enterprise | 64

Changes in Behavior and Syntax | 65

Junos Fusion Enterprise | 65

Known Behavior | 66

Junos Fusion Enterprise | 66

Known Issues | 67

Resolved Issues | 67

Resolved Issues: 17.4R1 | 68

Documentation Updates | 68

Migration, Upgrade, and Downgrade Instructions | 69

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 69

Upgrading an Aggregation Device with Redundant Routing Engines | 71

Preparing the Switch for Satellite Device Conversion | 72

Converting a Satellite Device to a Standalone Switch | 73

Upgrade and Downgrade Support Policy for Junos OS Releases | 75

Downgrading from Release 17.4 | 76

Product Compatibility | 77

Hardware and Software Compatibility | 77

Hardware Compatibility Tool | 77

Junos OS Release Notes for Junos Fusion Provider Edge | 78

New and Changed Features | 78

Hardware | 79

Changes in Behavior and Syntax | 79

Known Behavior | 80

Known Issues | 80

Resolved Issues | 81

Resolved Issues: 17.4R1 | 81

Documentation Updates | 81

Migration, Upgrade, and Downgrade Instructions | 82

Basic Procedure for Upgrading an Aggregation Device | 82

Upgrading an Aggregation Device with Redundant Routing Engines | 85

Preparing the Switch for Satellite Device Conversion | 85

Converting a Satellite Device to a Standalone Device	87
Upgrading an Aggregation Device	89
Upgrade and Downgrade Support Policy for Junos OS Releases	89
Downgrading from Release 17.4	89
Product Compatibility	90
Hardware Compatibility	90
Junos OS Release Notes for MX Series 5G Universal Routing Platforms	91
New and Changed Features	92
Hardware	93
Authentication, Authorization, and Accounting (AAA) (RADIUS)	93
Class of Service (CoS)	94
Dynamic Host Configuration Protocol (DHCP)	94
EVPNs	96
General Routing	98
High Availability (HA) and Resiliency	99
Interfaces and Chassis	99
Junos OS XML API and Scripting	101
Layer 2 Features	101
Logical Systems	102
Management	103
MPLS	107
Operation, Administration, and Maintenance (OAM)	108
Routing Protocols	109
Services Applications	112
Software Defined Networking (SDN)	115
Software Installation and Upgrade	117
Subscriber Management and Services	118
System Logging	122
User interface and Configuration	122
VPNs	124
Changes in Behavior and Syntax	124
Interfaces and Chassis	125
Management	126
MPLS	126

Multicast	127
Network Management and Monitoring	127
Routing Protocols	128
Security	128
Services Applications	129
Software Licensing	129
Subscriber Management and Services	129
Known Behavior	130
General Routing	131
Interfaces and Chassis	132
Layer 2 Ethernet Services	133
Subscriber Management and Services	133
Known Issues	133
Class of Service (CoS)	134
EVPNs	134
Forwarding and Sampling	135
General Routing	135
Infrastructure	138
Interfaces and Chassis	138
MPLS	139
Network Management and Monitoring	139
Platform and Infrastructure	139
Routing Protocols	140
Services Applications	142
VPNs	142
Resolved Issues	143
Resolved Issues: 17.4R1	143
Documentation Updates	158
Subscriber Management Provisioning guide	158
Migration, Upgrade, and Downgrade Instructions	159
Basic Procedure for Upgrading to Release 17.4	160
Procedure to Upgrade to FreeBSD 11.x-Based Junos OS	160
Procedure to Upgrade to FreeBSD 6.x-Based Junos OS	162
Upgrade and Downgrade Support Policy for Junos OS Releases	164

Upgrading a Router with Redundant Routing Engines	165
Downgrading from Release 17.4	165
Product Compatibility	166
Hardware Compatibility	166
Junos OS Release Notes for NFX Series	167
New and Changed Features	167
Changes in Behavior and Syntax	168
Known Behavior	168
Juniper Device Manager	169
Known Issues	169
Infrastructure	170
IPsec	170
Juniper Device Manager	170
Junos Control Plane	172
vSRX	173
Resolved Issues	173
Resolved Issues: 17.2R1	174
Documentation Updates	174
Migration, Upgrade, and Downgrade Instructions	175
Upgrade and Downgrade Support Policy for Junos OS Releases	175
Basic Procedure for Upgrading to Release 17.4	175
Product Compatibility	178
Hardware Compatibility	178
Junos OS Release Notes for PTX Series Packet Transport Routers	180
New and Changed Features	181
Hardware	181
High Availability (HA) and Resiliency	182
Interfaces and Chassis	183
IPv6	184
Junos OS XML API and Scripting	184
Layer 2 Features	184
Layer 3 Features	184
Management	185
MPLS	188

Routing Protocols	190
Security	192
Services Applications	192
Software Installation and Upgrade	193
Changes in Behavior and Syntax	194
Class of Service (CoS)	194
Interfaces and Chassis	194
Management	196
MPLS	196
Multicast	197
Network Management and Monitoring	197
Security	198
Software Licensing	198
Known Behavior	199
General Routing	199
Network Management and Monitoring	199
Known Issues	200
General Routing	200
Interfaces and Chassis	202
Multiprotocol Label Switching (MPLS)	202
Resolved Issues	203
Resolved Issues: 17.4R1	203
Documentation Updates	206
Migration, Upgrade, and Downgrade Instructions	206
Upgrade and Downgrade Support Policy for Junos OS Releases	206
Upgrading a Router with Redundant Routing Engines	207
Basic Procedure for Upgrading to Release 17.4	207
Product Compatibility	211
Hardware Compatibility	211
Junos OS Release Notes for the QFX Series	212
New and Changed Features	212
Hardware	214
Class of Service (CoS)	214
EVPNs	214

General Routing	216
Interfaces and Chassis	216
Junos OS XML API and Scripting	217
Management	217
Multicast	218
MPLS	219
Network Management and Monitoring	221
Port Security	222
Routing Protocols	222
Services Applications	223
Software Installation and Upgrade	224
Virtual Chassis	224
Changes in Behavior and Syntax	225
Class of Service	226
General Routing	226
Management	226
MPLS	226
Network Management and Monitoring	227
Security	229
Software Licensing	229
Virtual Chassis and Virtual Chassis Fabric (VCF)	229
Known Behavior	230
Hardware	230
EVPNs	230
High Availability (HA) and Resiliency	231
Interfaces and Chassis	231
Junos Fusion Provider Edge	232
Layer 2 Features	232
Routing Protocols	232
Storage and Fibre Channel	232
Known Issues	233
EVPNs	234
Hardware	235
High Availability (HA) and Resiliency	235

Infrastructure	235
Interfaces and Chassis	236
Junos Fusion Provider Edge	236
Layer 2 Features	238
Management	238
MPLS	238
Network Management and Monitoring	238
Port Security	239
Software Installation and Upgrade	239
Subscriber Management and Services	239
Virtual Chassis and Virtual Chassis Fabric (VCF)	239
VLAN Infrastructure	239
Resolved Issues	240
Resolved Issues: 17.4R1	240
Documentation Updates	244
Migration, Upgrade, and Downgrade Instructions	245
Upgrading Software on QFX Series Switches	245
Installing the Software on QFX10002 Switches	248
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	248
Installing the Software on QFX10008 and QFX10016 Switches	250
Performing a Unified ISSU	254
Preparing the Switch for Software Installation	255
Upgrading the Software Using Unified ISSU	255
Upgrade and Downgrade Support Policy for Junos OS Releases	258
Product Compatibility	258
Hardware Compatibility	259
Junos OS Release Notes for SRX Series	260
New and Changed Features	260
Release 17.4R1-S1 New and Changed Features	262
Release 17.4R1 New and Changed Features	264
Changes in Behavior and Syntax	273
Chassis Cluster	274
Installation and Upgrade	274

MPLS	274
Management	274
NAT	274
Security	275
Known Behavior	276
Authentication and Access	276
J-Web	276
Layer 2 Ethernet Services	277
Platform and Infrastructure	277
Software Installation and Upgrade	278
UTM	278
VPNs	278
Known Issues	279
Chassis Clustering	279
Flow-based and Packet-based Processing	279
Interfaces	280
J-Web	280
Routing Protocols	280
VPNs	280
Resolved Issues	281
Application Layer Gateways (ALGs)	282
Chassis Cluster	282
Class of Service (CoS)	282
Flow-Based and Packet-Based Processing	282
Interfaces and Chassis	284
J-Web	284
Layer 2 Ethernet Services	284
Network Address Translation (NAT)	284
Network Management and Monitoring	285
Platform and Infrastructure	285
Routing Policy and Firewall Filters	285
System Logging	286
Unified Threat Management (UTM)	286
VPNs	286

Documentation Updates	286
Migration, Upgrade, and Downgrade Instructions	287
Upgrade and Downgrade Scripts for Address Book Configuration	287
Product Compatibility	291
Hardware Compatibility	291
Upgrading Using Unified ISSU	292
Compliance Advisor	292
Finding More Information	292
Documentation Feedback	292
Requesting Technical Support	294
Self-Help Online Tools and Resources	294
Opening a Case with JTAC	295
Revision History	295

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 17.4R1 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 | 14
- Changes in Behavior and Syntax | 15
- Known Behavior | 16
- Known Issues | 17
- Resolved Issues | 18
- Documentation Updates | 18
- Migration, Upgrade, and Downgrade Instructions | 19
- Product Compatibility | 20

These release notes accompany Junos OS Release 17.4R1 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Management | 14](#)
- [Timing and Synchronization | 14](#)

This section describes the new features or enhancements to existing features in Junos OS Release 17.4R1 for ACX Series Universal Access Routers.

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 [Enterprise Profile for the Precision Time Protocol](#).]

SEE ALSO

Changes in Behavior and Syntax	 15
Known Behavior	 16
Documentation Updates	 18
Known Issues	 17
Resolved Issues	 18
Migration, Upgrade, and Downgrade Instructions	 19
Product Compatibility	 20

Changes in Behavior and Syntax

IN THIS SECTION

- [Management](#) | [16](#)
- [Security](#) | [16](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 17.4R1 for the ACX Series Universal Access 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.

SEE ALSO

New and Changed Features 14
Known Behavior 16
Documentation Updates 18
Known Issues 17
Resolved Issues 18
Migration, Upgrade, and Downgrade Instructions 19
Product Compatibility 20

Known Behavior

There are no known limitations in Junos OS Release 17.4R1 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	14
Changes in Behavior and Syntax	15
Documentation Updates	18
Known Issues	17
Resolved Issues	18
Migration, Upgrade, and Downgrade Instructions	19
Product Compatibility	20

Known Issues

IN THIS SECTION

- [Interfaces and Chassis | 17](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R1 for the ACX Series Universal Access 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	 14
Changes in Behavior and Syntax	 15
Known Behavior	 16
Documentation Updates	 18
Resolved Issues	 18
Migration, Upgrade, and Downgrade Instructions	 19
Product Compatibility	 20

Resolved Issues

There are no fixed issues in Junos OS 17.4R1 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	 14
Changes in Behavior and Syntax	 15
Known Behavior	 16
Documentation Updates	 18
Known Issues	 17
Migration, Upgrade, and Downgrade Instructions	 19
Product Compatibility	 20

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 for the ACX Series documentation.

SEE ALSO

New and Changed Features	 14
--	----------------------

Changes in Behavior and Syntax 15
Known Behavior 16
Known Issues 17
Resolved Issues 18
Migration, Upgrade, and Downgrade Instructions 19
Product Compatibility 20

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 19

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Universal Access 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 | 14](#)

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

[Known Behavior | 16](#)

[Documentation Updates | 18](#)

[Known Issues | 17](#)

[Resolved Issues | 18](#)

[Product Compatibility | 20](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 20](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 14
Changes in Behavior and Syntax 15
Known Behavior 16
Documentation Updates 18
Known Issues 17
Resolved Issues 18
Migration, Upgrade, and Downgrade Instructions 19

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- New and Changed Features | 22
- Changes in Behavior and Syntax | 29
- Known Behavior | 31
- Known Issues | 33
- Resolved Issues | 37
- Documentation Updates | 39
- Migration, Upgrade, and Downgrade Instructions | 40
- Product Compatibility | 41

These release notes accompany Junos OS Release 17.4R1 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Hardware | 23](#)
- [Authentication, Authorization and Accounting \(AAA\) | 23](#)
- [EVPNs | 23](#)
- [Junos OS XML API and Scripting | 25](#)
- [Layer 2 Features | 25](#)
- [Management | 25](#)
- [Multicast | 27](#)
- [Routing Protocols | 27](#)
- [Software Installation and Upgrade | 28](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R1 for the EX Series.

NOTE: The following EX Series switches are supported in Release 17.4R1: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.4R1, 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](#).

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.](#)]

•

NOTE: This feature is documented but not supported in Junos OS Release 17.4R1.

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 17.4R1, 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 both local customer edge (CE) devices only. Proxy ARP and ARP suppression are not supported on remote CE devices.

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](#).]

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

- **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 | 29](#)

[Known Behavior | 31](#)

[Known Issues | 33](#)

[Resolved Issues | 37](#)

[Documentation Updates | 39](#)

[Migration, Upgrade, and Downgrade Instructions | 40](#)

[Product Compatibility | 41](#)

Changes in Behavior and Syntax

IN THIS SECTION

- Management | 29
- Network Management and Monitoring | 29
- Security | 30
- Software Licensing | 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 17.4R1 for the EX Series.

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](#).]

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)**—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](#).]

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 (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](#).]

SEE ALSO

[New and Changed Features | 22](#)

[Known Behavior | 31](#)

[Known Issues | 33](#)

[Resolved Issues | 37](#)

[Documentation Updates | 39](#)

[Migration, Upgrade, and Downgrade Instructions | 40](#)

Known Behavior

IN THIS SECTION

- [Authentication, Authorization, and Accounting \(AAA\)](#) | 31
- [High Availability \(HA\) and Resiliency](#) | 31
- [Interfaces and Chassis](#) | 32
- [Junos Fusion Enterprise](#) | 32
- [Management](#) | 32
- [Platform and Infrastructure](#) | 32

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 for the EX Series.

