

Release Notes: Junos[®] OS Release 17.4R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

13 January 2021

Contents	Introduction 13
	New Features in 17.4R3 13
	Junos OS Release Notes for ACX Series 13
	New and Changed Features 14
	Release 17.4R3 New and Changed Features 14
	Release 17.4R2 New and Changed Features 15
	Release 17.4R1 New and Changed Features 15
	Changes in Behavior and Syntax 16
	Management 17
	Network Management and Monitoring 17
	Platform and Infrastructure 18
	Security 19
	Software Licensing 19
	Subscriber Management and Services 19
	Known Behavior 20
	General Routing 20
	Known Issues 21
	General Routing 22
	Interfaces and Chassis 24

Layer 2 Features	24
MPLS	24
Routing Protocols	24
Virtual Chassis	24
Resolved Issues	25
Resolved Issues: 17.4R3	25
Resolved Issues: 17.4R2	27
Resolved Issues: 17.4R1	29
Documentation Updates	29
Migration, Upgrade, and Downgrade Instructions	30
Upgrade and Downgrade Support Policy for Junos OS Releases	30
Junos OS Release Notes for EX Series Switches	31
New and Changed Features	32
Release 17.4R3 New and Changed Features	33
Release 17.4R2 New and Changed Features	33
Release 17.4R1 New and Changed Features	34
Changes in Behavior and Syntax	39
EVPNs	39
Interfaces and Chassis	40
Management	40
Multicast	40
Network Management and Monitoring	40
Platform and Infrastructure	41
Routing Protocols	42
Security	42
Software Licensing	42
Subscriber Management and Services	42
Virtual Chassis	43
Known Behavior	44
EVPN	45
High Availability (HA) and Resiliency	45
Infrastructure	45
Interfaces and Chassis	45
Junos Fusion Enterprise	45

Platform and Infrastructure | 46

Routing Protocols | 46

Virtual Chassis | 46

Known Issues | 47

Authentication and Access Control | 47

EVPN | 47

General Routing | 48

Infrastructure | 49

Interfaces and Chassis | 50

Junos Fusion Enterprise | 50

Layer 2 Ethernet Services | 50

Layer 2 Features | 51

Multicast | 51

Platform and Infrastructure | 51

Routing Protocols | 52

Subscriber Access Management | 52

Resolved Issues | 53

Resolved Issues: 17.4R3 | 53

Resolved Issues: 17.4R2 | 59

Resolved Issues: 17.4R1 | 63

Documentation Updates | 65

Migration, Upgrade, and Downgrade Instructions | 66

Upgrade and Downgrade Support Policy for Junos OS Releases | 66

Junos OS Release Notes for Junos Fusion Data Center | 67

New and Changed Features | 67

Changes in Behavior and Syntax | 68

Known Behavior | 68

Junos Fusion Data Center | 69

Known Issues | 69

Resolved Issues | 70

Resolved Issues: Junos OS Release 17.4R3 | 70

Resolved Issues: Junos OS Release 17.4R2 | 70

Resolved Issues: Junos OS Release 17.4R1 | 70

Documentation Updates | 71

Migration, Upgrade, and Downgrade Instructions | 71**Basic Procedure for Upgrading an Aggregation Device | 72****Preparing the Switch for Satellite Device Conversion | 74****Autoconverting a Switch into a Satellite Device | 76****Manually Converting a Switch into a Satellite Device | 79****Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 81****Configuring Satellite Device Upgrade Groups | 82****Converting a Satellite Device to a Standalone Device | 84****Upgrade and Downgrade Support Policy for Junos OS Releases | 84****Downgrading from Release 17.4 | 84****Junos OS Release Notes for Junos Fusion Enterprise | 86****New and Changed Features | 86****Release 17.4R3 New and Changed Features | 87****Release 17.4R2 New and Changed Features | 87****Release 17.4R1 New and Changed Features | 87****Changes in Behavior and Syntax | 88****Junos Fusion Enterprise | 89****Known Behavior | 89****Junos Fusion Enterprise | 89****Known Issues | 90****Junos Fusion Enterprise | 90****Resolved Issues | 91****Resolved Issues: 17.4R3 | 91****Resolved Issues: 17.4R2 | 92****Resolved Issues: 17.4R1 | 92****Documentation Updates | 93****Migration, Upgrade, and Downgrade Instructions | 93****Basic Procedure for Upgrading Junos OS on an Aggregation Device | 94****Upgrading an Aggregation Device with Redundant Routing Engines | 95****Preparing the Switch for Satellite Device Conversion | 96****Converting a Satellite Device to a Standalone Switch | 97****Upgrade and Downgrade Support Policy for Junos OS Releases | 97****Downgrading from Release 17.4 | 98**

Junos OS Release Notes for Junos Fusion Provider Edge | 99

New and Changed Features | 99

Release 17.4R3 New and Changed Features | 100

Release 17.4R2 New and Changed Features | 100

Release 17.4R1 New and Changed Features | 100

Changes in Behavior and Syntax | 100

Security | 101

Known Behavior | 101

Junos Fusion Provider Edge | 102

Known Issues | 102

Junos Fusion Provider Edge | 103

Resolved Issues | 104

Resolved Issues: 17.4R3 | 104

Resolved Issues: 17.4R2 | 105

Resolved Issues: 17.4R1 | 105

Documentation Updates | 105

Migration, Upgrade, and Downgrade Instructions | 106

Basic Procedure for Upgrading an Aggregation Device | 106

Upgrading an Aggregation Device with Redundant Routing Engines | 109

Preparing the Switch for Satellite Device Conversion | 109

Converting a Satellite Device to a Standalone Device | 111

Upgrading an Aggregation Device | 113

Upgrade and Downgrade Support Policy for Junos OS Releases | 113

Downgrading from Release 17.4 | 113

Junos OS Release Notes for MX Series 5G Universal Routing Platforms | 114

New and Changed Features | 115

Release 17.4R3 New and Changed Features | 116

Release 17.4R2-S2 New and Changed Features | 116

Release 17.4R2 New and Changed Features | 116

Release 17.4R1 New and Changed Features | 119

Changes in Behavior and Syntax | 149

Class of Service (CoS) | 150

EVPNs | 150

General Routing | 151

High Availability (HA) and Resiliency	151
Interfaces and Chassis	151
Management	154
MPLS	154
Multicast	157
Network Management and Monitoring	157
Routing Protocols	158
Security	159
Services Applications	159
Software Defined Networking	160
Software Installation and Upgrade	161
Software Licensing	161
Subscriber Management and Services	161
User Interface and Configuration	164
Known Behavior	164
EVPN	165
General Routing	166
Infrastructure	169
Interfaces and Chassis	169
Junos Fusion Provider Edge	170
Layer 2 Ethernet Services	170
Multiprotocol Label Switching (MPLS)	170
Platform and Infrastructure	170
Routing Protocols	170
Services Applications	171
Software Defined Networking (SDN)	172
Software Installation and Upgrade	172
Subscriber Management and Services	172
Known Issues	173
Class of Service (CoS)	174
EVPN	175
Forwarding and Sampling	176
General Routing	177
High Availability (HA) and Resiliency	189

Infrastructure	189
Interfaces and Chassis	189
Layer 2 Features	191
Layer 2 Ethernet Services	191
MPLS	192
Network Management and Monitoring	194
Platform and Infrastructure	194
Routing Policy and Firewall Filters	196
Routing Protocols	197
Services Applications	200
Subscriber Access Management	201
User Interface and Configuration	201
VPNs	201
Resolved Issues	202
Resolved Issues: 17.4R3	203
Resolved Issues: 17.4R2	234
Resolved Issues: 17.4R1	260
Documentation Updates	275
Subscriber Management Access Network Guide	275
Subscriber Management Provisioning guide	275
Subscriber Management VLANs Interfaces Guide	276
Migration, Upgrade, and Downgrade Instructions	276
Basic Procedure for Upgrading to Release 17.4	277
Procedure to Upgrade to FreeBSD 11.x-Based Junos OS	277
Procedure to Upgrade to FreeBSD 6.x-Based Junos OS	280
Upgrade and Downgrade Support Policy for Junos OS Releases	282
Upgrading a Router with Redundant Routing Engines	282
Downgrading from Release 17.4	282
Junos OS Release Notes for NFX Series	283
New and Changed Features	284
Release 17.4R3 New and Changed Features	284
Release 17.4R2 New and Changed Features	284
Release 17.4R1 New and Changed Features	284
Changes in Behavior and Syntax	285

Known Behavior | 285

Known Issues | 286

Virtual Network Functions | 286

Juniper Device Manager | 286

Junos Control Plane | 286

Resolved Issues | 287

Documentation Updates | 288

Migration, Upgrade, and Downgrade Instructions | 288

Upgrade and Downgrade Support Policy for Junos OS Releases | 288

Basic Procedure for Upgrading to Release 17.4 | 289

Junos OS Release Notes for PTX Series Packet Transport Routers | 291

New and Changed Features | 291

Release 17.4R3 New and Changed Features | 292

Release 17.4R2 New and Changed Features | 292

Release 17.4R1 New and Changed Features | 292

Changes in Behavior and Syntax | 304

Class of Service (CoS) | 305

General Routing | 305

Interfaces and Chassis | 305

Management | 308

MPLS | 308

Multicast | 309

Network Management and Monitoring | 309

Routing Policy and Firewall Filters | 310

Security | 311

Software Licensing | 311

Subscriber Management and Services | 311

Known Behavior | 312

General Routing | 313

Interfaces and Chassis | 313

Known Issues | 314

General Routing | 314

Infrastructure | 318

Interfaces and Chassis | 318

Layer 2 Features	318
MPLS	318
Routing Protocols	319
Resolved Issues	319
Resolved Issues: 17.4R3	320
Resolved Issues: 17.4R2	323
Resolved Issues: 17.4R1	327
Documentation Updates	329
Migration, Upgrade, and Downgrade Instructions	330
Upgrade and Downgrade Support Policy for Junos OS Releases	330
Upgrading a Router with Redundant Routing Engines	330
Basic Procedure for Upgrading to Release 17.4	331
Junos OS Release Notes for the QFX Series	334
New and Changed Features	335
Release 17.4R3 New and Changed Features	336
Release 17.4R2 New and Changed Features	336
Release 17.4R1 New and Changed Features	336
Changes in Behavior and Syntax	347
Class of Service (CoS)	348
EVPNs	348
General Routing	348
Interfaces and Chassis	348
Management	349
MPLS	349
Network Management and Monitoring	350
Routing Policy and Firewall Filters	352
Security	352
Software Licensing	353
Virtual Chassis	353
Known Behavior	354
Class of Service (CoS)	355
EVPN	355
General Routing	355
Interfaces and Chassis	356

Junos Fusion Satellite Software	356
Layer 2 Features	356
MPLS	356
Routing Protocols	357
Platform and Infrastructure	357
Virtual Chassis	358
Known Issues	358
EVPN	359
Forwarding and Sampling	360
General Routing	360
Interfaces and Chassis	365
Layer 2 Ethernet Services	365
Layer 2 Features	366
MPLS	366
Network Management and Monitoring	367
Platform and Infrastructure	367
Routing Protocols	367
Virtual Chassis	368
Resolved Issues	369
Resolved Issues: 17.4R3	369
Resolved Issues: 17.4R2	380
Resolved Issues: 17.4R1	387
Documentation Updates	391
Migration, Upgrade, and Downgrade Instructions	391
Upgrading Software on QFX Series Switches	392
Installing the Software on QFX10002 Switches	394
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	394
Installing the Software on QFX10008 and QFX10016 Switches	396
Performing a Unified ISSU	400
Preparing the Switch for Software Installation	401
Upgrading the Software Using Unified ISSU	401
Upgrade and Downgrade Support Policy for Junos OS Releases	404

Junos OS Release Notes for SRX Series | 405

New and Changed Features | 405

Release 17.4R3 New and Changed Features | 406

Release 17.4R2 New and Changed Features | 406

Release 17.4R1-S1 New and Changed Features | 406

Release 17.4R1 New and Changed Features | 408

Changes in Behavior and Syntax | 417

Release 17.4R3 Changes in Behavior and Syntax | 417

Release 17.4R2 Changes in Behavior and Syntax | 419

Known Behavior | 422

Authentication and Access | 422

Chassis Clustering | 422

J-Web | 423

Layer 2 Ethernet Services | 423

Platform and Infrastructure | 423

User Interface and Configuration | 424

VPNs | 424

Known Issues | 424

Chassis Clustering | 425

Flow-Based and Packet-Based Processing | 425

Intrusion Detection and Prevention (IDP) | 425

J-Web | 426

Platform and Infrastructure | 426

VPNs | 426

Resolved Issues | 427

Resolved Issues: 17.4R3 | 428

Resolved Issues: 17.4R2 | 436

Resolved Issues: 17.4R1 | 445

Documentation Updates | 449

Migration, Upgrade, and Downgrade Instructions | 449

Upgrade and Downgrade Scripts for Address Book Configuration | 450

Upgrading Using ISSU | 454

Compliance Advisor | 454

Finding More Information | 454

Documentation Feedback | 455

Requesting Technical Support | 456

Self-Help Online Tools and Resources | 456

Creating a Service Request with JTAC | 457

Revision History | 457

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 17.4R3 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.

New Features in 17.4R3

Feature	Release Note Section
Preventing validation of magic numbers in PPP peer-originated keepalive messages (MX Series)	“New and Changed Features” on page 115
LACP hold-up timer configuration support on LAG interfaces (PTX Series)	“New and Changed Features” on page 291

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [New and Changed Features | 14](#)
- [Changes in Behavior and Syntax | 16](#)
- [Known Behavior | 20](#)
- [Known Issues | 21](#)
- [Resolved Issues | 25](#)
- [Documentation Updates | 29](#)
- [Migration, Upgrade, and Downgrade Instructions | 30](#)

These release notes accompany Junos OS Release 17.4R3 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R3 New and Changed Features | 14](#)
- [Release 17.4R2 New and Changed Features | 15](#)
- [Release 17.4R1 New and Changed Features | 15](#)

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

Release 17.4R3 New and Changed Features

There are no new features or enhancements to existing features for ACX Series in Junos OS Release 17.4R3.

Release 17.4R2 New and Changed Features

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

Release 17.4R1 New and Changed Features

Management

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

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

Routing Protocols

- **Enhancements to BGP to support attribute transparency (ACX Series)**—Starting with Junos OS Release 17.4R1, BGP feature is enhanced to support attribute transparency for NEXT_HOP, AS_PATH, MULTI_EXIT_DISC, AIGP, and Communities attributes. This feature also provides BGP API enhancements (Add, Get, Modify, Update, Remove, Monitor APIs) to support EBGp and make the route server programmable.

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

Timing and Synchronization

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

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

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

SEE ALSO

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

[Known Behavior | 20](#)

[Documentation Updates | 29](#)

[Known Issues | 21](#)

[Resolved Issues | 25](#)

[Migration, Upgrade, and Downgrade Instructions | 30](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Management | 17](#)
- [Network Management and Monitoring | 17](#)
- [Platform and Infrastructure | 18](#)
- [Security | 19](#)
- [Software Licensing | 19](#)
- [Subscriber Management and Services | 19](#)

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 17.4R3 for the ACX Series Universal Metro Routers.

Management

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

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

Network Management and Monitoring

- **SNMP syslog messages changed (ACX 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 (ACX 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 (since this is an important message but less frequent)
- IFL LinkUp -> LOG_INFO (no change)
- IFD and IFL LinkDown -> LOG_WARNING (no change)

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

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

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

- **The NETCONF server omits warnings in RPC replies when the `rfc-compliant` statement is configured and the operation returns `<ok/>` (ACX Series)**—Starting in Junos OS Release 17.4R3, when you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an `<ok/>` element and one or more `<rpc-error>` elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the `rfc-compliant` statement is not configured, the NETCONF server might issue an RPC reply that includes both an `<rpc-error>` element with a severity level of warning and an `<ok/>` element.

Platform and Infrastructure

- **DMA recovery mechanism (ACX Series)**—Starting in Junos OS Release 17.4R3, a recovery mechanism has been introduced that is triggered in case the router enters an Idle state on any DMA channels. The recovery mechanism resets the PFE reboot to recover from Idle state.

The following recovery message is logged in the RE syslog message:

```
CHASSISD_FPC_ASIC_ERROR: <FPC 0> ASIC Error detected errorno 0x0000ffff FPC
restart initiated
CHASSISD_IFDEV_DETACH_FPC: ifdev_detach_fpc(0)
```

The following recovery message is logged in the PFE syslog message:

```
BCM DMA channel error detected
Resetting the PFE
```

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 (ACX 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 are correct to the numbers of days. Before this fix, for example, for a 1-year subscription license, the duration was generated as 364 days. After the fix, the duration of the 1-year subscription now shows as 365 days.

See [show system license](#).

Subscriber Management and Services

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

The new behavior is as follows:

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

In earlier releases:

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

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

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

SEE ALSO

[New and Changed Features | 14](#)

[Known Behavior | 20](#)

[Documentation Updates | 29](#)

[Known Issues | 21](#)

[Resolved Issues | 25](#)

[Migration, Upgrade, and Downgrade Instructions | 30](#)

Known Behavior

IN THIS SECTION

- [General Routing | 20](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 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.

General Routing

- Shared-buffer maximum default for IFL Queues is 66%, independent of the shared-buffer maximum knob under IFL scheduler configuration. [PR1275796](#)
- With enterprise profile, with multiple masters configured, PTP servo gets stuck in FREERUN state after the master is failed by disabling the IFL. [PR1281798](#)

- Error messages seen on loading basic iflset configuration on ACX5000 Junos routers. IFLSet in hierarchical-scheduler is not supported for HCOS in ACX5000. **ACX5000 fpc0 ACX_COS_HALP(acx_hqos_update_iflset_stats:xxxx): Invalid Queue index for iflset x ACX5k fpc0 ACX_COS_HALP(acx_hqos_update_iflset_stats:xxxx): Invalid Queue index for iflset y.** These log messages are harmless and there is no traffic impact. [PR1290166](#)

SEE ALSO

New and Changed Features	 14
Changes in Behavior and Syntax	 16
Documentation Updates	 29
Known Issues	 21
Resolved Issues	 25
Migration, Upgrade, and Downgrade Instructions	 30

Known Issues

IN THIS SECTION

- [General Routing](#) | [22](#)
- [Interfaces and Chassis](#) | [24](#)
- [Layer 2 Features](#) | [24](#)
- [MPLS](#) | [24](#)
- [Routing Protocols](#) | [24](#)
- [Virtual Chassis](#) | [24](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 for the ACX Series Universal Metro Routers.

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

General Routing

- Aggregate interface on ACX Series routers is permanently down after reboot, when **link-speed** is configured on 12.3X54-D10.6. [PR1022248](#)
- Forwarding when using non-existing SSM map source address in IGMPv3 instead of pruning. This is a day 1 design issue which needs to be redesigned. The impact is more, but definitely this needs some soaking time in DCB before it gets ported in previous versions. [PR1126699](#)
- When ACX 2100/2200 are used as ingress PE routers for Layer 2 circuit connections, and the PE-CE interface (UNI) is an aggregated Ethernet interface, then upon MPLS path switchover, the traffic can get blackholed. [PR1194551](#)
- ACX1000/ACX2000/ACX4000 does not support EVPN, therefore this PR removes EVPN CLI on these platforms. [PR1208248](#)
- Under certain scenarios, if VPLS instances and Layer 3 NNI interfaces are deleted together in the same commit, then a traffic duplication is observed for the VPLS traffic. To avoid such instances, it is recommended to delete or deactivate the Layer 3 NNI interfaces and VPLS instances in separate commits. [PR1260156](#)
- Junos CLI show class-of-service interfaces queue <ifl> does not display Queue buffer usage per logical interface. However the same can be viewed using PFE shell command. [PR1272822](#)
- In normal/software MAC learning mode, when incremental MAC traffic of higher range then the profile is received then after feb restart the MAC entries will not been seen in the software CLI alone though present in the hardware table. [PR1277436](#)
- OCM 100FX SFPs o with this part No. are not supported in this release. [PR1279202](#)
- On ACX5000, the buffer is corrupted on port 0 (*/*/0) and error message **MACDRAINTIMEOUT** and **dcbcm_check_stuck_buffers** are observed, which could eventually lead to port 0 (*/*/0) flapping. [PR1284590](#)
- There is a conflict when LACP packet comes in untagged/prio-tagged VPLS IFL. In the earlier stage of pipeline, filter entry to snoop LACP packet takes higher precedence over filter entry to assign SVP/SrcGport for untagged/prio-tagged VPLS IFL. Since the "interface-specific/input-list" firewall matches SVP/SrcGport in later stage of pipeline, the LACP packets are not hitting the firewall. [PR1346380](#)
- IFL classifier info should not be shown in output of **show class-of-service interface <ifd>** on ACX5000. [PR1353828](#)
- As part of the pic_periodic, before setting the port to master/slave mode, AN bit is checked if AN is complete and this would return if AN is still in progress. Since An was disabled, this port wasn't set to either mode and this was going on in a loop causing the CPU to go high. [PR1360844](#)
- The remote fault signalling is not supported for 1G fiber SFP during Auto-negotiation. Therefore in releases without the fix of this PR, we get cosmetic log error under **show interfaces extensive Link partner: Link mode: Full-duplex, Flow control: None, Remote fault: Down, Reason: Link partner offline. RFI ignored since AN is in default mode.** [PR1362490](#)

- Because of a race condition, in which the **class-of-service** configuration request for an interface is received before the e1-interface is created, a circuit with specified class-of-service parameters is created. Because of this, the interface creation fails resulting in traffic not flowing on the e1-interface and then (if e1-interfaces are further disabled or enabled) a core file is generated. [PR1378747](#)
- On Junos OS Release 17.3 and later releases, ACX5000, Packet Forwarding Engine syslog frequently shows the following errors messages: **acx_cos_tcp_bind_queues:736 parent acx_cos_tcp_ifd for ifd:ae0 doesn't exist for ifl:549 In 17.3R3-S1**. The error logs appear only from time to time, and this can be related with an interface flap. In Junos OS Release 18.1R3, the logs appear constantly, without any interface flap. This message is related to HCOS checking (even without HCOS configured). In software fix, we should check if the aggregate interface has HCOS configured or not. If not, we should return gracefully from this function without throwing this error. This is a harmless message. [PR1392088](#)
- On ACX1000/2000/4000/5048/5096 platforms, after a new child IFL with VLAN and filter is added on an aggregated Ethernet IFD or changing the VLAN ID of a child IFL with filter, traffic over the AE IFD might get filtered with that filter on the child IFL. For example: ae-0/0/0 is an IFD and ae-0/0/0.100 is an IFL. [PR1407855](#)
- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. [PR1411015](#)
- Interface with SFP-1FE-FX transceiver optic (740-021487) does not come UP on ACX series routers. [PR1439384](#)

Interfaces and Chassis

- When an unnumbered interface is binding to an interface which has more than one IP address and one of the IPs is deleted, the family inet of the unnumbered interface might get deleted. The issue results in traffic loss for all the services that rely on the family inet of the unnumbered interface. Configure **preferred-source-address** on the unnumbered interface will prevent deletion of the IP hence avoiding the deletion of the family inet of the unnumbered interface. [PR1412534](#)

Layer 2 Features

- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case the subscriber may need more time to get IP address assigned. The subscriber may remain in this state until it's lease expires if it has previously bound with the address in the option 50. [PR1435039](#)

MPLS

- Dynamically configured RSVP LSPs for LDP link protection might not come up after disabling/enabling protocol MPLS. [PR1432138](#)

Routing Protocols

- With IS-IS configured and in a very rare case, memory corruption might occur, this might cause rpd crash continuously. [PR1455432](#)

Virtual Chassis

- ACX5000 reports false parity error messages such as **soc_mem_array_sbusdma_read**. The ACX5000 SDK might raise false alarms for parity error messages such as **soc_mem_array_sbusdma_read**. This is a false positive error message. [PR1276970](#)

SEE ALSO

[New and Changed Features | 14](#)

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

[Known Behavior | 20](#)

[Documentation Updates | 29](#)

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.4R3 | 25
- Resolved Issues: 17.4R2 | 27
- Resolved Issues: 17.4R1 | 29

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: 17.4R3

Class of Service (CoS)

- CoS is incorrectly applied on Packet Forwarding Engine, leading to an egress traffic drop. [PR1329141](#)
- Firewall process crash might be seen with Multifield Classifier configuration. [PR1436894](#)

General Routing

- SNMP MIB walk/get/set on jnxDomCurrentTable and jnxDomNotifications might fail on ACX platforms. [PR1076943](#)
- On ACX5000 platform, if scaled logical interfaces exist, the logical interfaces might not all come up. [PR1229492](#)
- The 1G copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. [PR1286709](#)
- Incorrect packet statistics are reported in the ifHCInUcastPkts OID. [PR1306656](#)
- ACX Series routers support from dual-tagged through untagged packets Layer 3 traffic. [PR1307666](#)
- Port XE-0/3/0 did not turn up. [PR1328207](#)
- bcmDPC task is high eventhough Interuppt START_BY_START flag set to 0. [PR1329656](#)
- The fxpc process might use high CPU on ACX5000 after upgrade. [PR1360452](#)

- On a ACX ring topology, after link between ACX and MX flap, VPLS RI on PE (MX) have no MAC of CE over layer 2 circuit. [PR1360967](#)
- ARP reply drops when you add temporal buffer-size on the NNI interface. [PR1363153](#)
- Commit error is seen when configuring **mac-table-size** under bridge domain after the upgrade to Junos OS Release 15.1R7. [PR1364811](#)
- ACX5000: **fpc0 (acx_rt_ip_uc_lpm_install:LPM route add failed) Reason : Invalid parameter after configuring lpm-profile.** [PR1365034](#)
- VPLS with "vlan-id-list" is not working properly in some releases when the link between a PE device and a CE device is an aggregated Ethernet interface with a single member link and child physical interface flap. [PR1365894](#)
- The fxpc might crash after an interface is changed on ACX5000 routers. [PR1378155](#)
- The Layer 2 circuit might stop forwarding traffic when one core interface flapping happens. [PR1381487](#)
- The DMA failure errors might be seen when the cache flush or the cache is full. [PR1383608](#)
- On ACX led on GE interface goes down when speed 10M is added. [PR1385855](#)
- On ACX Series platforms the **forwarding-option dhcp-relay forward-only** knob stops working and the DHCP packets are dropped. [PR1392261](#)
- Certain builds of Junos OS do not allow you to upgrade or commit configuration changes when the SI service interface is used. [PR1393729](#)
- [ACX] MTU is not properly applied - and output of - ping mpls l2circuit sweep is giving lower values than expected. [PR1393947](#)
- ACX5048 rpm rfc2544-benchmarking test failing to start. [PR1395730](#)
- FPC might crash after offline or online MIC-3D-16CHE1-T1-CE-H. [PR1402563](#)
- ACX drops DNS responses which contain an underscore. [PR1410062](#)
- VPLS traffic might stop across ACX5000 with the aggregated Ethernet interface. [PR1412042](#)
- Junos PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- Number of inet-arp policers implemented on ACX 5000 has been increased from 16 to 64. [PR1413807](#)
- The swap memory is not initialized on boot on ACX5048/5096. [PR1415898](#)
- CoS table error can sometimes cause traffic outages and SNMP timeouts if the optic is plugged out and inserted back in. [PR1418696](#)
- High CPU usage on fxpc process might be seen on ACX5000 platform. [PR1419761](#)
- The FPC/fxpc crash might be observed on ACX platforms. [PR1427362](#)
- The l2cpd process might crash and generate a core dump when interfaces are flapping. [PR1431355](#)
- In ACX platforms, **no-vrf-propagate-ttl** might not work after activate or deactivate of CoS configuration. [PR1435791](#)

- In ACX Series, auto exported route between VRFs might not reply for ICMP echo requests. [PR1446043](#)
- 2circuit with a "backup-neighbor" (hot-standby) configured might stop forwarding traffic after failovers. [PR1449681](#)

Layer 2 Features

- The traffic with triple or more 802.1Q tags might fail to forward. [PR1415769](#)

Routing Protocols

- ACX5000: console management port device authentication credentials are logged in clear text (CVE-2019-0069). [PR1408195](#)
- Loopback address exported into other VRF instance might not work on ACX Series platforms. [PR1449410](#)
- MPLS LDP may still use stale MAC of the neighbor even if the LDP neighbor's MAC changes. [PR1451217](#)

Services Applications

- The spd might crash when **any-ip** is configured in the 'from' clause of the NAT rule with the static translation type. [PR1391928](#)

Resolved Issues: 17.4R2

Layer 2 Ethernet Services

- DHCPv6 relay ignores replies from server when renewing. [PR1354212](#)

Platform and Infrastructure

- On Junos OS, the **next-hop** index allocation fails and private index space get exhausted through incoming ARP requests to management interface (CVE-2018-0063). [PR1360039](#)
- DFW filter related errors seen while running tdm script. [PR1175190](#)
- MPLS LSP are being affected due to NH failed to be programmed. [PR1195419](#)
- Several error logs are seen on ACX Series router when link in primary path of LSP is flapped. [PR1204714](#)
- Transit ARP packets are being punted to the Routing Engine. [PR1263012](#)
- Common software fix for PR1204589 and PR1256073 that addresses Traceroute behavior while selecting the source address and adding CLI command for the same to configure the same. [PR1279191](#)
- ACX/AMX:fxpc core file is observed during unified ISSU. [PR1318771](#)
- On ACX platforms, network events might cause Layer2circuit traffic forwarding to fail with the "Table Full" message. [PR1319591](#)
- With **auto-installation** usb configured, interface related commits might not take effect due to a dcd error. [PR1327384](#)

- The major alarm about **Fan & PSU Airflow direction mismatch** might be seen by removing management cable. [PR1327561](#)
- The IPv6 service outage might occur after executing **clear ipv6 neighbor**. [PR1330791](#)
- ACX: Stale filter entries are present in TCAM. [PR1334784](#)
- The DHCP negotiations might fail and eventually cause outage if scaling number of DHCP clients reboot at the same time. [PR1335957](#)
- Unable to commit multiple ethernet-ring instances on ACX Series routers. [PR1337497](#)
- The Arp-reply packet might be dropped in a I2-circuit secondary path when using ieee-802.1 classifier. [PR1341126](#)
- [ACX5000] IPv6 /64 route is not installed in Packet Forwarding Engine for VRF **routing-instance** when lpm prefix-65-127 = disable. [PR1341714](#)
- PR to reduce **egress-vlan-xlate** entries in BD with **vlan-id-list**. [PR1343028](#)
- ACX5000: Traffic destined for specific ip within a subnet gets blackholed. [PR1345098](#)
- Filter is not working properly when applied using **input-list**. [PR1346380](#)
- NAT might not work and the spd might crash. [PR1346546](#)
- fxpc will crash on PFE command **show pfe context_vlan**. [PR1349721](#)
- ifl classifier info should not be shown in output of show class-of-service interface <ifd> on ACX5000. [PR1353828](#)
- On ACX Series routers, ARP policer for IFL is not working. [PR1356170](#)
- Memory leak is observed when ACX is under high traffic load. [PR1358127](#)
- ACX is incorrectly allowing to configure higher values in burst-size-limit than what the HW support. [PR1361482](#)
- [ACX5000] IPsec SA as OSPFv3 authentication is not working in Junos OS Release 16.2R2 and Release 17.3R2. [PR1363487](#)
- PCEP delegation-priority might not be honored. [PR1365560](#)
- The 'commit' or 'commit check' might fail due to the error **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)

Resolved Issues: 17.4R1

Layer 2 Ethernet Services

- JDHCPD memory leak during dhcp/pppoe login or logout loop. [PR1289780](#)

Platform and Infrastructure

- FAN on ACX Series routers intermittently gives **FAN Failure** alarms. [PR1127846](#)
- ACX1100 with midplane part no 650-062965 might fail to initialize FPGA. [PR1134335](#)
- Error messages **chassisd[1825]: pvidb_get_root_node: Error(2) retrieving rootnode value** might be seen. [PR1198817](#)
- High CPU utilization is seen due to clksyncd process. [PR1238067](#)
- ACX does not forward DHCP-RELAY requests with IRB interface after upgrade. [PR1243687](#)
- Tagged/untagged LLDP, LACP packets dropped on VPLS CE facing aggregate Ethernet interface. [PR1245242](#)
- The 1G copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. [PR1286709](#)
- ACX2x00-AC is reporting false PEM0 alarms periodically. [PR1310488](#)
- Error syslog on output/egress firewall filter on ACX Series routers. [PR1316588](#)
- ACX/AMX:fxpc core file is observed during unified ISSU. [PR1318771](#)

SEE ALSO

New and Changed Features 14
Changes in Behavior and Syntax 16
Known Behavior 20
Documentation Updates 29
Known Issues 21
Migration, Upgrade, and Downgrade Instructions 30

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 for the ACX Series documentation.

SEE ALSO

New and Changed Features 14
Changes in Behavior and Syntax 16
Known Behavior 20
Known Issues 21
Resolved Issues 25
Migration, Upgrade, and Downgrade Instructions 30

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 30

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

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

SEE ALSO

[New and Changed Features | 14](#)

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

[Known Behavior | 20](#)

[Documentation Updates | 29](#)

[Known Issues | 21](#)

[Resolved Issues | 25](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- [New and Changed Features | 32](#)
- [Changes in Behavior and Syntax | 39](#)
- [Known Behavior | 44](#)
- [Known Issues | 47](#)
- [Resolved Issues | 53](#)
- [Documentation Updates | 65](#)
- [Migration, Upgrade, and Downgrade Instructions | 66](#)

These release notes accompany Junos OS Release 17.4R3 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R3 New and Changed Features | 33](#)
- [Release 17.4R2 New and Changed Features | 33](#)
- [Release 17.4R1 New and Changed Features | 34](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R3 for the EX Series.

NOTE: Starting in Junos OS Release 17.4R1, MC-LAG is not supported on EX switches except for EX9200. Use the Virtual Chassis feature instead to provide equivalent functionality.

NOTE: The following EX Series switches are supported in Junos OS Release 17.4R3: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.4R3, J-Web is supported on the EX4300 and EX4600 switches in both standalone and Virtual Chassis setup.

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

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

For details about the J-Web distribution model, see [Release Notes: J-Web Application Package Release 17.4A1 for EX4300 and EX4600 Switches](#).

Release 17.4R3 New and Changed Features

- There are no new features or enhancements to existing features for EX Series Switches in Junos OS Release 17.4R3.

Release 17.4R2 New and Changed Features

EVPNs

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

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

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

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

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

Restoration Procedures and Failure Handling

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

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

Release 17.4R1 New and Changed Features

Hardware

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

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

Authentication, Authorization and Accounting (AAA)

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

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

EVPNs

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

The EVPN-MPLS interworking feature offers the following benefits:

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

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

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

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

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

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

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

Junos OS XML API and Scripting

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

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

Layer 2 Features

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

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

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

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

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

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

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

Management

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

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

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

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

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

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

Multicast

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

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

Routing Protocols

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

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

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

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

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

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

Software Installation and Upgrade

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

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

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

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

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

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

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

SEE ALSO

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

[Known Behavior | 44](#)

[Known Issues | 47](#)

[Resolved Issues | 53](#)

[Documentation Updates | 65](#)

[Migration, Upgrade, and Downgrade Instructions | 66](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [EVPNs | 39](#)
- [Interfaces and Chassis | 40](#)
- [Management | 40](#)
- [Multicast | 40](#)
- [Network Management and Monitoring | 40](#)
- [Platform and Infrastructure | 41](#)
- [Routing Protocols | 42](#)
- [Security | 42](#)
- [Software Licensing | 42](#)
- [Subscriber Management and Services | 42](#)
- [Virtual Chassis | 43](#)

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 17.4R3 for the EX Series.

EVPNs

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

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

Interfaces and Chassis

- **No support for performance monitoring on AE Interfaces (EX4300)**—Y.1731 performance monitoring (PM) over aggregated Ethernet interfaces is not supported on EX4300 switches. [See [sla-iterator-profile](#).]

Management

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

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

Multicast

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

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

Network Management and Monitoring

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

Before this change:

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

After this change:

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

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

- **SNMP syslog messages changed (EX Series)**—Starting in Junos OS Release 17.4R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD — AgentX master agent failed to respond to ping. Attempting to re-register
NEW — AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD — NET-SNMP version %s AgentX subagent connected
NEW — NET-SNMP version %s AgentX subagent Open-Sent!

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

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

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

- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (EX Series)**—Starting in Junos OS Release 17.4R3, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an <ok/> element and one or more <rpc-error> elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an <rpc-error> element with a severity level of warning and an <ok/> element.

Platform and Infrastructure

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the show interfaces mc-ae extensive command. The output now displays the following two additional fields:
 - Local Partner System ID?LACP partner system ID as seen by the local node.
 - Peer Partner System ID?LACP partner system ID as seen by the MC-AE peer node.

Previously, the show interfaces mc-ae extensive command did not display these additional fields.

[See [show interfaces mc-ae](#)]

Routing Protocols

- **Change in the default behavior of advertise-from-main-vpn-tables configuration statement**—BGP now advertises EVPN routes from the main bgp.evpn .0 table. You can no longer configure BGP to advertise the EVPN routes from the routing instance table. In earlier Junos OS Releases, BGP advertised EVPN routes from the routing instance table by default.

[See [advertise-from-main-vpn-tables](#).]

Security

- **Support for logging SSH key changes**—Starting with Junos OS Release 17.4R1, the configuration statement **log-key-changes** is introduced at the `[edit system services ssh]` hierarchy level. When **log-key-changes** configuration statement is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time **log-key-changes** was enabled. If **log-key-changes** was never enabled, then Junos OS logs all the authorized SSH keys.
- **Syslog or log action on firewall drops packets (EX4600 switches)**—Starting in Junos OS 17.4R3, if you configure a syslog or log action on an ingress firewall filter, control packets, and ICMP packets sent to the Routing Engine might be dropped.