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

Authentication, Authorization, and Accounting (AAA)

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

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)
- On an EX4300 Virtual Chassis, when you perform an NSSU, there might be more than 5 seconds of traffic loss for multicast traffic. [PR1125155](#)

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](#)

Junos Fusion Enterprise

- On a Junos Fusion Enterprise, **show ethernet-switching table** takes a few minutes to show entries when an extended port receives with MAC count set to 150K. [PR1117567](#)
- 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](#)

Management

- The **request system halt** command is not supported on EX4300-MP due to HW limitations. The **request system power-off** command is supported instead. [PR1271304](#)

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](#)

SEE ALSO

[New and Changed Features | 22](#)

[Changes in Behavior and Syntax | 29](#)

[Known Issues | 33](#)

[Resolved Issues | 37](#)

[Documentation Updates | 39](#)

[Migration, Upgrade, and Downgrade Instructions | 40](#)

[Product Compatibility | 41](#)

Known Issues

IN THIS SECTION

- Authentication, Authorization, and Accounting (AAA) | 34
- Hardware | 34
- High Availability (HA) and Resiliency | 34
- Infrastructure | 34
- Interfaces and Chassis | 35
- Junos Fusion Enterprise | 35
- Layer 2 Features | 35
- Management | 35
- MPLS | 35
- Network Management and Monitoring | 35
- Operation, Administration, and Maintenance | 36
- Platform and Infrastructure | 36
- Spanning-Tree Protocols | 36
- Virtual Chassis | 36

This section lists the known issues in hardware and software in Junos OS Release 17.4R1 for the EX Series.

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

Authentication, Authorization, and Accounting (AAA)

- On EX Series platforms, dot1x might stop authentication if continuous dot1x clients reauthentication requests cannot be processed. As a workaround, restart the dot1x process, restart the authd process, increase the EX Series platform/dot1x/authd (daemon) memory sizes, or reduce the system load by increasing the dot1x protocol reauthentication value. [PR1300050](#)

Hardware

- 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](#)

High Availability (HA) and Resiliency

- The issue will be visible during NSSU if two adjacent members have LACP links in one AE— for example, xe-2/0/1 (on member 2) and xe-3/0/1 (on member 3) are part of one AE, then traffic drop will be observed during NSSU, when member 3 will be rebooted. [PR1311977](#)

Infrastructure

- A memory leak issue occurs when DHCP snooping is enabled. [PR1303241](#)
- In JDM, (running on secondary server) jdmd daemon may core if GNF add-image is aborted by pressing CTRL-C. [PR1321803](#)

Interfaces and Chassis

- For "mge" (Multi-rate) interfaces, in the interface speed configuration, "auto" option will not be present. By default (without any configuration or in factory default mode), "mge" interface advertises all the speeds if AN is enabled. If a particular speed is set, that speed along with all the lower speeds will be advertised if AN is enabled. [PR1282915](#)

Junos Fusion Enterprise

- 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](#)

Layer 2 Features

- ERP route update fails during the addition of a new member to the ERP-configured VLAN. [PR1301595](#)

Management

- If there is need to change the MAC table size to higher value than default, then CPU spike is noticed when show command is executed and it take more time than normal condition to display the required details. There is no impact to the MAC learning and traffic. [PR1322041](#)

MPLS

- On EX Series switches, unified ISSU is not supported with MPLS configuration. [PR1264786](#)

Network Management and Monitoring

- EX Series switches configured with SFLOW and MAC RADIUS might incorrectly send MAC authentication requests for transit DHCPv6 traffic that is picked up by the SFLOW agent. [PR1298646](#)
- MACsec issue: the **show security macsec statistics** command does not show expected results. Statistics are incorrectly cleared for each physical interface once per second. [PR1283544](#)

Operation, Administration, and Maintenance

- Configuration statements that were allowed in Junos OS Release 12.3 are now invalid in Junos OS Release 14.1X53 and 15.1. As a result, when you upgrade an EX Series switch from Junos OS Release 12.3 to 14.1X53 or 15.1R1, the switch might lose its configuration and run in a line-card mode or go to "amnesiac" mode. [PR1281947](#)

Platform and Infrastructure

- 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](#)

Spanning-Tree Protocols

- When both RSTP and VSTP are enabled in the system and the VoIP interface data VLAN is not set under VSTP, the interface will be blocked for the data VLAN. [PR1306699](#)

Virtual Chassis

- In certain RTG configurations, packet drops are seen during the failover/switchover of the master Routing Engine in Junos OS Release 15.1R1 with GRES and NSR enabled. This issue occurred because of the delay in ARP update during the failover/switchover of the master Routing Engine. It is fixed in Junos OS Release 15.1.R7. [PR1278214](#)

SEE ALSO

[New and Changed Features | 22](#)

[Changes in Behavior and Syntax | 29](#)

[Known Behavior | 31](#)

[Resolved Issues | 37](#)

[Documentation Updates | 39](#)

[Migration, Upgrade, and Downgrade Instructions | 40](#)

[Product Compatibility | 41](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R1](#) | 37

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: 17.4R1

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 22
Changes in Behavior and Syntax 29
Known Behavior 31
Known Issues 33
Documentation Updates 39
Migration, Upgrade, and Downgrade Instructions 40
Product Compatibility 41

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 for the EX Series switches documentation.

SEE ALSO

New and Changed Features 22
Changes in Behavior and Syntax 29
Known Behavior 31
Known Issues 33
Resolved Issues 37

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 40

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 37
Documentation Updates 39
Product Compatibility 41

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 41

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 22
Changes in Behavior and Syntax 29
Known Behavior 31
Known Issues 33
Resolved Issues 37
Documentation Updates 39

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

- [New and Changed Features | 42](#)
- [Changes in Behavior and Syntax | 43](#)
- [Known Behavior | 43](#)
- [Known Issues | 44](#)
- [Resolved Issues | 44](#)
- [Documentation Updates | 45](#)
- [Migration, Upgrade, and Downgrade Instructions | 45](#)
- [Product Compatibility | 61](#)

These release notes accompany Junos OS Release 17.4R1 for the Junos Fusion Data Center. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

There are no new features in Junos OS Release 17.4R1 for Junos Fusion Data Center.

SEE ALSO

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

Documentation Updates	45
Migration, Upgrade, and Downgrade Instructions	45
Product Compatibility	61

Changes in Behavior and Syntax

There are no changes in behavior and syntax for Junos Fusion Data Center in Junos OS Release 17.4R1.

SEE ALSO

New and Changed Features	42
Known Behavior	43
Known Issues	44
Resolved Issues	44
Documentation Updates	45
Migration, Upgrade, and Downgrade Instructions	45
Product Compatibility	61

Known Behavior

There are no known limitations in Junos OS Release 17.4R1 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	42
Changes in Behavior and Syntax	43
Known Issues	44
Resolved Issues	44
Documentation Updates	45
Migration, Upgrade, and Downgrade Instructions	45

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.4R1 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	42
Changes in Behavior and Syntax	43
Known Behavior	43
Resolved Issues	44
Documentation Updates	45
Migration, Upgrade, and Downgrade Instructions	45
Product Compatibility	61

Resolved Issues

There are no fixed issues in Junos OS 17.4R1 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	42
Changes in Behavior and Syntax	43
Known Behavior	43
Known Issues	44
Documentation Updates	45

[Migration, Upgrade, and Downgrade Instructions | 45](#)

[Product Compatibility | 61](#)

Documentation Updates

This section lists the errata or changes in Junos OS Release 17.4R1 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 | 42](#)

[Changes in Behavior and Syntax | 43](#)

[Known Behavior | 43](#)

[Known Issues | 44](#)

[Resolved Issues | 44](#)

[Migration, Upgrade, and Downgrade Instructions | 45](#)

[Product Compatibility | 61](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 46](#)
- [Preparing the Switch for Satellite Device Conversion | 47](#)
- [Autoconverting a Switch into a Satellite Device | 50](#)
- [Manually Converting a Switch into a Satellite Device | 53](#)
- [Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 55](#)
- [Configuring Satellite Device Upgrade Groups | 56](#)
- [Converting a Satellite Device to a Standalone Device | 58](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 60](#)
- [Downgrading from Release 17.4R1 | 60](#)

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, replacing *n* with the spin number:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

NOTE: The QFX5100-48SH and QFX5100-48TH switch models are shipped from the factory with satellite device software. You cannot convert these switches to become standalone devices.

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 software image for your platform, using the following guidelines:
 - If the satellite device is a EX4300 switch, you install a standard, signed **jinstall** version of Junos OS.
 - If the satellite device is a QFX5100 switch that can be converted to a standalone device, you must install a Preboot eXecution Environment (PXE) version of Junos OS. The PXE version of Junos OS software supports the same feature set as the other Junos OS software packages for a release, but is specially engineered to install Junos OS onto a device running satellite software. The PXE Junos OS package name uses the format **install-media-pxe-qfx-5-version-domestic-signed.tgz**.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.
7. Copy the software to the routing platform or to your internal software distribution site.
8. Remove the satellite device from the automatic satellite conversion configuration.

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

by entering the **show** command at the [edit chassis satellite-management auto-satellite-conversion] hierarchy level.

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

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

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

9. Commit the configuration.

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

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

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

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

```
[edit]
user@aggregation-device> request chassis satellite install
/var/tmp/install-media-pxe-qfx-5-14.1X53D43.7.tgz fpc-slot 101
```

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

```
[edit]
user@aggregation-device> request chassis satellite install
/var/tmp/jinstall-ex-4300-14.1X53D43.7-domestic-signed.tgz fpc-slot 101
```

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

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 14.1, 14.2, 15.1, and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or even from Junos OS Release 14.1 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Downgrading from Release 17.4R1

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 42
Changes in Behavior and Syntax 43
Known Behavior 43
Known Issues 44
Resolved Issues 44
Documentation Updates 45
Product Compatibility 61

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 61

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 42
Changes in Behavior and Syntax 43
Known Behavior 43
Known Issues 44
Resolved Issues 44

Documentation Updates | 45

Migration, Upgrade, and Downgrade Instructions | 45

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 | 67
- Documentation Updates | 68
- Migration, Upgrade, and Downgrade Instructions | 69
- Product Compatibility | 77

These release notes accompany Junos OS Release 17.4R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION

- Junos Fusion Enterprise | 64

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R1 for Junos Fusion Enterprise.

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

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	 67
Documentation Updates	 68
Migration, Upgrade, and Downgrade Instructions	 69
Product Compatibility	 77

Changes in Behavior and Syntax

IN THIS SECTION

- [Junos Fusion Enterprise](#) | [65](#)

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 17.4R1 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	67
Documentation Updates	68
Migration, Upgrade, and Downgrade Instructions	69
Product Compatibility	77

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 17.4R1 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 Enterprise, when the satellite devices of a cluster are rebooted, the output of the CLI command **show chassis satellite** shows the Port State of the cascade ports as **Present**. [PR1175834](#)

SEE ALSO

New and Changed Features	63
Changes in Behavior and Syntax	65
Known Issues	67
Resolved Issues	67
Documentation Updates	68
Migration, Upgrade, and Downgrade Instructions	69
Product Compatibility	77

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.4R1 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	 67
Documentation Updates	 68
Migration, Upgrade, and Downgrade Instructions	 69
Product Compatibility	 77

Resolved Issues

IN THIS SECTION

- [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: 17.4R1

- 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 68
Migration, Upgrade, and Downgrade Instructions 69
Product Compatibility 77

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 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	 67
Migration, Upgrade, and Downgrade Instructions	 69
Product Compatibility	 77

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | [69](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | [71](#)
- [Preparing the Switch for Satellite Device Conversion](#) | [72](#)
- [Converting a Satellite Device to a Standalone Switch](#) | [73](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [75](#)
- [Downgrading from Release 17.4](#) | [76](#)

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 17.4R1:

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 commands, where *n* is the spin number:

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-17.4R1.n.tgz
```

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

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-17.4R1.n-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

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 the satellite device from the Junos Fusion topology.

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

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

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

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** from the menu and select the switch platform series and model for your satellite device.

4. Select the software image for your platform, using the following guidelines:

- If the satellite device is a EX4300 switch, you install a standard, signed **jinstall** version of Junos OS.
- If the satellite device is a QFX5100 switch that can be converted to a standalone device, you must install a Preboot eXecution Environment (PXE) version of Junos OS. The PXE version of Junos OS software supports the same feature set as the other Junos OS software packages for a release, but is specially engineered to install Junos OS onto a device running satellite software. The PXE Junos OS package name uses the format **install-media-pxe-qfx-5-version-domestic-signed.tgz**.

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

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

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

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

```
[edit]
```

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

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

```
[edit]
```

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

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

8. Commit the configuration.

To commit the configuration to both Routing Engines:

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

To commit the configuration to a single Routing Engine:

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

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

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

For example, to install a software package stored in the **/var/tmp** directory on the aggregation device onto a switch acting as the satellite device using FPC slot 102:

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

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

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

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases

provide direct upgrade and downgrade paths— you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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 | 67](#)

[Documentation Updates | 68](#)

[Product Compatibility | 77](#)

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 77](#)
- [Hardware Compatibility Tool | 77](#)

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 67
Documentation Updates 68
Migration, Upgrade, and Downgrade Instructions 69

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- New and Changed Features | 78
- Changes in Behavior and Syntax | 79
- Known Behavior | 80
- Known Issues | 80
- Resolved Issues | 81
- Documentation Updates | 81
- Migration, Upgrade, and Downgrade Instructions | 82
- Product Compatibility | 90

These release notes accompany Junos OS Release 17.4R1 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Hardware | 79

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R1 for Junos Fusion Provider Edge.

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	 79
Known Behavior	 80
Known Issues	 80
Resolved Issues	 81
Documentation Updates	 81
Migration, Upgrade, and Downgrade Instructions	 82
Product Compatibility	 90

Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 17.4R1.

SEE ALSO

New and Changed Features	 78
Known Behavior	 80
Known Issues	 80
Resolved Issues	 81
Documentation Updates	 81
Migration, Upgrade, and Downgrade Instructions	 82
Product Compatibility	 90

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 for Junos Fusion Provider Edge.

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

SEE ALSO

New and Changed Features 78
Changes in Behavior and Syntax 79
Known Issues 80
Resolved Issues 81
Documentation Updates 81
Migration, Upgrade, and Downgrade Instructions 82
Product Compatibility 90

Known Issues

There are no known issues in the Junos OS Release 17.4R1 for Junos Fusion Provider Edge.

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

SEE ALSO

New and Changed Features 78
Changes in Behavior and Syntax 79
Known Behavior 80
Resolved Issues 81
Documentation Updates 81
Migration, Upgrade, and Downgrade Instructions 82
Product Compatibility 90

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R1 | 81](#)

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: 17.4R1

Junos Fusion Provider Edge

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

SEE ALSO

New and Changed Features 78
Changes in Behavior and Syntax 79
Known Behavior 80
Known Issues 80
Documentation Updates 81
Migration, Upgrade, and Downgrade Instructions 82
Product Compatibility 90

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 for Junos Fusion Provider Edge documentation.

SEE ALSO

New and Changed Features 78

Changes in Behavior and Syntax 79
Known Behavior 80
Known Issues 80
Resolved Issues 81
Migration, Upgrade, and Downgrade Instructions 82
Product Compatibility 90

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 82
- Upgrading an Aggregation Device with Redundant Routing Engines | 85
- Preparing the Switch for Satellite Device Conversion | 85
- Converting a Satellite Device to a Standalone Device | 87
- Upgrading an Aggregation Device | 89
- Upgrade and Downgrade Support Policy for Junos OS Releases | 89
- Downgrading from Release 17.4 | 89

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

The download and installation process for Junos OS Release 17.4R1 is different that for earlier Junos OS releases.

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

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

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

For upgrades from Junos Release 14.2 and earlier:

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

All other upgrades:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.4R1 **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 17.4R1, 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	 78
Changes in Behavior and Syntax	 79
Known Behavior	 80
Known Issues	 80
Resolved Issues	 81
Documentation Updates	 81
Product Compatibility	 90

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility](#) | 90

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 78
Changes in Behavior and Syntax 79
Known Behavior 80
Known Issues 80
Resolved Issues 81
Documentation Updates 81
Migration, Upgrade, and Downgrade Instructions 82

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

IN THIS SECTION

- [New and Changed Features | 92](#)
- [Changes in Behavior and Syntax | 124](#)
- [Known Behavior | 130](#)
- [Known Issues | 133](#)
- [Resolved Issues | 143](#)
- [Documentation Updates | 158](#)
- [Migration, Upgrade, and Downgrade Instructions | 159](#)
- [Product Compatibility | 166](#)

These release notes accompany Junos OS Release 17.4R1 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Hardware | 93
- Authentication, Authorization, and Accounting (AAA) (RADIUS) | 93
- Class of Service (CoS) | 94
- Dynamic Host Configuration Protocol (DHCP) | 94
- EVPNs | 96
- General Routing | 98
- High Availability (HA) and Resiliency | 99
- Interfaces and Chassis | 99
- Junos OS XML API and Scripting | 101
- Layer 2 Features | 101
- Logical Systems | 102
- Management | 103
- MPLS | 107
- Operation, Administration, and Maintenance (OAM) | 108
- Routing Protocols | 109
- Services Applications | 112
- Software Defined Networking (SDN) | 115
- Software Installation and Upgrade | 117
- Subscriber Management and Services | 118
- System Logging | 122
- User interface and Configuration | 122
- VPNs | 124

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R1 for the MX Series routers.

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](#).]

-  **NOTE:** This feature is documented but not supported in Junos OS Release 17.4R1.

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 17.4R1, 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](#).]

- **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 MIC-MACSEC-MRATE MIC with MACsec support on MPC8E and MPC9E (MX2000 line of routers)**—Starting in Junos OS Release 17.4R1, the MIC-MACSEC-MRATE MIC extends Media Access Control Security (MACsec) capabilities on MPC8E and MPC9E MPCs installed in MX2010, MX2020, and MX2008 routers. Each MPC supports two MIC-MACSEC-MRATE 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](#).]

- **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](#)]

- **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-IPv6 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 NAPT](#).]

- **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 | 124](#)

[Known Behavior | 130](#)

[Known Issues | 133](#)

[Resolved Issues | 143](#)

[Documentation Updates | 158](#)

[Migration, Upgrade, and Downgrade Instructions | 159](#)

[Product Compatibility | 166](#)

Changes in Behavior and Syntax

IN THIS SECTION

• [Interfaces and Chassis | 125](#)

• [Management | 126](#)

• [MPLS | 126](#)

• [Multicast | 127](#)

- [Network Management and Monitoring | 127](#)
- [Routing Protocols | 128](#)
- [Security | 128](#)
- [Services Applications | 129](#)
- [Software Licensing | 129](#)
- [Subscriber Management and Services | 129](#)

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 17.4R1 for MX Series routers.

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**.

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.

- **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.
- **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](#).]

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.

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.

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 17.3R1, use the following command when configuring database memory for Enhanced Subscriber Management:

set system configuration-database max-db-size

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

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

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

SEE ALSO

[New and Changed Features | 92](#)

[Known Behavior | 130](#)

[Known Issues | 133](#)

[Resolved Issues | 143](#)

[Documentation Updates | 158](#)

[Migration, Upgrade, and Downgrade Instructions | 159](#)

[Product Compatibility | 166](#)

Known Behavior

IN THIS SECTION

- [General Routing | 131](#)
- [Interfaces and Chassis | 132](#)
- [Layer 2 Ethernet Services | 133](#)
- [Subscriber Management and Services | 133](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 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](#)
- 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](#)
- 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](#)
- 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](#)
- Support for enterprise profile support is with only 10G interfaces. 40G & 100G may result in phase alignment issue. [PR1310048](#)
- When changing encapsulation from VXLAN to MPLS or vice versa, need to deactivate and reactivate the instance. [PR1326430](#)

Interfaces and Chassis

- 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.

- In a node slicing context, issuing the command **edit 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](#).

- **Configuration not validated after interface is renamed or replaced (MX Series)**—On MX Series routers, after an existing interface in a configuration is renamed or replaced, the configuration is not validated during commit operation. The same configuration with the modified interface name, which might or might not be supported, is saved to the database without any commit errors. If the saved configuration is unsupported, then when an operation is later performed on it, the behavior or response is unknown.

For example, suppose the ge-1/0/0 interface supports the speed value to be configured (say, 1 Gbps) but the ae0 interface does not. You commit the following configuration on the ge- interface:

```
user@host# set interfaces ge-1/0/0 speed 1g
```

Later, you rename ge-1/0/0 to ae0 and commit the configuration, as shown below:

```
user@host# rename interfaces ge-1/0/0 to ae0
```

No validation is performed for the renamed interface ae0, and there are no commit errors. Although unsupported, the configuration is saved to the database.

This is a known issue.

Layer 2 Ethernet Services

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

Subscriber Management and Services

- The **all** option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the **all** option with the **clear services l2tp destination**, **clear services l2tp session**, or **clear services l2tp tunnel** statements in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

SEE ALSO

New and Changed Features 92
Changes in Behavior and Syntax 124
Known Issues 133
Resolved Issues 143
Documentation Updates 158
Migration, Upgrade, and Downgrade Instructions 159
Product Compatibility 166

Known Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 134](#)
- [EVPNs | 134](#)
- [Forwarding and Sampling | 135](#)
- [General Routing | 135](#)
- [Infrastructure | 138](#)
- [Interfaces and Chassis | 138](#)
- [MPLS | 139](#)
- [Network Management and Monitoring | 139](#)

- Platform and Infrastructure | 139
- Routing Protocols | 140
- Services Applications | 142
- VPNs | 142

This section lists the known issues in hardware and software in Junos OS Release 17.4R1 for MX Series routers.

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

Class of Service (CoS)

- A CoS scheduler update can fail when all of the following conditions are met: (1) Dynamic subscribers exist on an aggregated Ethernet bundle, (2) the 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, and (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 CoS update, (2) commit the configuration, (3) reapply the CoS update, and (4) commit the configuration. [PR1276459](#)

EVPNs

- In an EVPN scenario with static MAC configured in the EVPN instance, the static MAC route will be received and installed in the remote EVPN instance table. However, the static MAC route cannot be installed in the EVPN instance table after deactivating and activating the static MAC in the EVPN instance. [PR1193754](#)
- In an EVPN network with VXLAN encapsulation configured for direct-next hop 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](#)

- Core link flap might result in inconsistent global mac count. [PR1328956](#)
- Using the command **clear evpn mac-table** static pinned macs get deleted like other mac entries. This is an unexpected behaviour. [PR1329391](#)

Forwarding and Sampling

- In some stress test conditions, the sampled process crashes and generates a core file when connecting to L2BSA and EVPN subscribers aggressively. [PR1293237](#)
- When subscriber services that are enabled for interim volume accounting goes down it could rarely cause a core in PFED daemon with backtrace `pfed_timer_manager_c::remove_serv_id`. no workaround is required since PFED daemon auto recovers itself over the restart and no corrective action is required. [PR1296969](#)
- Heap memory leaks on DPC when the flow-specification route is changing. [PR1305977](#)

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 CoS-related configuration on MPC5E/MPC6E/MPC2E-NG/MPC3E-NG line cards, the following error messages might be seen: **trinity_insert_ifl_channel:6449 ifl 495 chan_index 495 NOENT** and **jnh_ifl_topo_handler_pfe(11591): ifl=495 err=1 updating channel table nexthop**. [PR1186645](#)
- Upgrading using unified ISSU might trigger a flap in the interfaces on MX Series routers. The following message might be seen: **SFP: pointer Null, sfp_set_present**. [PR1200045](#)
- On platforms with 64-bit X86 Routing Engines, if IPv6 is configured, then either IPv6 router advertisement or Multicast Listener Discovery (MLD) update can cause rpd to crash and generate a core file. [PR1224376](#)
- When virtual switch type is changed from IRB type to regular bridge, interfaces under the OpenFlow protocol are removed. The OpenFlow process (daemon) fails to program any flows. [PR1234141](#)
- On MX Series routers with XM chipset (for example: MPC3E/MPC4E/MPC5E/MPC6E/MPC2E-NG/MPC3E-NG), the MPC might reboot after ISSU completion. [PR1256145](#)

- In a node virtualization setup at high scale with nonstop active routing (NSR) configured, in rare occasions, the guest network function (GNF) might restart MPC9 linecards during a Routing Engine switchover. [PR1259910](#)
- Monitoring FPC temperature is not applicable on MX1RU platform, since MX1RU is single board design with logical FPC. [PR1263315](#)
- Because of transient hardware error conditions, only syslog events XMCHIP(x) FI: Cell underflow at the state stage - Stream 0, Count 65535 are reported, which is a sign of a fabric stream wedge. Additional traffic flow register pointers are validated and if stalled a new CMERROR alarm is raised: **XMCHIP(x) FI: Cell underflow errors with reorder engine pointers stalled - Stream 0, late_cell_value 65535, max_rdr_ptr 0x6a9, reorder_ptr 0x2ae.** [PR1264656](#)
- This issue occurs when an interface comes online and both OAM protocol and MKA protocol try to establish their respective sessions. Because of contention between these two protocols, OAM takes down the interface and MKA fails to establish connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
- On an MX Series Virtual Chassis system in a scaled subscriber management scenario, if unified ISSU is performed while the BGP protocol sessions are active and if the sessions are clients of BFD, then the BGP sessions might go down and come back up again, which might cause traffic loss. [PR1265407](#)
- This very specific issue occurs when the Packet Forwarding Engine is oversubscribed with unknown unicast flood with no MAC learning, which is not a common configuration. During unified ISSU, only the Packet Forwarding Engine gets wedged. However, this issue is not seen when the Packet Forwarding Engine is oversubscribed with L3 traffic or with L2 traffic with MAC learning. [PR1265898](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- In an L2BSA scaling scenario, after bringing up about 12,000 subscribers, one or more FPCs will reboot. [PR1273353](#)
- 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 via programs running on the device itself, and veriexec operating system restricts arbitrary programs from running on Junos OS. There are no known exploit vectors utilizing signed binaries shipped with Junos OS itself. [PR1282562](#)
- During re-injection of JSR data into the associated parent socket, if the socket has gone down by that time, a kernel crash may occur as a rare race condition. [PR1282573](#)
- MX Series Virtual Chassis only: 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](#)
- iLatency (calculated by differing producer timestamp and gRPC server timestamp) can sometimes be negative for Packet Forwarding Engine related telemetry packets due to drift in Routing Engine and Packet Forwarding Engine NTP servers. [PR1303376](#)

- This issue is applicable when using MPLS LSPs and resource RSVP-TE) self-ping. When rpd sends out a self-ping packet and an RSVP packet at the same time, these packets might overwrite the kernel's packet buffers, causing memory corruption and the kernel to panic. [PR1303798](#)
- If syslog errors "pfeman_inline_ka_steering_gencfg_handler: nh not found for nh=< pfh nhid>" are seen on the FPC after it reboots, it is likely that steering rules used for BFD packet redirection are not installed correctly. This may be caused due to unexpected replay order of IPC messages from kernel when the FPC reboots. It may be advisable to reconfigure the impacted BFD sessions that use the respective < pfh nhid> for the redirect rules. [PR1308884](#)
- Customers running 32-bit Junos OS might observe RPD core when traceoptions are enabled. [PR1305440](#)
- Subscriber management ISSU bug fix. [PR1309983](#)
- When upgrading JDM there is a possibility that the JDM daemon will not be running after the upgrade. No errors are reported during the upgrade. [PR1313964](#)
- **show version detail** CLI operational commands hangs for more than 120 seconds in master Routing Engine and more than 60 seconds in backup Routing Engine when extensible subscriber services related configuration is present in router. [PR1314242](#)
- Baseline Stats feature is not working correctly while verifying baseline-stats/Autologin feature for static subscribers over ipv6 demux interface in Routing-I. [PR1322132](#)
- A mobiled core will occur in systems where one Routing Engine is running Junos version 16.2R1 or 17.1R1 and the other Routing Engine is running version 16.1 or 17.2 or later. The core happens on the 16.2R1 or 17.1R1 slot when it is operating as the system's master Routing Engine. The cause is a message that is sent from the backup to the master that the master fails to understand. This situation can happen at various times during ISSU or when the system has GRES enabled with mixed Junos versions. This issue has been fixed starting with 16.1R2 and 17.1R2. [PR1322904](#)
- **show subscribers client-type vlan subscriber-state active logical-system default routing-instance < routing-instance name>** does not work in this release. This is only a display issue and not a functional issue. **>show subscriber** or **show subscriber detail** may be used instead to achieve the same result. [PR1322907](#)
- The CLI command **request vmhost halt routing-engine other** does not achieve the intended action. [PR1323546](#)
- For payload prefix resolve through SRTE color multi-path protocol-nexthop , initially route resolution all works correctly; thereafter due some network change events, the SRTE multiparth next hop updates may stuck in the async-ket io thread. To recover , flap the corresponding BGP session. [PR1324669](#)

Infrastructure

- The gstatd process for 64-bit Junos image does not get to the correct path in the code and as a result, the gstatd process fails to start. [PR1074084](#)
- The configuration statement "set system ports console log-out-on-disconnect", logs the user out from the console and closes the console connection. If the configuration statement "set system syslog console any warning" is used along with the earlier configuration and if there is no active telnet connection to the console, the daemons try to open the console and hang as they wait for a "serial connect" that is received only by doing a telnet to the console. As a workaround, remove the later configuration by using "set system syslog console any warning", which solves the issue. [PR1230657](#)
- When the **set system log-out-on-disconnect** command is enabled, the Junos OS eventd process (daemon) will block the console-open(). However, during this stage with the syslog console configured (always logs on console), any logging will continue even if console session has ended. While console logging is in wait state by eventd, syslog rotation freezes and some processes involved in logging in to the system would also go into the wait state, causing an undesirable behavior. [PR1253544](#)
- The syscalltrace.sh script gets installed as part of the Junos from 16.1R1 and above release, and it is triggered whenever there is replication error on the backup Routing Engine. It logs the system function call to the output file which provides additional debug information. But it might create large files due to a bug in this script. Juniper recommends to uninstall this script after Junos upgrade in production network. The uninstallation of this script will not have any functionality impact on router. [PR1306986](#)

Interfaces and Chassis

- On On IPV6 neighborship is not created on the IRB interface. [PR1198482](#)
- When configuring an aggregated Ethernet interface and after commit, some harmless log messages might appear. The MRU of the aggregated Ethernet interface might reset to the default value (for example,1522). The child links of the aggregated Ethernet interface get reset to the default MRU. [PR1261423](#)
- Junos OS upgrade involving Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later main releases with CFM configuration can cause cfmd to generate a core file after the upgrade. This is due to the old version of `/var/db/cfm.db`. [PR1281073](#)
- LAG member links running LACP in slow mode might get disassociated from the LAG bundle with a combination of restart interface-control and FPC offline/online trigger.

The issue was seen with scale configuration on DUT. The scale details are:

2800 CFM sessions

2800 BFD sessions

2043 BGP peers

3400 VRF instances [PR1298985](#)

- Some CFM sessions do not come up after DUT with the MPC9E line card line card is rebooted with scale config Scale details of the DUT ----- 2800 CFM sessions 2800 BFD sessions 2043 BGP peers 3400 VRF instances Hardware used: MPC9E line cards (MPC8/9E) PR scenario ----- Reboot DUT and peer nodes directly connected to it Symptoms of the PR ----- Some CFM sessions on peer routers are stuck in "start state" Error message "jnh_packet_get_host_headers_bypass" scroll on DUT. [PR1300515](#)
- Y.1731 Delay measurement is not supported on MPC6. [PR1303672](#)
- While executing query on PPP related Junos Telemetry Interface sensor, one may hit intermittent jpppd core. Core will end up restarting JPPPD daemon automatically without having any functional impact. [PR1311396](#)

MPLS

- Because of the current way of calculating bandwidth, you see a minimal discrepancy between MPLS statistics and adjusted bandwidth reported. The algorithm will be enhanced so that both values match 100 percent. [PR1259500](#)
- In an SRTE path, if the top label has explicit-v6-null (label 2), it is not removed even when there are real labels beneath. [PR1287337](#)
- On configuring SR-TE path with "0" explicit NULL as inner most label, SR-TE path doesn't get installed as with label "0". [PR1287354](#)
- When NG-MVPN is configured with RSVP provider tunnels and NSR is used, then the egress router for the tunnel might not correctly replicate some of the tunnel state to the backup routing engine, leading to temporary traffic loss during NSR failover for the affected tunnels. [PR1293014](#)
- Swapping the binding SID between colored and non-colored static SR LSPs might cause rpd to generate a core file. [PR1310018](#)
- The "show mpls container-lsp" output will not show any egress LSP until the Enhanced FRR is enabled for these egress LSPs. [PR1314960](#)

Network Management and Monitoring

- Syslog duplicate entries of hostname and timestamp are breaking the standard logging format. [PR1304160](#)

Platform and Infrastructure

- When using show | compare method to commit, part of configuration may be treated as noise and return syntax error. [PR1042512](#)
- This issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have XE/GE links toward the switch and LAG bundles as uplinks toward upstream routers. The XE/GE link is part of the

physical loop in the topology. Spanning tree protocols such as VSTP/RSTP/MSTP are used for loop avoidance. Some MAC addresses are not learned on DUT when LAG bundles that are part of such bridge domains are flapped and other events such as spanning tree root bridge change occur. [PR1275544](#)

- With 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 ISSU before all forwarding states are restored. This issue does not affect BUM traffic. As a workaround, nonstop bridging (NSB) can be configured at the [set protocols layer2-control nonstop-bridging]. This reduces traffic flood 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 MQCHIP raise a Major CM alarm. [PR1276132](#)
- MX-MPC1-3D, MX-MPC2-3D or MPC-3D-16XGE does not raise a major CMERROR alarm for a combination of link sanity interrupts and cell underflow drop. This might have a permanent impact on packet forwarding due to transient hardware failure. [PR1276144](#)
- When a user configures a firewall filter with one more more sampling actions on the MX204 performance can be degraded due to the use of the sampling action. [PR1303529](#)
- tcp-ping or rpm probe with probe interval of 1 sec fails due to race condition in tcp handler due to which the data packet is being treated as duplicate 3rd ack and dropped . any tcp packet received after connection is established will not have this issue. [PR1308952](#)
- Error message not observed during telnet with username longer than acceptable limit. [PR1312265](#)
- Performing an ISSU on a BSYS can result in a GNF FPC being stuck in the ISSU reboot state when an AFI is in use. [PR1318394](#)
- In JDM, (running on secondary server) jdmd daemon may core if GNF add-image is aborted by pressing CTRL-C. [PR1321803](#)

Routing Protocols

- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a non-default setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- Continuous soft core files might be generated due to bgp-path-selection code. The routing protocol process (rpd) forks a child and the child asserts to produce a core file. The problem is with route-ordering and it is auto-corrected after collecting the soft-assert-core file, without any impact to the traffic or service. [PR815146](#)
- LDP OSPF are 'in sync' state and the reason observed for this is "IGP interface down" with ldp-synchronization enabled for ospf. `user@host> show ospf interface ae100.0 extensive` Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93,

Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity As per the current analysis, "IGP interface down" is observed as the reason because although LDP notified OSPF that LDP sync was achieved, OSPF was not able to take note of the LDP sync notification as the OSPF neighbor was not up yet. The issue is under investigation. [PR1256434](#)

- Performance degradation occurs during computation of LFA and RLFA. This issue does not impact functionality. [PR1264564](#)
- Two multicast tunnel (mt) interfaces are seen for each of the PIM neighbors after VPN-Tunnel-Source activation or deactivation. However, ideally, the same tunnel source should be for both IPv4 and IPv6 address families, if both are using the same PIM tunnel. [PR1281481](#)
- In an IS-IS SR LAN scenario, advertising adj-sids might be missed for a few neighbors if the TLV length gets exhausted. This is not a common scenario. [PR1288331](#)
- When BGP Monitoring Protocol (BMP) sends out route monitoring messages for BGP routes that have unusable or unresolved next hops, the route monitoring messages for those routes might contain a BGP update with an MP_REACH_NLRI path attribute that specifies an incorrect path attribute length. This might occur for any of the address families, except IPv4 and Flowspec (for example, IPv6 can be impacted). This issue could result in unexpected behavior or failures in BMP station applications. [PR1292848](#)
- If a router works as a graceful restart helper during a peering establishment, the newly established peer might lose some of the negotiated capabilities and interpret the updates incorrectly. This can cause peer drops or invalid routes to be received. [PR1293174](#)
- The routing protocol process (rpd) might restart unexpectedly when configuring rib-groups and routing-instances with static routes in a certain order. [PR1298262](#)
- While the device is booting up with 17.4R1 image, error: channel 0: chan_shutdown_read: shutdown() failed for fd 10 [i0 o3]: Socket is not connected messages may show up. These are benign "error" messages and no functionality impact. [PR1300409](#)
- The mcsnoopd process is generating a core file in this scenario. When mcsnoopd tries to terminate gracefully, it tries to clean up all the resources it has used. For this cleanup to happen, the task infrastructure waits for 10 minutes. In these 10 minutes, the KRT task cleanup is not happening properly and it generates a core file. [PR1305239](#)
- Type codepoints for tunnel encapsulation attribute and its subTLVs are now aligned to temporarily assigned values assigned by IANA (From: <https://tools.ietf.org/html/draft-ietf-idr-segment-routing-te-policy-00#section-8>) - Subsequent Address Family Identifiers (SAFI) Parameters * SR Policy SAFI - 73 (From: <https://www.iana.org/assignments/bgp-parameters/bgp-parameters.xhtml#tunnel-sub-tlvs>) - BGP Tunnel Encapsulation Attribute Tunnel Types * SR Policy Type - 15 - BGP Tunnel Encapsulation Attribute TLV * Preference subTLV - 12 * Binding SID - 13 * Segment List subTLV - 128 * Remote Endpoint subTLV - 6 * Color subTLV - 4 (From: <https://tools.ietf.org/html/draft-ietf-idr-segment-routing-te-policy-00>) - Segment List Sub-TLV * Weight - 9. [PR1315486](#)

- Rpd generates a core file in the Backup Routing Engine due to Route Distinguisher (RD) clash between new RT instance updated by master and a deleting RT instance in the Backup Routing Engine. [PR1319587](#)
- RPT_MMX_REGRESSIONS: Rpd core file is observed 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](#)
- When BGP SRTE, static SRTE is configured on the router with both BGP SRTE, static SRTE having the supported 32k SRTE routes installed and then the full configuration is removed, then rpd might crash during the cleanup of the routes, other configuration. This crash is due to a BSID mpls.0 route associated with an SRTE policy remaining around without getting cleaned up. [PR1322133](#)

Services Applications

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

VPNs

- In a Next-generation MVPN scenario, the ipv6 RP bootstrap route type 3 (S-PMSI AD route) prefix for all PIM routers - ff02::d - persist after BSR data stop. [PR1269234](#)

SEE ALSO

[New and Changed Features | 92](#)

[Changes in Behavior and Syntax | 124](#)

[Known Behavior | 130](#)

[Resolved Issues | 143](#)

[Documentation Updates | 158](#)

[Migration, Upgrade, and Downgrade Instructions | 159](#)

[Product Compatibility | 166](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R1 | 143](#)

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: 17.4R1

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 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 struck 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 onlineing 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 (FRU upgarde) 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](#)
- I2tp 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 may 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 mLDP 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 may 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 Io0 is deleted. [PR1286351](#)
- An alarm-mgmt 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

- DTCP non-optimized trigger attributes can delay mirrored traffic forwarding in scaled environments. [PR1269770](#)
- 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

Changes in Behavior and Syntax	 124
Known Behavior	 130
Known Issues	 133
Documentation Updates	 158
Migration, Upgrade, and Downgrade Instructions	 159
Product Compatibility	 166

Documentation Updates

IN THIS SECTION

- [Subscriber Management Provisioning guide](#) | [158](#)

This section lists the errata and changes in Junos OS Release 17.4R1 documentation for MX Series.

Subscriber Management Provisioning guide

- The *Broadband Subscriber Sessions User Guide* did not report that you can suspend AAA accounting, establish a baseline of accounting statistics, and resume accounting. This feature was introduced in Junos OS Release 15.1R4.

SEE ALSO

New and Changed Features	 92
Changes in Behavior and Syntax	 124
Known Behavior	 130
Known Issues	 133
Resolved Issues	 143
Migration, Upgrade, and Downgrade Instructions	 159
Product Compatibility	 166

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.4 | 160](#)
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS | 160](#)
- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | 162](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 164](#)
- [Upgrading a Router with Redundant Routing Engines | 165](#)
- [Downgrading from Release 17.4 | 165](#)

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 | 92](#)

[Changes in Behavior and Syntax | 124](#)

[Known Behavior | 130](#)

[Known Issues | 133](#)

[Resolved Issues | 143](#)

[Documentation Updates | 158](#)

[Product Compatibility | 166](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 166](#)

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 | 92](#)

[Changes in Behavior and Syntax | 124](#)

[Known Behavior | 130](#)

[Known Issues | 133](#)

[Resolved Issues | 143](#)

[Documentation Updates | 158](#)

[Migration, Upgrade, and Downgrade Instructions | 159](#)

Junos OS Release Notes for NFX Series

IN THIS SECTION

- New and Changed Features | 167
- Changes in Behavior and Syntax | 168
- Known Behavior | 168
- Known Issues | 169
- Resolved Issues | 173
- Documentation Updates | 174
- Migration, Upgrade, and Downgrade Instructions | 175
- Product Compatibility | 178

These release notes accompany Junos OS Release 17.4R1 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

This section describes the new features or enhancements to existing features in Junos OS Release 17.4R1 for NFX Series devices.

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 | 168](#)

[Known Behavior | 168](#)

[Known Issues | 169](#)

[Resolved Issues | 173](#)

[Documentation Updates | 174](#)

[Migration, Upgrade, and Downgrade Instructions | 175](#)

[Product Compatibility | 178](#)

Changes in Behavior and Syntax

There are no changes in behavior and syntax for NFX Series in Junos OS Release 17.4R1.

SEE ALSO

[New and Changed Features | 167](#)

[Known Behavior | 168](#)

[Known Issues | 169](#)

[Resolved Issues | 173](#)

[Documentation Updates | 174](#)

[Migration, Upgrade, and Downgrade Instructions | 175](#)

[Product Compatibility | 178](#)

Known Behavior

IN THIS SECTION

- [Juniper Device Manager | 169](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 for the NFX Series.

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

Juniper Device Manager

- JDM shell configurations of interfaces override JDM CLI configurations. As a workaround, use the JDM CLI to configure interfaces. [PR1155749](#)
- SR-IOV interfaces do not support more than 64 VLANs on NFX250. [PR1156348](#)

SEE ALSO

[New and Changed Features | 167](#)

[Changes in Behavior and Syntax | 168](#)

[Known Issues | 169](#)

[Resolved Issues | 173](#)

[Documentation Updates | 174](#)

[Migration, Upgrade, and Downgrade Instructions | 175](#)

[Product Compatibility | 178](#)

Known Issues

IN THIS SECTION

- [Infrastructure | 170](#)
- [IPsec | 170](#)
- [Juniper Device Manager | 170](#)
- [Junos Control Plane | 172](#)
- [vSRX | 173](#)

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

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

Infrastructure

- You might not be able to upgrade from Junos OS Releases 15.1X53-D40 and 15.1X53-D45 to Junos OS Release 17.2R2. As a workaround, you can use the image file on a USB, configure the NFX device to boot from the USB, and install the upgrade. [PR1252323](#)

IPsec

- There is no CLI command to clear interface flow-statistics on ipsec-nm. [PR1216474](#)
- Initial allocation of hugepages is not guaranteed when the srpxfe is killed or restarted. [PR1233794](#)

Juniper Device Manager

- There might be no checks when you configure the IP address on different logical units of interfaces. The commit will go through, and will be displayed in the configuration. [PR1150512](#)
- The following commands are not supported:
 - **clear system reboot** and **clear system commit**
 - **restart gracefully**, **restart immediately**, **restart init**, and **restart soft**
 - **show ethernet-switching**, **show version brief**, **show version all members**, and **show system services service-deployment**

[PR1154819](#)

- When you use the NETCONF command to display system information details such as model and OS, the system OS is displayed as QFX. [PR1160055](#)
- Ubuntu package does not successfully install on the JDM container. As a workaround, install the package passwd by using the **sudo apt-get install passwd** command, which enables the **useradd** command again. [PR1168680](#)
- When you configure a static route on JDM in enhanced-orchestration disabled mode, there might not be an explicit check to validate the IP address. [PR1173039](#)
- System host bridge uses a default MTU of 1500 and does not support jumbo frames. Currently there is no CLI to configure the MTU on the host bridge. [PR1192169](#)
- The Network Service Orchestrator module commits the configuration on JDM, Junos Control Plane, and IPsec-NM sequentially. If the commit fails on any one of these system VNFs, the Network Service Orchestrator module automatically rolls back to the older configuration on the VNF where the commit

error is seen. But, all prior Network Service Orchestrator module configuration commits on the earlier VNFs continue to exist and is not reversed. [PR1196253](#)

- There is no commit check if the PCI address is reused for different interfaces in a VNF. As a workaround, we recommend that you stop the VNF and then add or delete interfaces. [PR1205497](#)
- Certain VNFs support hot plugging of virtio interfaces when the VNF is running. When a VLAN mapped interface is hot plugged to VNFs such as Centos, it is seen that the interface is not reachable from the vjunos0 VM. As a workaround, delete the VNF configuration and recommit the complete configuration along with the new interface. [PR1213451](#)
- After enabling or disabling the ipsec-nm service on the NFX250 platform, a warning message might not be displayed asking for a consent to reboot the device. The enabling or disabling action will be effective only after the device is rebooted. Similarly, no warning is displayed when enhanced orchestration is either enabled or disabled. [PR1213489](#)
- Pre-allocation of hugepages might not consider the available memory and proper commit check is required. It is advisable to use the feature based on free system memory availability. By default, the system requires up to 6 to 7 gigabytes of memory for various operations. The system might not function properly if more memory than what is available is allocated. [PR1213944](#)
- While spawning a VNF, there might not be a commit check for the valid image type supported. [PR1221642](#)
- If a VNF requests for more memory than the available system memory, commit might go through without any errors resulting in VNF going into a shut off state. As a workaround, use the show system visibility memory command to check the available free memory before spawning a VNF. Alternatively, check the log files and the VNF shut off reason will be captured in /var/log/syslog file. [PR1221647](#)
- The following commands are not supported:
 - **show host**
 - **request system software delete**
 - **request system software rollback**
 - **request system storage cleanup**

[PR1219972](#)

- DHCP service can be configured on custom system bridges for service chaining. There might be no commit check if the lower and higher values of the pool range are swapped. [PR1223247](#)
- If the configured TACACS+ server has an IP that can be accessed from JDM, the tacplus pam might not wait till timeout in case TACACS+ server is unreachable. [PR1224420](#)
- The Swap memory information displays incorrect values in the **show system visibility jdm** command output for NFX250 platforms with optimized SSD layouts. [PR1227528](#)
- With enhanced orchestration mode enabled and routing over management configured on vSRX for WAN redundancy for critical traffic, the system CPU utilization will reach 100% if the WAN link goes down,

and traffic routes through out-of-band management. vSRX might not respond to ping or management requests. Egress traffic through management might be throttled. [PR1233478](#)

- Removing the IRB configuration along with the DHCP configuration on JDM and rolling back the configuration might result in the DHCP service not functioning for service chaining of VNFs. [PR1234055](#)
- Hugepages that are preconfigured through the CLI are not used if a custom init-descriptor is used. [PR1245330](#)
- When a VLAN tag is configured through a JDM CLI on a VNF that is provisioned to a DPDK-enabled VM and the VM is spawned, the VLAN filtering or striping configuration on the VNF stops taking effect. Removing and recommitting the JDM VLAN ID configuration on the VNF can resolve the issue unless the system or the VNF is rebooted. [PR1251596](#)
- The **show system visibility cpu** command on JDM has the field values for IOWait and Intr always set to zero. [PR1258361](#)
- Configuring more than the available number of SR-IOV interfaces in enhanced mode might result in a state where the used MAC addresses for such interfaces are not released back to the system MAC pool on deletion of the VNF. [PR1259975](#)

Junos Control Plane

- The Alarm LED will be amber for a major alarm instead of red. In the NFX250-S1E model, the Alarm LED does not blink for any alarms. [PR1146307](#)
- When the option **accept-source-mac mac-address** is configured on an interface and then deleted, no additional MAC's will be learned on the interface. Only the MACs which were earlier configured will be available. [PR1168197](#)
- When LLDP is configured on vjunos0 on an NFX250 Network Services platform, the system name TLV(5) might not be advertised. [PR1169479](#)
- Configuring DSCP and DSCPv6 classifiers together on a Layer 2 interface is not supported. [PR1169529](#)
- There might a traffic drop in IPv4 multicast traffic on JCP when flow-control is configured on interfaces and multicast traffic is more than 400pps. [PR1191794](#)
- On an interface with family inet configured, you might not be able to configure a classifier or rewrite rules. [PR1262840](#)
- If the traffic in the out-of-band interface is heavy, the control plane connectivity might get blocked for some time while the packets are processed. If this interruption persists, the connection between the Packet Forwarding Engine (PFE) and control plane is cleared, which results in a PFE restart or shutdown. You must ensure that there is no heavy traffic flow in the management VLAN. [PR1270689](#)

vSRX

- On an NFX250-S1E platform running vSRX VNF, the performance of SR-IOV with UTM and IDP is lower than VirtIO with UTM and IDP. [PR1214118](#)
- If per-unit-scheduler is not configured, the IFD shaping fails and no packet is queued. [PR1264556](#)
- After configuring the IFD shaping, the ingress interface cannot receive packets. [PR1264850](#)
- The current maximum number of concurrent SIP calls is below the specified maximum limit. [PR1273356](#)

SEE ALSO

[New and Changed Features | 167](#)

[Changes in Behavior and Syntax | 168](#)

[Known Behavior | 168](#)

[Resolved Issues | 173](#)

[Documentation Updates | 174](#)

[Migration, Upgrade, and Downgrade Instructions | 175](#)

[Product Compatibility | 178](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.2R1 | 174](#)

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

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

Resolved Issues: 17.2R1

Juniper Device Manager

- User defined login class is not supported on JDM. [PR1155965](#)
- Ping with record route option will not work for virtio interfaces. [PR1162659](#)
- Default gateway assigned by phc for clients connected via front panel ports will be 10.10.10.254. [PR1168284](#)
- The CLI to configure the time zone is not functional. [PR1169675](#)
- SNMP Trap is not supported on JDM. [PR1173216](#)

Junos Control Plane

- Transmit rate of "0" cannot be configured on schedulers. [PR1158085](#)
- If a cable is not connected to the front panel RJ-45 ports, the status led will blink. [PR1168054](#)
- SFP-T transceivers are not supported. [PR1151575](#), [PR1166808](#), [PR1168203](#)

SEE ALSO

[New and Changed Features | 167](#)

[Changes in Behavior and Syntax | 168](#)

[Known Behavior | 168](#)

[Known Issues | 169](#)

[Documentation Updates | 174](#)

[Migration, Upgrade, and Downgrade Instructions | 175](#)

[Product Compatibility | 178](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 documentation for NFX Series.

SEE ALSO

[New and Changed Features | 167](#)

[Changes in Behavior and Syntax | 168](#)

[Known Behavior | 168](#)

[Known Issues | 169](#)

[Resolved Issues | 173](#)

[Migration, Upgrade, and Downgrade Instructions | 175](#)

[Product Compatibility | 178](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 175](#)
- [Basic Procedure for Upgrading to Release 17.4 | 175](#)

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 **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 17.4R1:

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.

NOTE: After you install a Junos OS Release 17.4R1 **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 for adding the software package to ensure that the device reboots successfully. This is the default behavior when the software package being added is for a different release. Adding the **reboot** command reboots the device after the upgrade is validated and installed. When the reboot is complete, the device displays the login prompt. The loading process might take 5 to 10 minutes. Rebooting occurs only if the upgrade is successful.

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