Software Licensing

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

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

Subscriber Management and Services

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

- Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

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

In earlier releases:

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

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

Virtual Chassis

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

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

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

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

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

SEE ALSO

New and Changed Features	 32
Known Behavior	 44
Known Issues	 47
Resolved Issues	 53
Documentation Updates	 65
Migration, Upgrade, and Downgrade Instructions	 66

Known Behavior

IN THIS SECTION

- [EVPN](#) | [45](#)
- [High Availability \(HA\) and Resiliency](#) | [45](#)
- [Infrastructure](#) | [45](#)
- [Interfaces and Chassis](#) | [45](#)
- [Junos Fusion Enterprise](#) | [45](#)
- [Platform and Infrastructure](#) | [46](#)
- [Routing Protocols](#) | [46](#)
- [Virtual Chassis](#) | [46](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 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.

EVPN

- 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](#)

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)

Infrastructure

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

Interfaces and Chassis

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

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 150000. [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](#)

Platform and Infrastructure

- On EX4300 and EX4600 switches, if a remote analyzer has an output IP address that is reachable through a route learned by BGP, the analyzer might be in a down state. [PR1007963](#)
- On an EX4300 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On EX4300 switches, when 802.1X single-suplicant authentication is initiated, multiple "EAP Request Id Frame Sent" packets might be sent. [PR1163966](#)
- On EX4300 10G links, preexisting MACsec sessions might not come up after the following events: Process (pfex, dot1x) restart or system restart link flaps [PR1294526](#)
- Repeated mode switching by enable/disable interface or setting and removing otn-options rate can cause dfe tuning to get stuck for a long time on the CFP2-DCO tunable DWDM optics resulting interfaces being down for around 30 minutes. [PR1452597](#)

Routing Protocols

- mcsnoopd might crash when all the core facing interfaces that are part of the L2 domain have flapped and it is attempting to flood a packet received over a CE interface, over the core-facing interfaces. [PR1329694](#)

Virtual Chassis

- Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on an QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur and its considered to be known behavior. [PR1347902](#)

SEE ALSO

[New and Changed Features | 32](#)

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

[Known Issues | 47](#)

[Resolved Issues | 53](#)

[Documentation Updates | 65](#)

[Migration, Upgrade, and Downgrade Instructions | 66](#)

Known Issues

IN THIS SECTION

- [Authentication and Access Control | 47](#)
- [EVPN | 47](#)
- [General Routing | 48](#)
- [Infrastructure | 49](#)
- [Interfaces and Chassis | 50](#)
- [Junos Fusion Enterprise | 50](#)
- [Layer 2 Ethernet Services | 50](#)
- [Layer 2 Features | 51](#)
- [Multicast | 51](#)
- [Platform and Infrastructure | 51](#)
- [Routing Protocols | 52](#)
- [Subscriber Access Management | 52](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 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 and Access Control

- The output of the command **show lldp detail** is not consistently displayed in any organized order. [PR1390760](#)

EVPN

- 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](#)
- In an EVPN environment, proxy ARP and ARP suppression is enabled on the PE device by default for reducing the flooding of ARP packets. However, in the case of ARP probe packets used in the process of Duplicate Address Detection (DAD), the client might treat the IP address that it is in use as duplicated address after receiving the proxied packets from PE device. [PR1427109](#)

General Routing

- From the code analysis, the CPU rate limiting and corresponding queue points to 100 pps in Junos OS Release 12.3 for ARP traffic. But in case of Junos OS Release 11.4, the rate limiter value was 3 Kpps. [PR1165757](#)
- On an EX9200-12QS line card, interfaces with the default speed of 10 Gigabit Ethernet are not brought down even when the remote end of a connection is misconfigured as 40 Gigabit Ethernet. [PR1175918](#)
- On EX4300 10-Gigabit links, preexisting MACsec sessions might not come up after the following events: process (pfex, dot1x) restart, system restart, or link flaps. [PR1294526](#)
- l2ald process may crash and generate a core file on EX Series VC when converted a trunk port to dot1x access port while tagged traffic is flowing. There might be a race-condition, where interface mode is being changed while traffic is running and l2ald has processed interface delete but dot1x has not. [PR1362587](#)
- Currently, other than QFX5100-24q and EX4600, PIC 1 is not supported on any other platforms inline with QFX5100. The **set chassis fpc 0 pic 1 port <x> channel-speed disable-auto-speed-detection** command cannot be used on PIC 1. This will result in a commit error **[edit chassis fpc 0 pic 1 port 2 channel-speed] channel-speed disable-auto-speed-detection PIC:1 not valid for Auto-speed-disable mode. error: configuration check-out failed**. So, if you want to disable auto-channelisation on PIC1, you have to disable auto-speed-detection for whole FPC. **set chassis fpc 0 auto-speed-detection disable**. [PR1362647](#)
- EX4300 Virtual Chassis systems might fail to register some jnxOperating SNMP OIDs related to the Routing Engines. This behavior is more likely if Virtual Chassis members 0 and 1 (FPC0 and FPC1) are not selected as Routing Engine. [PR1368845](#)
- Scale of 150 VRRP is not tested before, there are no issues observed for 100 VRRP groups. At the higher scale, there are no drops but traffic gets flooded for group beyond 100. [PR1371520](#)
- When **show** command is taking a long time to display results, the STP might change states as BPDUs are no longer processed and cause lots of outages. [PR1390330](#)
- On QFX5110 line of Series switches, uRPF check in strict mode will not work properly. [PR1417546](#)
- The issue is limited to the database related to MAC-MOVE scenario. When dhcp-security is configured, if multiple IPv4 and IPv6 client's MAC-MOVE occur, the jdncpd might consume 100 percent CPU and jdncpd crashes. [PR1425206](#)
- Multiple EX Series switches might be unable to commit baseline configuration after zeroize {master:0}[edit] **root# commit check Mar 26 05:50:48 mustd: UI_FILE_OPERATION_FAILED: File /var/run/db/enable-process.data doesn't exist Mar 26 05:50:48 mgd[1938]: UI_FILE_OPERATION_FAILED: Failed to open /var/run/db/enable-process.data+ file error: Failed to open /var/run/db/enable-process.data+ file error: configuration check-out failed: daemon file propagation failed**. [PR1426341](#)

- On EX9200 line of switches, when configuring too many VLANs and interfaces under VSTP a commit error might occur **xSTP:Trying to configure too many interfaces for given protocol**. [PR1438195](#)
- When **mac-table-aging-time** is configured, the bridge domain sequence get incremented unnecessarily. As a result, the MACs get flushed when the change message is received by l2-learning daemon with new sequence number. [PR1403358](#)
- Micro BFD session with timer configured with less than 3x500ms (such as 3x100ms) might flap upon inserting a QSFP to other port. [PR1435221](#)
- On EX Series platforms, if particular 100G port is used, CPU might hang or interface might be stuck down on the 100G port. This issue might cause traffic disruption in the network. [PR1440526](#)
- A sequence issue is observed when Virtual Chassis member is rebooted in aggregated interface. After rebooting VC member, the Routing Engine kernel injects MAC entry to FPC that rebooted. Because of the sequence issue, the Routing Engine added MAC entry, originally source MAC entry, to FPC as remote MAC entry. And MAC entry is never be aged out because it is a remote entry. [PR1440574](#)
- On the EX9214 device, if the MACsec-enabled link flaps after reboot, the error **errorlib_set_error_log(): err_id(-1718026239)** is observed. [PR1448368](#)
- The sFlow sample packets might stop on one aggregated Ethernet member link if ingress sFlow is configured on the member link. This might cause inaccurate monitoring on the network traffic. [PR1449568](#)
- The l2ald and eventd processes are hogging 100 percent after **clear ethernet-switching table** command is issued. As a result, continuous syslog errors **l2ald[18605]: L2ALD_IPC_MESSAGE_INVALID: Invalid message received (message type 0, subtype 0): null message** are observed. [PR1452738](#)
- On EX4300 and EX4600 Virtual Chassis or VCF scenario with VXLAN used, when configuring a firewall filter and commit, the firewall filter might not be able to be applied in a particular VC/VCF member for TCAM space running out. [PR1455177](#)
- Syslog **timeout connecting to peer database-replication** is generated when command **show version detail** is issued. [PR1457284](#)
- On QFX5100 and EX4600 platforms, the fxpc (Packet Forwarding Engine manager) process might crash when multiple BGP IPv6 sessions (for instance around 500) are flapped and then restored at the same time. [PR1459759](#)
- When tunnel-services are configured on a PIC, the optics measurements that subscribed through gRPC might not be streamed. [PR1468435](#)

Infrastructure

- The **set system ports console log-out-on-disconnect** command does not work. [PR1146891](#)
- When an SNMP poll is performed for the following OID's, the backup Routing Engine returns the value 6 (6=down) for the FAN and 1 (1=unknown) for the PSU's, even though the FAN and PSU's are UP. Fan: 1.3.6.1.4.1.2636.3.1.13.1.6 PSU: 1.3.6.1.4.1.2636.3.1.13.1.6.2. [PR1360962](#)

- On EX Series platforms, when you configure a large number of firewall filters on some interfaces, the FPC crashes generating core files. [PR1434927](#)
- Packet Forwarding Engine sometimes does not come up after system reboot. Timeout is required to handle the fifo tx/rx error. Debug sysctls are been removed. Mutex been added to handle to race condition. [PR1454950](#)

Interfaces and Chassis

- On GRES switchover, VSTP port cost on aggregated Ethernet interfaces might get changed, leading to topology change. [PR1174213](#)
- When dynamic DHCP sessions are existing in the device, if multiple commits in parallel are performed, the commit might hang up. [PR1470622](#)

Junos Fusion Enterprise

- On a Junos Fusion, when using LLDP, the "Power via MDI" and "Extended Power via MDI" TLVs are not transmitted. [PR1105217](#)
- On a Junos Fusion Enterprise, when 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](#)
- In Junos Fusion Enterprise environment, when EX3400 is being used as Satellite Device (SD), the cascade port on aggregation device (AD) might go down after it's connected SD reboot. [PR1382091](#)
- In Junos Fusion Enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, loop-detect feature does not work for ports 0-23, since the loop detect filter is not properly applied. [PR1426757](#)
- In a Junos Fusion Enterprise environment, when traffic originates from a peer device connected to the aggregation device and the ICL is a LAG, there might be a reachability issue if the cascade port is disabled and traffic has to flow through the ICL LAG to reach the satellite device. As a workaround, use single interface as the ICL instead of a LAG. [PR1447873](#)

Layer 2 Ethernet Services

- On EX4300, EX4600, switches with spine-leaf scenario, when two or more than two underlay interfaces with ECMP are brought down on leaf devices, the multihop BFD overlay sessions between spines and leafs might flap. And if BFD flaps, the protocols depending on the BFD (typically, IBGP protocols) might also flap, that leads to traffic impact. [PR1416941](#)
- On EX Series platforms with service dhcp enabled, the jdhcpd_era log files constantly consume 121M of space out of 170M, resulting into file system full and traffic impact. Memory usage of **/var/log/** will reach 100 percent. [PR1431201](#)

- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet might be dropped. In such a case, the subscriber might need more time to get IP address assigned. The subscriber might remain in this state until it's lease expires if it has previously bound with the address in the option 50. [PR1435039](#)

Layer 2 Features

- **eswd[1200]: ESWD_MAC_SMAC_BRIDGE_MAC_IDENTICAL: Bridge Address Add: XX:XX:db:2b:26:81 SMAC is equal to bridge mac hence don't learn** is seen in syslog every few minutes on ERPS owner. The logs occur during ERPS PDU in ERPS setup. [PR1372422](#)
- On EX4600 platforms, if copper base SFP-T is used, it might not get up on physical layer and the MAC/ARP learning might not work if it gets up. The PR fixes both layer-1 and layer-2 issues in this scenario. [PR1437577](#)
- On EX Series platforms with STP disabled, the LLDP function might fail when a Juniper Networks device connects to a non-Juniper one. In this scenario, the LLDP PDU with destination MAC 01:80:c2:00:00:00, which is one of the three reserved MAC addresses for LLDP in IEEE 802.1AB, will be ignored by Juniper LLDP process, and this causes the LLDP function failure. This issue has a service impact. [PR1462171](#)

Multicast

- IGMP query packets might be duplicated between L2 interfaces with IGMP Snooping is enabled. [PR1391753](#)

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](#)
- On EX4300 switches, the software upgrade in FIPS mode fails with the following error: **ERROR: py-base-powerpc-18.1R1.9.tgz: not a signed package.** [PR1371427](#)
- The first IRB stops working on adding the second IRB to an aggregated Ethernet and then removing it. [PR1423106](#)
- On all Junos OS platforms, when a device is upgraded to a newer version and **retry-options** statement exists in the configuration file, after the upgrade, the older version of the login-attempts and login-locks exist on an upgraded device. Under these circumstances, the device might not be accessible through ssh/telnet/console and the sshd process might crash. [PR1435173](#)

- On EX4300, and EX4600 line of Series switches, DMA buffer leaking might hit once the next-hop of received traffics is not resolved and eventually cause an FPC/pfex to crash if the DMA buffer runs exhaustion. [PR1436642](#)
- In EX4300 switches when 1G SFP is connected to 10G port, autonegotiation (AN) is enabled, many issues like ARP, link down might be caused. Therefore, when AN is disabled somehow corrupting the TX_DISABLE field resulting in Laser Tx remain enabled when disabling and plug-out - plug-in. [PR1445626](#)

Routing Protocols

- On EX4300 and EX4600 Series switches, if host destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log/syslog action (for example, 'filter <> term <> then log/syslog'), such packets should not be dropped and reach the Routing Engine. [PR1379718](#)
- If IGMP v2 is used and proxy mode is used for igmp-snooping, multicast traffic might be dropped because by default proxy sends queries/reports in IGMP v3 version, until the device receives new IGMP v2 query or report. [PR1425621](#)
- On EX4600 with service provider (SP) style VLAN configuration (in this method, each VLAN-ID is locally significant to a physical interface), if **interface-mac-limit** or **mac-table-size** is configured (that is, software MAC learning is enabled) and the scale of MAC addresses on the box is more than 2000, traffic might be dropped after Q-in-Q enabled interface is flapped or a change is made to the vlan-id-list. [PR1441402](#)

Subscriber Access Management

- The authd reuse address too quickly before jdhcpd completely cleanup the old subscriber with flooding error log. The log such as **jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815**. [PR1402653](#)

SEE ALSO

[New and Changed Features | 32](#)

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

[Known Behavior | 44](#)

[Resolved Issues | 53](#)

[Documentation Updates | 65](#)

[Migration, Upgrade, and Downgrade Instructions | 66](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues:17.4R3 | 53](#)
- [Resolved Issues: 17.4R2 | 59](#)
- [Resolved Issues: 17.4R1 | 63](#)

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:17.4R3

Authentication and Access Control

- Dot1xd core file might be observed dot1x interface is configured with EAP-PEAP as an authentication protocol. [PR1322891](#)
- Without dot1x configuration, the syslog **dot1xd[2192]: task_connect: task PNACAUTH./var/run/authd_control addr /var/run/authd_control: Connection refused** is generated repeatedly. [PR1406965](#)

Class of Service (CoS)

- CoS is incorrectly applied on Packet Forwarding Engine, leading to egress traffic drop. [PR1329141](#)

EVPN

- A few minutes of traffic loss might be observed during recovery from link failure. [PR1396597](#)
- The device might proxy the ARP probe packets in an EVPN environment. [PR1427109](#)
- ARP request/NS might be sent back to the local segment by DF router. [PR1459830](#)

Forwarding and Sampling

- The l2ald process might observe memory leak on Junos OS platforms. [PR1455034](#)

General Routing

- The RE-PFE **out-of-sync** errors might be seen in syslog. [PR1232178](#)
- Syslogs contain messages with %PFE-3: fpc0 ifd null, port 28 dc-pfe: %USER-3: ifd null, port 28 : %PFE-3: fpc0 ifd null, port 29 dc-pfe: %USER-3: ifd null, port 29. [PR1295711](#)
- EX4300-32F MACsec session stays down on 1G or 10G links after events when events are performed with running traffic. [PR1299484](#)
- MACsec causes dot1xd JTASK_SCHED_SLIP or FPC disconnect. [PR1322302](#)
- QFX5000 platforms might display fpc0 error requesting CMTFPC SET INTEGER, illegal setting 37 observed after upgrade. [PR1340897](#)
- The 40-Gigabit interfaces might not forward traffic. [PR1349675](#)
- When VOIP VLAN is set as NATIVE VLAN on the port, the interface still shows up as a tagged interface and drops all untagged traffic. [PR1349712](#)
- The l2ald process might crash and generate a core file on EX Series Virtual Chassis when converting a trunk port to dot1x access port with tagged traffic flowing. [PR1362587](#)
- FPM board status is missing in SNMP MIB walk result. [PR1364246](#)
- OAM Ethernet connectivity-fault-management configured on an aggregated Ethernet interfaces is not supported and no commit error is observed. [PR1367588](#)
- Unable to use Ansible to collect RSI from EX9200. [PR1367913](#)
- IPv6 router advertisement (RA) messages can increase internal Kernel memory usage. [PR1369638](#)
- The dot1xd might crash when dot1xd receives incorrect reply length from the authd. [PR1372421](#)
- The interface might not flap when both flap-on-disconnect and port-bounce are sent. [PR1372619](#)
- MAC refresh packet might not be sent out from the new primary link after RTG failover. [PR1372999](#)
- The rpd process might crash when route flap and LSP flap occur with CBF enabled. [PR1374558](#)
- FPC might crash when flapping the output interface of analyzer or sampling. [PR1374861](#)
- RIPv2 update packets might not send with IGMP snooping enabled. [PR1375332](#)
- Unable to commit with a configuration of packet-length in egress firewall filter on EX9200. [PR1378901](#)
- ARP request packets might be sent out with 802.1Q VLAN tag. [PR1379138](#)
- The dot1x does not work with Microsoft NPS server. [PR1381017](#)
- IRB interface does not turn down when master of VC is rebooted or halted. [PR1381272](#)
- Constant memory leak might lead to FPC memory exhaustion. [PR1381527](#)

- ARP/Ethernet-table is pointing to down aggregated Ethernet interface if MTU is changed. [PR1385199](#)
- On EX9200 platforms, the warning message **prefer-status-control-active is used with status-control standby** might be seen whenever you commit an operation. [PR1386479](#)
- MAC learning might stop working on some LAG interfaces. [PR1389411](#)
- The input rate statistics might not increase if there are non-standard packets flow. [PR1389908](#)
- The dhcp-security binding table might not be updated due to the renew request with '0.0.0.0' value in 'ciaddr'. [PR1394341](#)
- The subscriber bindings might not be successful on EX Series platforms. [PR1396470](#)
- The authd might stop when issuing **show network-access requests pending** command during the authd restarting. [PR1401249](#)
- The TCP connection between ppmdd and ppmann might be dropped due to a kernel issue. [PR1401507](#)
- The STP does not work when aggregated interfaces number is "ae1000" or above in QFX5000 and "ae480" or above in EX Series switches. [PR1403338](#)
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface. [PR1403528](#)
- In a very rare situation the router can crash with VMCore when there is a logical interface deletion. [PR1404507](#)
- PEM alarm for backup FPC remains on master FPC though backup FPC is detached from Virtual Chassis. [PR1412429](#)
- Virtual Chassis might become unstable and FXPC core files are generated when there are a lot of configured filter entries. [PR1422132](#)
- MACsec connection on EX4600 platforms might not come back up after interface disconnect/reconnect. [PR1423597](#)
- The jdncpd might consume 100 percent CPU and crash if **dhcp-security** is configured. [PR1425206](#)
- Rebooting or halting Virtual Chassis member might cause 30 seconds down on RTG link. [PR1427500](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- The mc-ae interface might get stuck in waiting state in dual mc-ae scenario. [PR1435874](#)
- Commit check error for VSTP on EX9200s: **xSTP:Trying to configure too many interfaces for given protocol**. [PR1438195](#)
- The DHCP Snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it. [PR1438351](#)
- The EX4600 and QFX5100 Virtual Chassis might not come up after replacing Virtual Chassis port fiber connection with DAC cable. [PR1440062](#)
- DHCP snooping static binding not take effect after deleting and re-adding the entries. [PR1451688](#)

- Configuration change in VLAN all option might affect the per-VLAN configuration. [PR1453505](#)
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used [PR1458559](#)

Infrastructure

- Packets with the DEI/CFI bit set to 1 in the L2 header might not be forwarded. [PR1326855](#)
- Traffic might silently get dropped or discarded with indirect next hop and load balancing. [PR1376057](#)
- The kernel crash when GRES configuration is enabled and committed. [PR1376362](#)
- The traffic to the NLB server might not be forwarded if the NLB cluster works on multicast mode. [PR1411549](#)
- Some of EX Series platforms might generate vmcore by panic and reboot. [PR1456668](#)

Interfaces and Chassis

- The logical interfaces in EVPN routing instances might flap after committing configurations. [PR1425339](#)
- The traffic might be forwarded to incorrect interfaces in MC-LAG scenario. [PR1465077](#)

Junos Fusion Enterprise

- PoE over LLDP negotiation is not supported on Junos Fusion Enterprise setup. [PR1366106](#)
- New satellite device cannot be added to the Fusion scenario. [PR1374982](#)
- The l2ald might crash while issuing **clear ethernet-switching table persistent-learning** command. [PR1409403](#)
- Extended ports in JFE do not adjust MTU when VoIP is enabled. [PR1411179](#)
- The traffic might get silently dropped or discarded in Junos Fusion Enterprise scenario with dual-AD. [PR1417139](#)
- Loop-detect feature is not working in Junos Fusion Enterprise. [PR1426757](#)

Layer 2 Ethernet Services

- Junos OS core file jdhcpd.core.0 is found in dhcpv6_packet_handle. [PR1329390](#)
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. [PR1373807](#)
- The malfunction of core isolation feature in EVPN-VXLAN scenarios causes traffic drop. [PR1417729](#)
- The DHCP DECLINE packets are not forwarded to DHCP server when forward-only is set within dhcp-reply. [PR1429456](#)
- On EX9200, DHCP-relay is stripping the 'GIADDR' field in messages towards the DHCP clients. [PR1443516](#)

Layer 2 Features

- RTG MAC refresh packets will be sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flap. [PR1389695](#)
- The traffic with triple or more 802.1Q tags might fail to forward. [PR1415769](#)

Layer 3 Features

- The l2ald might crash when issuing **clear ethernet-switching table persistent-learning**. [PR1381739](#)

Network Management and Monitoring

- Over temperature trap does not send out even though there is temperature hot alarm. [PR1412161](#)

Platform and Infrastructure

- Ping does not go through device after WTR timer expires in ERPS scenario. [PR1132770](#)
- Packet drop might be seen on the logical tunnel interfaces lt-x/2/x or lt-x/3/x. [PR1345727](#)
- Interface flapping is seen on EX4300 switch. [PR1361483](#)
- The LLDP TLV with the incorrect switch port capabilities might be sent. [PR1372966](#)
- On EX4300 switches, the software upgrade in FIPS mode fails and an error message **py-base-powerpc-18.1R1.9.tgz: not a signed package** is observed. [PR1371427](#)
- ECMP route installation failure with log messages like unilist install failure might be observed on EX4300 device. [PR1376804](#)
- Packet drops on interface if the statement **gether-options loopback** is configured. [PR1380746](#)
- Traffic loss seen in Layer 2 VPN with GRE tunnel. [PR1381740](#)
- On EX4300 loss-priority high set to multicast packets is overridden. [PR1382893](#)
- EX4300 device chooses incorrect bridge-id as RSTP bridge-id. [PR1383356](#)
- After EX4300 Virtual Chassis is upgraded to Junos OS Release 18.2R1 **jdhcpd: shmlog: shared log header is NULL** log message can be seen. [PR1387871](#)
- Unicast DHCP request might get misforwarded to backup RTG link. [PR1388211](#)
- ICMPV6 packets are not classified with static or multifield forwarding-class mapping. [PR1388324](#)
- Layer 3 IP route might be deleted after L2 next hop change is seen. [PR1389688](#)
- Continuous log messages get printed on EX4300 after upgrading to Junos OS Release 17.4 or later. [PR1391942](#)
- On EX4300 Series switches when a firewall filter is applied to a loopback interface, other firewall filters for multicast traffic might fail. [PR1392082](#)
- EX4300 OAM LFM might not work on **extended-vlan-bridge** interface with native VLAN configured. [PR1399864](#)

- Traffic drop is seen on EX4300 when 10G fiber port is using 1 Gigabit Ethernet SFP optics with autonegotiation enabled. [PR1405168](#)
- Untagged traffic is single-tagged in Q-in-Q scenario on EX4300 platforms. [PR1413700](#)
- In EX4300 few ports might remain in dot1x 'connecting' state and fail to transition to 'authenticated' state. [PR1417270](#)
- On EX4300 runt counter is never incremented. [PR1419724](#)
- EX4300 does not send fragmentation needed message when MTU is exceeded with DF bit set. [PR1419893](#)
- The pfex process might crash and core files might be generated when SFP is reinserted. [PR1421257](#)
- Traffic loss is seen when one of logical interfaces on LAG is deactivated or deleted. [PR1422920](#)
- Auditd crashes when accounting RADIUS server not reachable. [PR1424030](#)
- SNMP (ifHighSpeed) value does not appear properly for VCP interfaces only. It appears as zero. [PR1425167](#)
- Interface flapping scenario might lead to ECMP nexthop install failure on EX4300s. [PR1426760](#)
- IPv6 traffic might be dropped when static /64 Ipv6 routes are configured. [PR1427866](#)
- EX4300 does not drop FCS frames with CRC error on XE interfaces. [PR1429865](#)
- Unicast ARP requests are not replied with **no-arp-trap** option. [PR1429964](#)
- EX4300 enables the soft error recovery feature on the Packet Forwarding Engine, which can automatically detect the Packet Forwarding Engine parity error and recover by itself. [PR1430079](#)
- The ERPS failover does not work as expected on EX4300 device. [PR1432397](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- The PoE might not work after upgrading the PoE firmware on EX4300 platforms. [PR1446915](#)
- The firewall filters might not be created due to TCAM issues. [PR1447012](#)
- NSSU cause a traffic loss again after the backup to master transitions. [PR1448607](#)
- ERP might not revert back to IDLE state after reload/reboot of multiple switches. [PR1461434](#)

Routing Protocols

- The PPM mode for BFD session in EX4300 is centralized and not distributed by default. [PR1361800](#)
- EX4300 might drop incoming IS-IS hello packets when IGMP or MLD snooping is configured. [PR1400838](#)
- On Junos OS EX4600 switches, console management port device authentication credentials are logged in clear text. [PR1408195](#)
- ICMPv6 RA packets generated by Routing Engine might be dropped on the backup member of Virtual Chassis if **igmp-snooping** is configured. [PR1413543](#)

- Error message **RPD_DYN_CFG_GET_PROF_NAME_FAILED: Get profile name for session XXX failed: -7**, might be seen in syslog after restarting routing daemon. [PR1439514](#)
- The bandwidth value of the DDOS-protection might cause the packets loss after the device reboot. [PR1440847](#)
- Junos OS BFD sessions with authentication flaps after a certain time. [PR1448649](#)
- Loopback address exported into other VRF instance might not work on EX Series platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even the LDP neighbor's MAC changes. [PR1451217](#)

Spanning Tree Protocols

- The l2cpd might crash if the VSTP traceoptions and VSTP VLAN all commands are configured. [PR1407469](#)

User Interface and Configuration

- Switch might unable to commit baseline configuration after zeroize. [PR1426341](#)

Virtual Chassis

- Current MAC address might change when deleting one of the multiple L3 interfaces. [PR1449206](#)

Resolved Issues: 17.4R2

Authentication and Access Control

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

EVPN

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

High Availability (HA) and Resiliency

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

Infrastructure

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

- There is support for archiving dmesg file `/var/run/dmesg.boot`. [PR1327021](#)
- Enabling mac-move-limit stops ping on flexible-vlan-tagging enabled interface. [PR1357742](#)
- The dot1x filter might be removed from the Packet Forwarding Engine when **static-mac-address** ages out or is learned by eswd. [PR1335125](#)

Interfaces and Chassis

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

Platform and Infrastructure

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

- A memory leak is seen for dot1xd. [PR1313578](#)
- The Fan speed might frequently fluctuate between normal and full for MX Series platform. [PR1316192](#)
- The interface with 1G SFP might go down if no-auto-negotiation is configured. [PR1315668](#)
- Replace the **show vlans evpn** command to the **show ethernet-switching evpn** command for the EX9200 line of switches.. [PR1316272](#)
- IGMPv3 on EX4300 does not have the correct outgoing interfaces in the Packet Forwarding Engine that are listed in the kernel. [PR1317141](#)
- The L2cpd core files might be seen if the interface is disabled under VSTP and enabled under RSTP. [PR1317908](#)
- The vmcore might be seen and the device might reboot after the ICL is changed from an aggregated Ethernet to a physical interface. [PR1318929](#)
- High latency might be observed between a master Routing Engine and another FPC. [PR1319795](#)
- VLAN might not be processed, which leads to improper STP convergence. [PR1320719](#)
- Multicast traffic might not be forwarded to one of the receivers. [PR1323499](#)
- MAC learning issue and new VLANs creation failure might happen for some VLANs on an EX4300 platform. [PR1325816](#)
- The L2cpd might create a core file. [PR1325917](#)
- Extra EAP request packets might be sent unnecessarily. [PR1328390](#)
- EX4300 crashes when it receives more than 120kpps ARPs on me0 interface. [PR1329430](#)
- EX Series switches do not send RADIUS request after modifying the interface-range configuration. [PR1326442](#)
- The major alarm **Fan & PSU Airflow direction mismatch** might be seen by removing the management cable. [PR1327561](#)
- The SNMP trap message is always sent out with log about **Fan/Blower OK** on an EX4300 Virtual Chassis switch. [PR1329507](#)
- When exhausting a TCAM table, the filter might be incorrectly programmed. [PR1330148](#)
- The Rpd process crashed and generated core files on the new backup Routing Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler** after disabling NSR and GRES. [PR1330750](#)
- The dot1xd might crash if ports in multi-supplcant mode flaps. [PR1332957](#)
- The interface on which the VSTP is disabled by CLI might stay in the **Discarding** state after rebooting the device. [PR1333684](#)
- STP BPDUs are not sent out on the other active child when the anchor FPC has no active child. [PR1333872](#)

- MQSS errors and alarms might occur when the interface goes down. [PR1334928](#)
- EX9208: vstp vlan all statement has created L2CPD core files are generated during Routing Engine switchover or commit. [PR1341246](#)
- EX4300 storm control does not generate any action log after adding an RTG configuration. [PR1335256](#)
- IGMP packets are forwarded out of an RTG backup interface. [PR1335733](#)
- An L2cpd memory leak appears on EX Series platforms with VoIP configured. [PR1337347](#)
- The **show spanning-tree statistics bridge** command output gives 0 for all VLAN instance IDs. [PR1337891](#)
- MAC source address filter with the configuration statement **accept-source-mac**. does not work if MAC move limit is configured. [PR1341520](#)
- MSTP might not work normally after permitting a commit. [PR1342900](#)
- The filter might not be programmed in the Packet Forwarding Engine even though TCAM entries are available. [PR1345296](#)
- Statistics daemon PFED might generate core files on an upgrade between certain releases. [PR1346925](#)
- After the EX9200 FPC comes online, the other FPC CPU might use 100 percent and has traffic loss for about 30 seconds. [PR1346949](#)
- On EX4300 or EX4600 switches the VLAN translation feature does not work for the control plane traffic. [PR1348094](#)
- On EX4300 platforms, traffic drop might happen if LLC packets are received with DSAP and SSAP as 0x88 and 0x8e. [PR1348618](#)
- Running RSI via console port might cause system crash and reboot. [PR1349332](#)
- EX4600 detects a **LATENCY OVER-THRESHOLD** event with the incorrect value. [PR1348749](#)
- Commit error observed if box is downgraded from Junos OS 18.2/18.3 release to Junos OS Release 17.3R3. [PR1355542](#)
- On EX4300 platforms (Virtual Chassis and standalone) running Junos OS Release 16.1R5 or Junos OS Release 16.1R6, a firewall filter with a syslog option is unable to send syslog messages to the syslog server. [PR1351548](#)
- A high usage chassis alarm in "/var" does not clear from the EX4300 Virtual Chassis when a file is copied from fpc1 (master) to fpc0 (backup). [PR1354007](#)
- The ports using an SFP-T transceiver might still be up after system halt. [PR1354857](#)
- The FPC might crash due to the memory leak caused by the VTEP traffic. [PR1356279](#)
- Some interfaces cannot be added under STP configuration. [PR1363625](#)
- On EX4300/EX4600 platforms, the l2ald process might crash in dot1x scenario. [PR1363964](#)
- Packet Forwarding Engine might crash if encountering frequent MAC move. [PR1367141](#)
- The **request system zeroize** non-interactively might not erase the configuration on EX4300. [PR1368452](#)

Routing Protocols

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

Resolved Issues: 17.4R1

Authentication, Authorization, and Accounting (AAA)

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

Class of Service (CoS)

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

EVPNs

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

Infrastructure

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

Junos Fusion Enterprise

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

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

Layer 2 Features

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

Network Management and Monitoring

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

Platform and Infrastructure

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

Port Security

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

User Interface and Configuration

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

Virtual Chassis

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

VLAN Infrastructure

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

SEE ALSO

New and Changed Features 32
Changes in Behavior and Syntax 39
Known Behavior 44
Known Issues 47
Documentation Updates 65
Migration, Upgrade, and Downgrade Instructions 66

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 for the EX Series switches documentation.

SEE ALSO

New and Changed Features 32
Changes in Behavior and Syntax 39
Known Behavior 44
Known Issues 47
Resolved Issues 53
Migration, Upgrade, and Downgrade Instructions 66

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 66

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

New and Changed Features 32
Changes in Behavior and Syntax 39
Known Behavior 44
Known Issues 47
Resolved Issues 53

Junos OS Release Notes for Junos Fusion Data Center

IN THIS SECTION

- New and Changed Features | 67
- Changes in Behavior and Syntax | 68
- Known Behavior | 68
- Known Issues | 69
- Resolved Issues | 70
- Documentation Updates | 71
- Migration, Upgrade, and Downgrade Instructions | 71

These release notes accompany Junos OS Release 17.4R3 for the Junos Fusion Data Center. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

There are no new features in Junos OS Release 17.4R3 for Junos Fusion Data Center.

SEE ALSO

Changes in Behavior and Syntax 68
Known Behavior 68
Known Issues 69
Resolved Issues 70
Documentation Updates 71

Changes in Behavior and Syntax

There are no changes in behavior and syntax for Junos Fusion Data Center in Junos OS Release 17.4R3.

SEE ALSO

New and Changed Features	67
Known Behavior	68
Known Issues	69
Resolved Issues	70
Documentation Updates	71
Migration, Upgrade, and Downgrade Instructions	71

Known Behavior

IN THIS SECTION

- [Junos Fusion Data Center](#) | 69

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 for Junos Fusion Data Center.

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

Junos Fusion Data Center

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

SEE ALSO

New and Changed Features 67
Changes in Behavior and Syntax 68
Known Issues 69
Resolved Issues 70
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

Known Issues

There are no known issues in hardware and software in Junos OS Release 17.4R3 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 67
Changes in Behavior and Syntax 68
Known Behavior 68
Resolved Issues 70
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Junos OS Release 17.4R3 | 70](#)
- [Resolved Issues: Junos OS Release 17.4R2 | 70](#)
- [Resolved Issues: Junos OS Release 17.4R1 | 70](#)

This section lists the issues fixed in the Junos OS Release 17.4R3 for Junos Fusion Data Center.

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

Resolved Issues: Junos OS Release 17.4R3

- The ppmmd process on the aggregation device might generate a core file when using authentication key-chain with BFD. [PR1375647](#)
- BUM traffic may get dropped on peer Fusion aggregation device when the link between the satellite device and local aggregate device goes down. [PR1384440](#)

Resolved Issues: Junos OS Release 17.4R2

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

Resolved Issues: Junos OS Release 17.4R1

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

SEE ALSO

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

Documentation Updates

This section lists the errata or changes in Junos OS Release 17.4R3 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 67
Changes in Behavior and Syntax 68
Known Behavior 68
Known Issues 69
Resolved Issues 70
Migration, Upgrade, and Downgrade Instructions 71

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 72](#)
- [Preparing the Switch for Satellite Device Conversion | 74](#)
- [Autoconverting a Switch into a Satellite Device | 76](#)
- [Manually Converting a Switch into a Satellite Device | 79](#)
- [Configuring a Switch into a Satellite Device Before Connecting It to a Junos Fusion Topology | 81](#)
- [Configuring Satellite Device Upgrade Groups | 82](#)

- [Converting a Satellite Device to a Standalone Device | 84](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 84](#)
- [Downgrading from Release 17.4 | 84](#)

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

To download and install Junos OS:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

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

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

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

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

All other customers, use the following command:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Preparing the Switch for Satellite Device Conversion

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

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

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

Customers with EX4300 switches, use the following command:

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

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

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

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

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

```
[edit]
```

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

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

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

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

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

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

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

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

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

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

Autoconverting a Switch into a Satellite Device

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

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

Before you begin:

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

To autoconvert a switch into a satellite device:

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

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

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

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

2. Log in to the aggregation device.

3. Configure the cascade ports.

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

```
[edit]
```

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

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

Examples:

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

```
[edit]
```

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

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

```
[edit]
```

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

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

```
[edit]
```

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

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

```
[edit]
```

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

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

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

```
[edit]
```

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

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

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

```
[edit]
```

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

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

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

```
[edit]
```

```
user@satellite-device# commit
```

After committing the configuration, associate a satellite software image to the upgrade group:

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

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

7. Enable automatic satellite conversion:

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

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

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

8. Commit the configuration:

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

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

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

Manually Converting a Switch into a Satellite Device

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

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

Before you begin:

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

To manually convert a switch into a satellite device:

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

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

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

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

Examples:

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

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

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

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

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

```
[edit]
```

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

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

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

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

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

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

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

After committing the configuration, associate a satellite software image to the upgrade group:

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

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

6. Manually configure the switch into a satellite device:

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

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

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

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

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

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

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

Before you begin:

- Ensure that your switch that will become a satellite device is running a compatible conversion release of Junos OS. See [Junos Fusion Hardware and Software Compatibility Matrices](#).
- Ensure that you have copied the satellite software onto the device that will become a satellite device.

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

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

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

For instance, to install the satellite software package **satellite-3.1R1.n-signed.tgz** stored in the `/var/tmp/` directory on the switch, where *n* is the spin number:

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

- To install satellite software onto a QFX5100 switch, use the **satellite-3.1R1.n-signed.tgz** satellite software package.
 - To install satellite software onto a EX4300 switch, use the **satellite-ppc-3.1R1.n-signed.tgz** satellite software package.
2. The device will reboot to complete the satellite software installation.

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

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

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

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

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

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

4. Commit the configuration.

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

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

Configuring Satellite Device Upgrade Groups

To simplify the upgrade process for multiple satellite devices, you can create a software upgrade group at the aggregation device, assign satellite devices to the group, and install the satellite software on a groupwide basis.

To create a software upgrade group and assign satellite devices to the group, include the **satellite** statement at the **[edit chassis satellite-management upgrade-groups upgrade-group-name]** hierarchy level.

To configure a software upgrade group and assign satellite devices to the group:

1. Log in to the aggregation device.
2. Create the software upgrade group, and add the satellite devices to the group.

```
[edit]
user@aggregation-device# set chassis satellite-management upgrade-groups
upgrade-group-name satellite satellite-member-number-or-range
```

upgrade-group-name is the name of the upgrade group, and the **satellite-member-number-or-range** is the member numbers of the satellite devices that are being added to the upgrade group. If you enter an existing upgrade group name as the **upgrade-group-name**, you add new satellite devices to the existing software upgrade group.

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

[edit]

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Downgrading from Release 17.4

To downgrade from Release 17.4 to another supported release, follow the procedure for upgrading, but replace the 17.4 **jinstall** package with one that corresponds to the appropriate downgrade release.

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

[New and Changed Features](#) | 67

[Changes in Behavior and Syntax](#) | 68

[Known Behavior](#) | 68

Known Issues | 69

Resolved Issues | 70

Documentation Updates | 71

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- New and Changed Features | 86
- Changes in Behavior and Syntax | 88
- Known Behavior | 89
- Known Issues | 90
- Resolved Issues | 91
- Documentation Updates | 93
- Migration, Upgrade, and Downgrade Instructions | 93

These release notes accompany Junos OS Release 17.4R3 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

New and Changed Features

IN THIS SECTION

- Release 17.4R3 New and Changed Features | 87
- Release 17.4R2 New and Changed Features | 87
- Release 17.4R1 New and Changed Features | 87

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

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

Release 17.4R3 New and Changed Features

There are no new features in Junos OS Release 17.4R3 for Junos Fusion Enterprise.

Release 17.4R2 New and Changed Features

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

Release 17.4R1 New and Changed Features

Junos Fusion Enterprise

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

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

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

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

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

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

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

The EVPN-MPLS interworking feature offers the following benefits:

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

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

SEE ALSO

Changes in Behavior and Syntax	 88
Known Behavior	 89
Known Issues	 90
Resolved Issues	 91
Documentation Updates	 93
Migration, Upgrade, and Downgrade Instructions	 93

Changes in Behavior and Syntax

IN THIS SECTION

- [Junos Fusion Enterprise](#) | [89](#)

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 17.4R3 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 86
Known Behavior 89
Known Issues 90
Resolved Issues 91
Documentation Updates 93
Migration, Upgrade, and Downgrade Instructions 93

Known Behavior

IN THIS SECTION

- [Junos Fusion Enterprise | 89](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

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

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

SEE ALSO

[New and Changed Features | 86](#)

[Changes in Behavior and Syntax | 88](#)

[Known Issues | 90](#)

[Resolved Issues | 91](#)

[Documentation Updates | 93](#)

[Migration, Upgrade, and Downgrade Instructions | 93](#)

Known Issues

IN THIS SECTION

- [Junos Fusion Enterprise | 90](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- In Junos Fusion Enterprise environment, when EX3400 is being used as Satellite Device (SD), the cascade port on Aggregation Device (AD) might go down after its connected SD reboot. [PR1382091](#)
- In Junos Fusion Enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, loop-detect feature does not work for ports 0-23, since the loop detect filter is not properly applied. [PR1426757](#)

SEE ALSO

New and Changed Features	 86
Changes in Behavior and Syntax	 88
Known Behavior	 89
Resolved Issues	 91
Documentation Updates	 93
Migration, Upgrade, and Downgrade Instructions	 93

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R3](#) | [91](#)
- [Resolved Issues: 17.4R2](#) | [92](#)
- [Resolved Issues: 17.4R1](#) | [92](#)

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: 17.4R3

- PoE over LLDP negotiation is not supported on Junos Fusion Enterprise setup. [PR1366106](#)
- New satellite device can not be added to the Fusion scenario. [PR1374982](#)
- The l2ald process might generate a core file if issuing the **clear ethernet-switching table persistent-learning** command. [PR1409403](#)
- Extended ports in Junos Fusion Enterprise do not adjust MTU when VoIP is enabled. [PR1411179](#)
- The traffic might get blackholed in Junos Fusion Enterprise scenario with dual-AD. [PR1417139](#)

Resolved Issues: 17.4R2

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

Resolved Issues: 17.4R1

- On Junos Fusion Enterprise, traffic shaping is not supported on the extended ports. [PR1268084](#)
- On a Junos Fusion Enterprise with dual aggregation devices (ADs), if you apply Routing Engine loopback filters and bring down the cascade port on one of the ADs, the satellite device (SD) on the AD where the cascade port is down goes to ProvSessDown due to a TCP session drop over the ICL interface. [PR1275290](#)
- VRRP has a split-brain state in dual autodiscovery Junos Fusion. [PR1293030](#)
- An aggregation device without a cascade port cannot reach hosts over ICL link if they are authenticated by 802.1X in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)
- The 802.1X authentication might fail in a Junos Fusion setup. [PR1299532](#)
- IPv6 multicast is not forwarded over an MC-LAG ICL interface until the interface is toggled. [PR1301698](#)
- The l2ald process generates a core file with no apparent trigger. [PR1302344](#)
- All 802.1X authentication sessions are removed when the AUTO ICCP link is disabled. [PR1307588](#)
- The dot1x process might generate a core file in a Junos Fusion setup with dual aggregation devices. [PR1303909](#)
- LACP aggregated Ethernet interfaces go to down state when performing **commit synchronize**. [PR1314561](#)

SEE ALSO

New and Changed Features 86
Changes in Behavior and Syntax 88
Known Behavior 89
Known Issues 90
Documentation Updates 93
Migration, Upgrade, and Downgrade Instructions 93

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 for Junos Fusion Enterprise documentation.

SEE ALSO

New and Changed Features 86
Changes in Behavior and Syntax 88
Known Behavior 89
Known Issues 90
Resolved Issues 91
Migration, Upgrade, and Downgrade Instructions 93

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 94](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 95](#)
- [Preparing the Switch for Satellite Device Conversion | 96](#)
- [Converting a Satellite Device to a Standalone Switch | 97](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 97](#)
- [Downgrading from Release 17.4 | 98](#)

This section contains the procedure to upgrade or downgrade Junos OS and satellite software for a Junos Fusion Enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos Fusion Enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

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

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

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

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

To download and install Junos OS Release 17.4R2:

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

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

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

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

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

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

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

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

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

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

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

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

```
[edit]
```

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


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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Switch

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Downgrading from Release 17.4

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

NOTE: You cannot downgrade more than three releases.

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

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

SEE ALSO

[New and Changed Features | 86](#)

[Changes in Behavior and Syntax | 88](#)

[Known Behavior | 89](#)

[Known Issues | 90](#)

[Resolved Issues | 91](#)

[Documentation Updates | 93](#)

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- New and Changed Features | 99
- Changes in Behavior and Syntax | 100
- Known Behavior | 101
- Known Issues | 102
- Resolved Issues | 104
- Documentation Updates | 105
- Migration, Upgrade, and Downgrade Instructions | 106

These release notes accompany Junos OS Release 17.4R3 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R3 New and Changed Features | 100
- Release 17.4R2 New and Changed Features | 100
- Release 17.4R1 New and Changed Features | 100

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

Release 17.4R3 New and Changed Features

There are no new features in Junos OS Release 17.4R3 for Junos Fusion Provider Edge.

Release 17.4R2 New and Changed Features

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

Release 17.4R1 New and Changed Features

Hardware

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

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

SEE ALSO

Changes in Behavior and Syntax	 100
Known Behavior	 101
Known Issues	 102
Resolved Issues	 104
Documentation Updates	 105
Migration, Upgrade, and Downgrade Instructions	 106

Changes in Behavior and Syntax

IN THIS SECTION

- [Security](#) | [101](#)

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 17.4R3 for Junos Fusion Fabrics.

Security

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

SEE ALSO

[New and Changed Features | 99](#)

[Known Behavior | 101](#)

[Known Issues | 102](#)

[Resolved Issues | 104](#)

[Documentation Updates | 105](#)

[Migration, Upgrade, and Downgrade Instructions | 106](#)

Known Behavior

IN THIS SECTION

- [Junos Fusion Provider Edge | 102](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 for Junos Fusion Provider Edge.

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

Junos Fusion Provider Edge

- Layer 2 filter with multiple term containing mixed Layer 2 and Layer 3/Layer 4 match conditions do not get programmed on QFX10000 as AD. This is due to an ASIC limitation. [PR1286708](#)
- The FPCs were not online after an image upgrade due to a lack of space in `/var/tmp` directory. After ensuring enough space in `/var/tmp`, this issue was never seen. [PR1296082](#)
- The `no-mac-learning` and `interface-mac-limit` statements are not supported on extended ports or LAGs of extended ports. [PR1296731](#)
- The `interface-set` CLI command in firewall filter match condition is not supported on QFX10000 AD. [PR1298633](#)
- Filters: Policy route action is not supported on interfaces with `vxlan-vni` configuration along with routing instances. [PR1298683](#)
- Filters: Next-ip action for firewall filters is not supported with EVPN VXLAN vni configuration. [PR1298688](#)
- Configuration synchronization is not triggered when you issue the rollback command on the local aggregation device (AD). [PR1298747](#)

SEE ALSO

New and Changed Features 99
Changes in Behavior and Syntax 100
Known Issues 102
Resolved Issues 104
Documentation Updates 105
Migration, Upgrade, and Downgrade Instructions 106

Known Issues

IN THIS SECTION

- [Junos Fusion Provider Edge | 103](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 for Junos Fusion Provider Edge.

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

Junos Fusion Provider Edge

- The license installed will not be deleted, unless it is explicitly deleted using the **request** command. After disabling the cascade port, the license count will be marked as zero only after the satellite information is purged from the neighbor database. Previously this satellite neighbor information persisted for only 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: JUNOSxxxxxx License version: 4 Software Serial
Number: 99999B999999999 Customer ID: USER-SWITCH Features: SD-QFX5100-48SH-48TH-4PK -
SD 4 pack QFX5000-10-JFD permanent {master:0} user@host> show system license usage
Licenses Licenses Expiry Feature name used installed needed bgp 1 0 1 invalid SD-QFX5100-48SH-48TH
0 4 0 permanent {master:0} user@host> show system alarms
4 alarms currently active Alarm time Class
Description 2017-08-29 13:14:27 UTC Minor BGP Routing Protocol usage requires a license 2017-08-28
17:25:27 UTC Major FPC0: PEM 1 Not Powered 2017-08-28 17:25:27 UTC Major FPC Management 1
Ethernet Link Down [PR1294951](#)
- Configuration synchronization is not triggered when you issue the rollback command on the local aggregation device (AD). [PR1298747](#)
- When changing fpc slot-id, always delete the old configuration, commit, and then apply the new configuration. Otherwise, sdpd and mib2d might generate core files. Example: (1) delete chassis satellite-management fpc 101 cascade-ports et-0/0/11 (2) commit (3) set chassis satellite-management fpc 102 cascade-ports et-0/0/11 (4) commit . [PR1309080](#)
- In Junos Fusion scenario, if there are more than 12 cascade-ports configured to a satellite device (SD), the satellite discovery and provisioning process (sdpd) may continuously crash after committing, as a result, the SD cannot be managed from the aggregation device (AD). Traffic loss may not be observed

right after sdpd crash, and since it is continuous to crash if there is no interruption, the related FPCs on AD device may reach 100% CPU utilization. [PR1437387](#)

SEE ALSO

New and Changed Features		99
Changes in Behavior and Syntax		100
Known Behavior		101
Resolved Issues		104
Documentation Updates		105
Migration, Upgrade, and Downgrade Instructions		106

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R3](#) | [104](#)
- [Resolved Issues: 17.4R2](#) | [105](#)
- [Resolved Issues: 17.4R1](#) | [105](#)

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: 17.4R3

Junos Fusion Provider Edge

- Laser receive power of extended ports is higher than the output power of the peer link. [PR1358007](#)
- The ppmmd process on AD might crash when using **authentication key-chain** with BFD. [PR1375647](#)

- The spmd core file might be seen after executing **request support information** on Aggregation Device. [PR1375732](#)
- BUM traffic might get dropped on peer Fusion Aggregation Device when the link between Satellite Device and local Aggregate Device goes down. [PR1384440](#)

Resolved Issues: 17.4R2

Junos Fusion Provider Edge

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

Resolved Issues: 17.4R1

Junos Fusion Provider Edge

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

SEE ALSO

[New and Changed Features | 99](#)

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

[Known Behavior | 101](#)

[Known Issues | 102](#)

[Documentation Updates | 105](#)

[Migration, Upgrade, and Downgrade Instructions | 106](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 for Junos Fusion Provider Edge documentation.

SEE ALSO

[New and Changed Features | 99](#)

Changes in Behavior and Syntax 100
Known Behavior 101
Known Issues 102
Resolved Issues 104
Migration, Upgrade, and Downgrade Instructions 106

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 106
- Upgrading an Aggregation Device with Redundant Routing Engines | 109
- Preparing the Switch for Satellite Device Conversion | 109
- Converting a Satellite Device to a Standalone Device | 111
- Upgrading an Aggregation Device | 113
- Upgrade and Downgrade Support Policy for Junos OS Releases | 113
- Downgrading from Release 17.4 | 113

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 17.4R3 is different that for earlier Junos OS releases.

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

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

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

For upgrades from Junos OS Release 14.2 and earlier:

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

All other upgrades:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.4R3 **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 17.4R3, 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 | 99](#)

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

[Known Behavior | 101](#)

[Known Issues | 102](#)

[Resolved Issues | 104](#)

[Documentation Updates | 105](#)

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

IN THIS SECTION

- [New and Changed Features | 115](#)
- [Changes in Behavior and Syntax | 149](#)
- [Known Behavior | 164](#)
- [Known Issues | 173](#)
- [Resolved Issues | 202](#)
- [Documentation Updates | 275](#)
- [Migration, Upgrade, and Downgrade Instructions | 276](#)

These release notes accompany Junos OS Release 17.4R3 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R3 New and Changed Features | 116](#)
- [Release 17.4R2-S2 New and Changed Features | 116](#)
- [Release 17.4R2 New and Changed Features | 116](#)
- [Release 17.4R1 New and Changed Features | 119](#)

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R3 for the MX Series routers.

Release 17.4R3 New and Changed Features

Subscriber Management and Services

- **Preventing validation of magic numbers in PPP peer-originated keepalive messages (MX Series)**—Starting in Junos OS Release 17.4R3, you can include the **ignore-magic-number-mismatch** statement to disable the Packet Forwarding Engine from validating PPP magic numbers received during PPP keepalive (Echo-Request/Echo-Reply) exchanges. Because validation is not performed, the Packet Forwarding Engine does not detect whether the remote peer sends a magic number that does not match the number agreed upon during LCP negotiation. This prevents PPP from tearing down the session in the event of a mismatch. This capability is useful when the remote PPP peers include arbitrary magic numbers in the keepalive packets. Configuring this statement has no effect on LCP magic number negotiation or on the exchange of keepalives when the remote peer magic number is the expected negotiated number.

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

Release 17.4R2-S2 New and Changed Features

Routing Protocols

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

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

Release 17.4R2 New and Changed Features

EVPNs

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

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

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

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

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

Interfaces and Chassis

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

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

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

[See [Alarm Overview](#).]

MPLS

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

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

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

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

Restoration Procedures and Failure Handling

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

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

Software Installation and Upgrade

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

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

Subscriber Management and Services

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

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

Release 17.4R1 New and Changed Features

Hardware

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

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

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

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

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

- **Periodic refresh of authorization profile on TACACS+ server (MX Series)**—Starting with Junos OS Release 17.4R1, periodic refresh of the authorization profile that is received from the TACACS server is supported.

The authorization profile that is configured for the user on the TACACS server is sent to the Junos OS device after the user is successfully authenticated. The authorization profile is stored locally on the Junos OS device. With the periodic refresh feature, the authorization profile is periodically fetched from the TACACS server to refresh the authorization profile that is stored locally. User authorization is reevaluated using the refreshed authorization profile.

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

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

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

Class of Service (CoS)

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

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

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

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

Dynamic Host Configuration Protocol (DHCP)

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

Reauthentication is enabled by the following triggers:

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

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

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

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

Prior to 17.4R1 Release, subscriber entry creation constituted of Junos OS DHCPD (DHCP daemon) memory resources, session database resources, authentication procedure, accounting, dynamic profile instantiation, dynamic interface creation, firewall, CoS association, and more. if a customer network has

some non-customer controlled traffic for which a relay agent entry is created then it would be an unnecessary utilization of resources, and an incorrect association of profiles.

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

EVPNs

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

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

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

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

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

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

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

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

- **EVPN active standby multihoming to a single PE device (MX Series)**—Starting in Junos OS Release 17.4R1, Juniper supports EVPN active-standby multihoming. When you configure a protect (backup) interface for a primary interface on the same PE router, the protect interface becomes active when the primary interface fails and network traffic is switched to the protect interface.

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

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

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

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

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

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

The EVPN-MPLS interworking feature offers the following benefits:

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

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

General Routing

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

[See [Understanding Hybrid Mode](#).]

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

[See [phy-timestamping](#).]

- **Support for PTP over Ethernet, hybrid mode, and G.8275.1 profile (MPC7E-10G, MPC7E-MRATE, MPC8E, MPC9E)**—Starting in Junos OS Release 17.4R1, MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E support the following features:

- **PTP over Ethernet**— PTP over Ethernet enables effective implementation of packet-based technology that enables the operator to deliver synchronization services on packet- based mobile backhaul networks. PTP over Ethernet uses multicast addresses for communication of PTP messages between the slave clock and the master clock. The IEEE 1588 standard defines two types of multicast MAC addresses 01-80-C2-00-00-0E (link local multicast) and 01-1B-19-00-00-00 (standard Ethernet multicast) for PTP over Ethernet operations.
- **Hybrid mode**— In hybrid mode, the Synchronous Ethernet equipment clock (EEC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP.

[See [Understanding Hybrid Mode](#)]

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

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

High Availability (HA) and Resiliency

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

[See [show l2vpn connections](#)]

Interfaces and Chassis

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

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

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

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

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

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

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

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

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

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

The following features are supported on MX204 platform:

- Basic Layer 2 features including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
- Class of service (CoS)
- Firewall filters and policers
- Integrated routing and bridging (IRB)
- Layer 2 protocols
- Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs
- Layer 3 routing protocols and MPLS
- Layer 3 inline services
- Multicast forwarding
- Port mirroring
- Spanning-tree protocols, such as STP, MSTP, RSTP, and VSTP
- Synchronous Ethernet and Precision Time Protocol (IEEE 1588)
- Tunneling
- **Support for MACsec PSK keychain (MX2010, MX2020)**—Starting in Junos OS Release 17.4R1, MX2020 and MX2010 supports Key Agreement Protocol Fail Open mode. The MACsec PSK chains hitless rollover feature is documented in Junos OS Release 17.4R1, but not supported.
- **Strong encryption for configuration secrets (MX2020, MX2010, and MX2008 routers)**—Starting in Junos OS Release 17.4R1, the MX2020, MX2010 and MX2008 routers support strong encryption for configuration secrets. To use strong encryption for your configuration secrets, you need to configure a master password. The master password enables you to derive an encryption key that you use with the AES256-GCM standard to encrypt configuration secrets. This new encryption method uses the \$8\$ formatted strings.
[See [Hardening Shared Secrets in Junos OS](#).]
- **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-IPv4 NAT. This facilitates the phased introduction of IPv6 on the Internet by providing backward compatibility with IPv4.

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

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

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

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

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

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

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

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

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

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

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

given translated external IP address and port are always mapped to the same internal IP address. Deterministic NAT mapping eliminates the need for logging address translations.

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

- **Support for inline video monitoring for IPv6 flows (MX Series)**—Starting in Junos OS Release 17.4R1, MX Series routers support the inline video monitoring of IPv6 flows and IPv6-over-MPLS flows to measure media delivery index (MDI) metrics. MDI information enables you to identify devices that are causing excessive jitter or packet loss for streaming video applications.

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

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

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

Software Defined Networking (SDN)

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

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

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

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

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

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

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

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

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

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

[Configuring JDM on the x86 Servers](#)

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

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

[See [Abstracted Fabric Interface](#)]

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

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

Software Installation and Upgrade

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

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

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

- **Support for Zero Touch Provisioning (ZTP) (MX150)**—Starting in Junos OS Release 17.4R1, MX150 routers, powered with vMX, support zero touch provisioning. Zero touch provisioning enables you to provision new routers in your network automatically either by executing a script file or by loading a configuration file. In either case, the information is detected in a file on the Dynamic Host Control Protocol (DHCP) server. When you physically connect the MX150 router to the network and boot it with a default configuration, it attempts to upgrade the Junos OS Software automatically using information

detected on the DHCP server. If you do not configure the DHCP server to provide this information, the MX150 router boots with the pre-installed software and default configuration.

- **Support for unified ISSU on the CFP2-DCO-T-WDM-1 transceiver (MX Series)**—Starting in Junos OS Release 17.4R1, unified in-service software upgrade (unified ISSU) is supported on the CFP2-DCO-T-WDM-1 transceiver when the transceiver is installed on the MPC5E-100G10G MPC or the MIC6-100G-CFP2 MIC (installed on the MX2K-MPC6E MPC).

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

Subscriber Management and Services

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

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

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

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

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

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

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

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

DHCP server copies the RADIUS-sourced DHCP options, and also adds the DHCP server-sourced options to the packet and sends the response back to the client.

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

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

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

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

- **Enhancements to share CPE parameters between broadband network gateway (BNG) and RADIUS server (MX Series)**—Starting in Junos OS Release 17.4R1, the following enhancements are made to facilitate better communication between the broadband network gateway (BNG) and the RADIUS server:
 - CPE parameters such as DHCPv4 (VSA 26-208) and DHCPv6 (VSA 26-209) packet headers are shared between the broadband network gateway (BNG) and the RADIUS server.
 - A new VSA 26-207 is introduced that facilitates the exchange of DHCPv6 options with the RADIUS server, thereby ensuring that VSA 26-55 is dedicated to the exchange of DHCPv4 options.
 - A new statement, **family-state-change-immediate-update**. When configured at the **[edit access profile]** hierarchy level, the DHCP (both DHCPv4 and DHCPv6) server sends an immediate interim accounting report to the RADIUS server when the second family (IPv4 or IPv6) is activated or the first family gets deactivated.
 - A new VSA 26-210 is added to convey the reason for the accounting-request message in the start and interim accounting request packets sent to the RADIUS server. This helps the RADIUS server to determine the reason of the start and interim accounting that is being sent.

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

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

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

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

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

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

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

You must also enable enhanced IP networking services at the **[edit chassis network-services enhanced-ip]** hierarchy level.

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

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

[See [Understanding Distributed IGMP](#) .]

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

Layer 2 wholesale service connections are subject to the same class and quality-of-service management transformations.

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

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

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

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

- **Automatic 64-bit mode and maximum configuration database size (MX Series)**—Starting in Junos OS Release 17.4R1, when enhanced IP network services and enhanced subscriber management are enabled and a Routing Engine in the system has at least 32 GB of RAM, subscriber management daemons on that Routing Engine run in 64-bit mode. For consistent operation, all Routing Engines in the system must have the same amount of memory.

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

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

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

- **Enhancement to Gx-Plus Application (MX Series)**—Starting in Junos OS Release 17.4R1, the following enhancements to the Gx-Plus client application on the BNG are available:
 - When a monitored service is deactivated separate from a subscriber logout, the CCR-U indicates that the service is no longer active and includes the service's usage data.
 - The router updates the monitoring key and threshold values when they are received in a RAR message from the PCRF.

- A CCR-U is sent to the PCRF after the router sends an RAA message in response to an RAR message that requests service activations or deactivations.
- When the PCRF returns threshold values that are lower than the current values, the new threshold becomes the sum of the current value and the returned value.
- The PCEF has default minimum threshold values. If the change between the current value and the value returned by the PCRF is less than the minimum value, then the new value is adjusted to the minimum.
- The CCR-I message includes the Diameter AVP Subscription-Id attribute (443) with the Subscription-Id-Type Diameter AVP sub-attribute (450) set to 4 (END_USER_PRIVATE) and the Subscription-Id-Data Diameter AVP sub-attribute (444) set to **reserved**.

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

- **RADIUS attributes added to LNS messages (MX Series)**—Starting in Junos OS Release 17.4R1, the LNS includes the following RADIUS attributes when it sends an Access-Request message to the RADIUS server:
 - Tunnel-Type (64)
 - Tunnel-Medium-Type (65)
 - Tunnel-Client-Endpoint (66)
 - Tunnel-Server-Endpoint (67)
 - Acct-Tunnel-Connection (68)
 - Tunnel-Assignment-Id (82)
 - Tunnel-Client-Auth-Id (90)
 - Tunnel-Server-Auth-Id (91)

System Logging

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

[See [ukern-trace](#)]

User interface and Configuration

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

by bit error rate (BER) threshold levels) on Ethernet interfaces, and take corrective actions if the BER threshold value drops to a value in the range of 10^{-13} to 10^{-5} .

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

The following new configurations have been introduced at the **[edit interfaces *interface-name*]** hierarchy level to support the physical link degrade monitoring and recovery feature on Junos OS:

- To monitor physical link degrade on Ethernet interfaces, configure the **link-degrade-monitor** statement.
- To configure the BER threshold value at which the corrective action must be triggered on or cleared from an interface, use the **link-degrade-monitor thresholds (set *value* | clear *value*)** statement.

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

- To configure the link degrade interval value, use the **link-degrade-monitor thresholds interval *value*** statement. The configured interval value determines the number of consecutive link degrade events that are considered before any corrective action is taken.
- To configure link degrade warning thresholds, use the **link-degrade-monitor thresholds (warning-set *value* | warning-clear *value*)** statement.
- To configure the link degrade action that is taken when the configured BER threshold level is reached, use the **link-degrade action media-based** statement.
- To configure the link degrade recovery options, use the **link-degrade recovery (auto interval *value* | manual)** statement. The recovery mechanism triggers the recovery of a degraded link.

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

VPNs

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

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

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

SEE ALSO

[Changes in Behavior and Syntax | 149](#)

[Known Behavior | 164](#)

[Known Issues | 173](#)

[Resolved Issues | 202](#)

[Documentation Updates | 275](#)

[Migration, Upgrade, and Downgrade Instructions | 276](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 150](#)
- [EVPNs | 150](#)
- [General Routing | 151](#)
- [High Availability \(HA\) and Resiliency | 151](#)
- [Interfaces and Chassis | 151](#)
- [Management | 154](#)

- MPLS | 154
- Multicast | 157
- Network Management and Monitoring | 157
- Routing Protocols | 158
- Security | 159
- Services Applications | 159
- Software Defined Networking | 160
- Software Installation and Upgrade | 161
- Software Licensing | 161
- Subscriber Management and Services | 161
- User Interface and Configuration | 164

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 17.4R3 for MX Series routers.

Class of Service (CoS)

- **Junos commit notification of unsupported configuration**—Junos OS does not support changing the **hierarchical-scheduler** mode of a logical tunnel interface, or redundant logical tunnel interface, if an active pseudowire subscriber interface is attached to it. A commit error has now been added to provide the notification.

EVPNs

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

- EXP: unknown Oxtype:Oxsub-type:Oxvalue

where type, sub-type, and value are defined in RFC 4360 *BGP Extended Communities Attribute*, RFC7153 *IANA Registries for BGP Extended Communities*. Internet Assigned Numbers Authority (IANA) maintains a registry with information on the type and subtype field values at

<https://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xhtml>

- **Support for an VNI of zero**—Starting with Junos OS Release 17.4R3, Junos supports using a VXLAN Network Identifier (VNI)=0 when configuring a bridge domain or vlan in an EVPN-VXLAN network.

General Routing

- **User confirmation prompt for configuring the sub-options of request vmhost commands (MX Series and PTX series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm a [yes,no] for the sub-options also.
 - **request vmhost reboot**
 - **request vmhost poweroff**
 - **request vmhost halt**

In previous releases, the confirmation prompt was available for only the main options.

High Availability (HA) and Resiliency

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

Interfaces and Chassis

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

```

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

```

- **Modified output of the show chassis ethernet-switch command**—The ports 24 and 26 on the MX240, MX480, and MX960 routers with the RE-S-X6-64G Routing Engines are dedicated for external Ethernet connectivity. The **show chassis ethernet-switch** command on these routers displays the link status for these ports as **External Ethernet**.
- **Recovery of PICs that are stuck because of prolonged flow controls (MS-MIC, MS-MPC, MS-DPC, MS-PIC 100, MS-PIC 400, and MS-PIC 500)**—Starting in Junos OS Release 16.1R7, if interfaces on an MS-PIC, MS-MIC, MS-MPC, or MS-DPC are in stuck state because of prolonged flow control, Junos OS restarts the service PICs to recover them from this state. However, if you want the PICs to remain in stuck state until you manually restart the PICs, configure the new option **up-on-flow-control** for the **flow-control-options** statement at the **[edit interfaces mo-fpc/pic/port multiservice-options]** hierarchy level. In releases before Release 16.1R7, there is no action taken to recover service PICs from this state unless one of the options for the **flow-control-options** statement is configured, or service PIC is manually restarted.
- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the **show interfaces mc-ae extensive** command. The output now displays the following two additional fields:
 - Local Partner System ID—LACP partner system ID as seen by the local node.
 - Peer Partner System ID—LACP partner system ID as seen by the MC-AE peer node.

Previously, the **show interfaces mc-ae extensive** command did not display these additional fields.
- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (MX Series)**—In Junos OS Release 17.4R3, the **show lacp interfaces | display xml** command displays a new XML tag element **<lacp-hold-up-state>**. The **<lacp-hold-up-state>** displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces were in a single **<lacp-hold-up-information>** XML tag. Now, for each interface it is displayed in a separate **<lacp-hold-up-information>** XML tag.
- **No support for WAN-PHY mode on MX Series MPCs**—In Junos OS Releases 17.4R2, 17.4R3, and later, on the following MPCs or routers, you cannot configure **wan-phy** mode at 10-Gbps, 40-Gbps, and 100-Gbps on a per-port basis:
 - MPC7E-10G, MPC7E-MRATE, MX2K-MPC8E, and MX2K-MPC9E
 - MPC10003

- MX204 router
- JNP10K-LC2101 MPC
- **IRB not supported on Pseudowire Subscriber (PS) Logical Interface in bridge-domain (MX Series)**—In Junos OS Release 17.4R3, Integrated routing and bridging (IRB) is not supported on Pseudowire Subscriber (PS) Logical Interface. Hence you cannot add IRB to bridge domain with PS interface, that is, you cannot configure IRB and PS interface in the same bridge domain.

Note that adding IRB to a bridge-domain having Pseudowire Subscriber (PS) Logical Interface causes kernel crash and continuous reboot of the router until the configuration is rolled back.

NOTE: IRB is not supported on PS only in bridge-domain.

[See [bridge-domain](#).]

Management

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

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

MPLS

- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (MX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP ressignaling events. This is because the LSP is responsive to even minor increases or decreases in the utilization when current reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of around 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP ressignaling events for small LSPs of around 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP ressignaling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless they are huge, for example, 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value-based threshold along with the percentage-based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at the **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.
- **Support for label history for MPLS protocol (MX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at the **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entries per label . By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range. The **clear mpls label history** command clears the label history details.
- **Support for default time out duration for self-ping on an LSP instance (MX Series)**—Starting in Junos OS 17.4R1, the default time out duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also configure the self-ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is

enabled. The LSP types like CCC, P2MP, VLAN-based , and non-default instances do not support self-ping . You can configure **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.

- **Support for Flap and MBB counter for LSP (MX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on the master routing engine (RE) only:
 - Flap counter-- Counts the number of times a LSP flaps down or up.
 - MBB counter— Counts the number of times a LSP incurs MBB.

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

- **Support for inet.0 and inet.3 labeled unicast BGP route for protocol LDP (MX Series)**--- Starting in Junos OS Release 17.4R2, LDP egress policy is supported on both inet.0 and inet.3 routing Information bases (RIBs) also known as routing table for labeled unicast BGP routes. If a routing policy is configured with a specific (inet.0 and inet.3) RIB, the egress policy is applied on the specified RIB. If no RIB is specified and a prefix is present on both inet.0 and inet.3 RIBs for labeled unicast BGP routes, then inet.3 RIB is preferred. However, prior to Junos OS Release 12.3R1 and starting with Junos OS Release 16.1R1, LDP egress policy is always preferred on inet.0 RIB and support for inet.3 RIB egress policy for labeled unicast BGP routes was disabled. In Junos OS Release 12.3R1 and later releases up to Junos Release 16.1R1, LDP egress policy was supported in inet.3 RIBs, in addition to inet.0 RIBs, for labeled-unicast BGP routes.
- **New output fields to monitor LSP resigaling count**—Starting in Junos OS Release 17.4R1, the **show mpls lsp** command output displays the **Flap Count** and **MBB Count** output fields, that capture the historical count of the number of times a specific LSP has been resigaled because of autobandwidth-triggered reservation change, or other changes along the path. The flap count displays the number of times an LSP flaps down and up, and the MBB count displays the number of times an LSP incurred a make before break.
- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (MX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
- Starting in Junos OS Release 17.4R1, a new configuration statement - **adjust-threshold-absolute** - is introduced at the **[edit protocols mpls]** hierarchy level to specify the changes in the average label-switched path (LSP) utilization to trigger automatic bandwidth adjustment in bits per second (bps).
Currently, this change is specified as a percentage using the **adjust-threshold statement**. The **adjust-threshold-absolute** statement (bps) can be used in conjunction with the existing **adjust-threshold statement** (percent).
- Starting in Junos OS Release 17.4R1, the **spring-traffic-engineering** statement at the **[edit protocols]** hierarchy level is replaced with the **source-packet-routing** statement, although the support for the **spring-traffic-engineering** statement is provided as an alias. This replacement does not introduce any functionality change, and is intended for maintaining consistency across the terms used in Source Packet Routing in Networking (SPRING) or segment routing features.

- **Loss of traffic over bypass MPLS LSPs**—If RSVP link or node protection is enabled along with global RSVP authentication, there is loss of traffic over bypass MPLS LSPs at the time of local repair, when the point of local repair (PLR) and the merge point devices have different versions of the Junos OS software installed on them. That is, one device is running a release prior to Junos OS Release 16.1, and the other device is running a release starting with Junos OS Release 16.1R4-S12.
- **Bandwidth allocation**—For a label-switched path (LSP) that has both **bandwidth** and **minimum-bandwidth** for autobandwidth configured under the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level, the LSP bandwidth is adjusted differently.

The LSP is initiated with the bandwidth value configured under the **bandwidth** statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level. At the expiry of the **adjust-interval** timer, the LSP bandwidth gets adjusted based on the traffic flow.

If the bandwidth to be signaled is less than the value configured under the **minimum-bandwidth** statement at the `[edit protocols mpls label-switched-path lsp-name autobandwidth]` hierarchy level, then the LSP is signaled only using the minimum bandwidth.

If the bandwidth to be signaled is greater than the value configured under the **maximum-bandwidth** statement at the `[edit protocols mpls label-switched-path lsp-name autobandwidth]` hierarchy level, then the LSP is signaled only using the maximum bandwidth.

- Previously, when you configured zero (0) as the bandwidth of an RSVP interface, the bandwidth value was overwritten with the default interface bandwidth (raw hardware bandwidth), leading to unexpected behavior in the LSP setup. Starting with Junos OS Release 17.4R1-S5, when you configure zero as the bandwidth, 0 is applied as the RSVP bandwidth.

[See [bandwidth \(Protocols RSVP\)](#).]

Multicast

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

Network Management and Monitoring

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

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

- **SNMP syslog messages changed (MX Series)**—In Junos OS Release 17.4R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD —AgentX master agent failed to respond to ping. Attempting to re-register
NEW —AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD —NET-SNMP version %s AgentX subagent connected
NEW —NET-SNMP version %s AgentX subagent Open-Sent!