```
user@host> request system software add validate reboot source/jinstall-17.4
R1.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.4 R1.SPIN-export-signed.tgz
```

Replace the source with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 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.

SEE ALSO

[New and Changed Features | 167](#)

[Changes in Behavior and Syntax | 168](#)

[Known Behavior | 168](#)

[Known Issues | 169](#)

[Resolved Issues | 173](#)

[Documentation Updates | 174](#)

[Product Compatibility | 178](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 178](#)

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](#) | 167

Changes in Behavior and Syntax	168
Known Behavior	168
Known Issues	169
Resolved Issues	173
Documentation Updates	174
Migration, Upgrade, and Downgrade Instructions	175

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- New and Changed Features | 181
- Changes in Behavior and Syntax | 194
- Known Behavior | 199
- Known Issues | 200
- Resolved Issues | 203
- Documentation Updates | 206
- Migration, Upgrade, and Downgrade Instructions | 206
- Product Compatibility | 211

These release notes accompany Junos OS Release 17.4R1 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Hardware | 181](#)
- [High Availability \(HA\) and Resiliency | 182](#)
- [Interfaces and Chassis | 183](#)
- [IPv6 | 184](#)
- [Junos OS XML API and Scripting | 184](#)
- [Layer 2 Features | 184](#)
- [Layer 3 Features | 184](#)
- [Management | 185](#)
- [MPLS | 188](#)
- [Routing Protocols | 190](#)
- [Security | 192](#)
- [Services Applications | 192](#)
- [Software Installation and Upgrade | 193](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R1 for the PTX Series.

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 resignalling events for small LSPs of say 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 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 | 194](#)

[Known Behavior | 199](#)

[Known Issues | 200](#)

[Resolved Issues | 203](#)

[Documentation Updates | 206](#)

[Migration, Upgrade, and Downgrade Instructions | 206](#)

[Product Compatibility | 211](#)

Changes in Behavior and Syntax

IN THIS SECTION

- Class of Service (CoS) | 194
- Interfaces and Chassis | 194
- Management | 196
- MPLS | 196
- Multicast | 197
- Network Management and Monitoring | 197
- Security | 198
- Software Licensing | 198

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 17.4R1 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
```

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 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 (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](#).]

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 (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](#).]

SEE ALSO

[New and Changed Features | 181](#)

[Known Behavior | 199](#)

[Known Issues | 200](#)

[Resolved Issues | 203](#)

[Documentation Updates | 206](#)

[Migration, Upgrade, and Downgrade Instructions | 206](#)

[Product Compatibility | 211](#)

Known Behavior

IN THIS SECTION

- [General Routing | 199](#)
- [Network Management and Monitoring | 199](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 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

- 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](#)
- In the specific case of semi-graceful RCB reboot initiated by the internal shell command: 'vhclient init 0', GRES takes longer to complete i.e 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](#)
- After jacking-in the FPC, the output of "show chassis hardware" might indicate "No Power" for the FPC for initial 20 seconds; but will display the right status after that. [PR1319156](#)

Network Management and Monitoring

- **Unmatched SNMP traps in special scenarios (PTX10008 and PTX10016)**—Unmatched SNMP traps are generated for PTX10008 (starting in Junos OS Release 17.2R1) and for PTX10016 (starting in Junos OS Release 17.4R1) routers during the following two scenarios:
 - When you take the Routing Engine offline by issuing the **request vmhost power-off other-routing-engine** configuration statement, and then bring it back online by issuing the **request vmhost power-on other-routing-engine** statement.