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

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

Before this change:

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

After this change:

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

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

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

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

- **The NETCONF server omits warnings in RPC replies when the `rfc-compliant` statement is configured and the operation returns `<ok/>` (MX Series)**—Starting in Junos OS Release 17.4R3, when you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an `<ok/>` element and one or more `<rpc-error>` elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the `rfc-compliant` statement is not configured, the NETCONF server might issue an RPC reply that includes both an `<rpc-error>` element with a severity level of warning and an `<ok/>` element.
- A decrease in the MPLS label-switched path (LSP) statistics pauses the SNMP MIB `mplsLspInfoAggrOctets` count for one MPLS statistics gathering interval. In such cases, the `mplsLspInfoAggrOctets` value is updated only after completing one more interval of the MPLS statistics gathering.

Routing Protocols

- **Option to configure SPRING bandwidth utilization change threshold in percentage (MX Series)**—Starting in Junos OS Release 17.4R1, you can specify a change threshold in percentage beyond which RSVP triggers IGP updates. To configure the change threshold percentage, configure `percent percent` at the `[edit protocols rsvp interface update-threshold-max-reservable]` hierarchy level.
- **BGP enterprise trap `jnxBgpM2BackwardTransition` notification for IPv4 neighbors (MX Series)**—Starting in Junos OS Release 17.4R2, when an IPv4 BGP neighbor transitions from a higher state to a lower state, an enterprise trap `jnxBgpM2BackwardTransition` is sent in addition to an existing standard trap notification `bgpM2BackwardTransition`. In earlier Junos OS releases only `bgpBackwardTransition` trap notification was generated when a BGP IPv4 neighbor's state transitioned to a lower state.
- **Modified output of `show route forwarding-table`**—Starting in Junos OS Release 17.4R2, the output of `show route forwarding-table` command does not display the next-hop address for static routes that use point-to-point (P2P) interfaces.

[See [show route forwarding-table](#).]

- **MPLS configuration mandatory for indirect next-hop interfaces**—Starting in Junos OS Release 17.4R1, it is mandatory for an indirect next-hop's forwarding interface to have family MPLS configured. In a BGP network if the MPLS configuration for an indirect next-hop's forwarding interface is deleted or when the BGP labeled unicast interface is deactivated, all routes with indirect next hop undergo a route resolution again, which might impact traffic routing until the route resolution is completed. In earlier Junos OS releases when family MPLS was deleted, the indirect next-hop route was removed from the forwarding table and could not be recovered even when MPLS was reactivated.

Security

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

[See [protocol-version](#)]

Services Applications

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

JSERVICES_NAT_OUTOF_PORTS_APP

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

- **Support for host generated traffic on a GRE over GRE tunnel (MX Series)**—In Junos OS Release 17.4R3, you can send host generated traffic on a GRE over GRE tunnel. However, when path maximum transmission unit (PMTU) is updated for the outer GRE tunnel, MTU for inner GRE tunnel is not corrected.
- **Change in error message displayed while fragmenting or de-fragmenting IPv6 GRE tunnel interface (MX Series routers)**—In Junos OS Release 17.4R3, on a IPv6 GRE tunnel interface, when you enable fragmentation using the **allow-fragmentation** command or disable fragmentation using the **do-not-fragment** command, the following error message is displayed:

Fragmentation for V6 tunnels is not supported

In earlier Junos OS releases, the following message was displayed:

dcd_config_ifl_tunnel:Fragmentation for V6 tunnels is notsupported

Software Defined Networking

- **The 32-bit libstdc++ package no longer required for Junos Node Slicing setup**—Starting in Junos OS Release 17.4R2, you need not install the additional 32-bit **libstdc++** package for Red Hat Enterprise Linux (RHEL) or Ubuntu to set up Junos Node Slicing.
- **Installation or upgrade using remotely located installation package (MX480, MX960, MX2010, MX2020, MX2008)**—While performing Junos installation or upgrade on the base system (BSYS) or guest network function, if you provide a URL to the remotely located installation package (for example, an ftp file) in the command **request system software add *package-file-path***, the router locally copies the package, performs checks such as multi-version compatibility checks on the package, and then installs the package. The installation process is aborted if any errors are found during the checks. Previously, if you tried to perform installation or upgrade using a remotely located file, the router would skip multi-version checks and display an error message, but would not abort the installation process.

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

- The output of the **show mpls lsp ingress locally-provisioned** command is expected to display only label-switched paths (LSPs) that have been provisioned locally by the Path Computation Client (PCC). However, the **locally-provisioned** option was displaying all the LSPs, instead.

Starting in Junos OS Release 17.4R2, the **locally-provisioned** option in the **show mpls lsp ingress** command is behaving as expected.

Software Installation and Upgrade

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

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

Software Licensing

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

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

Subscriber Management and Services

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

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

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

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

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

```
set system configuration-database max-db-size
```

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

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

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

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

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

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

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

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

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

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

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

The new behavior is as follows:

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

In earlier releases:

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

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

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

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

- **Change to ICRQ message inclusion of the ANCP Access Line Type AVP (MX Series)**—Starting in Junos OS Release 17.4R2, the ICRQ message includes the ANCP Access Line Type AVP (145) when the received ANCP Port Up message includes a DSL-type of 0 (OTHER). In earlier releases, the AVP is not sent when the value is 0.
- **Out-of-address SNMP trap requires thresholds to be configured (MX Series)**—Starting in Junos OS Release 17.4R3, the behavior has changed for generating an out-of-address SNMP trap for an address pool configured at the **[edit access address-assignment]** or **[edit routing-instance name address-assignment]** hierarchy levels. You must now configure both the high-utilization and abated-utilization thresholds. When the number of assigned addresses surpasses the high-utilization threshold, a high-utilization trap is generated. If all the addresses are assigned from the pool, an out-of-address trap is generated and an out-of-address syslog message is sent.

In earlier releases, an out-of-address trap is generated when the address pool is exhausted, regardless of whether the thresholds are configured.

If the number of assigned addresses subsequently drops below the abated-utilization threshold, an abate-high-utilization trap is generated; this behavior is unchanged.

- **Disabling a pseudowire underlying interface (MX Series)**—Starting in Junos OS Release 17.4R3, you cannot disable the underlying logical tunnel (lt) interface or redundant logical tunnel (rlt) interface when

a pseudowire is anchored on that interface. If you want to disable the underlying interface, you must first deactivate the pseudowire.

[See [Configuring a Pseudowire Subscriber Logical Interface Device](#).]

User Interface and Configuration

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

SEE ALSO

[New and Changed Features | 115](#)

[Known Behavior | 164](#)

[Known Issues | 173](#)

[Resolved Issues | 202](#)

[Documentation Updates | 275](#)

[Migration, Upgrade, and Downgrade Instructions | 276](#)

Known Behavior

IN THIS SECTION

- [EVPN | 165](#)
- [General Routing | 166](#)
- [Infrastructure | 169](#)
- [Interfaces and Chassis | 169](#)
- [Junos Fusion Provider Edge | 170](#)
- [Layer 2 Ethernet Services | 170](#)
- [Multiprotocol Label Switching \(MPLS\) | 170](#)
- [Platform and Infrastructure | 170](#)
- [Routing Protocols | 170](#)
- [Services Applications | 171](#)

- [Software Defined Networking \(SDN\) | 172](#)
- [Software Installation and Upgrade | 172](#)
- [Subscriber Management and Services | 172](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 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.

EVPN

- Routing instances of type **evpn** configured with a **vlan-id** will advertise MAC (type 2) routes with the VLAN value in the Ethernet tag field of the MAC route. Advertising MAC routes with a nonzero VLAN is incompatible with the EVPN VLAN-based service type. To enable interoperability between a Junos OS routing instance of type **evpn** and a remote EVPN device operating in VLAN-based mode, the Junos routing instance should be configured with **vlan-id none** so that the Ethernet tag in advertised MAC routes is set to zero. [PR945247](#)
- A provider edge (PE) device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE device. The IGP instance running in the VRF on the PE might be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- In a scaled up EVPN VPWS configurations (approximately 8000 EVPN VPWS), during Routing Engine switchover, rpd scheduler slip messages might be seen. [PR1225153](#)
- EVPN remote MAC may not be installed in bridge-domain / ethernet switching table if load-balance per-packet is not configured in multi-homing scenario remote esi would be shown as **unresolved** in **show evpn instance extensive** output From EVPN traceoptions: **evpn_mac_msg_send_to_l2ald:1172 EVPN MAC instance::vlan::mac [Flags: 0x0] Instance ESI xx:xx:xx:xx:xx:xx:xx:xx:xx:xx not yet resolved**[PR1295846](#)
- In an EVPN network with VXLAN encapsulation configured for direct-nexthop mode ("pure type 5" mode without overlay gateway addresses), at least one type 5 route per VRF from a remote endpoint must be received and installed in the local routing table of a device, to enable the local device to forward inbound type 5 traffic received from the remote endpoint. If the local device has not installed at least

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

- When changing encapsulation from VXLAN to MPLS or vice versa, need to deactivate and reactivate the instance. [PR1326430](#)

General Routing

- On MX Series routers with MS-MPC or MS-MIC, memory leaks can be seen with `jnx_msp_jbuf_small_oc` object, upon sending millions of Point-to-Point Tunneling Protocol control connections (3 through 5 million) 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 through 30,000 cps. [PR1087561](#)
- Source-prefix filtering and protocol filtering of the CGNAT sessions are incorrect. For example, **show services sessions extensive protocol udp source-prefix <0:7000::2>** displays incorrect filtering of the sessions. [PR1179922](#)
- Chef for Junos OS supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev resources. The netdev resource developed for interface configuration has a limitation to configuring the XE interface. The netdev interface resource determines that speed is a configurable parameter that is supported on a GE interface but not on an XE interface. Hence, the netdev interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages `chef-11.10.4_1.1.*.tgz` `chef-11.10.4_2.0.*.tgz` in all platforms `{i386/x86-32/powerpc}`. [PR1181475](#)
- 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 with NPU image size less than 60MB to Junos OS Release with NPU image size great than 60MB. [PR1222540](#)
- This issue has not been addressed, and it is probably not easy to address either. The problem is, when some route/NH has been created by the app, it's assumed that it can propagate to the rest of the system. KRT asynchronously picks up this state for propagation. There is no reverse indication to the app, if there was an error in propagating the state. The system is supposed to eventually reconcile. So, if SPRING-TE produces a `<route, NH>` pair that looks legal from the app standpoint, but KRT is not able to download it to the kernel, because kernel rejected the NH, the `<route, NH>` sort of gets stuck in RPD. In the meantime, the previous version of the route (L-ISIS in this case) that was downloaded still lingers in the kernel and Packet Forwarding Engine. [PR1253778](#)
- On a Junos-based platform, sometimes it might occur that FPC is stuck in offline state with the reason **Restarted by cli command** after restarting the FPC immediately after restarting chassisd. This occurs

due to the fact that it takes sometime for the system to stabilize after chassisd restarts. Though chassisd would provide the FPC status and be able to accept the commands but in the back end it would be doing many initializations. So wait until all the PIC status are also available before issuing any command that makes fpc online, offline, and restart. [PR1275530](#)

- CFM is not supported for L2-over-GRE tunnel. CCM can pass through as transit traffic via GRE interfaces transparently using data path. Link trace functionality uses mac-learning and re-injecting LTM on GRE interface in case the bridge is configured with CFM. This is not a supported feature. [PR1275833](#)
- On MX104 **JTASK_SCHED_SLIP** is seen on commit randomly. [PR1281016](#)
- With Junos Os Releases 16.2R1, 16.1R4 releases or above, the error message about **jlaunchd, "jlaunchd: %AUTH-1: commit-batch is thrashing, not restarted**, might be seen after system reboot or Routing Engine switchover. [PR1284271](#)
- At reboot RHEL 7.3 servers report libvirtd[6282]: segfault at 10 ip 00007f87eab09bd0. No core file is left and no operational impact is known. [PR1287808](#)
- When LLDP is configured on multihomed extended ports, the peer might have duplicate entries for a duration of the hold timer (default: 120 seconds) during catastrophic configuration events such as redundancy group ID change and redundancy group name change. The duplicate entry would be deleted after the LLDP hold timer expired on the peer. [PR1291519](#)
- Race condition where on Ubuntu based external servers G-ARP may not be sent from jmgmt0 interface, resulting in loss of connectivity to management IP of JDM. [PR1291836](#)
- This is a limitation or expected behavior for smart SFPs. When you insert a smart SFP, 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 the af interfaces statically configured bandwidth. Also, the fabric capacity of the Packet Forwarding Engine is slightly higher than that of the WAN interface of same bandwidth. Because 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 Packet Forwarding Engine and kernel for sometime. 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 output for next **show dynamic-tunnels database** command. [PR1297913](#)
- For CFP2-DCO-T-WDM-1 pluggable, Rx payload type shown incorrectly (shown 0 vs 7). [PR1300423](#)
- We do see the underflow error during FPC cold boot and initial traffic start cases. But these error are limited and should not appear once traffic is stabilized. [PR1306280](#)

- UDP Setup rate for DetNat64 is approx. 10% less than setup rate of stateful-nat64 for 15M sessions on single NPU. DetNat64 needs extra processing while creating sessions and hence it's setup rate is 10% less than setup rate of stateful-nat64. [PR1307451](#)
- Support for enterprise profile is only provided for 10-Gigabit Ethernet interfaces. Use of 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. [PR1310048](#)
- A mobiled core will occur in systems where one RE is running Junos version 16.2R1 or 17.1R1 and the other RE 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 RE. 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](#)
- Parametrized (aka converged) HTTP redirect/rewrite services (CPCD) are not supported on Mx104 platform with MS-MIC. please note that other flavors of CPCD continue to work fine with this combination, Mx104 platform with MS-MIC. [PR1330340](#)
- When a new instance of Virtual Route Reflector - vRR - is launched, the factory default configuration has dhcp client and auto image turned on. Even after DHCP config is removed, access-internal default routes installed by DHCP client may persist and cause reachability problem. This typically happens during initial installation, and restart routing immediately can clear the problem. [PR1335925](#)
- Forward filter with log for Inline NAT is not supported. [PR1385843](#)
- Newer B0 DCO modules(740-087314) HGFEC implementation is different and standardized vs. A0 (740-072229) which has different implementation causing link not to come up for interop between B0 and A0. [PR1394134](#)
- IDS aggregate config knob will not be considered for the installation of the IDS dynamic filter. [PR1395316](#)
- On vMX platforms, the link flapping for the ixgbe interface might trigger PF (Physical Function) to reset for ixgbe, but the VF (Virtual Function) reset will not be done. The issue results in traffic drop for the interface. [PR1424626](#)
- HQoS configuration on ps interface anchored to logical-tunnel will fail to commit with the following error: [edit class-of-service interfaces ps0 unit 10]'output-traffic-control-profile'cannot configure traffic control profile (pic has no CoS queuing) error: configuration check-out failed. [PR1429927](#)

Infrastructure

- Junos boots from OAM volume after shutdown. The root cause is the file system super block is corrupted, but what caused the super block corruption is unknown. [PR1296861](#)

Interfaces and Chassis

- The same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating incorrect configuration. [PR1221993](#)

- **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.

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

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

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

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

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

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

Junos Fusion Provider Edge

- The FPCs were not online after an image upgrade due to lack of space in /var/tmp directory. After ensuring enough space in /var/tmp , this issue was never seen. [PR1296082](#)

Layer 2 Ethernet Services

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

Multiprotocol Label Switching (MPLS)

- When Flow-Label (FL) is enabled for PW, the OAM packets are not sent with Flow-Label because RPD is not aware of the Flow-Label values assigned by PFE software. Hence the packets are getting dropped by PFE at the tail-end PE. The remote PE is expecting the packet with FL and PW label. [PR1217566](#)
- An SR-TE path with "0" explicit NULL as inner most label, SR-TE path does not get installed with label "0". [PR1287354](#)

Platform and Infrastructure

- On all Junos OS platforms, execution of Python scripts through enhanced automation does not work on veriexec images. [PR1334425](#)

Routing Protocols

- This is not a functionality impacting BUG, Issue here is BGP NSR replication starts after some delay in certain cases. [PR1256965](#)
- RPD-Packet Forwarding Engine out-of-sync during MoFRR convergence. [PR1284463](#)
- The mcsnoopd process might crash when all the core-facing interfaces that are part of the Layer 2 domain have flapped and the mcsnoopd is attempting to flood a packet received over a CE interface, over the core-facing interfaces. [PR1329694](#)

Services Applications

- Session counters for cleartext traffic are not updated after decryption. Decrypted packet count can, however, be obtained by running the following command. `show security group-vpn member ipsec statistics`. [PR1068094](#)
- Hide ha detail if its not configured on a particular interface. [PR1383898](#)
- Broadband-edge platforms do not support service-set integration with dynamic profiles when the service set is representing a carrier-grade NAT configuration. As a workaround, you can use next-hop service set configurations and routing options to steer traffic to a multiservices (ms) interface where NAT functionality can be exercised. The following configuration snippet shows the basics of statically configuring the multiservices interface next hop and a next-hop service set. Traffic on which the service is applied is forced to the interface inside the network by configuring that interface as the next hop. This configuration does not show other routing-options or NAT configurations relevant to your network.

```

routing-options {
  static {
    route 0.0.0.0/0 {
      next-hop ms-3/0/0.1;
      preference 0;
    }
  }
  ...
}
services {
  service-set CGN {
    nat-rules CGN_SAMPLE;
    next-hop-service {
      inside-service-interface ms-3/0/0.1;
      outside-service-interface ms-3/0/0.2;
    }
  }
  nat {
    ...
  }
}

```

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

Software Defined Networking (SDN)

- **JDM restart failure in the first attempt**—In some cases, after the Juniper Device Manager (JDM) is stopped, attempts to restart it may fail and result in the following error message: **Job for jdm.service failed because the control process exited with error code. See "systemctl status jdm.service" and "journalctl -xe" for details.** The message may further contain the following: **error: Cannot set interface flags on 'macvlan8': Device or resource busy.** Note that the actual mac vlan value may vary. As a workaround, you can reboot the server. Before initiating the reboot, ensure that the mastership is switched to the guest network functions (GNFs) on the other server to minimize disruption on the affected server.

Software Installation and Upgrade

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

Subscriber Management and Services

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

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

SEE ALSO

New and Changed Features	 115
Changes in Behavior and Syntax	 149
Known Issues	 173
Resolved Issues	 202
Documentation Updates	 275
Migration, Upgrade, and Downgrade Instructions	 276

Known Issues

IN THIS SECTION

- [Class of Service \(CoS\)](#) | [174](#)
- [EVPN](#) | [175](#)
- [Forwarding and Sampling](#) | [176](#)
- [General Routing](#) | [177](#)
- [High Availability \(HA\) and Resiliency](#) | [189](#)
- [Infrastructure](#) | [189](#)
- [Interfaces and Chassis](#) | [189](#)
- [Layer 2 Features](#) | [191](#)
- [Layer 2 Ethernet Services](#) | [191](#)
- [MPLS](#) | [192](#)
- [Network Management and Monitoring](#) | [194](#)
- [Platform and Infrastructure](#) | [194](#)
- [Routing Policy and Firewall Filters](#) | [196](#)
- [Routing Protocols](#) | [197](#)
- [Services Applications](#) | [200](#)
- [Subscriber Access Management](#) | [201](#)
- [User Interface and Configuration](#) | [201](#)
- [VPNs](#) | [201](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 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 might fail when all of the following conditions are met:
 - Dynamic subscribers exist on an aggregated Ethernet bundle.
 - CoS traffic-control-profile or scheduler-map (or both) applied to these dynamic subscribers is from a static configuration.
 - The relevant static CoS is modified in the same configuration commit as a modification to the aggregated Ethernet bundle (either a leg add or leg remove) containing the subscribers.
 - The leg add or leg remove in the commit is the first or last leg to be added or removed from a line card.

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

subscriber cos update not applied to interface <interface-name>status <id> or subscriber cos update not applied to interface-set <interface-set-name> status <id>. These messages indicate that the last update to the subscriber or interface set was not applied.

As a workaround:

- (1) Remove the last class-of-service update.
- (2) Commit the configuration.
- (3) Re-apply the class-of-service update.
- (4) Commit the configuration. [PR1276459](#)
- Configuration of hidden statement **rate-limit-burst** in the class-of-service hierarchy. The commit needs to push an update for CoS code handling on all the Packet Forwarding Engines and during this time, if an interface settings (Internal attributes for an interface) was found to be NULL. The Interface settings are usually stored in a memory location and the pointer to it became NULL because CoSd did not check for the NULL values and resulted in segmentation fault. Channelized interface setting was found to be NULL for channelized interfaces, but the CoS code handling the configuration rate-limit-burst in Packet Forwarding Engine de-referenced the setting without doing NULL check, resulting in generating core files. [PR1425667](#)

EVPN

- The Layer 2 address learning process (l2ald) might generate a core file in a scaled Layer 2 setup, including bridge domain, VPLS, EVPN, and so on. The l2ald core file usually follows a kernel page fault that recovers on its own. In some cases, a manual restart of the process is needed to recover logs: **/kernel: %KERN-3-BAD_PAGE_FAULT: pid 69719 (l2ald), uid 0: pc 0x88beb5ce got a read fault at 0x6ca, x86 fault flags = 0x4 /kernel: %KERN-6: pid 69719 (l2ald), uid 0: exited on signal 11 (core dumped) init: %AUTH-3: l2-learning (PID 69719) terminated by signal number 11.** A core file is generated. [PR1142719](#)
- In an EVPN scenario with static MAC configured in the EVPN instance, the remote EVPN instance can see the MAC route information. However, after deactivating and activating the static MAC in the EVPN instance, and then checking the MAC route information in the remote EVPN instance, no such MAC route is found in the EVPN route table. [PR1193754](#)
- In an EVPN network with VXLAN encapsulation configured for direct-nexthop mode ("pure type 5" mode without overlay gateway addresses), at least one type 5 route per VRF from a remote endpoint must be received and installed in the local routing table of a device, 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](#)
- The issue is applicable to mac-in-mac PNN-EVPN and does not affect any other scenario. When PBB EVPN configuration is reloaded on MX Series routers, error logs are seen while deleting interfaces related to backbone bridge component. These errors does not result in any functional issues. [PR1323275](#)
- PBB EVPN cannot flood traffic towards a core layer. Traffic recovers by performing **restart l2-learning**. In addition to this, there is a limitation in PBB EVPN active/active (A/A) unicast traffic forwarding. If entropy in the traffic is not sufficient, then uneven load balancing causes a problem on MH peer A/A routers. This causes a drop for return traffic. These issues are applicable to mac-in-mac private network-to-network (PNN)-EVPN and does not affect any other scenario. [PR1323503](#)
- In an EVPN-VXLAN deployment, the rpd process might crash on the new master Routing Engine after performing a GRES. [PR1333754](#)
- In Junos OS platforms, the l2ald daemon might crash during MAC address processing. The MAC learning process will be impacted during the period of l2ald crash. The l2ald recovers itself. [PR1347606](#)
- Bidirection Layer 2 traffic floods for around 5 seconds for streams from SH to MH, when the **clear mac table** command is executed on the MX Series router because MAC addressing takes time to develop in the system. The **clear mac table** is a disruptive command that deletes all dynamic MAC addresses in the system. [PR1360348](#)
- In the scenario of EVPN Type-5 Route with MPLS encapsulation for EVPN-MPLS on MX platforms, if statement **chained-composite-next-hop ingress no-evpn** is configured, the EVPN type-5 route might be lost in the EVPN routing table. [PR1362222](#)
- Type 2 EVPN routes are missing after deactivating/activating protocol EVPN. [PR1362598](#)

- When EVPN is configured with class-of-service-based forwarding (CBF), traffic might be lost for the CBF services. [PR1374211](#)
- On EVPN-VXLAN scenario, during BGP flapping, the next-hop towards a VTEP (Virtual Tunnel End Point) might not be programmed properly, so if the traffic (especially inter-VNI traffic) destination is hashed through this Leaf/VTEP node, traffic loss might be seen. The reason is that due to BGP flap, the **route delete and route add request to rpd** might get compressed, which results in VXLAN database not getting updated with right unicast next hop to stitch it with VXLAN Encapsulation nexthop (VENH). Hence, VENH will not have unicast next hop to forward the traffic. [PR1415450](#)
- On an MX104 router, the **chained-composite-next-hop ingress evpn** is missing in junos-defaults group, this configuration statement has to be configured to make EVPN work Junos OS Releases prior to 18.3R1-S1/18.3R2, otherwise EVPN does not work as expected. [PR1415466](#)
- When DHCP is used with EVPN, Layer 2 learning daemon adds a destination route to kernel with a "permanent remote" flag while dhcp daemon adds a destination route with "permanent" flag. There could be a race condition where the layer 2 learning destination route gets overwritten by dhcp route, causing the remote flag to get deleted. This subsequently leads to the ARP route to age out in kernel. To ensure that dhcp routes are not added to kernel, **forward-only** statement must be configured under forwarding-options dhcp-relay. [PR1439568](#)

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](#)
- Heap memory leaks occur on DPC when the flow specification route is changed. [PR1305977](#)
- Fusion: Firewall filter not applied as input filter to extended port when used for Layer 2VPN. [PR1311013](#)
- This PR should fix some hints for the CLI commands to avoid confusion. With the fix, it will be like this:

```
{master}[edit] user@router# set firewall flexible-match source-ipv6-match bit-length
```

Possible completions: Length of integer input (1..32 bits), Optional length of string input (1..128 bits)

```
<<<< added information that for integer the limit is 32bit {master}[edit] labroot@beltway-re1# set
firewall flexible-match source-ipv6-match bit-length 120
```

```
{master}[edit] user@router# commit check re1: commit-check failed
commit-check failed error:
configuration check-out failed
```

for range, added the syntax check that no ", " or " is supported.

```
{master}[edit] user@router# set firewall family inet6 filter flex-match-v6 term source-ipv6 from
flexible-match-range range 0x00000001-0x00010001, 0x00010001-0x00010070^ syntax error.
```

```
{master}[edit] user@router# set firewall family inet6 filter flex-match-v6 term source-ipv6 from
flexible-match-range range 0x00000001-0x00010001
```


Possible completions:<[Enter]> Execute this command + **apply-groups** Groups from which to inherit configuration data + **apply-groups-except** Don't inherit configuration data from these groups **bit-length** Length of the data to be matched in bits (1..32) **bit-offset** Bit offset after the (match-start + byte) offset (0..7) **byte-offset** Byte offset after the match start point **flexible-range-name** Select a flexible match from predefined template **fieldmatch-start** Start point to match in packet| Pipe through a command {master}[edit] user@router# set firewall family inet6 filter flex-match-v6 term source-ipv6 from flexible-match-range range 0x00000001-0x00010001 or^ syntax error. [PR1389103](#)

- The error of traffic not getting policed as expected is seen after local switching for VLAN 100 AND 10. While verifying Selective Local-Switching functionality with 4000 VLANs. [PR1436343](#)
- ARP packets are getting dropped by the Packet Forwarding Engine after **chassis-control** is restarted. [PR1450928](#)
- Commit failure with error might be seen and the dfwd crashes when applying a firewall filter with action **then traffic-class** or **then dscp** to an interface. [PR1452435](#)
- On MX platforms, for an aggregated Ethernet bundle of at least two members hosted at two different FPCs, if the aggregated Ethernet interface is with CoS output-traffic-control-profile of shaping-rate and with the output filter of policer with logical-bandwidth-policer and bandwidth-percent, the aggregated Ethernet interface might have incorrect effective output policing rate. [PR1466698](#)
- On the MX platform with MPC line card (except DPC line card) used, if an input firewall filter is configured at the ingress VPLS interface, the packet with a VLAN priority of five, with three or more VLAN tags might be forwarded into the wrong queue. When this occurs, it might cause traffic loss due to congestion as all traffic is forwarded into the default queue. [PR1473093](#)
- If the **policy-map xx** option is configured under **family mpls** for filters, then the filters might not take effect after committing. [PR1478964](#)

General Routing

- If a Layer 3 interface is receiving a GRE encapsulated packet and interface has two filters attached in ingress as follows:

- a. **Family any** with action as mirror
- b. **Family inet** with action as **decapsulate gre**

then the expected behavior is that mirrored copy must have the GRE headers as well. However, that is not working as expected (and a bug) due to presence of filter (b). If the customer is interested in mirroring entire packet that came on the interface (that includes GRE header as well), then the workaround is to deactivate/disable the "decapsulate gre action of filter (b). [PR1090854](#)

- 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](#)

- While upgrading from Junos OS Release 15.1F based images to Junos OS 16.x and later releases or downgrading from Junos OS Release 16.x to Junos OS Release 15.1F images, if the **validate** option is enabled, chassisd might crash and upgrade or downgrade will fail. This issue should not be seen if both base and target images are from Junos OS Release 15.1F or Junos OS Release 16.x and later. [PR1171652](#)
- When same UID objects are used in both inet and inet6 services of the same subscriber session, deactivation the first session cause conditions which avoid releasing UID entry after deactivation second service session. This leads to having stale UID entry and can cause subscriber's connection problem in the future when UID pool would be completely exhausted. The probability of hitting the issue increases if amount subscriber to amount of unique services ratio is approaching 1 (that is, when almost every subscriber has a service with unique service objects). [PR1188434](#)
- 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](#)
- Source-address based Filter Based Forwarding is used under **forwarding-options** to steer the packets towards AMS bundle in the Vodafone configuration. When you remove the from **source-address** condition from the filter, the reverse traffic gets looped back into the AMS bundle. Under this condition, Prolonged Flow Control generates core files are seen. We do have from source-address configured in the SFW rule, which should have dropped the packets, which are getting looped back into the AMS bundle, but, this is not happening, even though SFW functionality works as expected for other packets. [PR1192184](#)
- With MPC8/9 MRATE MIC and plug-in optics module(QSFP28-100GBASE-LR4), bit errors might be seen. [PR1200010](#)
- Upgrading using unified ISSU might trigger a flap in the interfaces on MX Series routers. The following message might be seen: **SFP: pointer Null, sfp_set_present**. [PR1200045](#)
- After system boot up or after PSM reset, you might see the PSM INP1 or INP0 circuit Failure error message. [PR1203005](#)
- SMID daemon has stopped responding to the management requests after a jl2tpd (L2TP daemon) crash on an MX960 BNG. [PR1205546](#)
- 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](#)
- After configuring PCEP following log seen - pccd: [89798] Could not decode message from rpd. This might impact in growth of memory of pccd process over time, which can be cleared by restarting the process. [PR1235692](#)
- Aggregated Ethernet interface link remains down when it is configured with link-speed mixed mode. The issue was seen depending on the order of configuration. Creating aggregated Ethernet interface first, then add child interface to the aggregated Ethernet, aggregated Ethernet never becomes up. [PR1241275](#)

- Continuous logging as **PEM power status has changed, run power budget again** at chassisd logs is noticed. [PR1242847](#)
- When gRPC subscription for telemetry data with 2 seconds frequency, the jsd process might crash. [PR1247254](#)
- Load Balancing is uneven across Aggregate Ethernet member links when the AE bundle is part of an Equal Cost Multi-Path (ECMP) path. The AE member-links needs to span Virtual Chassis members. [PR1255542](#)
- On MX Series routers with XM chipset (for example, MPC3E, MPC4E, MPC5E, MPC6E, MPC2E-NG, and MPC3E-NG), the MPC might reboot when the unified ISSU completes. [PR1256145](#)
- The following cosmetic error is observed as the output: **mshpmand[190]: msvcs_session_send: Plugin id 3 is not present in the svc chain for session.** Please open a JTAC case to confirm. [PR1258970](#)
- After router reboot or JSD (JET service process) process crash, sometimes the listening socket for JSD (JET service process) is not operational. [PR1263748](#)
- The issue occurs when an interface comes online and both the OAM protocol and the MKA protocol try to establish their respective sessions. Because of contention between these two protocols, OAM takes down the interface and MKA fails to establish connection (because the interface is down, it cannot send out MKA packets). [PR1265352](#)
- On an MX Series Virtual Chassis system in a scaled subscriber management scenario, if a unified ISSU is performed while the BGP protocol sessions are active and such BGP sessions are clients of BFD, then these BGP sessions might go down and come back up again, causing traffic loss. [PR1265407](#)
- 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](#)
- GNFs in a node-slicing setup currently do not support Junos snapshot/recovery mechanisms. [PR1268943](#)
- DEP does not support dh group group19, encryption algorithm aes-256-cbc and hash sha-384 in its list of default proposals. These must be configured explicitly in the configuration. [PR1269160](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- Incorrect counters for output packets on child links ae0 interface when configured with new feature 'revertive'. [PR1273983](#)
- For inline jflow when configured **template-referesh-rate** and **option-refresh-rate** with both packets and seconds interval configuration options, the packets interval configuration is not working. [PR1274206](#)
- On vMX platform, performance of the Intel X710 NIC is lower compared to the performance of Intel 82599 NIC. Because, 10G line rate can be achieved at 512 byte packet size for X710 NICs where as same can be achieved at 256 bytes for 82599 NICs. [PR1281366](#)
- A routine within an internal Junos OS sockets library is vulnerable to a buffer overflow. Malicious exploitation of this issue might lead to a denial-of-service (kernel panic) or be leveraged as a privilege

escalation through local code execution. The routines are only accessible through programs running on the device itself, and veriexec restricts arbitrary programs from running on Junos OS. There are no known exploit vectors utilizing signed binaries shipped with Junos OS itself. Refer to [JSA10792](#) for more information. [PR1282562](#)

- MX Series Virtual Chassis only: When using a channelized configuration on MPC7, MPC8, or MCP9 MRATE PIC QSFP interfaces for VCP connections between members, a VCP interface needs to be configured on channel 0 of each QSFP to activate the port. [PR1283283](#)
- Due to vendor code limitation, ungraceful removing of summit MACsec TIC from chassis might cause a crash or unpredictable result. [PR1284040](#)
- This is in an internal change as Syslog usage is deprecated. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful/relevant messages. [PR1284625](#)
- This is in an internal change as Syslog usage is deprecated, however, there may be customer impact due to syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful/relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. [PR1284643](#)
- TVP platforms do not support chassisd hard restart command due to infra limitation. FPC power off does not happen cleanly as the old chassisd process initiates FPC power off command and exits. restart chassisd hard with GRES on MX10003 causes new chassisd process to open reconnect window and wait for connection. RE and FPC goes out of sync and FPC reconnect is not handled which causes FPC to be restarted multiple times. Finally, FPC comes online. [PR1293314](#)
- Junos OS releases with a fix committed in Junos OS Releases 15.1R5-S4, 16.1R4-S3, 16.1R5, and 17.3R1 with XM-based linecards (MPC3E, MPC4E, MPC5E, MPC6E, MPC2E-NG, or 3E-NG) might report **DDR3 TEMP ALARM** chassisd's error log message. [PR1293543](#)
- In some Junos MX deployments, random syslog messages are observed as below for FPC cards **fpcx ppe_img_ucose_redistribute Failed to evict needed instr to GUMEM - xxx left**. These messages are not an issue and would not cause any service impact. These messages will be addressed as "INFO" level messages. On Junos Packet Forwarding Engine there are dedicated UMEM and shared GUMEM memory blocks. This informational message indicates some evicting events between UMEM and GUMEM and can be safely ignored. [PR1298161](#)
- When a GRES or NSR is performed on a BSYS the master Routing Engine on the GNFs (virtual nodes/network slices) will detect the BSYS chassisd restart and enter a NSR hold down delay. During this time CLI commands to evoke a switchover on the master Routing Engine will indicate the system is not NSR ready. This is similar to a stand alone MX if chassisd is restarted on the master Routing Engine. Note that a CLI command on the BU RE will succeed. This too is in keeping with standalone MX behavior. [PR1298571](#)
- 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](#)

- Support for enterprise profile is only provided for 10-Gigabit Ethernet interfaces. Use of 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. [PR1310048](#)
- Alarm is raised if Mixed AC PEMs are present. Changed the criteria to check whether mixed AC is present. If the PEM is AC(HIGH) first bit of **pem_voltage** is set and if it is AC(LOW) second bit of **pem_voltage** is set. So if both first and second bit is set then MIXED AC is present. [PR1315577](#)
- Making changes in services traffic-load-balance instance for one instance, can lead to refresh of existing instances. [PR1318184](#)
- In JDM, (running on secondary server) jdmd daemon might generate core files if GNF add-image is aborted by pressing CTRL+C. [PR1321803](#)
- With regards fpc restarts/Virtual Chassis splits, the design of MX Series Virtual Chassis infra relies on the integrity of the TCP connections and the reactions to failure situation might not handle in graceful way : tcp connection timeout because of jlock hog crossing boundary value (5 seconds) causing bad consequences in MX Series Virtual Chassis currently no other easy solutions that would be able to reduce this jlock hog besides enable marker infra in MX Series Virtual Chassis setup. [PR1332765](#)
- USB is not pass through, hence you cannot access USB in Junos VM. [PR1333201](#)
- The output of the CLI command **show class-of-service fabric statistics** now calculates traffic that was dropped because of internal errors in the fabric forwarding path. [PR1338647](#)
- In Multiprotocol Label Switching (MPLS)/Resource Reservation Protocol (RSVP) environment, when the label-switched path (LSP) flapping cause RSVP LSP reroute, LSP might stick in Dn state with **Record route: <self>...incomplete**. [PR1343289](#)
- On MX platform with 100M SFP used on MIC-3D-20GE-SFP-E/MIC-3D-20GE-SFP-EH, SFP might not work if it is not from Fiberxon or Avago. [PR1344208](#)
- There is a possibility of MACSEC sessions not establishing if FPCs go through continuous cycles of offline/online (more than 10 times) followed by restarting dot1xd. [PR1344358](#)
- On Next Generation Routing-Engine (NG-RE), a failure of the Hardware Random Number Generator (HWRNG) will leave the system in a state where there is not enough entropy available to operate. [PR1349373](#)
- In some cases, Online insertion and removal (OIR) of a MIC on an FPC can lead to black-holing of traffic destined to the FPC. The only way to recover from this is to restart the FPC. The issue will not be seen if use the corresponding CLI commands to offline and then online the MIC. [PR1350103](#)
- On all Junos platforms, licenses might not take effect after successfully committing a license key configuration. [PR1350302](#)
- During stress conditions error log messages regarding route add, change, or delete might be incorrect. [PR1350713](#)
- When ephemeral DB instance is configured, if committing changes which are unrelated to IGMP/MLD (such as **set interfaces ge-0/0/1.0 description**), and the number of ephemeral commits reaches to ephemeral DB maximum size, the ephemeral DB purge might happen. Then it would purge all the commits

and rollover. On this purge the mgd gives all the applications a FULL COMMIT view. And on this FULL COMMIT view IGMP/MLD deletes all configurations and adds it back again. This might cause PIM to prune the groups on those interfaces and send join messages again. Finally, the multicast traffic flapping and drop might be seen. [PR1352499](#)

- On MX platform with the subscriber-management feature enabled, if the combination of an Ascend-Data-Filter (ADF) and a redirect filter is applied to the subscribers, it may cause a leak in the Broadband Edge (BBE) filter index. The index is not released when the subscriber logs out. Due to this issue, new subscribers are not able to connect when all the indexes are used up. [PR1353672](#)
- The log of **SMART ATA Error Log Structure error: invalid SMART checksum.** might be seen on FPC with WINTEC mSata SSD. [PR1354070](#)
- BGP IPv4 PIC: Packet Forwarding Engine Selector stuck in a rerouted state on the Unilist NH after Primary aggregated Ethernet link deactivate or activate. [PR1354786](#)
- If the packets are destined to specific MAC address (such as last two octets are 0x1101, 0x1102, 0x1103, 0x1104, 0x1106, 0x1108, 0x1109, 0x110a, and so on), they might be dropped on the remote-end device when going through MX104 built-in xe(10GE) ports. [PR1356657](#)
- The bbe-smgd process may restart unexpectedly. This issue is found while performing subscriber service's GRES test suite under heavy load. [PR1359290](#)
- Craftd messages are generated on Summit MX 3RU (mx10003) and Summit MX 1RU (mx204) platforms. Summit platforms do not have Craft Interface. Hence these errors are expected, and can safely be ignored. When Craftd daemon tries to open the device, it fails with a junk char in the fatal error message because the error no is not mapped to a string in the kernel code. **Feb 20 01:49:38 MX craftd[xxxx]: craftd detected platform mx10002 Feb 20 01:49:38 MX craftd[xxxx]: LIBJSNMP_SA_IPC_REG_ROWS: ns_subagent_register_mibs: registering 1 rows Feb 20 01:49:38 MX craftd[xxxx]: fatal error, failed to open smb device: ,JlÈ.** [PR1359929](#)
- With MPC5E, MPC2E-NG, and 3E-NG and large-scale configurations along with large amount of traffic causing non-zero stats on multiple queues, when executing unified ISSU, the ISSU prepare stage might take longer time than usual because PR 1283850 introduced a bug which could cause the stats disable to take longer. As a result, the chassisd triggers restart/crash of the MPC and the ISSU completes after the crash. [PR1369635](#)
- After successfully delegating a locally configured LSP to a PCE, the router still displays 0 as the "Delegated" counter value under the output of CLI command **show path-computation-client status**. [PR1369929](#)
- The voltage high alarm might not be cleared when voltage level comes back to normal for MIC on MPC5E. [PR1370337](#)
- On MX platform enabled with enhanced subscriber management, if the subscriber profile initiates a filter service for each subscriber, and there are large scale of Broadband Edge (BBE) subscribers (for example, 10000) logging in and out repeatedly, the filter service might fail to get installed for the subscriber due to this issue. In some rare condition, it might also lead to the Flexible PIC Concentrator (FPC) crash. [PR1374248](#)

- I/O session used for communicating between threads is freed due to FSM state transition. After freeing the memory, the fields of the I/O session are used for tracing causing RPD core files. [PR1374759](#)
- In subscriber scenario, if the "service-accounting-deferred" is configured on dynamic-profile, and there is multicast to a large number of destinations on the same physical port, the FPC Errors might be seen. [PR1380566](#)
- **set vmhost <>** configuration command needs to be available user who have **system-control permission** (in order to be in line with **set system <>** command). But **set vmhost <>** is available user who has **system-control permission** corrected the same. [PR1383706](#)
- It is possible to configure the purge timeout of programmable RPD clients to 'never'. This will mean that the routes added by PRPD clients will not be deleted when client disconnects. They will stay until routing daemon restarts or it is deleted by the client that added the route. This can be configured using following CLI command. Note the programmable API for setting purge timeout does not support this feature yet. Set **routing-options programmable-rpd purge-timeout never**. [PR1384303](#)
- Due to transient Hardware condition single-bit error (SBE) event are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary Hardware replacements. [PR1384435](#)
- On MX platform enabled with subscriber scenario, if large scale of subscribers (for example, more than 1000 subscribers) set up connections simultaneously, the setup rate might be 30 percent lower than expected. [PR1384722](#)
- MPC2E NG/MPC3E NG card will go in error with error id XM Chip Error code: 0x701ca [PR1384830](#)
- When traceoptions are enabled with a lot of trace flags or 'flag all', the rpd might crash due to buffer overflow issue. This is a timing issue. [PR1387050](#)
- In low end 32-bit systems, rpd has a lower level of available memory. It is desired to have a log message to alert customer when the average memory usage or transient memory usage exceeds thresholds. [PR1387465](#)
- During Zero Touch Provisioning (ZTP) process, default route is being cleaned up by code. Due to this if a static default route is configured in the initial configuration (configuration file downloaded from the file server for ZTP), the route will fail to work. This might lead to ZTP failure or device access issue after ZTP. [PR1387724](#)
- Bbe-smgd core files when MTU configuration is changed with subscribers are still logged in on the ifd. MTU configuration change should only be done when there are no subscribers logged in on the ifd. Catastrophic configuration changes should be done only in maintenance mode, when no subscribers are on the ifd. [PR1389611](#)
- If the statement **persist-groups-inheritance** is configured, when trying to add additional sites to existing group and routing-instance configuration, error might be observed and it leads to fail to commit after issuing "commit check". [PR1391668](#)
- On MX2008 platform with MPC9E, in line rate traffic with a redundant SFB2 scenario, if offline one redundant SFB2, there might be tail or sometimes WRED drops in MPC9E, resulting in partial traffic

loss. Under normal circumstances, the SFBs should be auto fail-over if one of them fails, and there should be only a little packet dropped momentarily. [PR1395591](#)

- MPC 7, MPC 8, and MPC9 cards have a local disk which they keep a copy of the software image. The cards boot from the disk when an image is there, and boot from the chassis network (through BOOTP) when an image is not there. Presumably, new MPC 7, MPC 8, and MPC9 cards do not have an image on the disk and would require a network boot. On single chassis, there is no problem. But on MX Series Virtual Chassis, the network boot does not work. [PR1396268](#)
- On MX Series platforms, if Channelized OC MIC (such as 1xCOC12/4xCOC3 CH-CE) is used, the MPC card/AFEB/TFEB (Forwarding Engine Board) might crash with core files. This is not easily reproducible. The traffic through the MIC would be impacted. [PR1396538](#)
- The Junos RPD daemon has facilities to attempt to trap certain classes of non-fatal bugs by continuing to run, but leaving a "soft" core file. Leaving a soft core is intended to be non-disruptive to routing and forwarding. This PR implements a mechanism by which users may disable soft core files being generated. [PR1396935](#)
- Router is advertising the ESMC QL of PRC even though the Current clock status is holdover. This behavior is addressed in this PR and will be applicable to all platforms. [PR1398129](#)
- In MPLS over UDP or MPLS over GRE scenario, if the next hop type of the MPLSoUDP/MPLSoGRE tunnel is interface route, the tunnel may not come up. [PR1398362](#)
- The authentication module for JET RPCs and Telemetry fails in authenticating usernames or passwords of certain lengths. Hence the users will be unable to execute JET APIs or Junos Streaming Telemetry. [PR1401854](#)
- After upgrading Junos to Junos OS Release 17.2 or later, the statement **chained-composite-next-hop ingress l3vpn extended-space** cannot be configured any longer on a Logical system. [PR1402390](#)
- On MX Series platform with MS-MPC card used, in race condition, if the MS-MPC is used on HA (High Availability) scenario (the **set interfaces ms-x/x/x redundancy-options redundancy-peer/redundancy-local** knob and GRES is configured), the FPC might crash due to the bus error (segmentation fault). The reason is that when two CPUs simultaneously access the same session-extension memory in the session structure, one for writing, the other for reading. A reading CPU gets an incorrect value and uses that as the memory address. This causes the bus error (segmentation fault). [PR1405917](#)
- The process rpd might crash after a non-forwarding route (that is, a route to an indirect next-hop association is non-forwarding indirect next-hop) which is received from multiple protocols is resolved again by using the non-forwarding path. [PR1407408](#)
- On MX platforms using MPC7E, MPC8E, MPC9E, MX10k-LC2101 or MX10003, when inline-jflow application is used, Fatal error on Hybrid Memory Cube (HMC) will perform **disable-pfe** action. Since Jflow records are hosted on the HMC memory partition, reading and writing to the HMC memory might trigger FPC crash and high FPC CPU utilization, causing slow convergence (adding/deleting routes or nexthops) for other Packet Forwarding Engines on the same FPC carrier. [PR1407506](#)

- Configuration database can remain locked after the ssh session is halted. [PR1410322](#)
- On MX2020 and MX2010 platform, traffic traversing MPC8E or MPC9E may be discarded after one of SFB2s goes offline and it is requested online. This is a timing issue as it is not reproducible all the time. [PR1410813](#)
- In MPC8 line card, enabling both bandwidth knob along with flex-flow-sizing knob may result in Jflow service getting disabled due to not able to allocate the memory requested by flex-flow-sizing knob. [PR1413513](#)
- PCE initiated LSPs get deleted from PCC if PCEP session goes down and gets re-established within **delegation-cleanup-timeout** period. [PR1415224](#)
- In Virtual Private LAN Service (VPLS) multihoming with Label-switched Interface (LSI) interfaces used scenario, if the IPv6 neighbor is established via the VPLS, the IPv6 neighbor might become unreachable after the primary link of the VPLS multihoming goes down. The issue results in traffic loss for the IPv6 neighbor. [PR1417209](#)
- With Netconf the xmlns attribute is printed twice for **rpc <get-arp-table-information>** to the router. [PR1417269](#)
- Certain JNP10008-SF and JNP10016-SF manufactured between July 2018 to March 2019 may have incorrect core voltage setting. The issue can be corrected by re-programmed the core voltage and updated the setting in nvram memory. [PR1420864](#)
- If HyperText Transfer Protocol (HTTP) Header Enrichment function is used, the traffic throughput decreases when traffic passes through Header Enrichment. [PR1420894](#)
- On MX platform, with 1xCOC12 or 4XCOC3 used, if channelized interfaces are configured, FPC CPU overuse might be seen. [PR1420983](#)
- On all Junos platforms, when the file system gets into full state and there is no enough spare disk space, it might get into a problematic system condition in some corner case while doing configuration commit. After that, if consecutive commits are still done in such a problematic status, commit-check failure logs might be seen eventually. Due to this issue, some process might be not running even if its configuration is present. [PR1423500](#)
- Even though **disk-failure-action reboot** or **disk-failure-action halt** is configured, the system does not reboot or halt as expected when it encounters the disk error. [PR1424187](#)
- The issue is limited to DB related to MAC-MOVE scenario. When dhcp-security is configured, if multiple IPv4 and IPv6 client's MAC-MOVE happens, the jdhcpd might consume 100% CPU and jdhcpd will crash afterwards. [PR1425206](#)
- On all junos platforms running 64-bit mode rpd, the rpd will crash continuously if MD5 authentication on any protocols (like MD5 authentication for BGP/ISIS/OSPF) is used along with master-password. [PR1425231](#)
- On some fixed MPCs with builtin PICs, the ENTITY MIB has incorrect contained in values for PICs when doing snmp mib walk for oid .1.3.6.1.2.1.47 .[PR1427305](#)

- On MX platforms with ppp configured, when something abnormal happens such as the user's dialup router is abnormally powered off, or the keepalive packet is dropped due to network problem, the ppp session will ageout, while in a rare case, the ppp session is not getting deleted accordingly, which make the new session cannot be created. So new session is not able to log in. The ppp traffic might be dropped since duplicate-protection feature on the interface. And the IP address of the ppp interface cannot be pingable. [PR1428212](#)
- MX is discarding the traffic coming from framed-route hosts even if IPoE subscriber installed with valid IPv4 framed routes. In Customer scenario, **demux-source** is configured with variable **\$junos-subscriber-ip-address**. [PR1429743](#)
- On MX platforms, in a subscriber management configuration, if VPLS encapsulation is only configured under user-facing interface while it is not configured under the core-facing interface, when incorrect configuration checks is being performed, this configuration is prevented from being applied. [PR1430360](#)
- Multiple delete of a non existing config statements produces errors via rpc load-configuration. [PR1431198](#)
- Dual Stack Subscriber Accounting Statistics are not baselined when one stack logs out. [PR1432163](#)
- When SSH keys are generated during downgrade or upgrade of an image (usually on the first boot), <output> XML tags are visible in the messages. Taking out the xml tags will cause issues in netconf session. This is a minor cosmetic issue, hence does not have impact on the functionality. [PR1432464](#)
- Digital Optical Monitoring MIB (jnxDomMib) currently not supported on MX150 and VMX platform. [PR1432982](#)
- On MX platform with Trio based FPCs, if **sa-multicast** ' is in the configuration, all traffic will be dropped. [PR1433306](#)
- URI portion in URL will become case-sensitive through a hidden configuration statement **url-case-sensitive** under **url-filter-template**. Existing behavior is the default that is, URL is case-insensitive including URL. **url-filter[web-filter] {profile <name> {????????url-case-sensitive;????????}} }**. [PR1434004](#)
- On MX series routers with MPC7E, MPC8E, or MPC9E installed, if optics QSFP-4X10GE-LR (Part number 740-054050) is used, the link might flap. [PR1436275](#)
- The CPU utilization on mib2d daemon might keep at high level in race conditions (it may get hit or triggered at times by some churn in the system, no specific trigger). [PR1437762](#)
- On all Junos platforms, if hash-key is enabled, packets might be dropped due to chassisd crash, even packets on other FPCs which the hash-key is disabled. [PR1437855](#)
- In VSTP scenario, if flexible vlan tagging is configured on the interface and multiple IFLs are configured for the interface, if **vlan all interface all** is configured under VSTP, not all interfaces are enabled for this protocol. [PR1439583](#)
- Egress stream flush failure and traffic blackhole could occur on a rare occasion for a repeatedly flapping link on MPC7, MPC8, and MPC9E cards. [PR1441816](#)

- When MX configured with **route-modify-same-preference** statement, if the RADIUS returned framed-route is incorrect, such as "192.168.1.4/24", it will leave that route in the bbe-smgd and rpd not cleared. [PR1445155](#)
- The multiservices PIC manager daemon (mspmand) runs on service PIC (MS-MPC/MS-MIC) and is responsible for managing URL Filtering service if URL filtering feature is configured. The mspmand process might crash if URL filtering is configured and one blacklisted domain name is a sub-string of another blacklisted domain name in URL filter database file. This would be continuous crashes until all the sub-string entries are removed from the system. [PR1445751](#)
- On the platforms that do not support Router Advertisement Guard (RA Guard), such as PTX, after issuing the command **show access-security router-advertisement-guard ??**, the process jdhcpd may crash. [PR1446034](#)
- On MX platform, when switchover a service interface that has NAT and GR configuration, the static route for NAT might never come up. [PR1446267](#)
- Currently ISIS is sending system host-name instead of system-id in OC paths in Isdb or Adjacency xpaths in periodic streaming and on-change notification. [PR1449837](#)
- In subscriber scenario, when a new burst-size of traffic-control-profiles (TCP) is configured under dynamic-profile, the new burst-size can not take effect, instead, the old burst-size is still activated actually. In the corner case, this will cause packets to drop. [PR1451033](#)
- On the MX10003 platform, the alarmd wont write the alarm messages to the syslog. [PR1453533](#)
- IPV6 accounting stop attributes are not correct for MLPPP subscribers. [PR1455175](#)
- In the previous Junos version, the CLI command combination of "invoke-on" and "display xml rpc" may give incorrect RPC command, because this combination is not supported in Junos. E.g., issuing the command "show version invoke-on all-routing-engines | display xml rpc". [PR1456578](#)
- MX-flexible-vlan-tagging change in MTU behavior from 15.1R [PR1456809](#)
- When VRRP (virtual router redundancy protocol) is configured on MC-LAG (multichassis link aggregation groups), traffic destined to VRRP virtual MAC address might get dropped because the virtual MAC is not correctly programmed in PFE (packet forwarding engine). [PR1459692](#)
- When labeled-IPv6 and non-labeled IPV6 prefixes are received with the same protocol nexthop and the outgoing interface does not have MPLS family enabled, the IPv6 non-labeled route will be in inactive state and remains in hidden state. [PR1460786](#)
- Fabric hardening (FH) is the process of controlling bandwidth degradation to prevent traffic black hole. When FH is processing, if SFB/SCB get failure, FH process will be stuck, which will get traffic lost. [PR1461356](#)
- With 17.4R2/17.4R2-S2+ release, if any MX10003 FPC restart or is set to be offline after GRES, the other FPC might coredump and all PICs might get stuck at offline state. Release in 18.x and above are not affected. [PR1462686](#)

- The DFE tuning enabled interfaces on certain MX platform might get stuck in down state, if the remote interface sends invalid code to the local interface. Link might not come up even after the remote peer has begun sending a good signal. [PR1463015](#)
- If any MIC of MIC-3D-2XGE-XFP / MIC-3D-4XGE-XFP / MIC-3D-20GE-SFP-E / MIC-3D-20GE-SFP-EH / MIC-MACSEC-20GE is installed in MPC2E-NG/MPC3E-NG card, the Microkernel (uKern) might hog for CPU on Packet Forwarding Engine (PFE) when there is a high rate of interface flaps (~30/40 flaps per second). This will eventually trigger the MPC2E-NG/MPC3E-NG card crash with an NGMPC core file. Normally the excessive interface flapping won't happen frequently in real world and it may be caused due to external environment. This fix will reduce the impact and prevent the uKern hog when having such conditions. [PR1463859](#)
- On MX platforms with MS-MIC/MS-MPC, when stateful firewall is configured with "application junos-dce-rpc-portmap" and RPC ALG is enabled (both Sun RPC and MS-RPC), the mspmand might crash continuously (about every 15 or 20 minutes). [PR1464020](#)
- If a netconf session is initiated over inband connection, the CPU utilization on mgd daemon might be stuck at 100% after the netconf session which is executing an RPC call for some commands gets interrupted by flapping interface. There is no impact observed to control-plane or forwarding-plane, the subsequent netconf session will continue to function. [PR1464439](#)
- Traceroute generates ICMP error message like Destination Host unreachable, Time exceeded etc which actually helps in identifying the intermediate hops. Code Logic for handling ICMP errors was not there as part of asymmetric processing. [PR1466135](#)
- In the PPPoE subscriber management environment, due to the PPPoE inline keepalives timeout events may get dropped by the RE (routing engine), the PPPoE subscribers get stuck. This issue may cause the PPPoE subscribers are unable to reconnect. [PR1467125](#)
- Crypto library shim memory utilization performance improvement by using data shim instead of control shim. [PR1467874](#)
- On Junos from 16.2R1 onwards, if "commit" is executed after "commit check", the daemon (e.g. dhcpcd, sampled) might not be started even the related configuration is successfully committed. [PR1468119](#)
- When tunnel-services are configured on a PIC, the optics measurements that subscribed via gRPC might not be streamed. [PR1468435](#)
- On all Junos platforms with l2cpd (Layer-2 control protocols) daemon, committing configuration changes which are processed by l2cpd (e.g., flexible-vlan-tagging, stacked-vlan-tagging, vlan-tagging, family ethernet-switching) might cause marginally memory leak. Committing the l2cpd processed configuration changes in a successive manner might cause the memory resource exhaustion (Some operations have the same effect as the committing action, e.g., bouncing a vlan-tagged interface in a successive way). Eventually, it could result in the l2cpd process crash. [PR1469635](#)
- VMCORE-../src/junos/bsd/sys/netjsr/jsr_prl.c:2128 [PR1472519](#)

High Availability (HA) and Resiliency

- The following error is seen during early unified ISSU validation phase: **error: not enough space in /var on re1**. As a workaround, make sure that the space available in /var is twice the size of the target image. This is the basic requirement for unified ISSU to proceed. [PR1354069](#)

Infrastructure

- The /var/run is in storage file system but it should be in memory file system. [PR1198395](#)
- The configuration statement **set system ports console log-out-on-disconnect**, logs the user out from the console and closes the console connection. If the configuration statement **set system syslog console any warning** is used with the earlier configuration and when there is no active telnet connection to the console, the process tries to open the console and hangs as it waits for a "serial connect" that is received only by 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](#)
- On MX Series, if GRES is not configured, while "master-only" IP is configured on fxp0.0, the IP address might not be applied to the interface after reboot. [PR1341325](#)
- If you pulled out a USB from the system while files are being copied, the kernel will panic and the system will restart. [PR1425608](#)
- On all Junos platforms that are upgraded to Junos OS Release 15.1 onward, when the duplex setting is changed on the management interface (for example, fxp0/em0), the duplex status of the management interface might not be updated in the output of the "show interface <>". [PR1427233](#)
- The service utility "cron" runs in the background and regularly checks /etc/crontab for tasks to execute and searches /var/cron/tabs for custom crontab files. These files are used to schedule tasks which "cron" runs at the specified times. "cron" daemon is started during boot. If for some reason, the "cron" process exits, the scheduled tasks will not be executed. "cron" was not restarted automatically and had to be started manually. [PR1463802](#)

Interfaces and Chassis

- Junos now checks IFL information under the ae interface and prints only if it is part of it. [PR1114110](#)
- In Junos BNG solutions, after commit event, when configuration contains duplicate vlan-id configured on aggregate and demux interfaces, Junos MX Series routers may go into db prompt mode and kernel core files asserted. [PR1274038](#)
- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after upgrade. This is because of 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](#)
- Y.1731 Dwlpy measurement is not supported on MPC6. [PR1303672](#)
- In MX Virtual chassis, flooding of the error message **?CHASSISD_CONFIG_ACCESS_ERROR: pic_parse_ifname: Check fpc rnage failed** can be seen with LACP enabled aggregated Ethernet interfaces on MPC7, MPC8, and MPC9 cards. The errors will only have impact for DWDM pics, which does not effect on the MPC7, MPC8, and MPC9 cards. Hence this syslog message can be safely suppressed. [PR1349277](#)
- In L2VPN (Layer 2 Virtual Private Network) scenario with sonet interface which is used for PE-CE link, that sonet interface might go down after enabling "keep-address-and-control" knob on it. [PR1354713](#)
- With ppp-service traceoptions configured as: **user@router> show configuration protocols ppp-service traceoptions file jtac-jpppd.log size 1g files 10; level all; flag all; filter {user {"subscriber@domain.com";}}**. It is expected to see only PPP negotiation events belong to subscriber defined in filter section. However in releases affected by this issue several stings of logs related to other (non interested) subscriber may be seen. [PR1370994](#)
- In large scale subscriber environment, changing ae member link configuration may cause two REs coredump. [PR1375638](#)
- Static demux0 logical interfaces do not come up after configuration change if underlying interface is et (100 GE). After the configuration change et interface gets flushed in order to reparse the configuration. During this DCD miss to create the dependency between demux0 logical interfaces and underlying et interface which results in flushing off the demux0 logical interfaces. This issue will be seen only if underlying interface is et. For all other interfaces this has been already taken care. This is day one issue. As a workaround, restart DCD (or the entire RE reboot) to clears the problem or else use 'commit full' instead of commit while committing new configuration. [PR1401026](#)
- On MX Series platforms, EX-SFP-1FE-LX SFP does not initialize with MIC-3D-20GE-SFP-E(EH). [PR1405271](#)
- On all Junos platforms, if this is any protocol running upon aggregated Ethernet interfaces, while committing any configuration changes related to aggregated Ethernet interfaces, unrelated aggregated Ethernet interfaces might go down. [PR1409535](#)
- When an unnumbered interface is binding to an interface which has more than one IP address and one of the IPs is deleted, the family inet of the unnumbered interface might be getting deleted. The issue results in traffic loss for all the services that rely on the family inet of the unnumbered interface. Configure preferred-source-address on the unnumbered interface will prevent deletion of the IP hence avoiding the deletion of the family inet of the unnumbered interface. [PR1412534](#)
- If aggregated interface(ae) has vrrp configuration, in following use cases, member IFLs will not be created after member IFD comes up and ae will be in down state.

1. fpc restart (request chassis fpc restart slot <>)
 2. chassis-control restart (restart chassis-control)
 3. reboot both RE (request system reboot both-routing-engines). So before performing above operations, it is advisable to remove vrrp configuration from aggregated interface(ae). [PR1429045](#)
- Customer need two knobs for EOAM CFM interoperability between MX10003 and Ciena CPE Two knobs: 1. primary-vid - this allows interop with Ciena CPE - which is used at every tower site to est. EOAM CFM session 2. enhanced-cfm-mode - provides required scale needed for EOAM for CBH and METROE services [PR1465608](#)
 - When EVPN setup in MPLS Active/Active (A/A) or VxLAN A/A environment, if Ethernet Segment Identifier (ESI) is configured on a physical interface (IFD) of multi-homed PE, Designated Forwarder (DF) election will not happen when the logical interface (IFL) under the IFD is disabled. As a result, this issue will cause traffic drop. [PR1467855](#)
 - When dynamic DHCP sessions are existing in the device, if multiple commits in parallel are performed, the commit might hang up. [PR1470622](#)

Layer 2 Features

- For router equipped with following line cards:
T4000-FPC5-3DMX-MPC3E-3DMPC5E-40G10GMPC5EQ-40G10GMPC6E MX2K-MPC6E.
If the router is working as VPLS PE, due to MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)
- With VPLS being configured, after upgrade to 15.1/16.1/17.x releases, in some circumstances VPLS LSI interface are not correctly created, causing remote MACs not being learnt and L2 VPLS outage. The issue is not reproduced and the code change is not a fix but add instrumentation using a hidden command 'show vpls ipc-history', which should be captured right away when the issue is seen on latest releases. show vpls ipc-history <<<<< show vpls connections show krt queue show route forwarding-table extensive /var/log/messages [PR1295664](#)
- On all Junos platforms with NSR enabled, under EVPN-VPLS scenario, the VPLS neighbors might stay in down state after configuration changes in vlan-id. [PR1428862](#)

Layer 2 Ethernet Services

- This is an internal change as syslog usage is deprecated, however, there may be customer impact due to syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful or relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. [PR1284592](#)
- In MC-LAG with force-up scenario, the LACP PDU loop might be seen when both MC-LAG nodes and access device use same admin key. [PR1379022](#)

- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case the subscriber may need more time to get IP address assigned. The subscriber may remain in this state until its lease expires if it has previously bound with the address in the option 50. [PR1435039](#)
- When DHCP is configured, if subscribers are moved from one routing-instance to another or if the subscriber is deleted and re-added, the jdhcpd process might go into infinite loop and cause 100% CPU usage. [PR1442222](#)
- There are two options to configure DHCP relay, one is using **dhcp-relay** knob, the other is using helper bootp. On MX10000 platform, neither of DHCP-Relay nor helper bootp function can work. [PR1447323](#)

MPLS

- When using **mpls traffic-engineering bgp-igp-both-ribs** with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit area border router (ABR) when there are two or more such area border routers (ABRs). This causes interdomain RSVP LSPs to break. RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. [PR1048560](#)
- The issue occurs when graceful Routing Engine switchover (GRES) is done between the master and backup Routing Engines of different memory capabilities. For example, one Routing Engine has only enough memory to run routing protocol process (rpd) in 32-bit mode while the other is capable of 64-bit mode. The situation could be caused by using Junos OS Release 13.3 or later with the configuration statement **auto-64-bit** configured, or, by using Junos OS Release 15.1 or later even without the configuration statement. Under these conditions, the rpd might crash on the new master Routing Engine. As a workaround, this issue can be avoided by using the CLI command **set system processes routing force-32-bit**. [PR1141728](#)
- The routing protocol process (rpd) might crash in the backup Routing Engine when LSP tunnels are present with an NSR configuration. [PR1186292](#)
- In a CE-CE setup, traffic loss might be observed over the secondary LSP on primary failover. [PR1240892](#)
- If the primary link goes down immediately after bypass (for example, FPC containing both primary and bypass or, both primary and bypass FPCs go down simultaneously) such that primary link goes down even before the PLR sends out any Path message after bypass down, then the nodes downstream of the PLR along the LSP path will be left with stale LSP state until refresh timeout. This condition will not result in any traffic loss. [PR1242558](#)
- 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](#)
- It takes longer to set-up Layer 3 VPN egress protection starting from Junos version 16.1R1. [PR1278535](#)

- In case of CSPF-disabled LSPs, if the primary path ERO is changed to an unreachable strict hop, sometimes the primary path stays up with the old ERO. The LSP does not switch to standby secondary. [PR1284138](#)
- An SR-TE path with "0" explicit NULL as inner most label, SR-TE path does not get installed with label "0". [PR1287354](#)
- Swapping the binding SID between colored and non colored static SR LSPs might cause rpd to generate a core file. [PR1310018](#)
- The Packet Forwarding Engine on Trio platform or PTX/QFX10000 follows a certain conversion logic to convert MPLS-VPN labels to certain channel values, and then back to MPLS-VPN labels. VPN labels having values 0x7FFFF and above (524287 and above) are affected by this conversion logic. [PR1323496](#)
- If inet address is not configured for the gr- interface, the gr- interface will borrow address from loopback interface. From 16.1R1, the RSVP creates a node-neighbor by default. There are duplicate neighbors with the same IP address since the gr- interface is borrowing address from loopback interface. The RSVP path lookup will fail because it gets confused with the node neighbor presence. So the RSVP LSP will not come up when it goes through the gr- interface which is borrowing address from the loopback interface. [PR1340950](#)
- Executing a **restart chassisd** in a router with scaled configuration might result in rpd core files. [PR1352227](#)
- When 'tunnel-services' is configured under 'chassis fpc <> pic <>', the vt-x/y/z physical interface (IFD) is created for the corresponding FPC. If 'protocols rsvp' is configured, RSVP will create a default vt-x/y/z.u logical interface (IFL) under the corresponding vt-x/y/z IFD. After applying a configuration change that will remove RSVP and trigger FPC restart, the vt-x/y/z.u IFL is not cleaned up due to a code issue. Hence the corresponding vt-x/y/z IFD cannot be cleaned up during the corresponding FPC coming up. The IFD cleaning keeps retrying which cause the corresponding FPC to be stuck in 'Ready' state. [PR1359087](#)
- When traceroute to a remote host for an MPLS LSP using the command **traceroute mpls bgp**, in very rare cases, it is possible that mplsoam daemon is holding the stale BGP instance handle in the query to the rpd process to get the information for the Forwarding Equivalence Class (FEC), hence rpd crash might occur because of the invalid instance. It may cause traffic impact till rpd comes back up. [PR1399484](#)
- On Junos platforms with scaled MPLS labels used, when the system is already running with high load, inefficient labels allocation might cause even higher CPU utilization at 100 percent for hours. The issue might affect traffic. [PR1405033](#)
- The LDP transit egress route for a BGP route has an indirect nexthop. In NSR and GRES scenario, after Routing Engine switchover, in some cases, LDP might fail to receive route flash for a BGP route from inet.0 and would not update the inet.3 route for the BGP route. As a result, the nexthop for LDP transit egress route will become unusable and the LDP transit egress route will get deleted. It could cause BGP sessions to go down and cause traffic drop. [PR1420103](#)
- Dynamically configured RSVP LSPs for LDP link protection may not come up after disabling/enabling protocol mpls. [PR1432138](#)
- In inter-domain RSVP (Resource Reservation Protocol) LSP (Label-switched Path) scenario, the rpd memory leak might be seen when the CSPF (Constrained Shortest Path First) tries to recompute the

path for the "down" LSP which is due to no route or ERO is incorrectly configured. The issue might lead to rpd crash when the rpd is out of memory and results in traffic loss. [PR1445024](#)

- In RSVP environment with link or node protection deployed, if two consecutive PLRs (point of local repair) along the LSP perform local repair simultaneously and if backup LSP signaling between the downstream PLR & MP (merge point) pair fails due to any reason, then the backup LSP signaling between the upstream PLR & MP pair also does not succeed. Then due to a software defect the upstream PLR does not correctly clean up the LSP state and continues to send traffic into the backup LSP, resulting in traffic blackhole at the downstream PLR. [PR1445994](#)
- In Link Aggregation Control Protocol (LACP) with Unilist next-hop scenario, when Resource Reservation Protocol (RSVP) protection or BGP Prefix-Independent Convergence (PIC) is used, if the LACP interface flapping happens fast enough, which might cause traffic blackhole. Due to a delay which causes the first "link down message" arriving at Packet Forwarding Engine (PFE) after the "link up message" already being received. So that PFE marks both of the primary and backup next-hop as unusable. (This is a timing issue) [PR1452866](#)
- On all platforms with BGP PIC configured, if doing some commit operations where RSVP ingress routes are affected, the rpd crash might be seen. [PR1471281](#)
- In a corner case on Junos platform, where the family ccc is configured along with any other existing family within the same interface, like inet, inet6, etc. (basically, Junos never allows to do so, but somehow a customer did it). And if the family ccc is deleted from the interface, which might cause kernel crash and the device reboot automatically, so all the traffic will be interrupted. [PR1478806](#)

Network Management and Monitoring

- Issue: snmpd daemon leaks memory in snmpv3 query path and crashes. Cause: The issue is caused by a memory leak when the request PDU is dropped by snmp when the configuration - **snmp filter-duplicates** is enabled. Each request PDU has a structure pointer for the SNMPv3 security details. This is allocated when the pdu is created or cloned. But while dropping the duplicate requests the corresponding free for this structure is not done, which causes the memory leak. [PR1392616](#)

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as `nh_ucast_change:291Referenced l2ifl not found`. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- The **login_getclass: unknown class 'j-idle-timeout** error is getting displayed when the user has not configured timeout value for root user. If the user has not configured timeout value, j-idle-timeout entry is not present in login.conf file and error message is displayed because j-idle-timeout class is not found. Steps to Reproduce:

1) Login to router as root.

2) Clear log messages.