- When you pull out the whole Routing Engine and Control Board FRU from the chassis, and then insert it back into the chassis.

SEE ALSO

[New and Changed Features | 181](#)

[Changes in Behavior and Syntax | 194](#)

[Known Issues | 200](#)

[Resolved Issues | 203](#)

[Documentation Updates | 206](#)

[Migration, Upgrade, and Downgrade Instructions | 206](#)

[Product Compatibility | 211](#)

Known Issues

IN THIS SECTION

- [General Routing | 200](#)
- [Interfaces and Chassis | 202](#)
- [Multiprotocol Label Switching \(MPLS\) | 202](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R1 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

- On platforms with 64-bit X86 Routing Engines, if IPv6 is configured, then either IPv6 router advertisement or Multicast Listener Discovery (MLD) update can cause rpd to crash and generate a core file. [PR1224376](#)
- PTX Series FPC3 might receive noise on the FPC console port and interpret it as valid signals. This might cause login fails on the console port, core files to be generated, or even reloads. [PR1224820](#)

- Management Daemon (MGD) might crash if the Openconfig package is installed immediately or within minutes of Network Agent (NA) 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. The crash will not occur if Network Agent (NA) package is installed after the openconfig package. [PR1265815](#)

- When you offline or restart an FPC 'x' that is sending traffic to FPC 'y', the following error messages are seen on the destination FPC:

```
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.
```

- It also results in a corresponding alarm being set on a destination FPC.
- Specific to PTX10000 is the transient alarm that gets set when this condition occurs by CMERROR infra.
- The alarm clears off later because the source FPC is turned offline.

[PR1268678](#)

- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- With non-enhanced mode, traffic loss is seen on v4 static-lsp with stitch operation not working on PTX Series with paradise. [PR1290942](#)
- When CFP2-DCO-T-WDM-1 is plugged into PTX PIC, repeated configuration rollback can cause the link to take a long time to come-up. [PR1301462](#)
- When CFP2-DCO-T-WDM-1 is plugged in PTX PIC, carrier frequency offset tca is raised after FPC restart, even when tca is not enabled. [PR1301471](#)
- iLatency (calculated by differing producer timestamp and gRPC server timestamp) can sometimes be negative for PFE related telemetry packets due to drift in RE and PFE NTP servers. [PR1303376](#)
- In PTX10008/PTX10016, the FPD LED for SIB might turn to a "GREEN-STEADY" state even before any of the SIBs can come online. [PR1311632](#)
- After deleting firewall and re-configuring again, the firewall filter counter values might not be as expected. [PR1319664](#)
- IPv6 payload prefix resolving over an SR-TE policy results in traffic being silently discarded or dropped if the **extended-nexthop-color** command is used. [PR1319273](#)
- When the command **set forwarding-options l2circuit-control-passthrough** is configured on a working LACP bundle, the interface will go down and no traffic will pass. [PR1320407](#)
- When DCO hot-plug is done before link is up, the FPC might crash. [PR1322260](#)

- When a URCB is inserted in the PTX10008/PTX10016 chassis and it becomes a standby CB, then the STS LED for that CB does not glow GREEN and remains unlit. [PR1325498](#)
- When the telemetry subscription is ended from the collector side, network agent fails to delete the configuration from ephemeral database. Hence the sensors will continue to be provisioned and data will be produced by the publishers without any subscribers. There are two workarounds for this issue:
 - 1. Editing the ephemeral database: This option should be exercised only when the collector(s) have been stopped. From the operational prompt, execute the following commands in sequence: **foo> edit ephemeral foo# delete This will delete the entire configuration Delete everything under this level? [yes,no] (no) yes foo# exit .**
 - 2. Restarting network agent: This option should be exercised only when the collector(s) have been stopped. From the operational prompt, execute following command: **foo> restart na-grpc-server Please verify that the configurations are no longer seen in ?show ephemeral-configuration? command output.** [PR1329134](#)

Interfaces and Chassis

- Junos OS upgrade involving Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later main releases with CFM configuration might cause cfmd to generate a core file after the upgrade. This is due to the old version of `/var/db/cfm.db`. [PR1281073](#)

Multiprotocol Label Switching (MPLS)

- LDP to BGP stitching with eBGP indirect nexthop having implicit null label had never worked on PTX Series routers. It works only when BGP indirect next hop has a real label. Workaround: (1) Ensure the peer advertises real label by adding another router between the egress and ingress PE device. (2) Use IBGP that gets resolved over LDP or RSVP-TE LSPs. This will ensure that the BGP indirect next hop has a real label. [PR1254702](#)
- When NG-MVPN is configured with RSVP provider tunnels and NSR is used, then the egress router for the tunnel might not correctly replicate some of the tunnel state to the backup routing engine, leading to temporary traffic loss during NSR failover for the affected tunnels. [PR1293014](#)
- For a static-label-switched-path configured with "stitch" knob, when resolving over an IP route whose nexthop contains no label operations, there might be packet loss. [PR1307938](#)

SEE ALSO

[New and Changed Features | 181](#)

[Changes in Behavior and Syntax | 194](#)

[Known Behavior | 199](#)

[Resolved Issues | 203](#)

[Documentation Updates | 206](#)

[Migration, Upgrade, and Downgrade Instructions | 206](#)

[Product Compatibility | 211](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R1 | 203](#)

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: 17.4R1

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](#)
- 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](#)
- PTX10016 : 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 | 181](#)

[Changes in Behavior and Syntax | 194](#)

[Known Behavior | 199](#)

[Known Issues | 200](#)

[Documentation Updates | 206](#)

[Migration, Upgrade, and Downgrade Instructions | 206](#)

[Product Compatibility | 211](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 documentation for PTX Series.

SEE ALSO

[New and Changed Features | 181](#)

[Changes in Behavior and Syntax | 194](#)

[Known Behavior | 199](#)

[Known Issues | 200](#)

[Resolved Issues | 203](#)

[Migration, Upgrade, and Downgrade Instructions | 206](#)

[Product Compatibility | 211](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 206](#)
- [Upgrading a Router with Redundant Routing Engines | 207](#)
- [Basic Procedure for Upgrading to Release 17.4 | 207](#)

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 **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 17.4R1:

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 17.4R1 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-17.4
R1.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.4
R1.SPIN-export-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 17.4 **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	181
Changes in Behavior and Syntax	194
Known Behavior	199
Known Issues	200
Resolved Issues	203
Documentation Updates	206
Product Compatibility	211

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 211](#)

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 181
Changes in Behavior and Syntax 194
Known Behavior 199
Known Issues 200
Resolved Issues 203
Documentation Updates 206
Migration, Upgrade, and Downgrade Instructions 206

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- New and Changed Features | 212
- Changes in Behavior and Syntax | 225
- Known Behavior | 230
- Known Issues | 233
- Resolved Issues | 240
- Documentation Updates | 244
- Migration, Upgrade, and Downgrade Instructions | 245
- Product Compatibility | 258

These release notes accompany Junos OS Release 17.4R1 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Hardware | 214
- Class of Service (CoS) | 214
- EVPNs | 214
- General Routing | 216
- Interfaces and Chassis | 216
- Junos OS XML API and Scripting | 217
- Management | 217
- Multicast | 218
- MPLS | 219

- Network Management and Monitoring | 221
- Port Security | 222
- Routing Protocols | 222
- Services Applications | 223
- Software Installation and Upgrade | 224
- Virtual Chassis | 224

This section describes the new features for the QFX Series switches in Junos OS Release 17.4R1.

NOTE: The following QFX Series platforms are supported in Release 17.4R1: QFX5100, QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016.

Hardware

- **QFX10000-30C-M line card (QFX100008 and QFX100016 switches)**—Starting in Junos OS Release 17.4R1-S2, the QFX10000-30C-M line card 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 17.4.1, 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 time-out duration for self-ping on an LSP instance (QFX 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 65535 (runs until success) seconds using the **self-ping-duration value** command at **[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 **[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 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 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 | 225](#)

[Known Behavior | 230](#)

[Known Issues | 233](#)

[Resolved Issues | 240](#)

[Documentation Updates | 244](#)

[Migration, Upgrade, and Downgrade Instructions | 245](#)

[Product Compatibility | 258](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service | 226](#)
- [General Routing | 226](#)
- [Management | 226](#)
- [MPLS | 226](#)
- [Network Management and Monitoring | 227](#)
- [Security | 229](#)
- [Software Licensing | 229](#)
- [Virtual Chassis and Virtual Chassis Fabric \(VCF\) | 229](#)

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 17.4R1 for the QFX Series.

Class of Service

- 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.

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 time out duration for self-ping on an LSP instance (QFX 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 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 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.

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](#).]

- **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](#).]

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 and Virtual Chassis Fabric (VCF)

- **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](#).]

SEE ALSO

[New and Changed Features | 212](#)

[Known Behavior | 230](#)

[Known Issues | 233](#)

[Resolved Issues | 240](#)

[Documentation Updates | 244](#)

[Migration, Upgrade, and Downgrade Instructions | 245](#)

[Product Compatibility | 258](#)

Known Behavior

IN THIS SECTION

- [Hardware | 230](#)
- [EVPNs | 230](#)
- [High Availability \(HA\) and Resiliency | 231](#)
- [Interfaces and Chassis | 231](#)
- [Junos Fusion Provider Edge | 232](#)
- [Layer 2 Features | 232](#)
- [Routing Protocols | 232](#)
- [Storage and Fibre Channel | 232](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 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.

Hardware

- On a QFX5110-32C switch, if a splitter cable is connected to a Spirent 10G CV/MX card, ports will not come up due to varied pre-empt settings for the splitter and DAC cables. There is a hardware limitation where we have no way in EEPROM to differentiate between splitter and DAC cable to apply different settings. As a workaround, use a 40G Spirent card with internal channelization on the Spirent side and manual channelization on the QFX5110-32C side. [PR1280593](#)

EVPNs

- A QFX10000 switch running Junos OS Release 17.4R1 or later software might experience a small and continuous traffic loss under the following conditions:
 - The switch is configured as a Layer 2 and/or Layer 3 VXLAN gateway in an EVPN-VXLAN topology with either a two-layer or collapsed IP fabric.
 - The switch has default ARP and MAC aging timer values.

Under these conditions, the following types of traffic flows might be impacted:

- Bidirectional Layer 3 traffic in a multihomed topology.
- Unidirectional Layer 3 traffic in a single-homed topology.

Note that this issue does not impact bidirectional Layer 3 traffic in a single-homed topology.

To prevent loss in these traffic flows, you must set the **aging-timer** configuration statement in the **[edit system arp]** hierarchy level so that the value is less than the value of the **global-mac-table-aging-time** configuration statement in the **[edit protocols l2-learning]** hierarchy level. [PR1309444](#)

- 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 before Junos OS Release 17.4R1, *permanent remote* means that the ARP route was learned from a remote PE device 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](#)

High Availability (HA) and Resiliency

- On a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than 5 seconds of traffic loss for multicast traffic. [PR1125155](#)

Interfaces and Chassis

- On QFX Series, the logical interface (IFD) and the physical interface (IFL) go down when traffic exceeds rate-limit. Storm control is supported only on interfaces configured in family Ethernet-switching. Moreover, in this family, we support only one IFL per IFD. Due to this, bringing down the IFD is acceptable. Flexible VLAN tagging is not supported on the interfaces enabled for storm control. [PR1295523](#)
- 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](#)

Junos Fusion Provider Edge

- The **no-mac-learning** and **interface-mac-limit** statements are not supported on extended ports or LAGs of extended ports. [PR1296731](#)
- Configuration synchronization is not triggered when you issue the rollback command on the local aggregation device (AD). [PR1298747](#)

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](#)

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](#)
- ERPS convergence takes time after GRES switchover and hence traffic loss is observed for a brief period. [PR1290161](#)
- For the QFX10002 and QFX10008 switches, you might observe an increase in the convergence time of OSPF routes when compared to Junos OS 17.3 releases. An average increase of 1.5 seconds is seen for 100,000 OSPFv3 routes. [PR1297541](#)

Storage and Fibre Channel

- If the configuration changes or any aggregation devices (AD) restart, you might see inconsistency in the output of **show ethernet-switching table** and **show fip snooping satellite** on different ADs for some time. It takes time for the ADs to completely restart and hence MAC addresses might be learned over EVPN (DRP flag). When AD restart is complete, MAC addresses should be learned locally and hence the DRP flag moves to the S flag. It can take up to 10 minutes to get consistent output for show commands.

The output for **show ethernet-switching table** on all ADs will show all the MAC addresses. However, the flags against the MAC addresses might be different on the ADs because the MAC addresses might be learned statically on some ADs and dynamically on others. The flag against the dynamic MAC addresses will be changed from D to S once those MAC addresses are relayed from the satellite device (SD) to the AD, which can take up to 10 minutes. However, there should not be any traffic drop. Traffic drop is expected only initially, when the AD has just been restarted. [PR1304173](#)

SEE ALSO

New and Changed Features	212
Changes in Behavior and Syntax	225
Known Issues	233
Resolved Issues	240
Documentation Updates	244
Migration, Upgrade, and Downgrade Instructions	245
Product Compatibility	258

Known Issues

IN THIS SECTION

- EVPNs | 234
- Hardware | 235
- High Availability (HA) and Resiliency | 235
- Infrastructure | 235
- Interfaces and Chassis | 236
- Junos Fusion Provider Edge | 236
- Layer 2 Features | 238
- Management | 238
- MPLS | 238
- Network Management and Monitoring | 238
- Port Security | 239
- Software Installation and Upgrade | 239
- Subscriber Management and Services | 239
- Virtual Chassis and Virtual Chassis Fabric (VCF) | 239
- VLAN Infrastructure | 239

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.4R1.