3) Exit and go to CLI mode and type **show log messages**. The login error should be logged in the messages. [PR1097799](#)

- The error messages about **LUCHIP(5) GUMEM1[77a0] mismatch** might be seen after MX Series MPC card with LU chipset goes offline or online. [PR1221195](#)
- With unified ISSU, momentary traffic loss is expected. In EVPN E-Tree, in addition to traffic loss, the known unicast frames can be flooded for around 30 seconds during ISSU before all forwarding states are restored. This issue does not affect BUM traffic. As a workaround, nonstop bridging (NSB) can be configured at **[set protocols layer2-control nonstop-bridging]**. This reduces traffic flood to around 10 seconds in a moderate setup. [PR1275621](#)
- Due to a transient hardware error condition the **CPQ Sram parity error** and **CPQ RDRAM double bit ECC error** syslog errors on MQCHIP raise a major CM alarm. [PR1276132](#)
- An accuracy issue occurs with three-color policers of both type single rate and two rate in which the policer rate and burst-size combination of the policer accuracy vary. This issue is present starting in Junos OS Release 11.4 on all platforms that use MX Series ASIC. [PR1307882](#)
- Traffic statistics may not match on PS after clearing interface statistics. [PR1328252](#)
- On all JunOS platforms, execution of Python scripts through enhanced automation does not work on verixec images. [PR1334425](#)
- Provides ability to configure host rsyslog from Junos guest. HOST side: The facility is one of the following keywords: auth, authpriv, cron,daemon, kern, lpr, mail, mark, news, security (same as auth), syslog, user, uucp and local0 through local7. The keyword security should not be used anymore and mark is only for internal use and therefore should not be used in applications. Anyway, you may want to specify and redirect these messages here. The facility specifies the subsystem that produced the message, that is, all mail programs log with the mail facility (LOG_MAIL) if they log using syslog. The priority is one of the following keywords, in ascending order: debug, info, notice, warning, warn (same as warning), err, error (same as err), crit, alert, emerg, panic (same as emerg). The keywords error, warn and panic are deprecated and should not be used anymore. The priority defines the severity of the message. Guest side: https://www.juniper.net/documentation/en_US/junos/topics/reference/general/syslog-facilities-severity-levels.html

remote : sync syslog server config from Junos to Linux & modify rsyslog.conf

set vmhost/app-engine syslog host <ip/ip6>any any

set vmhost/app-engine syslog host <ip/ip6>match xxx [PR1341549](#)

- In filter list (input-list/output-list) scenario, when the filters in the same filter list refer to a same nested filter, the FPC might crash continuously. The issue results in traffic loss during FPC crash and reboot. [PR1357531](#)
- In a Layer 3 VPN topology, traceroute to a remote PE device for a CE-facing network results in an ICMP TTL expired reply with a source address of only one of the many CE-facing networks. In Junos OS Releases 15.1R5, 16.1R3, and 16.2R1 and later releases, there is a kernel sysctl value,

icmp.traceroute_l3vpn. Setting this to 1 will change the behavior to select an address-based on the destination specified in the traceroute command. This PR adds the option to the configuration. [PR1358376](#)

- Sometimes OSPF flapping during unified ISSU from Junos OS Release 16.2R2 to Release 17.2R3. [PR1371879](#)
- One single port with Dual stack subscribers pppoe/dhcpv6 drop all the connections and no subscribers seen now. [PR1382288](#)
- In Junos Fusion provider edge setup, if CoS (class-of-service) is configured in the cascade port, when doing some CoS configurations changes, such as deactivating or activating CoS configurations on the cascade port, the traffic on this port would be silently dropped due to Packet Forwarding Engines mis programming for CoS queue of the cascade port. [PR1408159](#)
- On MX Series routers with MS-MPC cards, when FPC restart or routing-instance type is changed (for example, virtual-router to vrf), or RD is changed, traffic from a Group virtual private network (GVPN) tunnel to MPLS over UDP tunnel may fail to get decrypted on the MS-MPC, this will cause complete service loss. [PR1422242](#)
- On all Junos platforms with NSR enabled, the BGP session with hold-time 6 seconds or smaller flaps after the backup RE is pulled out ungracefully. [PR1428518](#)
- For the bridge-domains configured under an EVPN instance, the ARP suppression is enabled by default. This enables the EVPN to proxy the ARP, and reduces the flooding of ARP in the EVPN networks. Because of that, the storm-control is not taking effect to the ARP packets on the ports under such bridge-domain. [PR1438326](#)
- When executed over Junos CLI, Python op script is started as a separate process with the same user as the user which started the script. However, when the python op script is started from NETCONF session, the script started as a process from user "nobody". If the script is using PyEZ session to connect to the device and execute RPC commands, it will return the following error from Pyez: `ConnectError(host: None, msg: user "nobody" does not have access privileges.)`. This is fixed by executing with the python op script with the same user as the user from the NETCONF session which invoked op script. This means that the behavior from CLI and NETCONF sessions are the same. [PR1445917](#)
- In NTP with the boot-server scenario, when the router or switch boots, the NTP daemon will send a ntpdate request to poll the configured NTP boot-server to determine the local date and time. If the ntpdate is not be activated correctly while the device booting, the ntpdate might not work successfully. Then some cosmetic error messages of time synchronization might be seen, but there is no impact with time update since ntp daemon will update the time eventually. [PR1463622](#)
- On MX-VC setup with bridge-domains configured, if ae interface is used within bridge-domain, and if the ingress ae and egress ae interface host in different VC members, the Layer-2 traffic over ae sent from one member to another is getting corrupted. [PR1467764](#)

Routing Policy and Firewall Filters

- The rpd might crash during the policy configuration changes. [PR1357802](#)

- If a policy-option with only conditions **from route-distinguisher** and **then next-hop a.b.c.d** is applied to BGP, the next-hop for routes in the inet.0 might be set to this next-hop a.b.c.d, even though these routes do not carry any route-distinguisher value (l3vpn.inet.0 is unaffected). [PR1433615](#)

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](#)
- When only default routing-instance is present, the Junos command **show bgp summary** does not show the BGP ESTABLISH state. If the BGP state is not an ESTABLISHED state, then it shows the states as design (that is, Active, Idle, Connect). If there is a routing-instance configured (apart from master routing-instance inet.0), the BGP ESTABLISH state is showed properly. Issue happens for IPv4 BGP sessions only, on IPV6 we always see all the BGP states as default. [PR600308](#)
- 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](#)
- In rare cases, rpd might generate a core file with error **rt_notbest_sanity: Path selection failure on** The core is 'soft', which means there should be no impact to traffic or routing protocols. [PR946415](#)
- For single-hop eBGP session, upon interface down event, do not do GR helper logic. In problem state
Peer: 8.3.0.2 AS 100 Local: 8.3.0.1 AS 101Group: EBGp Routing-Instance: masterForwarding
routing-instance: masterType: External State: Active Flags: <>Last State: Idle Last Event: StartLast Error:
CeaseImport: [reject]Options: Holdtime: 90 Preference: 170 Local AS: 101 Local System AS: 0Number
of flaps: 2Last flap event: StopError: 'Cease' Sent: 1 Recv: 0NLRI we are holding stale routes for:
inet-unicastTime until stale routes are deleted or become long-lived stale: 00:01:54 >>>>>>>>Time
until end-of-rib is assumed for stale routes: 00:04:54Table inet.0RIB State: BGP restart is completeSend
state: not advertisingActive prefixes: 14Received prefixes: 21Accepted prefixes: 15Suppressed due to
damping: 0Stale prefixes: 21 [PR1129271](#)
- JTASK_SCHED_SLIP for rpd may be seen on doing restart routing or ospf protocol disable with scaled bgp routes in MX104 router. [PR1203979](#)
- Certain BGP traceoption flags (for example, "open", "update", and "keepalive") might result in (trace) logging of debugging messages that do not fall within the specified traceoption category, which results in some unwanted BGP debug messages being logged to the BGP traceoption file. [PR1252294](#)
- 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 0.0.0.01Type: P2P, Address: 10.0.60.93,**

Mask: 255.255.255.252, MTU: 9100, Cost: 1050Adj count: 1Hello: 10, Dead: 40, ReXmit: 2, Not StubAuth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTCProtection type: NoneTopology default (ID 0) -> Cost: 1050LDP sync state: in sync, for: 00:04:03, reason: IGP interface downconfig holdtime: infinity As per the current analysis, **IGP interface down** is observed as the reason because although LDP notified OSPF that LDP synchronization was achieved, OSPF was not able to take note of the LDP synchronization notification, because the OSPF neighbor was not up yet. The issue is under investigation. [PR1256434](#)

- When generating SNMP traps/notifications for BGP events from the jnxBgpM2 MIB, Junos was not properly emitting OBJECTS of type InetAddress with the expected length field. This will cause compliant SNMP tools to be able to parse the contents of those OBJECTS properly. In particular, the length field for the InetAddress OBJECT-TYPE was omitted. Using the set protocols bgp snmp-options emit-inet-address-length-in-oid command will cause these OBJECTS to be emitted in a compliant fashion. Given the length of time that this error has been in place, it was decided to leave the existing non-compliant behavior in place to avoid breaking tools that had accommodated the existing behavior as the default. [PR1265504](#)
- Two multicast tunnel (mt) interfaces are seen for each of the PIM neighbors after VPN-Tunnel-Source activation or deactivation. However, ideally, the same tunnel source should be used for both IPv4 and IPv6 address families, if both are using the same PIM tunnel. [PR1281481](#)
- This is in an internal change as Syslog usage is deprecated, however, there may be customer impact due to syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful/relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate [PR1284621](#)
- In rare cases RIP replication may fail as a result of performing NSR Routing Engine switch overs when the system is not NSR ready. [PR1310149](#)
- Rpd core file is observed at 0x094680ac in task_reconfigure_complete (ctx=0x9dfe940 <task_args>, seqnum=570) at ../../../../src/junos/lib/libtask/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 route target filtering (RTF) is configured for Virtual Private Network (VPN) routes and multiple BGP session flap, there is a slight chance that some of the peers might not receive the VPN routes after the flapped sessions come up. [PR1325481](#)
- In a large-scale OSPF network (for example, there are more than 500 devices in an area), OSPF remote loop free alternate (rLFA) default PQ node selection algorithm does not provide proper protection paths. [PR1335570](#)
- In JunOS 16.1 or higher, during BGP convergence, the input/output thread constructing the outgoing BGP PDU and manipulating the path attributes before hand-off the data to the socket. If this PDU length is zero, it will trigger an assertion and routing-protocol demon is restarting. [PR1351639](#)

- When **clear validation database** was issued back to back multiple times, we ended up with partial validation database (some validation entries were missing). This eventually recovered after up to 30 minutes (half of the Record Lifetime) when we did periodical full updates. [PR1326256](#)
- In a large-scale OSPF network (for example, there are more than 500 devices in an area), OSPF remote loop free alternate (rLFA) default PQ node selection algorithm does not provide proper protection paths. [PR1335570](#)
- When a BGP import policy changes IPv6 routes to have IPv4 nexthop, rpd might crash during route resolution. With the fix, changing route to have nexthop with different address family will not be allowed, if the route table does not have that resolution family configured. [PR1389557](#)
- In BGP scenario with multipath enabled, if applying import/export policy of IPv6 routes with IPv4 next hop to a BGP neighbor, the rpd might crash continuously. [PR1390428](#)
- If an import policy is applied to a BGP neighbor and the policy has indirect IPv4 next hop for IPv4 and IPv6 routes (IPv6 routes resolved over IPv4), when BGP unresolved route is withdrawn, rpd crash might be seen. [PR1391568](#)
- When 'as-path-group' is configured under BGP, if a configuration with a large scale as-path regex is committed, the route protocols flap might be seen. [PR1396344](#)
- When NSR (nonstop-routing) is enabled in local device and BGP GR (Graceful-Restart) is enabled in peer device, if the peer triggers a GR restart (it is usually caused by some failure in peer or the peer restarts rpd, etc), some BGP sessions might stuck in Idle state. The reason is that when the GR restart happens, the device is still doing the initial sync to the backup RE of the previous sessions, so some BGP sessions might stuck in Idle state because the router does not complete the process (the initial sync of the data set to the backup). [PR1412538](#)
- Change in route selection process. To select the better route between a non-BGP and BGP route, if you are at Step 7 of the route selection process (https://www.juniper.net/documentation/en_US/junos/topics/reference/general/routing-protocols-address-representation.html), then the BGP route is always the better one. [PR1415468](#)
- If IGMP v2 is used and proxy mode is used for igmp-snooping, multicast traffic might be dropped because by default proxy sends queries/reports in IGMP v3 version, until the device receives new IGMP v2 query or report. [PR1425621](#)
- In a scenario with IS-IS running single spf (shortest-path-first) for IPv4 and IPv6, that is, multi-topology is not enabled, when a new IS-IS link comes up, IFA (interface address) for IPv4 comes up quickly and the route is installed, but IFA for IPv6 is not up quickly because DAD (Duplicate Address Detection) is enabled by default. Therefore, after spf calculation, the next-hop list for IPv6 remains empty for about 11 seconds, and IS-IS ends up deleting the route. [PR1430581](#)
- By default, BGP multipath is for load balance with BGP neighbors in the same AS. For load balance with BGP neighbors in a different AS, the statement "multiple-as" is further needed. However if the statement "multiple-as" is only configured in some BGP groups but not in all BGP groups, the expected load balance will not work. [PR1430899](#)

- On all Junos platforms working as the source node (e.g. node S) where Per-Prefix Loop Free Alternate (PP-LFA) is configured for Open Shortest Path First (OSPF) routing protocol, if the destination prefix is learned from two originator nodes (e.g. node E and node F) with different costs, and both originator nodes E and F are directly connected with the source node S, PP-LFA might not work as expected in such scenario where the last hop needs to be protected on the penultimate hop. Due to this issue, an improper backup nexthop might be selected which couldn't handle node failure case and micro-loop might be seen. [PR1432615](#)
- In BGP segment routing traffic engineering (SRTE) scenario, process rpd might crash when knob "extended-nexthop-color" is added or removed from the BGP configuration. [PR1442952](#)
- When configuring an alternate incoming interface for a PIM RPF check using rpf-selection, you might find that additional groups outside the configured range switch to the alternate incoming interface. [PR1443056](#)
- On Junos platforms with BGP-PIC (protect core) and "add-path" enabled scenario, the rpd CPU utilization gets 100% due to incorrect path-selection. This issue may impact route update convergence or even cause routing protocols to flap. [PR1446861](#)
- If multipath is enabled, in some certain conditions, the rpd core might be seen while secondary route resolution. [PR1454951](#)
- With IS-IS configured and in a very rare case, memory corruption may occur, this may cause rpd crash continuously. [PR1455432](#)
- The rpd crash might be observed due to modification of router-id in OSPF NSSA with area-range configured. [PR1459080](#)
- On all Junos platforms running with Border Gateway Protocol (BGP), if both BGP multipath and BGP damping are configured, it might happen that, when the active route, for example r1, is withdrawn but it is not really deleted due to damping, then BGP might be unable to find its original gateway when the route r1 is relearned and becomes the best route again. It will lead to the rpd process crash. [PR1472671](#)

Services Applications

- It is not recommended to configure **ms- interface** when ams bundle in one-to-one mode has the same member interface. [PR1209660](#)
- Calling station was getting truncated after 64 bytes. As part of fix we are supporting calling-station till 128 bytes. [PR1462689](#)
- On MX platforms with MS-MPC/MS-MIC, after the IPsec VPN tunnel is up, if the NATted remote peer's IP address has been changed (e.g. NAT pool changed on peer), IKE SA might establish with an incorrect gateway, and kmd might crash frequently during this IKE SA IP migration. [PR1477181](#)

Subscriber Access Management

- Sometimes, when PPPoE subscribers login and logout from Junos OS 16.1 releases, the following messages are generated: `user@devcie> show log messages | match authd authd[5208]:`
`sdb_app_access_line_entry_read_by_uifl: uifl key 'demux0.xxxxxxxx': snapshot failed (-7) authd[5208]:`
`sdb_app_access_line_entry_read: uifl key 'demux0.xxxxxxxx': read failed` These messages indicate that authd daemon for subscriber authentication is attempting to read private data for an underlying interface which no longer exists (-7 = SDB_DATA_NOT_FOUND). These messages have no impact and can be safely ignored, where authd daemon is asking sdb for record that no longer exists. [PR1236211](#)
- authd re-use address too quickly before jdhcpd completely cleanup the old subscriber which flooding error log . The log such as `jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3` as it is already used by 1815. [PR1402653](#)
- Subscriber filtering for General Authentication Services traceoptions will report debug messages for other users. [PR1431614](#)
- The output of "test aaa ppp" is missing "<radius-server-data>" tag. [PR1444438](#)
- In Gx-Plus for Provisioning Subscribers scenario, when the PCRF (Policy and Charging Rules Function) server is unreachable or the diameter protocol is down, the subscriber login might fail to successfully establish a session or the subscribers might fail to bind a service policy by Gx-Plus after the PCRF Server connectivity is restored. [PR1449064](#)

User Interface and Configuration

- Junos Fusion `show chassis hardware satellite` command is not available on Junos OS Release 17.3.
`root@mx104> show chassis satellite detail` Satellite Alias: fusion FPC Slot: 101 Operational State: Online <...> Below, you can see no "show chassis hardware satellite" option: `root@MX104> show chassis hardware ?` Possible completions:<[Enter]> Execute this command clei-models Display CLEI barcode and model number for orderable FRUsdetail Include RAM and disk information in outputextensive Display ID EEPROM informationmodels Display serial number and model number for orderable FRUs| Pipe through a command. [PR1388252](#)
- Changing nested apply groups does not take affect. [PR1427962](#)

VPNs

- When switching from l2ckt to evpn vpws, deactivate and activate the instance. [PR1312043](#)
- JDI-RCT:Summit:Rpd core@ rtbit_reset, rte_tgtexport_rth [PR1379621](#)
- On all Junos platforms, if there are multiple interfaces configured under a single l2circuit/local-switching, and each of these interfaces has a description field configured under them, when l2circuit/local-switching

connections flapping continuously, memory usage increment might happen, eventually, it will result in rpd crash because of running out of memory. [PR1418870](#)

- In MVPN scenario with ingress replication selective provider tunnel used, if the knob "link-protection" is added/deleted from the LSP for MVPN, rpd crash might be seen. The reason is that when link-protection is deleted, the ingress tunnel is not deleted, and when link link-protection is added back, it tries to add same tunnel, hence the rpd asserts as same tunnel exists. Finally the rpd core might be seen. [PR1469028](#)

SEE ALSO

[New and Changed Features | 115](#)

[Changes in Behavior and Syntax | 149](#)

[Known Behavior | 164](#)

[Resolved Issues | 202](#)

[Documentation Updates | 275](#)

[Migration, Upgrade, and Downgrade Instructions | 276](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R3 | 203](#)
- [Resolved Issues: 17.4R2 | 234](#)
- [Resolved Issues: 17.4R1 | 260](#)

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: 17.4R3

Application Layer Gateways (ALGs)

- DNS requests with EDNS options might be dropped by DNS ALG. [PR1379433](#)

Authentication and Access Control

- MAC move might occur in DHCP security scenario. [PR1369785](#)
- Push-to-JIMS now supports push auth entry to all online jims servers. [PR1407371](#)

Class of Service (CoS)

- CoS is incorrectly applied on Packet Forwarding Engine, leading to egress traffic drop. [PR1329141](#)
- Configuring host-outbound-traffic under class-of-service may cause certain devices to crash. [PR1359767](#)
- The 802.1P rewrite may not work on inner VLAN. [PR1375189](#)
- FPC card might reboot when changing CoS mode from hierarchical-scheduler to per-unit-scheduler. [PR1387987](#)
- The cosd process might crash during committing configuration change through netconf. [PR1403147](#)
- Traffic drop occurs when deleting MPLS family or disabling interface which has non-default EXP rewrite-rules. [PR1408817](#)
- The host-inbound packets might be dropped if configuring host-outbound FC. [PR1428144](#)
- Firewall process crash might be seen with Multifield Classifier configuration. [PR1436894](#)

EVPN

- EVPN/VXLAN: MAC entry incorrectly programmed in Packet Forwarding Engine, leading to some traffic getting silently dropped or discarded. [PR1231402](#)
- L2ALD restarts when changing "protocols" related configuration. [PR1357911](#)
- EVPN: Last designated forwarder update time is not in sync with system time [PR1362997](#)
- Packet drop in EVPN stitching with IRB configured. [PR1363935](#)
- The EVPN implementation does not follow RFC-7432. [PR1367766](#)
- Small rpd memory leak when configuring EVPN. [PR1369705](#)
- EVPN active/active multihomed PE device occasionally prefers to route to a directly connected prefix using LSPs towards the multihomed peer instead of going directly out the IRB interface (which is up). [PR1376784](#)
- MAC addresses might disappear if the interface MTU of EVPN PE is changed. [PR1382966](#)
- The RA packets might be sent out without using the configured virtual gateway address. [PR1384574](#)
- EVPN-VXLAN: VTEP tunnel does not get deleted when EVPN peer goes down. [PR1390965](#)

- A few minutes of traffic loss might be observed during recovery from link failure [PR1396597](#)
- The BUM traffic might not be flooded in EVPN-MPLS scenario. [PR1397325](#)
- IPv6 link-local address for virtual-gateway address is marked as duplicate in EVPN. [PR1397925](#)
- EVPN Type 2 MAC+IP route is stuck when the route advertisement has two MPLS labels and withdrawal has one label. [PR1399726](#)
- ARP refresh functionality may fail in an EVPN scenario. [PR1399873](#)
- RPD core files are seen upon Routing Engine switchover with scaled EVPN configuration. [PR1401669](#)
- The rpd crashes due to memory corruption in EVPN. [PR1404351](#)
- EVPN database and bridge mac-table are out of sync due to the interface's flap [PR1404857](#)
- The rpd might crash on a leaf node when handling the withdrawal of remote or local MAC address in an EVPN-VXLAN scenario. [PR1405681](#)
- Local L2ALD proxy MAC+IP advertisements accidentally delete MAC+IP EVPN database state from remotely learned type 2 routes [PR1415277](#)
- The device may proxy the ARP Probe packets in an EVPN environment [PR1427109](#)
- Incorrect MAC count with "show evpn/bridge statistics". [PR1432293](#)
- Stale MAC addresses are present in the bridge mac-table in EVPN/MPLS scenario [PR1432702](#)
- Restarting I2-learning might cause some remote MAC addresses to move into forwarding 'dead' state [PR1441565](#)
- Traffic drop might be seen at EVPN Layer3 Gateway scenario [PR1442319](#)
- The bridge mac-table age timer does not expire for rbeb interfaces [PR1453203](#)
- ARP request/NS might be sent back to the local segment by DF router [PR1459830](#)

Forwarding and Sampling

- The kernel crash might be observed when there is a firewall filter modification. [PR1365265](#)
- In EVPN A-A scenario with MX acting as PE device, flood NHs to handle BUM traffic might not get created or miss certain branches when the configuration is performed in a particular sequence. [PR1377749](#)
- LTS Subscriber statistics reporting to RADIUS. [PR1383354](#)
- The lsi binding for the IPv6 neighbor is missing. [PR1388454](#)
- The filter counter is not written to the accounting file when accounting is enabled on the bridge firewall filter. [PR1392550](#)
- The srrd process might stuck at 100 percent usage when Jflow is used. [PR1393696](#)
- Junos OS: Firewall filter terms named "internal-1" and "internal-2" being ignored (CVE-2019-0036). [PR1394922](#)

- In some newer releases firewall filter action "decapsulate gre" cannot decapsulate ip-over-ip and ipv6-over-ip traffic. [PR1398888](#)
- The SRRD might crash when memory corruption occurs [PR1414568](#)
- The firewall filter configuration change might not be applied after software upgrade to Junos release 16.1R1 or later [PR1419438](#)
- EVPN enhancement for MAC flush mechanism in JUNOS [PR1421018](#)
- The device is in Amnesiac mode after ISSU with "mgd: error: configuration check-out failed" generate [PR1432664](#)
- High CPU utilization of l2ald is seen after replacing EVPN configuration. [PR1446568](#)
- [MX204] Input/Output counters of AE bundle/member links configured on non-default logical systems are not updated [PR1446762](#)
- The pfed might crash and not be able to come up on the PTX or TVP platforms [PR1452363](#)
- The l2ald process might observe memory leak on Junos platforms [PR1455034](#)

General Routing

- The command "show configuration | compare" shows unchanged configuration after deleting part of the configuration under the firewall section. [PR1042512](#)
- TACACS access does not work after upgrade. [PR1220671](#)
- Routing Engine-Packet Forwarding Engine out-of-sync errors might be seen in syslog [PR1232178](#)
- Mspmand core observed while sending TFO packets with high rate using tcpreplay tool [PR1253862](#)
- Error messages might be seen if the aggregated Ethernet interface hosted on MPC-3D-16XGE card flaps. [PR1279607](#)
- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/mspsmd. [PR1284654](#)
- The fault "PCI Device missing" alarm might be observed during upgrade [PR1301191](#)
- The rpd might crash by executing the command "show route extensive" during deletion of IS-IS configuration. [PR1301849](#)
- The mgd might crash when Ephemeral DB is used [PR1305424](#)
- The "LIBJNX_REPLICATE_RCP_ERROR" error message might be seen when the backup Routing Engine is not present in the dual Routing Engine scenario. [PR1305660](#)
- Incorrect packet statistics are reported in the ifHCInUcastPkts OID. [PR1306656](#)
- MACSEC causes dot1xd JTASK_SCHEDULE_SLIP or FPC disconnect [PR1322302](#)
- Potential heap leak at tcp_conn_create under UKERN TCP stack. [PR1326746](#)
- GRE interface might not come up after deactivating or activating the routing instances. [PR1327099](#)
- AI-script does not get automatic upgrade unless it is manually done after a Junos OS upgrade. [PR1337028](#)

- The rpd might crash when high-priority routes flap [PR1338895](#)
- fpc temperature mismatch for mpc8/9 on mx2k platform [PR1339077](#)
- Error message "RE does not have MAC map for mac type 7" might be seen on Summit platforms [PR1345637](#)
- Linux-based FPC should close chassis TCP control connection immediately when J-UKERN is crashing. [PR1347536](#)
- JSA10914 Junos OS: QFX10K and PTX Series: FPC process crashes after J-Flow processes a malformed packet (CVE-2019-0014) [PR1348417](#)
- The MPC might crash when the MIC is removed. [PR1350098](#)
- Large-scale users' login and logout may cause mgd memory leak. [PR1352504](#)
- Traffic loss might be seen on new master Routing Engine after the interface flaps followed by Routing Engine switchover in VRRP scenario [PR1353583](#)
- Trinity JNH memory leak when adding and removing unicast NH [PR1354225](#)
- Traffic drop might be seen after GRES if uRPF is configured. [PR1354285](#)
- Traffic might be blocked on MX with MS-MPC/MS-MIC [PR1358019](#)
- MPC/FPC might be unable to reply request messages to Routing Engine in a high subscriber scale scenario. [PR1358405](#)
- The "show chassis fpc" might show "Bad Voltage" for FPC powered off by configuration or CLI command after the command "show chassis environment fpc" is executed [PR1358874](#)
- FPC core files might be observed after GRES switchover [PR1361015](#)
- IP over VPLS traffic is affected by EXP rewrite rule on the core-facing MPLS interface. [PR1361429](#)
- MX Series router functioning as a BNG does not generate ESMC/SSM Quality Level failed SNMP trap. [PR1361430](#)
- The MS-MPC might reset continuously on MX Series platform. [PR1362271](#)
- JDI-RCT:M/Mx: Traffic loss of 1% is seen during GRES phase of ISSU from 17.3-20180527.0 to 17.3-20180527.0 [PR1362324](#)
- Executing "show route prefix proto ip detail" during route churn in a route scale scenario may lead to FPC crash. [PR1362578](#)
- The inline J-Flow sampling configuration might cause FPC crash on MX Series platforms. [PR1362887](#)
- Streaming telemetry data might not be received by one client when two clients subscribe to the same path with same frequency at the same time. [PR1363199](#)
- MX Series Virtual Chassis: Request to record VCCP heartbeat state change in syslog by default. [PR1363565](#)
- FPM board status is missing in SNMP MIB walk result. [PR1364246](#)

- The kernel might crash after repeatedly deactivating/activating interfaces/filter/class-of-services configurations due to accessing stale memory entry [PR1364477](#)
- Configuration commit might be delayed by 30 seconds. [PR1364621](#)
- The rpc command about interface unit might fail. [PR1365151](#)
- Extended traffic loss when performing ISSU/GRES with aggregated Ethernet interface configured with LACP. [PR1365316](#)
- MS-MPC/MS-PIC might be crash in NAT scenario. [PR1366259](#)
- JDI-RCT:M/Mx: Syslog errors seen " LOG : Err] Failed to allocate 2 jnh-dwords for encap-ptr(ether-da)!,LOG: Err] gen_encap_common: jnh-alloc failed! 8" [PR1366811](#)
- I2C error logs are seen when configuring wavelength on tunable SFP+. [PR1367224](#)
- The bbe-smgd process might crash during the authentication phase for L2BSA subscriber. [PR1367472](#)
- RTG interface status will be shown as incorrect status with show interface. [PR1368006](#)
- Junos OS:set system ports console insecure allows root password recovery on OAM volumes (CVE-2019-0035) [PR1368998](#)
- Error messages about mic_sfp_phy_mdio_sgmii_lnk_op might be seen after FPC is booting up on MX Series or EX9200 platform. [PR1369382](#)
- when you configure vrrp delegate-processing with tomcat enabled, the Packet Forwarding Engine dropped VRRP packets and count sw error. [PR1369503](#)
- SNMP mib walk causes KMD errors. [PR1369938](#)
- The rpd might crash after Routing Engine switchover is performed or the rpd is restarted if interface-based dynamic GRE tunnel is configured. [PR1370174](#)
- The bbe-smgd might crash when FPC is restarted. [PR1371926](#)
- The IPv6 routed packet might be transmitted through an interface whose VRRP state is in non-master. [PR1372163](#)
- The dot1xd might crash when dot1xd receives incorrect reply length from the authd. [PR1372421](#)
- Image installation on SD fails with the error " Unable to read reply from software add command to re1; error 1". [PR1372877](#)
- The Routing Engine might crash after non-GRES switchover [PR1373079](#)
- URL filtering might not work when the data interfaces move from one VRF to another. [PR1373582](#)
- LDP convergence delay might be seen after IGP metric change with the statement bgp-igp-both-ribs configured. [PR1373855](#)
- Cosmetic log "warning: [---] is protected, '---' cannot be deleted" is seen after commit using "configure private" in a configuration with "protect" flag present [PR1374244](#)

- FPC might be unable to work properly if one child interface is removed from an aggregated Ethernet bundle in dynamic VLAN subscriber scenario. [PR1374478](#)
- The bbe-smgd might crash continuously in centralized IGMP scenario. [PR1374530](#)
- The rpd process might crash when route flap and LSP flap occur with CBF enabled. [PR1374558](#)
- PCE initiated LSPs remain "Control status became local" after removing PCE configuration. [PR1374596](#)
- Few L2BSA subscribers might be stuck in init/terminating/terminated status after previous logout. [PR1375070](#)
- SFB and PDM/PSU related information is missing in jnxBoxAnatomy MIB on high end MX Series routers (MX2010/2020). [PR1375242](#)
- The bbe-smgd core files might be seen after doing GRES. [PR1376045](#)
- MS-MPC might have performance degradation under scaled fragmented packets. [PR1376060](#)
- PFE wedge may be observed if there are interfaces going to down state [PR1376366](#)
- Interface optic output power is not zero when the port has been disabled [PR1376574](#)
- MX Series BNG Node Slicing - DHCP Relay - AF Interfaces snoop and drop DHCP replies from DHCP server. [PR1377358](#)
- Packets might be dropped on data plane in the inline J-Flow scenario. [PR1377500](#)
- bbe-smgd cores trying to scale to 2000 BGP peers with 2000 dynamic PPPoE clients enabled w/routing-services. [PR1378542](#)
- After NAT64 router (with MS-MPC) translates an IPv6 fragment to IPv4 fragment, the router is not inserting the right value in the identification field of IPv4 Header. [PR1378818](#)
- ICMPv6 packets larger than 1024 might be dropped if "icmp-large-packet-check" is configured on IDS service. [PR1378852](#)
- ARP request packets might be sent out with 802.1Q VLAN tag [PR1379138](#)
- Traffic might get silently dropped or discarded when CoS configuration is changed on a PS interface. [PR1379530](#)
- Protocol adjacency might flap and FPC might reboot if jlock hog happens. [PR1379657](#)
- Remove the chassisd alarms for FPCs exceeding 90 percent of power budget and exceeding 100 percent of power budget [PR1380056](#)
- MQSS errors might cause FPC restart. [PR1380183](#)
- The Routing Engines might crash with various core files due to the deadlock issue on the SDB STS. [PR1380231](#)
- The rpd might crash on the new master Routing Engine when performing GRES. [PR1380298](#)
- IPv6 ping might fail for spine node in EVPN scenario. [PR1380590](#)

- The routes learned over an interface will be marked as "dead" next hop after changing the prefix length of an IPv6 address on that interface [PR1380600](#)
- Layer 3 VPN traffic might be dropped because one core-facing interface goes down. [PR1380783](#)
- Daemon dfwd might crash with DFWD_TRASHED_RED_ZONE log messages. [PR1380798](#)
- FPC might crash on PTX or QFX10000 after lo0 filter change [PR1380917](#)
- IRB interface does not turn down when master of Virtual Chassis is rebooted or halted. [PR1381272](#)
- Traffic is silently dropped or discarded when FPC is taken offline in an MC-LAG scenario. [PR1381446](#)
- Memory leak observed in MS-MPC card. [PR1381469](#)
- Constant memory leak might lead to FPC memory exhaustion. [PR1381527](#)
- The unicast traffic from IRB interface towards LSI might be dropped due to Packet Forwarding Engine mismatching at egress processing. [PR1381580](#)
- SSD lifetime might be shortened in OVSDDB environment. [PR1381888](#)
- All type of subscribers might not be able to log in after double GRES operations. [PR1382050](#)
- The MPC6E might crash while fetching PMC device states. [PR1382182](#)
- Flows are getting exported before the expiry of the configured active timeout value. [PR1382531](#)
- The chassisd might crash due to HW-DB errors on TVP based platforms [PR1383246](#)
- Domain name is not reported as part of the LLDP system name in "show lldp neighbor" command. [PR1383295](#)
- The configuration performed through NETCONF session might fail. [PR1383567](#)
- Adjusting mac-table-size configuration might cause l2ald crash. [PR1383665](#)
- The VC could not come up after upgrading to QFX5E platforms (TVP-based platforms for QFX5100 or QFX5200 switches) [PR1383876](#)
- The kmd crashes with core file after bringing up IPSec connection. [PR1384205](#)
- CoS attachment might be mistakenly removed for DHCPv4 stack when DHCPv6 stack fails to be brought up for a single session dual stack subscriber. [PR1384289](#)
- Missing statement "interface-description" for static subscribers. [PR1384421](#)
- MBFD flaps because clksync congests the scheduler for 100ms. [PR1384473](#)
- Multiple bbe-smgd core files with reference to bbe_mcast_vbf_dist_policy_service_encoder(). [PR1384491](#)
- BFD sessions might flap consistently [PR1384601](#)
- SNMP MIB walk returns unexpected data. [PR1384807](#)
- ARP and ethernet-table entry is pointing to an aggregated Ethernet interface whose state is down if MTU is changed. [PR1385199](#)

- On vMX systems, when you configure large number of interfaces, the vFPC CPU utilization might go very high periodically because of interface statistics collection running repeatedly. [PR1385853](#)
- The device with more than five IP addresses configured in the DHCP server group goes into amnesiac mode after reboot [PR1385902](#)
- IPSec VPN traffic might fail when passing through MS-MPC of MX Series routers with CGNAT enabled. [PR1386011](#)
- ALB-ECMP may not work as expected for LDP tunnels [PR1386061](#)
- The rpd process might end up with stuck krt queue entries in a VRF scenario [PR1386475](#)
- In subscriber management environment, DHCP subscriber might get stuck in terminated state. [PR1386662](#)
- In case an LSP is locally configured without an explicit path, ERO object remains empty in the PCRpt generated by PCC. [PR1386935](#)
- Uninitialized EDMEM[0x400094] Read (0x6db6db6d6db6db6d) logs are seen with sampling applied to a subscriber with routing-service applied. [PR1386948](#)
- On MX2000 platforms, backup CB's chassis environment status keeps 'Testing' after backup CB becomes online by removal/insert operation [PR1387130](#)
- The pccd might crash when changing delegation-priority [PR1387419](#)
- The bbe-smgd process might crash when two subscribers log in with the same framed-route prefix and preference values. [PR1387690](#)
- Output of the "show class-of-service interface" command incorrectly shows adjusting application as PPPoE IA tags for DHCP subscribers. [PR1387712](#)
- Some SFBs might go down when one of the PSMs in the chassis generates a bad output voltage which is out-of-range [PR1387737](#)
- IPsec IKE keys are not cleared when delete/clear notification is received [PR1388290](#)
- The bbe-smgd might not respond to the NS message for the SLAAC client on dynamic VLAN. [PR1388595](#)
- Fabric drops might be seen if using a newer generation of MPC with SFB2 [PR1388780](#)
- Incorrect value for flow packets/octetets fields might be seen in inline J-Flow scenario. [PR1389145](#)
- IGMP group threshold exceed log message prints a wrong demux logical interface. [PR1389457](#)
- BFD flaps are seen on PTX or QFX10K platforms with inline BFD [PR1389569](#)
- MX204 - Excluding "speed" CLI option under the interface level [PR1389918](#)
- The jnxFruState might show incorrect PIC state after replacing an MPC with another MPC having less PICs [PR1390016](#)
- Class of service adjustment-control-profile configuration for application DHCP tags does not get applied. [PR1390101](#)
- Traffic destined to VRRP VIP gets dropped as filter is not updated to related logical interface [PR1390367](#)

- The vmcore might be seen when routing changes are made on the peer spine in an EVPN-VXLAN scenario. [PR1390573](#)
- The statement routing-engine-power-off-button-disable does not work on MX204 and MX10003. [PR1391548](#)
- The bbe-smgd process might crash after committing configuration changes. [PR1391562](#)
- The bbe-smgd process might crash in a corner case if family inet6 is used in dynamic profile [PR1391845](#)
- The Packet Forwarding Engine might not respond with ICMP time exceeded error when packet arrives from the subscriber. [PR1391932](#)
- On ACX-Series platforms the 'forwarding-option dhcp-relay forward-only' knob stops working and the DHCP packets are dropped. [PR1392261](#)
- third-generation FPC reboot loop because of having internal intf issues [PR1393643](#)
- JUNOS enhancement configuration knob to modify mcontrol watchdog timeout [PR1393716](#)
- IPV6 Next-Hop programming issue might be observed on QFX10K/PTX1K/PTX10K devices [PR1393937](#)
- WITHDRAWN: Junos OS: gRPC hardcoded credentials may allow unauthorized access to systems with Junos Network Agent installed (REJECTED) [PR1394927](#)
- The l2ald process might crash when doing "commit check" for some specific configurations [PR1395368](#)
- The minor alarm of "Bottom Fan Tray Pred Fail" might be wrongly raised when the fan speed is at high speed on MX960 [PR1395539](#)
- The subscriber bindings might not be successful on QFX/EX platforms [PR1396470](#)
- Adding IRB to bridge domain with PS interface causes kernel crash. [PR1396772](#)
- The MS-MPC might core when mspmand receives a non-syn packet of TCP [PR1396785](#)
- Subscriber flapping might cause SMID resident memory leak. [PR1396886](#)
- The PPPoE subscribers are unable to reconnect after FPC reboot [PR1397628](#)
- Confirmation message is missing when issuing "request vmhost reboot re* " [PR1397912](#)
- The CLI command "show system firmware" gets hidden on MX platforms [PR1398022](#)
- On vMX platform, kernel core files are generated when the kernel state (ifstates) exceptions occur. [PR1398320](#)
- Junos OS: NFX150 Series, QFX10K Series, EX9200 Series, MX Series, PTX Series: Path traversal vulnerability in NFX150 and NG-RE leads to information disclosure (CVE-2019-0074) [PR1398333](#)
- IPSEC tunnel can not be established because the tunnel SA and rule are not installed in the PIC. [PR1398849](#)
- The bbe-smgd process might crash when executing "show pppoe logout" [PR1398873](#)
- CPU hog may be observed on PTX/QFX10000 Series platform [PR1399369](#)