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

EVPNs

- In EVPN-VXLAN setup having IRB interfaces with Virtual Gateway Address(VGA) configured, there is a possibility that the longest prefix match with destination address in Ipv6 source address logic selects a VGA for ipv6 neighborhood establishment. This leads to IPV6 packet loss as ND entries are not getting refreshed when VGA is greater than interface address is selected as source address. [PR1267830](#)
- On QFX10000 switches implementing EVPN/VXLAN, if the routing engine is repeatedly restarted on redundant gateways, then inter-vrf traffic will be dropped without notification. [PR1289091](#)
- 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 nexthop 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](#)
- If a host is multi-homed to a set of PEs for redundancy, when the host's MAC/IP is learned by one of these PEs, all PEs belong to this redundant set will install the /32 host route pointing to its local IRB interface in the tenant's IP-VRF table as long as its local multihoming ES interface connecting to this host is up - this is the optimized behavior that can be achieved with the knob "routing-option forwarding-table chained-composite-next-hop ingress evpn" on QFX5110 platform unless this knob is a part of Junos default configuration. Otherwise, without enabling this knob, if a PE attached the multi-homed ES learned this host's MAC/IP only from the control plane through EVPN, the PE will install the /32 host route pointing to the remote PE where it learned host's MAC/IP from. For a PE attached to the multi-homed ES and learned this host's MAC/IP locally through the data plane, the PE always installs the /32 host route pointed to its local IRB interface.[PR1321187](#)
- On a QFX5100 switch implementing EVPN-VXLAN, if the link partner is rebooted or power-cycled, the switch may generate a core file. Reboot of the affected device is recommended to aid in traffic convergence.[PR1321606](#)
- On QFX10002 switches implementing EVPN/VXLAN, clearing ARP on the device may impact L3 traffic.[PR1324228](#)
- EVPN-VXLAN Inter-VLAN traffic stream may gets black-holed when traffic was continuous for longer duration.[PR1324229](#)
- On QFX5110-48S switches implementing EVPN/VXLAN, when the remote Provider Edge (PE) leaf is powered off, the corresponding VTEP interface on the local PE leaf may not be deleted.[PR1324495](#)
- Even though an ARP route is learned locally, the output for the show arp command on the provider edge (PE) device on which the route was learned might display the route as permanent remote, which means that the route was learned from a remote PE device that advertised the route without a MAC address. This issue might occur under the following conditions: 1) in an EVPN-VXLAN topology with a two-layer or collapsed IP fabric in which a customer edge (CE) device is multihomed to QFX10000 switches that function as Layer 3 only or Layer 2/Layer 3 PE devices and run Junos OS Release 17.4R1 or later, 2) IRB

interfaces are configured on each of the PE devices with the same set of IP and MAC addresses.

[PR1324824](#)

Hardware

- When an QSFP+4x10G-IR (PSM4 optical transceiver) is connected between a QFX5200 and a PTX5000, the interfaces do not link because of a timing issue. When a port is channelized, the link goes down and the optical speed is set before the interface comes up. [PR1307400](#)

High Availability (HA) and Resiliency

- ERPS convergence takes time after GRES switchover and hence traffic loss is observed for a brief period. [PR1290161](#)
- When link-protection with the backup port state “down” and LACP are configured, if backup state “down” is removed from the configuration, what should happen is that both ports will be up and the primary port should pass all egress traffic. In some instances, however, traffic might pass through the backup port instead of the primary port. [PR1297597](#)
- The issue will be visible during NSSU if two adjacent members have LACP links in one AE— for example, xe-2/0/1 (on member 2) and xe-3/0/1 (on member 3) are part of one AE, then traffic drop will be observed during NSSU, when member 3 will be rebooted. [PR1311977](#)
- When LACPD system ID mismatch happens, then system ID can be explicitly configured on aggregate interface. Ensure that system IDs should be different on both the boxes. [PR1322067](#)

Infrastructure

- When the **set system ports console log-out-on-disconnect** command is enabled, the Junos OS eventd process (daemon) will block the console-open(), but during this stage with syslog console configured (always logs on console), any logging will continue even if the console session has ended. While console logging is in wait state by eventd, syslog rotation freezes and some daemons directly attached to logging in the system would also get into the wait state, causing an undesirable behavior. [PR1253544](#)
- There is a chance that Management Daemon (MGD) may crash if the Openconfig package is installed immediately or within minutes of Network Agent(NA) 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. The Crash will not happen if Network Agent(NA) Package is installed after the openconfig package [PR1265815](#)
- Support for enterprise profile support is with only 10G interfaces. 40G & 100G may result in phase alignment issue. [PR1310048](#)

- On Junos Automation Enhancement images there is a way to use 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. The regular script run is not impacted. [PR1324124](#)
- In JDM, (running on secondary server) jdmd daemon may core if GNF add-image is aborted by pressing CTRL-C. [PR1321803](#)

Interfaces and Chassis

- When a 100G port is channelized to either 4 x 25G or 2 x 50G speed, in the CLI command to display interface details, "show interface et-x/y/z" we see the "Speed" being shown as 100G but not the actual speed of 25G or 50G. This does not impact the physical speed capability of the ports when used in channelized mode. [PR1319884](#)

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 for 8 minutes; now neighbor information is being held for 8 hours. This time delay is introduced to avoid repeating the initial recognition of the satellite device for interface-down events.

```

user@host> show configuration | display set | grep et-0/0/30
set groups user-host-grp interfaces et-0/0/30 cascade-port
set chassis satellite-management fpc 101 cascade-ports et-0/0/30
set interfaces et-0/0/30 disable
{
master:0}
user@host> show chassis satellite terse
Device                               Extended Ports
Slot  State      Model              Total/Up    Version
100   Online     EX4300-48T         50/1        17.4-20170726_common_xxx.0
102   Online     QFX5200-32C-32Q    2/1         17.4-20170726_common_xxx.0
103   Online     QFX5110-48S-4C     3/2         17.4-20170726_common_xxx.0
{
master:0}
user@host> show chassis satellite neighbor
Interface  State      Port Info    System Name  Model              SW Version
et-0/0/30  Dn
et-0/0/18  Two-Way    et-0/0/18    sd102        QFX5200-32C-32Q    17.4-20170726_common_xxx.0
et-0/0/12  Two-Way    et-0/0/50    sd103        QFX5110-48S-4C     17.4-20170726_common_xxx.0

```

```

et-0/0/6    Two-Way    et-0/1/3    sd100      EX4300-48T
17.4-20170726_common_xxx.0
{
master:0}
user@host> show system license
License usage:
Licenses      Licenses      Licenses      Expiry
Feature name          used    installed    needed
bgp                    1         0         1    invalid
SD-QFX5100-48SH-48TH    0         4         0    permanent
Licenses installed:
License identifier: JUNOSxxxxxxx
License version: 4
Software Serial Number: 99999B999999999
Customer ID: USER-SWITCH
Features:
SD-QFX5100-48SH-48TH-4PK - SD 4 pack QFX5000-10-JFD
permanent
{
master:0}
user@host> show system license usage
Licenses      Licenses      Licenses      Expiry
Feature name          used    installed    needed
bgp                    1         0         1    invalid
SD-QFX5100-48SH-48TH    0         4         0    permanent
{
master:0}
user@host> show system alarms
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 getting 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, sdpc 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](#)

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](#)
- When per-packed load balancing is removed/deleted, next hop index might change. [PR1198092](#)
- On QFX5100 switches, static LAG link protection switchover/revert is not working consistently. [PR1286471](#)

Management

- If there is need to change the MAC table size to higher value than default, then CPU spike is noticed when show command is executed and it take more time than normal condition to display the required details. There is no impact to the MAC learning and traffic. [PR1322041](#)

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. Workaround: (1) Ensure the peer advertises a real label by adding another router between the egress and ingress PE devices. (2) Use IBGP that gets resolved over LDP or RSVP-TE LSPs. This will ensure that the BGP indirect next hop has a real label. [PR1254702](#)
- On QFX5100 switches, unified ISSU is not supported with MPLS configuration. [PR1264786](#)
- When NG-MVPN is configured with RSVP provider tunnels and NSR is used, then the egress router for the tunnel might not correctly replicate some of the tunnel state to the backup routing engine, leading to temporary traffic loss during NSR failover for the affected tunnels. [PR1293014](#)
- Traffic drop occurs on sending L3 traffic across mpls label switched path [PR1313977](#)

Network Management and Monitoring

- On QFX5200 switches, input error counters are incremented on the local physical interface, which does not carry any traffic for this platform. These errors are incremented only when the system is coming up; after the system is up, these counters are not incremented. The error messages have no impact on switch functionality. [PR1148794](#)
- MACsec issue: The **show security macsec statistics** command does not show expected results. Statistics are incorrectly cleared for each physical interface (IFD) under eth periodic (1 second). [PR1283544](#)

Port Security

- On the QFX10000-12C-DWDM Coherent Line Card, it is possible that sometimes the link flaps when MACsec is enabled on Ethernet interfaces. [PR1253703](#)

Software Installation and Upgrade

- On QFX-Series platforms with the ZTP feature enabled, the DHCP clients are not getting an IP address if the DHCP pool with /31 subnet is configured. However, if the DHCP pool with /30 or /24 is configured, it works fine. With /31 configured, the DHCP client state remains as “requesting” :