- The unexpected alarm might be shown on NG-RE [PR1399654](#)
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- The authd might stop when issuing show network-access requests pending command during the authd restart [PR1401249](#)
- Command "show | compare" output on global group changes lose the diff context after a rollback or 'load update' is performed [PR1401505](#)
- The subscriber route installation failed because some interfaces states are not properly installed. [PR1401506](#)
- The TCP connection for external or internal might be dropped due to a kernel issue [PR1401507](#)
- FPC core files due to a corner case scenario (race condition between RPF, IP flow). [PR1401808](#)
- The na-grpcd log file is not rotated and keeps growing until Routing Engine is out of disk space. [PR1401817](#)
- The mspmand process might crash with lots of error logs seen in high scaled MX platforms with MS-MPC/MS-MIC [PR1402260](#)
- The MPC might crash due to CPU overuse by dfw thread. [PR1402345](#)
- Some error logs might be seen on FPC when reading is attempted from Uninitialized memory location. [PR1402484](#)
- FPC might crash after offline/online MIC-3D-16CHE1-T1-CE-H. [PR1402563](#)
- DHCP subscriber cannot reconnect over dynamic VLAN demux interfaces due to RPF check failure [PR1402674](#)
- Host outbound traffic might be dropped on MPC7, MPC8, and MPC9. [PR1402834](#)
- Smg-service could become unresponsive when doing some GRE-related CLI operations. [PR1403480](#)
- The time synchronization through PTPoE might not work when Enhanced Subscriber Management is enabled on MX Series routers. [PR1404002](#)
- Continuous kernel crashes might be observed in backup Routing Engine or VC-BM. [PR1404038](#)
- With MS-MPC and MS-MIC service cards, SYSLOG messages for port block interim might show 0.0.0.0 for the private IP address and PBA release messages might show the NAT'd IP address as the private IP address. [PR1404089](#)
- The FPC might crash in a CoS scenario [PR1404325](#)
- repd continue core on VC-Bm when there are too many IPv6 address on one session (hit PR1384889) [PR1404358](#)
- Incorrect display of assigned prefixes to a subscriber in the output of 'show interface < dynamic demux interface>' [PR1404369](#)

- In a very rare situation Router can crash with VMCore when there is a IFL deletion [PR1404507](#)
- Configuration load override or load replace resets ANCP neighbors. [PR1405318](#)
- MX Series: Denial of Service vulnerability in MS-PIC component on MS-MIC or MS-MPC (CVE-2019-0065) [PR1405423](#)
- FPC crash might be seen when adding or deleting a leg to an AE bundle or FPC restarts in subscriber scenario [PR1405876](#)
- NAT64 translation issues of ICMPv6 Packet Too Big message with MS-MPC/MS-PIC [PR1405882](#)
- The FPC crash might be observed in MS-MPC HA environment [PR1405917](#)
- Fabric performance drop on MPC7, MPC8, or MPC9E and SFB2 based MX2000 platform. [PR1406030](#)
- The rpd might crash due to a race condition with the combination of community actions done at both a BGP import policy and a forwarding-table policy [PR1406357](#)
- Traffic impact might be seen if auto-bandwidth is configured for RSVP LSPs [PR1406822](#)
- FPC might crash during the subscriber related stress tests. [PR1407285](#)
- Layer 2 VPN might flap repeatedly after the link up between PE and CE devices. [PR1407345](#)
- The rpd might crash when a commit check is executed on LDP trace options filtering [PR1407367](#)
- The PFE might get disabled unexpectedly due to a auto correctable non-fatal hardware error on PTX or QFX10002/QFX10008/QFX10016 [PR1408012](#)
- Traffic forwarding failed when crossing VCF members [PR1408058](#)
- The ToS/DSCP and TTL fields might not be copied into the outer IP header in group VPN scenario. [PR1408168](#)
- The alarm 'Mismatch in total memory detected' is observed after issuing "request reboot vmhost routing-engine both". [PR1408480](#)
- The MPC line cards might crash when performing ISSU to Junos OS Release 19.1R1 or later [PR1408558](#)
- Syslog flooded with "Limit check for pppoe subscriber failed" messages [PR1408833](#)
- The misconfiguration of dynamic profile might cause the login issues of the subsequent subscribers [PR1409398](#)
- MX-MPC2-3D-EQ and MPC-3D-16XGE-SFPP will now show "Exhaust A" temperature, rather than Intake temperature. [PR1409406](#)
- Indirect-next-hop pointing to unknown unilist stuck with weight 65535 may occur after a link flap [PR1409632](#)
- The non-existent subscribers might appear in the 'show system resource-monitor subscribers-limit chassis extensive' output. [PR1409767](#)
- FPC might crash during next-hop change when using MPLS inline J-Flow [PR1409807](#)

- On MX10003 platform, after removing the FPC from a slot, when a new FPC is plugged in, chassis was showing old serial for this new FPC. [PR1409930](#)
- ACX drops DNS responses which contain an underscore [PR1410062](#)
- When using SFP+, the interface optic output might be non-zero even when the interface has been disabled. [PR1410465](#)
- Packets might be dropped if the traffic is forwarded through an LT interface [PR1410970](#)
- Kernel replication failure might be seen if an IPv6 route next-hop points to an ether-over-atm-llc ATM interface. [PR1411376](#)
- A steady increase of the PFE heap memory utilization may happen when PPPoE subscribers are flapping [PR1411389](#)
- Parity error might cause FPC alarm [PR1411610](#)
- JTASK_SCHED_SLIP error might be observed on VRR platform during NTP synchronization [PR1411679](#)
- GRE over GRE might not work for host-generated traffic [PR1411874](#)
- MX10003: The rpd crash with switchover-on-routing-crash does not trigger Routing Engine switchover and the rpd on master Routing Engine goes into STOP state. [PR1412322](#)
- Junos PCC might reject PCUpdate/PCCreate message if the metric type is other than type 2 [PR1412659](#)
- PPPoE subscribers might not be able to log in after ISSU. [PR1413004](#)
- The rpd memory leak might be seen due to an incorrect processing of a transient event. [PR1413224](#)
- During ISSU or merge virtual-chassis member back to the VC, CoS GENCFG writes failures may be observed [PR1413297](#)
- JFLOW: To Reduce max flow table Size when using Flex-flow-sizing [PR1413513](#)
- The support of inet6 filter attribute for ATM interface is broken in Junos OS Release 17.2R1 onwards. [PR1413663](#)
- The services load balance might not be effective for AMS if the hash key under the forwarding-options hierarchy is configured [PR1414109](#)
- FPC crash might be observed if it reaches heap utilization limit. [PR1414145](#)
- The PTX1000/PTX10002/QFX10002 may stop forwarding packets after the "chassis-control" process restarts [PR1414434](#)
- NPC might not apply configured resource-monitor thresholds after NPC restart [PR1414650](#)
- Firewall filters are not getting programmed into Packet Forwarding Engine. [PR1414706](#)
- The user might not enter configure mode due to mgd is in lockf status [PR1415042](#)
- ICMP MTU exceeded error generated from Packet Forwarding Engine does not reach the expected source. [PR1415130](#)
- The IRB interface might flap after committing configuration change on any interface [PR1415284](#)

- The bbe-smgd process might have memory leak while running "show system subscriber-management route route-type <> routing-instance <>" [PR1415922](#)
- Some IPsec tunnels might fail to pass traffic after GRES on MX platform [PR1417170](#)
- The ECMP fast reroute protection feature might not work on MX5, MX10, MX40, MX80, and MX104. [PR1417186](#)
- The IPv6 neighbor might become unreachable after the primary link goes down in VPLS multihoming scenario [PR1417209](#)
- An IPv4 packet with a zero checksum might not be translated to IPv6 packet properly under NAT64 scenario. [PR1417215](#)
- Some subscribers might be offline when doing GRES or daemon restart [PR1417574](#)
- No message indicates the failure when subscribers request NAT port failure happens under CGNAT with MS-MPC [PR1418128](#)
- The rpd core might be seen after changing the OSPF/OSPF3 interface cost [PR1418152](#)
- there is no SNMP Trap message generated for jnxHardDiskMissing/jnxHardDiskFailed on Summit MX [PR1418461](#)
- MX-GX+ Services are not synced up to the BACKUP RE with GRES/NSR enabled [PR1418594](#)
- Adding two or more ps interface may cause traffic drop in l2circuit scenario [PR1418610](#)
- lsp-cleanup-timer is not being honored when lsp-cleanup-timer is configured to be greater than 2147483647 [PR1418937](#)
- The PPPoE negotiation of subscriber connection might fail when 65535 is assigned as session ID. [PR1418960](#)
- RX alarms are not set as according to the threshold value configured for the DCO Tunable Optics. [PR1419204](#)
- A PPP session under negotiation might be terminated if another PPPoE client bears the same session ID. [PR1419500](#)
- CPU usage on Service PIC might spike while forming an IPsec tunnel under DEP/NAT-T scenario. [PR1419541](#)
- A new tunnel could not be established after changing the NAT mapping IP address until the IPEC SA Clear command is run [PR1419542](#)
- rtsock_peer_unconsumed_obj_free_int: unable to remove node from list logged extensively [PR1419647](#)
- The IPsec tunnel might get down when the Junos platforms and the peer both act as the initiator and try to bring an IPsec tunnel up at the same time [PR1420293](#)
- The bbe-smgd process might crash and might not recover in a rare scenario [PR1420376](#)
- MX: PTP phase aligned but TE/cTE is not good. [PR1420809](#)

- An interface may go to downstate on QFX10000/PTX10000 platform [PR1421075](#)
- MX LNS might fail to forward the traffic on the subscriber access route. [PR1421314](#)
- Failed to reload keyadmin database for /var/etc/keyadmin.conf. [PR1421539](#)
- MX Series Virtual Chassis: VCP port reports MTU value 9152 in the ICMP MTU exceeded message while the VCP port MTU is set to 9148. [PR1421629](#)
- After control plane event few ipsec tunnels failed to send traffic through the tunnel [PR1421843](#)
- The changed value of "remote-gateway" does not take effect when the router acts as an initiator of IPsec-VPN tunnel. [PR1421977](#)
- The CoS IEEE-802.1 classifier might not get applied when it is configured with service activation on underlying interface. [PR1422542](#)
- While committing huge configuration, customer is seeing the error "error: mustd trace init failed" [PR1423229](#)
- "set forwarding-options enhanced-hash-key symmetric" is not effective on MX10003 [PR1423288](#)
- Traffic is dropped after FPC reboot with aggregated Ethernet member links deactivated by remote device [PR1423707](#)
- The bbe-smgd process might crash after executing the command "show system subscriber-management route prefix <>" [PR1424054](#)
- The system does not reboot or halt as configuration when encountering the disk error [PR1424187](#)
- Interface with FEC disabled might flap after Routing Engine mastership switchover. [PR1425211](#)
- Soft GRE tunnel route lost after reboot/GRES or upgrade in WAG scenario. [PR1425237](#)
- The mspmand process might crash and restart with a mspmand core file created after doing a commit change to deactivate and activate service-set [PR1425405](#)
- All interfaces creation failed after NSSU [PR1425716](#)
- MPC reboot or RE mastership switchover might occur on MX204/MX10003 [PR1426120](#)
- IFL Targeting: 18k phantom distributed interfaces are displayed for AE interface with the targeted distribution enabled on it, when there are no active subscribers [PR1426157](#)
- Some CFM and BFD sessions might flap while collecting MPLS statistics [PR1426727](#)
- Traffic loss might be seen when multiple IPsec tunnels are established with the remote peer [PR1426975](#)
- Traffic might not flow through MACsec interface even after an unsupported cipher-suite is removed. [PR1427294](#)
- ENTITY MIB has incorrect containedIn values for some fixed MPCs with builtin PICs [PR1427305](#)
- Rebooting or halting Virtual Chassis member might cause traffic on RTG link to be down for about 30 seconds. [PR1427500](#)

- The subscriber IP route might get stuck in bbe-smgd if the subscriber IP address is the same as local IP address. [PR1428428](#)
- Incorrect IGMP interface counter for dynamic PPP interfaces. [PR1429018](#)
- L2TP subscriber and MPLS pseudowire subscriber volume accounting statistics value remains unchanged post ISSU. [PR1429692](#)
- The AE interface does not come up after rebooting the FPC/device though the physical member link is up [PR1429917](#)
- Configuration is prevented from being applied on MX in subscriber scenario [PR1430360](#)
- Inline LSQ might not work when it is configured on the same FPC where MIC-3D-16CHE1-T1 is slotted [PR1431069](#)
- Error might be observed when using a script to load configuration. [PR1431198](#)
- The l2cpd process might crash and generate a core dump when interfaces are flapping [PR1431355](#)
- During the stress tests, bbe-smgd process might crash on backup Routing Engine when performing GRES. [PR1431455](#)
- The bbe-smgd might crash if subscribers are trying to log in or log out and a configuration commit activity happens at the same time. [PR1431459](#)
- Subscribers coming from new IFDs might not login in due to 512 entries limit in the subscriber-limit table [PR1431566](#)
- Allow installation of three identical framed-routes in the same routing-instance. [PR1431891](#)
- MX10003 - PEM not present alarm raised when minimum required PEM exists in the system. [PR1431926](#)
- Traffic might be sent on the standby link of aggregated Ethernet bundle and be lost with LACP fast-failover enabled. [PR1432449](#)
- Change to in-use parameterized filter prefix-list could result in bbe-smgd core on backup Routing Engine. [PR1432655](#)
- Traffic will be dropped if 'sa-multicast' is in the configuration [PR1433306](#)
- RSI and RSI brief should not include "show route forwarding-table" when tomcat enabled. [PR1433440](#)
- MX URLF: URL case sensitivity support [PR1434004](#)
- The repd process might crash after booting first time with a newly installed Junos release [PR1434363](#)
- PFE memory leak might be seen if MLPPP links are flapped [PR1434980](#)
- MPC7/8/9/MX10003 MPC/EX9200-12QS/EX9200-40XS line card might crash in a scaling setup [PR1435744](#)
- The mc-ae interface may get stuck in waiting state in dual mc-ae scenario [PR1435874](#)
- The static PPP/PPPoE subscribers stuck in "init" state permanently and error message "Failed to create client session, err=SDB data corrupted" might be seen [PR1436350](#)

- LNS router might send the router-advertisement packet with NULL source link-layer option field [PR1437847](#)
- Subscriber flows might not be synchronized between AE members on MX-VC platforms [PR1438621](#)
- FPC on Virtual Chassis backup router might reboot in MX Series Virtual Chassis scenario. [PR1439170](#)
- The "vlan all interface all" combination not working as expected under VSTP [PR1439583](#)
- The bbe-smgd core dumps is seen after restarted [PR1439905](#)
- CoS related errors are seen and subscribers could not get service [PR1440381](#)
- DHCP offer packets towards IRB over LT interface getting dropped in DHCP relay environment [PR1440696](#)
- The layer2 dynamic VLAN might be missed when an interface is added or removed for an AE interface [PR1440872](#)
- For a route received via EBGp the AIGP value may not be considered as expected [PR1441438](#)
- The rpd may crash or consume 100% of CPU after flapping routes [PR1441550](#)
- On PTX/QFX AE outgoing traffic might be dropped after changes are made to AE [PR1441772](#)
- Egress stream flush failure and traffic blackhole might occur [PR1441816](#)
- The packets originating from the IRB interface might be dropped in VPLS scenario [PR1442121](#)
- The chassisd is unable to power off a faulty FPC after RE switchover which leading to chassisd restart loop [PR1442138](#)
- In "enhanced-ip" or "enhanced-ethernet" mode with DCU (destination-class-usage) accounting enabled, MS-DPC may drop all traffic that should egress via ae interface [PR1442527](#)
- The kmd process might crash and restart with a kmd core file created if IP of NAT mapping address for IPsec-VPN remote peer is changed. [PR1444183](#)
- Inline-keepalive might stop working for LNS subscribers if the knob "routing-services" is enabled [PR1444696](#)
- access route might stuck in bbe-smgd and RPD not cleared [PR1445155](#)
- Detached LACP member link gets LACP state as enabled in Packet Forwarding Engine during switchover because of device reboot [PR1445428](#)
- The mspmand process might crash if URL filtering is configured and one blacklisted domain name is a sub-string of another blacklisted domain name in URL filter database file [PR1445751](#)
- NAT service-set in certain scale might fail to get programmed [PR1446931](#)
- The J-Flow version 5 stops working after changing input rate value. [PR1446996](#)
- Interface attributes might cause high CPU usage of dcd [PR1448858](#)
- Interfaces might flap forever after deleting the interface disable configuration [PR1450263](#)
- VLAN config change with l2ald restart might cause Kernel sync issues and impact forwarding [PR1450832](#)

- IPSec[SNMP]: Snmp query for IPSec Decrypted/Encrypted packets does not fetch right values; observing KMD_SNMP_FATAL_ERROR [PR1451324](#)
- [MX] Error dropped packets seen on MQ/XM based MPC cards though there is no traffic flowing through the system [PR1451958](#)
- There is high temperature from "show chassis environment" output after MPC4E insert to slot 5 [PR1456457](#)
- The subscriber routes are not cleared from backup RE when session is aborted [PR1458369](#)
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used [PR1458559](#)
- The traffic might be stuck on MS-MPC/MS-MIC with sessions receiving huge number of affinity packets [PR1459306](#)
- The PPTP doesn't work with destination NAT [PR1460027](#)
- In EVPN scenario memory Leak might be observed when proxy-macip-advertisement is configured [PR1461677](#)
- The subscribers might not pass traffic after doing some changes to the dynamic-profiles filter [PR1463420](#)
- SMGD generated core files after essmd restart with reference to mmf_ensure_mapped (mmf=0xe8f0200, offset=4294967295, len=108) at ../src/junos/lib/libmmf/mmf.c:1972. [PR1372223](#)
- The subinfo process might crash with core file if flapping all the subscribers [PR1379482](#)
- Error messages "shmlog: argcnt 309 not enough memory" are generated every hour. [PR1384371](#)
- The bbe-smgd process generates repeated core files and stops running as a result of long-term session database shared memory corruption. [PR1388867](#)
- Delay in CLI output with second or more "show subscriber <> extensive" queries when first session is sitting at -(more)- prompt displaying "show subscribers extensive". [PR1390762](#)
- The framed route beyond the first might not be installed in a DHCP subscriber management environment. [PR1401148](#)
- Traffic loss seen in IGMP subscribers after GRES. [PR1402342](#)
- Change the default parameters for resource-monitor rtt-parameters [PR1407021](#)
- bbemg_smgd_lock_cli_instance_db should not log as error messages [PR1421589](#)
- A stuck lock in shared memory might prevent subscribers from logging in again after daemon crash [PR1424607](#)
- The jdhcpd might consume 100% CPU and crash if dhcp-security is configured [PR1425206](#)
- Show subscriber extensive incorrectly displays DNS (domain-name-server) address provided to DHCP clients [PR1457949](#)

High Availability (HA) and Resiliency

- FPCs rebooted although the MXVC ISSU output looks like successful [PR1376774](#)
- If FPGA on the new master CB has a specific hardware failure, the chassis might keep crashing after GRES switchover. [PR1393884](#)

Infrastructure

- Cleanup at thread exit in FreeBSD kernel causing memory leaks. [PR1328273](#)
- The command **show system virtual-memory | display xml validate** displays errors. [PR1356423](#)
- The error **jlaunchd: disk-monitoring is thrashing, not restarted** might be seen. [PR1380032](#)
- The alarm might be seen if the PEM's serial number starts with "1F1" [PR1398128](#)
- SNMP OID IFOutDiscards not updated when drops increasing. [PR1411303](#)
- The traffic to the NLB server may not be forwarded if the NLB cluster works on multicast mode. [PR1411549](#)

Interfaces and Chassis

- Subscribers might fail to access the device after deleting the needless aggregated Ethernet configuration. [PR1322678](#)
- Momentary traffic loss might happen when a GRES is performed. [PR1336455](#)
- Native-vlan-id support on ps-interface [PR1352933](#)
- Error messages like "ifname [ds-5/0/2:4:1] is chan ci candidate" are seen during a commit operation [PR1363536](#)
- In case of MPLS ,DMR packets are sent with different mpls exp bits if MX receives CFM DMM packets with varying exp values on MPLS header [PR1365709](#)
- In rare case, there might be L2TP subscribers stuck in terminated state. [PR1368650](#)
- Constant dcpfe process crash might be seen if using an unsupported GRE interface configuration [PR1369757](#)
- Unified ISSU could be aborted at "Timed out Waiting for protocol backup chassis master switch to complete" with MX Series Virtual Chassis configuration. [PR1371297](#)
- JDI-RCT: QFX5200 MCLAG: parse_remove_ifl_from_routing_inst() ERROR : No route inst on et-0/0/16.16386, errors seen after restart l2cpd daemon [PR1373927](#)
- The dcd process might go down when 'vlan-id none' is configured for the interface. [PR1374933](#)
- "PE Chip:pe0[0]: IPW: oversize_drop error" causes Major error on FPC [PR1375030](#)
- PPP Chap Challenge-Length Is Not Initialized With Default Value [PR1375145](#)
- Race condition during Routing Engine mastership transition might cause improper deletion/recreation of logical interface em0.0 Interface family address. [PR1376216](#)

- Duplicate IP cannot be configured on both SONET (so-) interface and other interfaces [PR1377690](#)
- Some error logs (Tx unknown LCP packet) might be reported by the bbe-smgd daemon on MX-Series platforms [PR1378912](#)
- The pfe_disable action does not disable the logical tunnel interfaces belonging to the affected Packet Forwarding Engine [PR1380784](#)
- Higher level OAM CFM between CE might not work in VPLS scenario [PR1380799](#)
- The dcd restarted unexpectedly after committing a configuration with static demux interface stacking over PS interface. [PR1382857](#)
- The jpppd process might crash if the EPD value contains a format specifier [PR1384137](#)
- Changing the value of mac-table-size to default might lead all FPC to reboot. [PR1386768](#)
- DCD core files can be seen after FPC restart if channelized interfaces are configured. [PR1387962](#)
- All DPCs might crash while adding or deleting a logical interface from the aggregated Ethernet bundle. [PR1389206](#)
- The interface-control process thrashes and dcd does not restart after adding an invalid demux interface to the configuration [PR1389461](#)
- Interim accounting updates might not be sent for subscribers after Junos OS selective update [PR1391011](#)
- The dcd memory leak might be seen when committing configuration change on static route tag [PR1391323](#)
- The dcd crash might be seen after deleting the subinterface from VPLS routing-instance and mesh-group. [PR1395620](#)
- "MIC Error code: 0x1b0002" alarm might not be cleared for MIC on MPC6 when the voltage has returned to normal [PR1398301](#)
- The backup Routing Engine might get stuck in amnesiac mode after reboot. [PR1398445](#)
- The transportd might consume 100% CPU for a prolonged period [PR1398967](#)
- All dcd operations might be blocked if profile-db is corrupt [PR1399184](#)
- [Cordoba] incorrect Lane chromatic dispersion values and false positive RX power high alarm [PR1400190](#)
- Certain otn-options cause interface flapping during commit. [PR1402122](#)
- The configuration "targeted-broadcast" does not work on IRB interface [PR1404442](#)
- The subscriber might not be able to access the device due to the conflicted assigned address. [PR1405055](#)
- The cfmd might fail to start after it is restarted [PR1406165](#)
- Inline periodic packet management (PPM) adjacency (rx) session might be programmed with the incorrect packet template. [PR1417707](#)
- The monitor ethernet loss-measurement command returns invalid ETH-LM request for unsupported outgoing logical interface. [PR1420514](#)

- Invalid speed value on an interface might cause other interface configuration loss [PR1421857](#)
- The syslog message "/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex" upon LFM related configuration commit on aggregated Ethernet interfaces [PR1423586](#)
- The logical interfaces in EVPN routing instances might flap after committing configurations. [PR1425339](#)
- The configuration statement "flexible-queuing-mode" is not working on FPCs of Virtual Chassis member 1. [PR1425414](#)
- CFM message is flooding. [PR1427868](#)
- The vrrpd process might crash after deleting VRRP sessions for several times [PR1429906](#)
- The NCP session might be brought down after IPCP Configure-Reject is sent [PR1431038](#)
- Mixed link-speed AE bundle could not add new sub-interface successfully [PR1437929](#)
- Targeted-distribution for static demux interface over aggregate Ethernet interface does not take correct LACP link status into consideration when choosing primary and backup links. [PR1439257](#)
- The cfmd process might crash after a restart on Junos 17.1R1 and above [PR1443353](#)
- Enhancement of add/delete a single vlan in vlan-id-list under interface family bridge [PR1443536](#)
- The ifinfo daemon might crash on the execution of "show interface extensive" command [PR1448090](#)
- Mismatched MTU value causes the RLT interface to flap [PR1457460](#)

J-Web

- Junos OS: J-Web Denial of Service due to multiple vulnerabilities in Embedthis Appweb Server. [PR1345330](#)
- Junos OS: Persistent XSS vulnerability in J-Web (CVE-2019-0047) [PR1410400](#)
- Junos OS: Session fixation vulnerability in J-Web (CVE-2019-0062) [PR1410401](#)

Layer 2 Features

- The backup VPLS router might still have MAC addresses after the primary router is rebooted and recovered in VPLS environment. [PR1356726](#)
- The traffic might not be transmitted correctly in a large scale of VPLS scenario. [PR1371994](#)
- Flow label is still used by ingress PE though the Egress PE is not configured/supporting for Flow label in a vpls multihomed Scenario. [PR1393447](#)
- The rpd crashes after iw0 interface is configured under a VPLS instance. [PR1406472](#)
- In a Layer 2 domain, there might be unexpected flooding of unicast traffic at every 32-40 seconds interval towards all local CE-facing interface. [PR1406807](#)
- Broadcast traffics may be discarded in a VPLS local-switching scenario. [PR1416228](#)

- Commit error will be seen but the commit is processed if adding more than o. [PR1420082](#)
- In VC scenario traffic drop might be seen when one VC member reboots and rejoins the VC [PR1453430](#)

Layer 2 Ethernet Services

- Junos core file `jdhcpd.core.0` found in `dhcpv6_packet_handle` is seen. [PR1329390](#)
- ZTP infra scripts are not included for MX PPC routers. [PR1349249](#)
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. [PR1373807](#)
- RADIUS accounting statistics are not cleared after subscriber logout. [PR1383265](#)
- The subscriber's authentication might fail when the link-layer address encoded in the DHCPv6 DUID is different from the actual link-layer hardware address. [PR1390422](#)
- The SNMP query on LACP interface might lead to lacpd crash. [PR1391545](#)
- Junos OS: jdhcpd crash upon receipt of crafted DHCPv6 solicit message (CVE-2019-0037). [PR1391983](#)
- On EVPN setups, incorrect wrong destination MAC addresses starting with 45 might show up when using the `show arp hostname` command. [PR1392575](#)
- After GRES switchover, LACP will be down on peer device and never been recovered automatically. [PR1395943](#)
- Log messages `dot1xd[]: task_connect: task ESP CLIENT:...: Connection refused` might be reported in Junos OS Release 17.4 or later. [PR1407775](#)
- jdhcpd becomes aware about some of the existing configuration only after 'commit full' or jdhcpd restart. [PR1419437](#)
- The jdhcpd process might consistently run at 100% CPU and not provide service if the `delay-offer` is configured for DHCP local server. [PR1419816](#)
- jdhcpd daemon might crash during continuous stress test. [PR1421569](#)
- The jdhcpd memory leak might happen on MX5, MX10, MX40/MX80, and MX104 when testing DHCP subscribers log-in and out. [PR1432162](#)

MPLS

- The command of "ping mpls l2circuit" might fail if the flow-label is enabled for l2circuit [PR1217566](#)
- MPLS routes might be dead if MPLS is disabled or deleted from the interface and enabled/added back to the same interface or the maximum number of labels is changed in quick succession without a delay [PR1355878](#)
- The LSP might remain UP even if no path is acceptable due to CSPF failure [PR1365653](#)
- RPD might restart after an MPLS LSP flap if "no-cspf" and "fast-reroute" are configured in an LSR ingress router. [PR1368177](#)

- RSVP authentication might fail between some Junos releases and cause traffic loss during local repair. [PR1370182](#)
- DSCP bit marking of LSP self-ping is not compliant with RFC7746 [PR1371486](#)
- The next hop of static LSP for MPLS might get stuck in dead state after changing the network mask of the outgoing interface. [PR1372630](#)
- The traceroute MPLS might fail when traceroute is executed from a Juniper device to another device not supporting RFC6424. [PR1372924](#)
- The rpd might crash when executing Routing Engine switchover under BGP environment and route churn occurs [PR1373313](#)
- The traffic might not be load-balanced equally across LSPs with ldp-tunneling configured [PR1373575](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition [PR1374102](#)
- The rpd process might crash continuously if nsr-synchronization or all flag is used in RSVP traceoptions [PR1376354](#)
- JSA10883: Junos OS: Receipt of a specifically crafted malicious MPLS packet leads to a Junos kernel crash (CVE-2018-0049) [PR1380862](#)
- The rpd might crash on backup Routing Engine after switchover [PR1382249](#)
- An RSVP-signaled LSP might stay in down state after a link in the path flaps. [PR1384929](#)
- Ingress LSPs down due to CSPF failure [PR1385204](#)
- Configured bandwidth 0 does not get applied on RSVP interface [PR1387277](#)
- The bypass LSP might pass through unexpected path that includes the same SRLG as the down protected TE link [PR1387497](#)
- The rpd process might keep crashing repeatedly if the LSP destination address is set to be 0.0.0.0 [PR1397018](#)
- The rpd might crash when LDP route with indirect next hop is deleted. [PR1398876](#)
- A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based-stats are used. [PR1401152](#)
- MPLS LSP traffic loss might be seen under rare conditions if CSPF is enabled [PR1402382](#)
- The L2circuit information is not advertised over the LDP session if "ldp dual-transport inet-lsr-id" is different from the router-id [PR1405359](#)
- The rpd might crash when RSVP bypass path flaps [PR1406400](#)
- LDP tunneling config triggers huge scheduler slips causing IGP flapping [PR1410827](#)
- Resources might be reserved for stale RSVP LSP when RSVP is disabled on the interface [PR1410972](#)
- The rpd might crash if longest-match is configured for LDP. [PR1413231](#)
- LDP route is not present in inet6.3 if IPv6 interface address is not configured [PR1414965](#)

- LDP routes might flap if committing any configuration changes. [PR1416032](#)
- Traffic might be silently dropped or discarded due to a long LSP switchover duration in RSVP-signaled LSP scenario [PR1416487](#)
- LDP route might be missing in inet.3 when enabling sr-mapping-client on LDP-SR stitching node [PR1416516](#)
- RSVP LSP might get stuck in down state in OSPF multiarea topology. [PR1417931](#)
- LDP might not update the LDP ingress route metric when inet.3 route flash happens before inet.0 [PR1422645](#)
- MPLS LSP auto-bandwidth statistics miscalculations might lead to high bandwidth reservation. [PR1427414](#)
- Traffic loss might be observed after changing configuration under "protocols mpls" in ldp-tunneling scenario [PR1428081](#)
- The LDP might withdraw a label for an FEC once the IGP route is inactive in inet.0 [PR1428843](#)
- When MBB for P2MP LSP fails, it is stuck in old path [PR1429114](#)
- SRLG entry shows Unknown after removing it from configuration in show mpls lsp extensive output or show mpls srlg. Shows Unknown-0xXX (XX will vary) [PR1433287](#)
- The P2MP LSP branch traffic might be dropped for a while when the Sender PE is doing switchover [PR1435014](#)
- The rpd will crash continuously if RSVP LSP link/node protection is configured [PR1435019](#)
- Traffic loss might be seen after LDP session flaps rapidly [PR1436119](#)
- The flow label is not pushed when "chained-composite-next-hop ingress l2ckt/l2vpn" is enabled [PR1439453](#)
- The LDP route and LDP output label are not showing in the inet.3 table and LDP database respectively if enable OSPF rib-group [PR1442135](#)
- RSVP Path message with long refresh interval is dropped between Junos pre-16.1 and 16.1+ nodes [PR1443811](#)
- The transit packets might be dropped if an LSP is added or changed on MX/PTX device [PR1447170](#)
- The LDP route timer is reset when committing unrelated configuration changes [PR1451157](#)
- High CPU usage and rpd core dump might be observed if "ldp track-igp-metric" is configured and IGP metric is changed [PR1460292](#)
- The device may use the local-computed path for the PCE-controlled LSPs after link/node failure [PR1465902](#)

Network Address Translation (NAT)

- The nsd process might crash during SNMP query for deterministic NAT pool information. [PR1436775](#)

Network Management and Monitoring

- The backup routing engine sends syslog messages to the syslog server with master fxp0 instead of lo0. [PR1341938](#)
- Child link missed from mib id dot3adAggPortAttachedAggID (OID - 1.2.840.10006.300.43.1.2.1.1.13). [PR1410439](#)
- The AGENTX session timeout between master (snmpd) and subagent triggers some daemon crash [PR1396967](#)
- The snmp query may not get data in scaled Layer 2 circuits environment. [PR1413352](#)
- Syslog match filtering does not work if single line of `/etc/syslog.conf` is over 2048 bytes. [PR1418705](#)

Platform and Infrastructure

- MAC addresses might not be learnt on bridge-domains after XE/GE interface flap [PR1275544](#)
- Distributed multicast might not be forwarded to a subscriber interface. [PR1277744](#)
- Junos OS: Login credentials are vulnerable to brute force attacks through the REST API (CVE-2019-0039) [PR1289313](#)
- The "show igmp statistics" command not including any statistics under interface aggregate for distributed multicast interfaces. [PR1289415](#)
- The dcd Micro BFD seems to be failing in dcd_commit_check log file even when BFD is not configured [PR1300796](#)
- The "Platform failed to bind rewrite" message can be seen when chassis control restart is done with the COS rewrite rule configured on aggregated Ethernet interface [PR1315437](#)
- Inline keepalive session might be down due to lcp-keepalive-failure on MPC5E/MPC6E PICO interfaces [PR1343687](#)
- The rpd might crash when doing Routing Engine switchover with NSR and logical-system configurations. [PR1345720](#)
- Packet drop might be seen on the logical tunnel interfaces lt-x/2/x or lt-x/3/x [PR1345727](#)
- RLT subinterfaces not reporting statistics. [PR1346403](#)
- lt- interface gets deleted with tunnel-services configuration still present. [PR1350733](#)
- Some line cards might crash in subscriber scenario enabled with distributed IGMP. [PR1355334](#)
- When forwarding-class-accounting statement is enabled on an interface, inside of a routing-instance of instance-type vrf, aggregate input forwarding-class statistics do not increment (egress statistics work fine). [PR1357965](#)
- Traffic might drop on new added interfaces on MX Series after unified ISSU [PR1371373](#)

- The logical tunnel interface might be unable to send out control packets generated by Routing Engine. [PR1372738](#)
- kernel and ksyncd core files are generated after dual CB flap at rt_nhfind_params: rt_nhfind() found an nh different from that onmaster 30326. [PR1372875](#)
- JNH memory leaks in multicast scenario with MoFRR enabled [PR1373631](#)
- The traffic traversing an IRB interface might not be tagged with a VLAN if the packets go through an additional routing-instance [PR1377526](#)
- FPC crash might be seen after FPC restarts [PR1380527](#)
- Packet drops on interface if the statement "gether-options loopback" is configured. [PR1380746](#)
- Traffic loss seen in Layer 2 VPN with GRE tunnel. [PR1381740](#)
- MAC learning might get stuck on MX Series router with DPC and MPC [PR1383233](#)
- Jlock hog might be reported at restart routing [PR1389809](#)
- Individual command authorization might cause mgd crash. [PR1389944](#)
- Traffic is dropped when passing through MS-DPC to MPC. [PR1390541](#)
- The command "commit synchronize" might fail because several internal connections are stuck. [PR1394370](#)
- When using ifconfig utility to bring down the PS logical interface, its Admin status is not going down as expected. [PR1396335](#)
- The packet might drop in tunnel interface with a checksum error [PR1396372](#)
- All FPC cards might restart after Layer 3 VPN routes churn. [PR1398502](#)
- RVT interface might get flapping [PR1399102](#)
- Syslog error message: [LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_pfe:3211): pfe_id 0 cchip 0[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_pfe:3272): No PFE found for pfe_id_start 0 is seen. [PR1402377](#)
- Some files are missing during log archiving [PR1405903](#)
- Abnormal queue-depth counters in "show interface queue" output on interfaces which associated to XM2 and 3 [PR1406848](#)
- IPv6 traffic might be dropped between VXLAN bridge-domain and IP/MPLS network [PR1407200](#)
- Traffic is getting dropped when there is a combination of DPC/MX-FPC card and MPC card on egress PE router in Layer 3 VPN [PR1409523](#)
- Junos OS: Insufficient validation of environment variables in telnet client may lead to stack-based buffer overflow (CVE-2019-0053) [PR1409847](#)
- FPC crash may be observed with scaled subscribers login attempts [PR1409879](#)
- The VLAN tag is incorrectly inserted on the access interface if the packet is sent from an IRB interface. [PR1411456](#)

- The MPC might crash when one MIC is pulled out during this MIC is booting up [PR1414816](#)
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10 (based system). [PR1417321](#)
- The op url command cannot run a script with libs from /config/scripts. [PR1420976](#)
- The ARP request might not be replied although "proxy-arp" is configured [PR1422148](#)
- show jnh trap-info with incorrect LU instance caused a crash and core files on FPC. [PR1423508](#)
- The policer bandwidth might be incorrect for the aggregate interface after activating the command 'shared-bandwidth-policer'. [PR1427936](#)
- The FPC might crash when the firewalls filter manager deals with the firewall filters [PR1433034](#)
- The device might not be accessible after the upgrade [PR1435173](#)
- The BGP session might flap after performing RE switchover simultaneously on both end of BGP peers [PR1437257](#)
- The next-hop MAC address in the output from "show route forwarding-table" command might be wrong [PR1437302](#)
- The multicast traffic is dropped while multicast ingress replication is configured with "local-latency-fairness" [PR1438180](#)
- The RPM udp-ping probe does not work in a multiple routing instance scenario. [PR1442157](#)
- Packets drop due to misssing destination MAC in the Packet Forwarding Engine. [PR1445191](#)
- Some hosts behind unnumbered interface are unreachable after the router/FPC restarts [PR1449615](#)
- The RE originated IPv6 packets might be dropped when interface-group rule is configured under IPv6 filter [PR1453649](#)