```
user@host> show dhcp client binding
```

IP address	Hardware address	Expires	State	Interface
0.0.0.0	00:00:5E:00:53:00	0	SELECTING	irb.0
0.0.0.0	00:00:5E:00:53:01	0	SELECTING	vme.0
10.160.136.65	00:00:5E:00:53:03	0	REQUESTING	et-0/0/0.0

[PR1298234](#).

Subscriber Management and Services

- Family Ethernet-switching can't be used when flexible-vlan-tagging is configured. The behavior is non-deterministic with this configuration and there is a possibility of seeing dcpfe core. [PR1316236](#)

Virtual Chassis and Virtual Chassis Fabric (VCF)

- If the **reboot** option is used for **request system software add** in a QFX5110 Virtual Chassis or VCF, the master Routing Engine may not get upgraded due to members rebooting before the master receives their responses for the upgrade. [PR1309103](#)

VLAN Infrastructure

- When a VLAN uses an IRB interface as the routing interface, the vlan-id parameter must be set to "none" to ensure proper traffic routing. This issue is platform independent. [PR1287557](#)

SEE ALSO

New and Changed Features	212
Changes in Behavior and Syntax	225
Known Behavior	230
Resolved Issues	240
Documentation Updates	244
Migration, Upgrade, and Downgrade Instructions	245
Product Compatibility	258

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.4R1 | 240

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: 17.4R1

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](#)
- OVSDb and Openflow have some limitations on QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016 switches running Junos OS Releases 17.1R1, 17.1R2, and 17.2R1. [PR1288227](#)
- 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 and Virtual Chassis Fabric (VCF)

- 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 212
Changes in Behavior and Syntax 225
Known Behavior 230
Known Issues 233
Documentation Updates 244
Migration, Upgrade, and Downgrade Instructions 245
Product Compatibility 258

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.4R1.

SEE ALSO

New and Changed Features 212
Changes in Behavior and Syntax 225
Known Behavior 230
Known Issues 233
Resolved Issues 240
Migration, Upgrade, and Downgrade Instructions 245
Product Compatibility 258

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 245
- Installing the Software on QFX10002 Switches | 248
- 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 | 248
- Installing the Software on QFX10008 and QFX10016 Switches | 250
- Performing a Unified ISSU | 254
- Preparing the Switch for Software Installation | 255
- Upgrading the Software Using Unified ISSU | 255
- Upgrade and Downgrade Support Policy for Junos OS Releases | 258

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 255](#)
- [Upgrading the Software Using Unified ISSU on page 255](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:

NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software— Junos OS, the active configuration, and log files— on the switch to an external storage device with the `command`.

Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, *jinstall-host-qfx-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 | 212](#)

[Changes in Behavior and Syntax | 225](#)

[Known Behavior | 230](#)

[Known Issues | 233](#)

[Resolved Issues | 240](#)

[Documentation Updates | 244](#)

[Product Compatibility | 258](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 259](#)

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	 212
Changes in Behavior and Syntax	 225
Known Behavior	 230
Known Issues	 233
Resolved Issues	 240
Documentation Updates	 244
Migration, Upgrade, and Downgrade Instructions	 245

Junos OS Release Notes for SRX Series

IN THIS SECTION

- New and Changed Features | 260
- Changes in Behavior and Syntax | 273
- Known Behavior | 276
- Known Issues | 279
- Resolved Issues | 281
- Documentation Updates | 286
- Migration, Upgrade, and Downgrade Instructions | 287
- Product Compatibility | 291

These release notes accompany Junos OS Release 17.4R1 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/software/junos/>.

New and Changed Features

IN THIS SECTION


- Release 17.4R1-S1 New and Changed Features | 262
- Release 17.4R1 New and Changed Features | 264

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R1 and Junos OS Release 17.4R1-S1 for the SRX Series devices.

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 17.4R1.



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:

Release 17.4R1-S1 New and Changed Features

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 \(MACsec\)](#).

GPRS

- **Support for GTP handover group (SRX4600)**— Starting with Junos OS Release 17.4R1-S1, 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

- **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

- **SRX4600 Services Gateway**— 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

Although the Junos OS SRX4600 device supports the same services that run on the Junos OS SRX5000 Series devices, it differs in its infrastructure implementation. The SRX4600 device is built on the X86 multi-core processor with the Eagle chip and its flow architecture has been modified to maximize use of that processor. 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-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 type **uname -a** in a terminal on your host OS to verify that the host OS is using the latest kernel version.

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

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 dl0 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 dl0 interface for 4G wireless modems: behavior aggregate classifiers, multifold classifiers, policers, shapers, schedulers, and rewrite rules. The dialer interface, dl0, 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 due to a 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, 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, SRX1500, SRX4100, and SRX4200 devices, and vSRX instances]**— Starting with Junos OS Release 17.4, UDP flood whitelist mechanism is implemented. 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 [Network DoS Attacks](#)]

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](#) | 273

[Known Behavior](#) | 276

[Known Issues](#) | 279

[Resolved Issues](#) | 281

[Documentation Updates | 286](#)

[Migration, Upgrade, and Downgrade Instructions | 287](#)

[Product Compatibility | 291](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Chassis Cluster | 274](#)
- [Installation and Upgrade | 274](#)
- [MPLS | 274](#)
- [Management | 274](#)
- [NAT | 274](#)
- [Security | 275](#)

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 17.4R1.

Chassis Cluster

- **ISSU chassis cluster hold timer (SRX1500, SRX4100, SRX4200, SRX4600)**—Starting with Junos OS Release 17.4R1, the hold timer for the initial reboot of the secondary node during the ISSU process is extended from 15 minutes (900 seconds) to 45 minutes (2700 seconds) in a chassis clusters

Installation and Upgrade

- **Factory-default configuration changes**— Starting in Junos OS Release 17.4R1, on SRX300, SRX320, SRX340, SRX345, and SRX550M devices, the **telnet** and **xnm-clear-text** options are not part of the system services configuration in the factory-default configuration of the device.

MPLS

- **Support for Path Computation Element Protocol** — Starting with Junos OS Release 17.4R1, the Path Computation Element Protocol (PCEP) is supported on SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, and SRX4200 devices, and vSRX instances.

Management

- **Changes to Junos OS YANG module naming conventions (SRX 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](#).]

NAT

- **Error message when active source NAT is configured**— When zones are not configured in the rule set and when active source NAT is configured with missing mandatory statement “from”, then the “**Missing mandatory statement: 'from' error: configuration check-out failed**” message is displayed when performing the commit and the configuration checkout fails.
- **Configuring source NAT rule or pool**— On SRX4600 devices, when you configure source NAT rule or pool with rule name or pool name as interface or service set, you will receive the following error message: **syntax error, expecting < data>**.

- If there is a source NAT rule named **interface**, the rule cannot be viewed using the **show security nat source rule interface** command.
- If there is a source NAT rule named **service-set**, the rule cannot be viewed using the **show security nat source rule service-set** command.
- If there is a source NAT pool named **interface**, the pool cannot be viewed using the **show security nat source pool interface** command.
- If there is a source NAT pool named **service-set**, the pool cannot be viewed using the **show security nat source pool service-set** command.
- If there is a source NAT pool named **interface**, the paired-address cannot be viewed using the **show security nat source paired-address pool-name interface** command.
- If there is a source NAT pool named **service-set**, the paired-address cannot be viewed using the **show security nat source paired-address pool-name service-set** command.

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.
- **TPM Firmware Update (SRX300, SRX320, SRX340, and SRX345)** – Starting with Junos OS Release 17.4R1, Trusted Platform Module (TPM) firmware has been updated. The upgraded firmware version provides additional secure cryptography and improves security. Updated TPM firmware is available along with the Junos OS package. For updating TPM Firmware, see [Upgrading TPM Firmware on SRX-Devices](#).

To confirm the TPM firmware version, use the **show security tpm status** command. The following additional new output fields are introduced:

- **TPM Family**— Displays Trusted Computing Group' s (TCG) TPM family version.
- **TPM Firmware version**— Displays the firmware version loaded in TPM.

SEE ALSO

[New and Changed Features | 260](#)

[Known Behavior | 276](#)

[Known Issues | 279](#)

[Resolved Issues | 281](#)

[Documentation Updates | 286](#)

[Migration, Upgrade, and Downgrade Instructions | 287](#)

[Product Compatibility | 291](#)

Known Behavior

IN THIS SECTION

- [Authentication and Access | 276](#)
- [J-Web | 276](#)
- [Layer 2 Ethernet Services | 277](#)
- [Platform and Infrastructure | 277](#)
- [Software Installation and Upgrade | 278](#)
- [UTM | 278](#)
- [VPNs | 278](#)

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.4R1 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](#)

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. The same DHCP relay can be configured using the CLI. [PR1205911](#)
- On SRX Series devices, the DHCP client bindings option under Monitor is removed. The same bindings can be seen in CLI using the **show dhcp client binding** command. [PR1205915](#)
- On SRX Series devices, you cannot view the custom log files created for event logging in J-Web.
- 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, 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](#)

Layer 2 Ethernet Services

- PPPoE + DHCPv6 cannot work in all SRX platforms with 15.1X49 and later versions. [PR1229836](#)

Platform and Infrastructure

- Starting in Junos OS Release 17.4R1, on SRX300, SRX320, SRX340, SRX345, and SRX550M devices, telnet and xnm-clear-text are not part of system services in factory-default configurations.
- The **log-out-on-disconnect** statement is not operational on SRX1500, SRX4100, SRX4200, and SRX4600 devices; on these devices, you must manually log out from the console with the **request system logout** command.
- On SRX4600 devices, USB disk is not made available to JUNOS. However, USB disk is available for Host OS (Linux) with full access. USB is still used in the booting process (install and recovery functions). [PR1283618](#)

Software Installation and Upgrade

- On SRX1500, SRX4100, and SRX4200 devices, ISSU is not supported for upgrading to 17.4 releases from previous Junos OS releases. ISSU is supported for upgrading from Junos OS 17.4 to successive 17.4 releases.

On SRX5400, SRX5600 and SRX5800 devices, ISSU is not supported for upgrading to 17.3 and higher releases from earlier Junos OS releases. ISSU is supported for upgrading from Junos OS 17.3 to Junos 17.4 releases.

NOTE: SRX300 Series devices and SRX550M devices do not support ISSU.

For more details on ISSU support on SRX devices, see <https://kb.juniper.net/KB17946>

UTM

- On SRX Series devices, if the category file transfer fails between the primary and secondary devices, then the file transfer results in an upgrading error and an error log is generated.

On SRX Series devices, during new category file installation, if the category filename is changed, then the new category file overwrites the old category file in the internal system and all related output information is replaced with the new category name.

- If the user profile has the same name as the base filter, then the Web filter uses the wrong profile.

VPNs

- On SRX5400, SRX5600, and SRX5800 devices, when CoS on st0 interface is enabled and the incoming traffic rate destined for st0 interface is higher than 300000 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. It is recommended that you configure appropriate policer on the ingress interface to limit the traffic below 300000 pps per SPU. [PR1239021](#)

SEE ALSO

[New and Changed Features | 260](#)

[Changes in Behavior and Syntax | 273](#)

[Known Issues | 279](#)

[Resolved Issues | 281](#)

[Documentation Updates | 286](#)

[Migration, Upgrade, and Downgrade Instructions | 287](#)

[Product Compatibility | 291](#)

Known Issues

IN THIS SECTION

- [Chassis Clustering | 279](#)
- [Flow-based and Packet-based Processing | 279](#)
- [Interfaces | 280](#)
- [J-Web | 280](#)
- [Routing Protocols | 280](#)
- [VPNs | 280](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R1.

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

Chassis Clustering

- On SRX Series devices with virtual private LAN service (VPLS) configurations, when both nodes are rebooted at the same time and FPC0 cold sync completes (when FPC nodes come online), secondary node stops working when you perform RGO failover (within 5 minutes). As a workaround, perform RGO failover after waiting for ~30 minutes, after system boots up. [PR1327357](#)

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 all SRX Series devices, filter-based forwarding 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 SRX Series devices, packet capture does not work after you change, delete, or add maximum capture size. [PR1304723](#)

Interfaces

- 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 SRX1500, SRX4100, SRX4200, and vSRX chassis clusters, invalid MAC address might be allocated for certain AE interfaces, which results in the interfaces not coming up. [PR1270166](#)
- On SRX4600 devices, the 1GE interface is not supported in Junos OS Release 17.4R1. [PR1315073](#)

J-Web

- On SRX Series devices, uploading certificate using the browse button stores the certificate in the device at `/jail/var/tmp/uploads/`, which will be deleted upon executing the **request system storage cleanup** CLI command. [PR1312529](#)
- On SRX Series devices, the values of address and address-range are not displayed in the inline address-set creation pop-up window of JIMS [PR1312900](#)

Routing Protocols

- On SRX Series devices, RIP is supported in packet-to-packet DC mode on st0 interfaces. [PR1141817](#)

VPNs

- When the SRX Series device is an initiator behind NAT, disabling NAT on the middle router causes an immediate new negotiation failure due to a failed attempt with port 4500. The next attempt will succeed by using port 500. Disabling NAT and bringing down all the existing tunnels and reestablishing the tunnels with port 500 is the expected behavior. [PR1273213](#)

SEE ALSO

New and Changed Features	260
Changes in Behavior and Syntax	273
Known Behavior	276
Resolved Issues	281
Documentation Updates	286
Migration, Upgrade, and Downgrade Instructions	287
Product Compatibility	291

Resolved Issues

IN THIS SECTION

- Application Layer Gateways (ALGs) | 282
- Chassis Cluster | 282
- Class of Service (CoS) | 282
- Flow-Based and Packet-Based Processing | 282
- Interfaces and Chassis | 284
- J-Web | 284
- Layer 2 Ethernet Services | 284
- Network Address Translation (NAT) | 284
- Network Management and Monitoring | 285
- Platform and Infrastructure | 285
- Routing Policy and Firewall Filters | 285
- System Logging | 286
- Unified Threat Management (UTM) | 286
- VPNs | 286

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.

Application Layer Gateways (ALGs)

- On SRX Series devices SIP packet might drop when SIP traffic performs destination NAT. [PR1268767](#)
- The pfed process crashes 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 due to incorrect device temperature readings which reportedly is a too high temperature, leading to an immediate proactive power-off of the device to protect the device from overheating. When this condition occurs, the following log message is shown

in file `/var/log/hostlogs/lcmd.log`: Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor.[PR1241061](#)

- 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 crashes 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 due to 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 260
Changes in Behavior and Syntax 273
Known Behavior 276
Known Issues 279
Documentation Updates 286
Migration, Upgrade, and Downgrade Instructions 287
Product Compatibility 291

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R1 for the SRX Series documentation.

SEE ALSO

New and Changed Features 260
--

Changes in Behavior and Syntax	273
Known Behavior	276
Known Issues	279
Resolved Issues	281
Migration, Upgrade, and Downgrade Instructions	287
Product Compatibility	291

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Scripts for Address Book Configuration | 287

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 | 288
- Running Upgrade and Downgrade Scripts | 289
- Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 290

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 289](#)).

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:

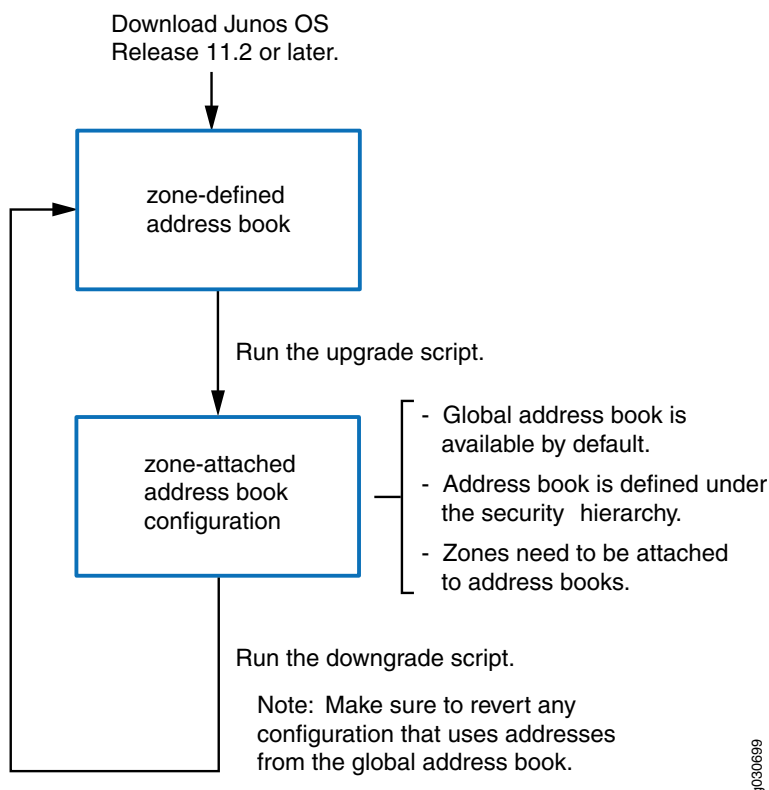
- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.

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 | 260](#)

[Changes in Behavior and Syntax | 273](#)

[Known Behavior | 276](#)

[Known Issues | 279](#)

[Resolved Issues | 281](#)

[Documentation Updates | 286](#)

[Product Compatibility | 291](#)

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 291](#)

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 | 260](#)

[Changes in Behavior and Syntax | 273](#)

[Known Behavior | 276](#)

[Known Issues | 279](#)

[Resolved Issues | 281](#)

[Documentation Updates | 286](#)

[Migration, Upgrade, and Downgrade Instructions | 287](#)

Upgrading Using Unified ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [High Availability User Guide for Routing Devices](#).

For information about ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\)](#) web application.

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information

For the latest, most complete information about known and resolved issues with the Junos OS, see the Juniper Networks Problem Report Search application at <https://prsearch.juniper.net>.


For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at <https://pathfinder.juniper.net/feature-explorer/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies— For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/assets/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties— For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC Hours of Operation — The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) tool located at <https://entitlementsearch.juniper.net/entitlementsearch/> .

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at <https://www.juniper.net/support/requesting-support.html>.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to support@juniper.net. For documentation issues, fill out the bug report form located at <https://www.juniper.net/cgi-bin/docbugreport/>.

Revision History

17 September 2021—Revision 20, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 October 2020—Revision 19, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 18, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 February 2020—Revision 17, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 May 2019—Revision 16, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 January 2019—Revision 15, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 November 2018—Revision 14, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 October 2018—Revision 13, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 August 2018—Revision 12, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 May 2018—Revision 11, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

19 April 2018—Revision 10, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 April 2018—Revision 9, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 March 2018—Revision 8, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 March 2018—Revision 7, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 February 2018—Revision 6, Junos OS Release 17.4R1— 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 © 2018 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