Routing Policy and Firewall Filters

- set metric multiplier offset may overflow/underflow. [PR1349462](#)
- MX Series Router: CLI statement **as-path-expand last-as** commit failure. [PR1388159](#)
- The rpd process might crash when **routing-options flow** configuration is removed. [PR1409672](#)
- Policy matching RD changes next-hop of the routes which do not carry RD [PR1433615](#)
- Routes resolution might be inconsistent if any route resolving over the multipath route [PR1453439](#)

Routing Protocols

- BGP might not advertise routes on the existing BGP peer after adding Layer 3 VPN instance [PR1237006](#)
- Multihop eBGP peering session exchanging EVPN routes can result in rpd core files when BGP updates are sent. [PR1304639](#)
- With Resource Certification (RPKI) enabled, RPD successive crashes during route validation DB processing [PR1309944](#)

- The BGP session might be stuck with high BGP OutQ value after GRES on both sides [PR1323306](#)
- Junos OS: RPD process crashes when BGP peer restarts (CVE-2019-0049) [PR1337304](#)
- The VRF static route might not be exported when route-distinguisher-id is used on RR in BGP Layer 3 VPN scenario. [PR1341720](#)
- The bfd process memory leak might be observed if enabling multi-hop BFD session for a static route with multiple qualified-next-hop [PR1345041](#)
- vFPC may continuously crash on vMX platform. [PR1364624](#)
- Ukern memory leak and core crash in BGP environment [PR1366823](#)
- Qualified next hop of static route might not be withdrawn when BFD is down [PR1367424](#)
- About 10 minutes traffic loss is caused by BGP flap during unified ISSU. [PR1368805](#)
- RE-based micro BFD packets do not go out with configured source IP when the interface is in logical-system [PR1370463](#)
- TCP sessions might be taken down during RE switchover [PR1371045](#)
- Route entry might be missing when IS-IS shortcut is enabled and MPLS link flaps. [PR1372937](#)
- static route age is the same as last commit. [PR1377279](#)
- The rpd process might crash after executing commit the configuration related to mapping-server-entry [PR1379558](#)
- 2019-01 Security Bulletin: Junos OS: OpenSSL Security Advisories [16 Apr 2018] and [12 June 2018] [PR1380686](#)
- The sshd authentication logs includes one syslog in UTC time [PR1382786](#)
- The rpd might crash under a rare condition if GR helper mode is triggered [PR1382892](#)
- Polling interface statistic and status becomes very slow when MPC CPU goes to 100% [PR1383373](#)
- The static route might persist even after its BFD session goes down [PR1385380](#)
- The rpd might crash after issuing operational command "show route detail" for RIP route [PR1386873](#)
- BGP sessions might keep flapping on backup Routing Engine if proxy-macip-advertisement is configured on IRB interface for EVPN-VXLAN. [PR1387720](#)
- Penultimate-hop router does not install BGP LU label causing traffic to be silently dropped or discarded. [PR1387746](#)
- IGMPv3/MLD membership requests could not work normally [PR1389119](#)
- Unexpected packet loss might be seen for some multicast groups during failure recovery with both MoFRR and PIM automatic MBB join load-balancing features enabled [PR1389120](#)
- In rare cases rpd might crash after Routing Engine switchover when BGP multipath and Layer 3 VPN vrf-table-label are configured [PR1389337](#)

- FPC might crash when BGP multipath is configured with protection [PR1389379](#)
- Race condition causes all the BGP sessions to flap after NSR switchover [PR1391084](#)
- Non-BGP protocol route with an AS PATH might cause inappropriate route selection [PR1391767](#)
- The pcmd on the Routing Engine might run with high CPU utilization after Routing Engine switchover. [PR1392704](#)
- RPD core files on backup Routing Engine during neighborship flap when using authentication-key with size larger than 20 character. [PR1394082](#)
- Multicast traffic might be interrupted in H-VPLS scenario [PR1394213](#)
- The rpd process might crash when rp-register-policy is configured with more than 511 terms [PR1394259](#)
- The best and the second-best routes might have the same weight value if BGP PIC is enabled [PR1395098](#)
- BGP DMZ LINK BANDWIDTH - not able to aggregate bandwidth, when applying the policy. [PR1398000](#)
- The rpd soft core files and inappropriate route selection might be seen when Layer 2 VPN is used [PR1398685](#)
- The process rpd might crash in BGP setup with NSR enabled. [PR1398700](#)
- Junos OS: BGP packets can trigger rpd crash when BGP tracing is enabled. (CVE-2019-0019) [PR1399141](#)
- The UHP behavior is not supported for LDP to SR stitching scenario [PR1401214](#)
- There might be unexpected packets drop in MoFRR scenario if active RPF path is disabled [PR1401802](#)
- The rpd might be stuck at 100% when auto-export and BGP add-path are configured [PR1402140](#)
- On the multi-access/broadcast network, third party BGP router might unexpectedly select RR router as next-hop to forward the IPv6 traffic. [PR1402255](#)
- M/Mx/QFX:mcsnoopd core generated immediately after the commit change related to VXLAN-EVPN configuration [PR1408812](#)
- The L3VPN link protection doesn't work after flapping the CE facing interface [PR1412667](#)
- The unexpected AS prepending action for AS path might be seen after the no-attrset statemnt is configured or deleted with vrf-import/vrf-export configuration. [PR1413686](#)
- The rpd gets stuck in a loop while doing the multipath calculation which leads to the high CPU usage [PR1414021](#)
- Dynamic routing protocol flapping with VM host Routing Engine switchover on NG-RE. [PR1415077](#)
- Junos OS: OpenSSL Security Advisory [26 Feb 2019] [PR1419533](#)
- A memory leak in rpd might be seen if source packet routing is enabled for IS-IS protocol [PR1419800](#)
- The bfdd process might crash on old master RE during GRES [PR1420694](#)
- IPv6 IS-IS routes might be deleted and not be reinstalled when MTU is changed under the logical interface level for family inet6 [PR1420776](#)

- Route churn might be seen after changing maximum-prefixes configuration from value A to vlaue B [PR1423647](#)
- The rpd might crash if no-propagate-ttl is configured in BGP multipath scenario [PR1425173](#)
- The rpd might crash in PIM scenario with auto-rp enabled [PR1426711](#)
- The rpd might crash while removing multicast routes that do not have an associated (S,G) state or activating the "accept-remote-source" knob on PIM upstream interface [PR1426921](#)
- The rpd might crash while handling the withdrawal of an imported VRF route [PR1427147](#)
- The rpd generates core file due to improper handling of graceful restart stale routes. [PR1427987](#)
- IPv6 aggregate routes are hidden [PR1431227](#)
- PIM-SM join message might be delayed with MSDP enabled [PR1433625](#)
- Removing SSH Protocol version 1 from configuration [PR1440476](#)
- RIP routes might be discarded by Juniper device over a /31 subnet interface [PR1441452](#)
- The rpd process might crash in inter-AS option B Layer 3 VPN scenario if CNHs is used [PR1442291](#)
- The rpd crash might be seen after configuring OSPF nssa area-range and summaries [PR1444728](#)
- The rpd might crash in OSPF scenario due to invalid memory access [PR1445078](#)
- JUNOS BFD sessions with authentication flaps after a certain time [PR1448649](#)
- The connection between ppmmd(RE) and ppmman(FPC) might get lost due to session timeout [PR1448670](#)
- The rpd scheduler slip for BGP GR might be up to 120s after the peer goes down [PR1454198](#)
- Prefix SID conflict might be observed in ISIS [PR1455994](#)
- The rpd scheduler slips might be seen on RPKI route validation enabled BGP peering router in a scaled setup [PR1461602](#)
- Install all possible next-hops for OSPF network LSAs [PR1463535](#)
- BGP peers might flap if the parameter of hold-time sets small [PR1466709](#)

Services Applications

- IPsec-VPN IKE security-associations might get stuck in "Not Matured" state. [PR1369340](#)
- Inline Service interface may not UP when bandwidth is configured. [PR1370405](#)
- NAT64 does not translate ICMPv6 Type 2 packet (packet is too big) correctly when MS-DPC is used for NAT64. [PR1374255](#)
- Twice NAT not supported on FTP ALG causes MS-PIC crash. [PR1383964](#)
- L2TP subscribers might be stuck in init state in a corner case. [PR1391847](#)
- The spd might crash when **any-ip** is configured in the 'from' clause of the NAT rule with the static translation type. [PR1391928](#)

- IP ToS bits are not copied to outer IPSec header. [PR1398242](#)
- Invalid Layer 4 checksum might be observed on IPv4 packets generated by NAT64 with MS-DPC after translating fragmented IPv6 UDP/TCP packets. [PR1398542](#)
- The ICMPv6 packet with embedded IPv6 fragment might not be translated correctly to IPv4 ICMP packet in a NAT64 with MS-DPC deployment. [PR1402450](#)
- The stale si- IFL might be seen when L2TP subscribers with duplicated prefixes or framed-route login. [PR1406179](#)
- The kmd process might crash on MX/ACX platforms when IKEv2 is used. [PR1408974](#)
- jpppd core files on LNS. [PR1414092](#)
- L2TP LAC might fail to tunnel static pp0 subscriber to the desired LNS. [PR1416016](#)
- IPsec SA may not come up when the Local gateway address is a VIP for a VRRP configured interface. [PR1422171](#)
- In subscriber with L2TP scenario, subscribers are stuck in INIT state forever. [PR1425919](#)
- The kmd process may crash when DPD timeout for some IKEv2 SAs happens. [PR1434521](#)
- Traffic might be dropped in IPsec VPN scenario when the VPN peer is behind a NAT device. [PR1435182](#)

Software Installation and Upgrade

- JSU might be deactivated from FPC in case of power cycle. [PR1429392](#)

Subscriber Access Management

- The authd process might not be started after executing RE switchover on backup Routing Engine or without GRES enabled. [PR1368067](#)
- Address pool does not correctly cycle to the beginning of the pool when linked-pool-aggregation parameter is defined. [PR1374295](#)
- The subscribers might be stuck in terminating state if radius redirect is used. [PR1376265](#)
- Radius VSA's, Actual-Data-Rate-Downstream and Actual-Data-Rate-Upstream values are not complaint with RFC 4679. [PR1379129](#)
- CoA updates subscriber with original dynamic-profile if radius has returned different dynamic-profile name. [PR1381230](#)
- Some subscribers fail to get SRL service as provided in Radius accept message even though the Radius messages can be sent and received. [PR1381383](#)
- The value of 'predefined-variable-defaults routing-instances' overrides the RADIUS-supplied VSA (26-1 Virtual-Router). [PR1382074](#)
- The RAA message may consist of additional AVP "Destination-Host" even it is not configured for Gx-Plus session. [PR1384011](#)

- Log Message: **authd: gx-plus: logout: wrong state for request session-id <xyz>**. [PR1384599](#)
- Multiple IPv6 IANA addresses assigned for one session in IPv6 PD binding failure scenarios. [PR1384889](#)
- Usage-Monitoring-Information AVP maybe activate service accounting. [PR1391411](#)
- The DHCPv6-PD client connection might be terminated after commit when RADIUS assigned address is not defined within the range of a local pool. [PR1401839](#)
- The authd crash might be seen due to a memory corruption issue. [PR1402012](#)
- JSRC used Radius Service accounting protocol instead of JSRC for SRC installed service. [PR1403835](#)
- Some continuous log messages could be seen. [PR1407923](#)
- Subscribers might not be able to re-login in Gx-plus provisioning scenario. [PR1418579](#)
- Address allocation issue with linked pools when using linked-pool-aggregation. [PR1426244](#)
- RADIUS authentication server might always be marked with DEAD. [PR1429528](#)
- On MX platforms a false error might be received for SAE policy activation/deactivation failure [PR1447632](#)

User Interface and Configuration

- The **show configuration** and **rollback compare** commands causing high CPU usage. [PR1407848](#)

VPNs

- Non-optimal route to source might be selected for NG-MVPN with unicast-umh-election enabled. [PR1315011](#)
- The process rpd may crash after configuration change in an Layer 2 VPN scenario. [PR1351386](#)
- In dual-homed NG-MVPN the receipt of type 5 withdrawal removes downstream join states for some routes. [PR1368788](#)
- The receivers belonging to a routing instance may not receive multicast traffic in an Extranet next-generation MVPN scenario. [PR1372613](#)
- The **accept-remote-source** knob configured on the core interface might cause traffic outage. [PR1375716](#)
- High rpd CPU utilization on the backup Routing Engine might be observed in MVPN+NSR scenario. [PR1392792](#)
- The rpd process crashes when LSP template for a provider tunnel is changed [PR1395353](#)
- Downstream interface is not removed from multicast route after getting PIM prune. [PR1398458](#)
- Dvaita JDI-RCT: NGMVPN Traffic drops seen for multicast groups with "selective" provider tunnels [PR1406757](#)
- The multicast traffic drop might be seen when **static-umh** is configured in NGMVPN scenario. [PR1414418](#)
- The deletion of (S,G) entry might be skipped after the PIM join timeout. [PR1417344](#)
- The rpd process might crash in rare conditions when Extranet NG-MVPN is configured. [PR1419891](#)

- MPLS LSP ping over I2circuit might not work when flow-label is enabled [PR1421609](#)
- The resumed multicast traffic for certain groups might be stopped in overlapping MVPN scenario [PR1441099](#)
- Memory leak might happen if PIM messages received over an MDT (mt- interface) in Draft-Rosen MVPN scenario [PR1442054](#)
- The rpd process might crash due to memory leak in "MVPN RPF Src PE" block [PR1460625](#)
- The I2circuit connections might be stuck in OL state after changing the I2circuit community and flapping the primary LSP path [PR1464194](#)

Resolved Issues: 17.4R2

Application Layer Gateways (ALGs)

- IKEv2 negotiation might fail with IKE ESP ALG enabled in an IKEv2 redirection scenario. [PR1329611](#)

Authentication and Access Control

- The client moves back to connecting state when VSTP is enabled along with dynamic vlan assigned once port get authenticated by dot1x [PR1304397](#)

Class of Service (CoS)

- CoS wildcard configuration is applied incorrectly after a router restart. [PR1325708](#)
- Remove CoS IDL from the jet IDL package. [PR1347175](#)
- The Routing Engine might get into amnesiac mode after restarting if **excess-bandwidth-share** is configured. [PR1348698](#)
- The aggregated Ethernet link-protection feature is not supported. [PR1355498](#)

EVPN

- EVPN traffic mapping to specific LSPs is not working. [PR1281415](#)
- The rpd might crash on platform using junos with evpn and nsr enabled after restarting the rpd process in EVPN environment [PR1320408](#)
- An EVPN discard route is installed on the local provider edge (PE) device after connection flaps on a remote PE device in a multihome EVPN topology. [PR1321125](#)
- If host is multihomed then all PEs should install the /32 host IP address pointing to its local IRB interface as long as its local multihomed ES interface is up. [PR1321187](#)
- The rpd crash might happen during EVPN/VXLAN configuration changes. [PR1321839](#)
- RPD crash on backup Routing Engine if NSR and IS-IS SR enabled. [PR1323980](#)
- The FPC might crash after deleting the VPLS configuration. [PR1324830](#)

- A core link flap might result in an inconsistent global MAC count. [PR1328956](#)
- On a deactivated end system identifier (ESI) for PS at a physical interface level, the rpd process generates core files for EVPN VPWS PWHT. [PR1332652](#)
- On doing **restart routing**, the rpd process might generate core files on a PE router that has a EVPN-VXLAN configuration. [PR1333331](#)
- MPLS label leak leads to label exhaustion and the rpd process crash [PR1333944](#)
- In an EVPN scenario with nonstop active routing (NSR) enabled, the rpd crashes and generates core files on the backup Routing Engine while any configuration changes on the master Routing Engine. [PR1336881](#)
- The rpd process might crash when executing CLI command "show route evpn-ethernet-tag-id" [PR1337506](#)
- In an EVPN-VXLAN environment, the BFD flap causes the VTEP to flap, causing the Packet Forwarding Engine to crash. [PR1339084](#)
- Traffic loss might be observed in an EVPN-VPWS scenario if the remote PE's interface comes down. [PR1339217](#)
- On EVPN-VXLAN scenarios, the traffic might get black-holed to interfaces that are down, but LACP is up. [PR1343515](#)
- The rpd might crash if the IRB interface and routing instance are deleted together in the same commit. [PR1345519](#)
- Traffic might be lost on a Layer 2 and Layer 3 spine node in a multihome EVPN scenario. [PR1355165](#)
- EVPN IRB configured with **no-gratuitous-arp-request** is still sending gratuitous ARP. [PR1356360](#)
- The rpd might crash if the EVPN instance refers to a vrf-export policy which doesn't have 'then community'. [PR1360437](#)
- Proxy ARP may not work as expected in an EVPN environment. [PR1368911](#)

Forwarding and Sampling

- The pfd process generates a core file in `pfed_process_session_state_notification_msg, pfed_timer_manager_c::remove_serv_id, pfed_delete_timer_id_by_serv_sid (serv_sid=0, serv_info=0x0)` at `../../../../src/junos/usr.sbin/pfed/pfed_timer.cc:16`. [PR1296969](#)
- Remote CE1 MAC address might take more time to clear after clearing MAC. [PR1304866](#)
- The dfwd process might crash during execution of **show firewall templates-in-use** command. [PR1305284](#)
- The second archive site in the accounting-file configuration is not used when the first one uses SFTP and is not reachable. [PR1311749](#)
- Accounting files with no records might be unexpectedly uploaded to the archive site. [PR1313895](#)
- The FPC CPU might reach 100 percent constantly if shared bandwidth policer is configured. [PR1320349](#)
- The error messages about `dfw_gencfg_handler` might be seen during a unified ISSU. [PR1323795](#)

- Ukernel leaks 6x40 bytes heap nodes upon each IPC path when handshaking or establishment occurs between l2alm and l2ald. [PR1326921](#)
- DHCP service crashes after the device is set to factory default by zeroize. [PR1329682](#)
- Some firewall filter counters might not be created in SNMP. [PR1335828](#)
- The error logical interface under VPLS might be blocked after MAC moving if the logical interfaces are on the same physical interface. [PR1335880](#)
- In EVPN-VXLAN **clear ethernet-switching table** might not work correctly. [PR1341328](#)
- Junos allows firewall filters with the same name under **edit firewall** and **edit firewall family inet** hierarchy levels [PR1344506](#)
- Commit failed when attempting to delete any demux0 unit numbers that are greater or equal to 1000000000. [PR1348587](#)
- The remote MAC might not be added in the forwarding table, which will cause a traffic drop in an EVPN scenario with RSVP and CBF configured. [PR1353555](#)
- The backup Routing Engine is writing dummy interface accounting records. [PR1361403](#)

General Routing

- In timing hybrid mode, MX Series MPC2 cards are not working with ACX with VLAN (native-vlan-id). [PR1076666](#)
- An rpd memory leak is caused by repeated RSVP reservation state block (RSB) deletes. [PR1115686](#)
- No warning is raised when the bridge family is configured with interface-mode trunk but without vlan-tagging or flexible-vlan-tagging. [PR1154024](#)
- An unexpected **MobileNext Gateway Activation license** alarm is observed when TDF gateway is configured. [PR1162518](#)
- The replacement PIC might bounce when PIC PB-4OC3-4OC12-SON-SFP (4x OC-12-3 SFP) is replaced with PB-4OC3-1OC12-SON2-SFP (4x OC-3 1x OC-12 SFP) and a CLI commit is made. [PR1190569](#)
- Agentd process crashes with core-dump [PR1197608](#)
- The **Unable to deregister sub error (131072) for error(0x1b0001) for module MIC** error messages are seen on the MPC5E card. [PR1221337](#)
- The error log **cc_mic_irq_status: CC_MIC(5/2) irq_status(0x1d) does not match irq_mask(0x20), enable(0x20), latch(0x1d)** is seen continuously for MIC-3D-4OC3OC12-1OC48. [PR1231084](#)
- The **chassisd[9132]: LIBJSNMP_NS_LOG_NOTICE: NOTICE: netsnmp_ipc_client_connection: unix connection error: socket(-1) main_session(0x9812f80)** error messages are seen after a chassis-control restart. [PR1243364](#)
- The GNF sometimes resets its MPC type 9 at NSR at a high scale. [PR1259910](#)

- On a vMX FPC, the software FPC might restart unexpectedly with the following message: **panic (format_string=format_string@entry=0x9e509c4 "Thread %s attempted to %s with irq priority at %d\n").** [PR1263117](#)
- The **show chassis FPC** command does not show temperature. [PR1263315](#)
- The load-based throttling functionality is not enabled by default. [PR1271739](#)
- Flexible PIC concentrator (FPC) crash/reboot is observed when bringing up about 12K Layer 2 Bit Stream Access(L2BSA) subscribers simultaneously. [PR1273353](#)
- Error messages observed on vty session while running script for IGMP Snooping over EVPN-VXLAN. [PR1276947](#)
- On an MX104 platform with GRES enabled, the chassis network-services might not get set as "Enhanced-IP". [PR1279339](#)
- BSYS logs messages are reporting that GNF owned PICs do not support power off configuration at commit when no such configuration is present. [PR1281604](#)
- The kernel might crash when an NSR enabled device has BGP peer flapping. [PR1282573](#)
- The enhancement of reporting total SBE errors when the corrected single-bit errors threshold of 32 is exceeded for MPC7E/MPC8E/MPC9E. [PR1285315](#)
- The LC, PFH, and Packet Forwarding Engine interfaces do not come up on Routing Engine 1. [PR1285606](#)
- The missing statement **Shared bandwidth policer not supported for interface ge-x/x/x** is seen during a failed commit in Junos OS Release 16.1R3. [PR1286330](#)
- The oneset or leaf-list configuration might not get deleted with the delete operation through JSON. [PR1287342](#)
- PPPoE cannot dial in due to all padi dropped as "unknown iif" when deactivated/activated AE configuration. [PR1291515](#)
- During PPPoE subscriber login errors like **vbf_flow_src_lookup_enabled** and **Failed to find iif structure, iifl** were seen on FPC. [PR1294710](#)
- The KRT queue might be stuck with the **RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0** error. [PR1295756](#)
- Some random number of ports on a 10-Gigabit MPC7E card might not come up after the remote system or line card restarts or interface flaps. [PR1298115](#)
- The log message about the shutdown time is incorrect when the system exceeds chassis over the temperature limit. [PR1298414](#)
- When the subscriber limit feature is configured, any new login request after the maximum number of subscribers is denied. [PR1298924](#)
- The error messages about PEM might be seen in the MX Series platform with AC PEM. [PR1299284](#)

- A chassisd core file is seen after the insertion of REMX2K-X8-64 in MX2000 line routers with the older RE-S-1800x4. [PR1300083](#)
- The ICMP/ICMPv6 error messages might be discarded while forwarding through an AMS interface. [PR1301188](#)
- Reported same IFD KV by two different sensors. [PR1301858](#)
- The rpd might crash when NSR is enabled and routing-instance specific configurations are committed. [PR1301986](#)
- Continuous interface flapping might lead to an unwanted MIC reset. [PR1302246](#)
- The multicast resolve-rate value might go back to default after system upgrade or reboot. [PR1303134](#)
- Internal latency is high during the initial subscription of sensors. [PR1303393](#)
- Fan speed changes frequently on MX Series after an upgrade to Junos OS software. [PR1303459](#)
- The fabric planes might go into "check" state after restarting the line cards with SFB2 used on the MX2010 or the MX2020. [PR1304095](#)
- The **start shell pfe network fpc** command is not working on the MX960. [PR1306236](#)
- /Frame: messages might be seen with Telemetry enabled. [PR1308513](#)
- FPC syslog errors with **pfeman_inline_ka_steering_gencfg_handler: nh not found** could mean that steering rules are not installed correctly. [PR1308884](#)
- After a smooth upgrade from SFB to SFB2, if one plane/SFB is restarted, link training fails between those planes and MPC6 cards. [PR1309309](#)
- First access-request is failing for L2BSA subscribers when changing the MTU of LACP aggregate Ethernet A10NSP interface. [PR1309599](#)
- Subscribers might not be able to access the device if dynamic VLAN is used. [PR1309770](#)
- Ninety percent of subscribers might go down after a unified ISSU from Junos OS Release 16.1 to Junos OS Release 17.3. [PR1309983](#)
- Local IPv6 interface address from the NDRA prefix is not removed from the service interface when the subscriber dual-stack session is removed. [PR1310752](#)
- The utilization of "commit check" just after setting the master-password can trigger an improper decoding of configuration secrets. [PR1310764](#)
- After guest network functions (GNFs) Routing Engine switches mastership as expected, the rpd might be unresponsive. [PR1310765](#)
- The incorrect error number might be reported for syslog messages with a prefix of %DAEMON-3-RPD_KRT_Q_RETRIES. [PR1310812](#)
- Fragmented UDP packet might be incorrectly parsed as a uBFD packet and dropped. [PR1311134](#)
- Suppress chassis alarm for switched off PEMs. [PR1311574](#)

- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- The rpd process generates a core file after multiple session flaps on a scale setup. [PR1312169](#)
- The PEM alarms and I2C failures are observed on MX240, MX480, MX960, EX92, and SRX5K. [PR1312336](#)
- A false over temperature SNMP trap could be seen when using MPC5/6/7/8/9 on an MX2020. [PR1313391](#)
- The IPv6 router-solicit (RS) packets are dropped in nondefault RI, but for default RI it is working. [PR1313722](#)
- The **show version detail** command gives severity **error log traffic-dird[20126]: main: swversion pkg: 'traffic-dird' name: 'traffic-dird' ret: 0**. [PR1313866](#)
- The jdmd subsystem is not responding after an upgrade. [PR1313964](#)
- The mspmand process generates a core file because of a flow-control seen while clearing CGNAT+SFW sessions. [PR1314070](#)
- When ccc is configured on a umic interface, ARP is not resolving and observing traffic loss. [PR1314149](#)
- The JDM link is incorrectly shown to be up when the underlying physical link is down. [PR1314180](#)
- The **show version detail | no-more** command hangs for more than 120 seconds in the master Routing Engine and more than 60 seconds in the backup Routing Engine. [PR1314242](#)
- The smgd process generates a core file with reference to bbe_cos_ifl_publish() bbe_cos_if.c:6543. [PR1314651](#)
- The rpd process might crash in a MoFRR scenario. [PR1314711](#)
- For MPC7E, there is an **IR-mode** commit failure. [PR1314755](#)
- The L2TP LAC might drop packets that have an incorrect payload length while sending packets to the LNS. [PR1315009](#)
- Continuous logs from vhlclient are seen for all the commands executed. [PR1315128](#)
- FPC crash is observed when a route has unilist next-hops in a RSVP scenario. [PR1315228](#)
- The **show version detail** command gives severity **error log mobiled: main Neither BNG LIC nor JMOBILE package is present,exit mobiled**. [PR1315430](#)
- The **show version detail** command might generate severity error log **main: name: SRD ret: 0**. [PR1315436](#)
- Sensors belong to the same producer with identical reporting interval are not streamed in parallel [PR1315517](#)
- The rpd process generates a core file when a **show route inetcolor.0** command is executed from the CLI. [PR1316078](#)
- The fan speed might frequently keep changing between normal and full for the MX Series platform. [PR1316192](#)

- The demux interface sends a neighbor solicitation with source link-address of all zeros 00:00:00:00:00:00 MAC. [PR1316767](#)
- The **show configuration <> | display json** command might not be properly enclosed in double quotes. [PR1317223](#)
- Linux-based microkernel might panic due to a concurrent update on mutable objects. [PR1317961](#)
- CoA shaping rate is not applied successfully after a unified ISSU from Junos OS Release 15.1R6.7 to Release 16.1R6.2. [PR1318319](#)
- The rpd process might crash when the link flaps on an adjacent router. [PR1318476](#)
- The bbe-smgd process might crash after performing GRES. [PR1318528](#)
- The FPC crashes on a configuration change for the Packet Forwarding Engine sensors. [PR1318677](#)
- Changed text reported in the **show chassis hardware** output for CFP2-DCO optical transceivers. [PR1318901](#)
- MS-MPC and MS-MIC might crash after a new IPsec tunnel is added. [PR1318932](#)
- The MPC with specific failure hardware might impact other MPCs in the same chassis. [PR1319560](#)
- The kernel might generate a core file if the number of routing instances created are more than 256. [PR1319781](#)
- The task replication might not be complete to certain network protocols after multiple GRES. [PR1319784](#)
- The error log message of **MIB2D_COUNTER_DECREASING: pfes_stats_delta: counter** might be seen on VMX. [PR1319996](#)
- Loading xmlproxy YANG files cause telemetry session and some daemons to restart. [PR1320211](#)
- The chassis MIB SNMP OIDs for VC-B member chassis are not available after an MX Series Virtual Chassis unified ISSU. [PR1320370](#)
- The **show subscriber summary** command displays an incorrect terminated subscriber count. [PR1320717](#)
- The PPP inline keepalive does not work as expected when CPE aborts the subscriber session. [PR1320880](#)
- The rpd process crashes during the BGP configuration change and telemetry streaming with OpenConfig. [PR1320900](#)
- MX Series routers send the IPv6 router advertisements and the DHCPv6 advertisements before sending IPCPv6 ACK from CPE. [PR1321064](#)
- CoS is not applied to the Packet Forwarding Engine when the VCP link is added. [PR1321184](#)
- The bbe-smgd process generates core files after massive clients log out and log in, in a PPPoE dual stack subscriber scenario. [PR1321468](#)
- A CoA-NAK with "Error-Cause = Invalid-Request" is sent back to the RADIUS server when a drop policy is applied under radius-flow-tap in an L2TP subscriber scenario. [PR1321492](#)
- The **show system schema module hierarchy** command is broken in the CLI. [PR1321682](#)

- In commit fast-synchronize mode, the commit operation might get stuck after the commit check is performed. [PR1322431](#)
- The rpd process might crash when two next hops are installed with the same next-hop index. [PR1322535](#)
- The rpd process might crash when the OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)
- MS-MIC interface IFLs remain down after many iterations of offline/online. [PR1322854](#)
- An incorrect output is observed while verifying the command **show subscribers client-type vlan subscriber-state active logical-system default routing-instance default**. [PR1322907](#)
- NCP Conf-Ack/Conf-Req packets might be dropped constantly from the MLPPP client. [PR1323265](#)
- CLI commands in **show system subscriber-management route routing-instance <XXX>** hierarchy show unexpected outputs. [PR1323279](#)
- JDM Management is unreachable after flapping physical JDM and GNF/VNF management interfaces. [PR1323519](#)
- The **request vmhost halt routing-engine other** command does not halt the backup Routing Engine. [PR1323546](#)
- Memory leaks in the MGD-API process during Get API Requests and Error Handling during Set API Request. [PR1324321](#)
- Subscribers might fail to log in after the interface is deactivated or activated. [PR1324446](#)
- A memory leakage is seen in the mosquito-nossl process. [PR1324531](#)
- The SNMP interface filter does not work when "interface-mib" is part of the dynamic-profile. [PR1324573](#)
- KRTQ entries are waiting in an async queue. [PR1324669](#)
- The VLAN rewrite function might put the wrong VLAN ID when an Ethernet OAM is configured on DPCE cards. [PR1325070](#)
- The SNMP values might not be increased monolithically. [PR1325128](#)
- The MPC cards might drop traffic under a high temperature. [PR1325271](#)
- Non-MACsec interfaces are impacted when first time MACsec is configured on one of the interfaces or respective FPC is rebooted. [PR1325282](#)
- IS-IS adjacency fails to establish because packets drop on Packet Forwarding Engine. [PR1325311](#)
- MACsec session might fail to establish on MX10003. [PR1325331](#)
- The VLAN demux interface does not respond to the ARP request in a subscriber scenario with an MX Series router after Junos OS Release 15.1 with subscriber-management enabled. [PR1326450](#)
- MACsec MKA transmit Interval is changed to the upper limit. [PR1326526](#)
- In an MX Series BNG, the CoS service object is not deleted properly for TCP and scheduler. [PR1326853](#)
- Some **show** commands were issued twice when a **request support information** is executed. [PR1327165](#)

- With auto-installation USB configured, interface related commits might not take effect due to a dcd error. [PR1327384](#)
- Minor alarm **LCM Peer Connection un-stable** is observed on an MX150 after the chassisd process startup or restart. [PR1328119](#)
- Only 5.5M TCP sessions can be established for a NAPT44_SFW_APP_EIM/EIF configuration on an MS-MIC. [PR1328510](#)
- The following message is constantly logged: **fm_feacap_sys_feature_get:Attribute DB init not yet done, reading from pvid (id: 18)**. [PR1328868](#)
- For the **show class-of-service interface demux0 <demux interface>** command, the Adjustment overhead-accounting mode does not provide the expected output. [PR1329212](#)
- When an AMS bundle has a single MAMs added to it, the subinterfaces do not recover after the subinterface has been disabled. [PR1329498](#)
- Host-outbound traffic is not rewriting IEEE-801.pbits for a dynamic subscriber IFL over a PS interface. [PR1329555](#)
- SNMP walks of Interfaces related MIB objects are slower than expected in a scaled configuration. [PR1329931](#)
- The **show services nat mappings address-pooling-paired** command times out and fails. [PR1330207](#)
- The **Too many supplies missing in Lower/Upper zone** alarm flaps (set/clear) every 20 seconds if a zone does not have the minimum required PSMs. [PR1330720](#)
- The packets might be dropped if one route is adverted by BGP, where the session is established through the subscriber interface. [PR1330737](#)
- The rpd process generates core files on the new backup Routing Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler** after disabling NSR+GRES [PR1330750](#)
- The FPC might be wedged when the LSQ interface receives fragmented packets. [PR1330998](#)
- Under very high scale, replication is not started for BGP and is stuck in progress for RIP and LDP after a NSR. [PR1331145](#)
- Chassis FPC temperature with non-NEBS optics is higher after a software upgrade. [PR1331186](#)
- The bbe-smgd process might crash after executing the **clear ancp access-loop circuit-id <circuit id of interface set>** command. [PR1332096](#)
- Inaccurate Jflow records might be seen for an output interface and a next hop. [PR1332666](#)
- On an MX150 platform, the **set chassis alarm management-ethernet link-down ignore** command is not ignoring the alarm for the FPC Mgt 0 interface. [PR1332799](#)
- The subinfo process might crash and it might cause the PPPoE subscribers to get disconnected. [PR1333265](#)

- JDID thrashes continuously and continuous log messages are observed in syslog. [PR1333632](#)
- Active/active (A/A) Multihoming EVPN VXLAN in some race conditions can trigger constant high CPU usage on the backup Routing Engine. [PR1334235](#)
- Two subscribers cannot reach the online state at the same time if they have an identical Frame-Route attribute value. [PR1334311](#)
- MPC8E or MPC9E reports high temperature alarms and fan speed changing continuously through full and normal speed iterations. [PR1334750](#)
- The rpd process crashes when performing the BGP configuration change. [PR1334846](#)
- The UID limit is reached in a large-scale subscriber scenario. [PR1334886](#)
- When using the **show subscribers** command and when the FPC number has two digits, the interface and IPv6 address get connected together for DHCPv6 PD. [PR1334904](#)
- The IPsec rule might not work if both IPv4 ANY-ANY term and IPv6 ANY-ANY term are configured for it. [PR1334966](#)
- Traffic drops on the MX Series LNS because of software error/unknown family exception when traffic goes to or comes from an MLPPP subscriber if 'routing-services' is presented in the dynamic-profile used by this subscriber. [PR1335276](#)
- The master LED glows on the master and the backup RCB, while performing the image upgrade on the master with GRES/NSR enabled. [PR1335514](#)
- There are hitless key chain rollover feature limitations on MIC-MACSEC-MRATE. [PR1335644](#)
- The RIP route updates might be partially dropped when NSR is enabled. [PR1335646](#)
- The MAC_STUCK might be seen on the MS-MPC or the MS-MIC. [PR1335956](#)
- JET application might not respawn after a normal exit. [PR1336107](#)
- Subscriber might experience SDB DOWN event and drop the clients' connections when issuing the **show subscribers** commands. [PR1336388](#)
- On an MX2000 with an SFB card installed, high traffic volume on an MPC7E, MPC8E or MPC9E might cause traffic drops with cell underflow messages. [PR1336446](#)
- The bbe-smgd might crash when doing a CoS configure of the interface set. [PR1336852](#)
- The **set protocols lldp neighbour-port-info-display port-id** command might not take effect. [PR1336946](#)
- The error log message **sdb_db_interface_remove: del ifl:si-<index> with licnese cnt non zero on** can be seen on LTS during a subscriber logout. [PR1337000](#)
- AI-script does not get an auto reinstall upon a Junos OS upgrade on a next-generation Routing Engine. [PR1337028](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- The MX10003 MPC offline button is not effective. [PR1340264](#)

- The CLI shows CB states online after pressing RCB offline button for 4 seconds or more. [PR1340431](#)
- Upon a reboot from a cold state (or after a Junos OS software upgrade), MX150 might not forward multicast traffic, including VRRP packets, from the Packet Forwarding Engine to the Routing Engine. [PR1341044](#)
- There might be traffic loss on some subscriber sessions when more than 32,000 L2TP subscriber sessions are anchored in the ASI interface. [PR1341659](#)
- The reboot of the Routing Engine might occur if the PPPoE interface is configured over an aggregated Ethernet or RETH interface. [PR1341968](#)
- With discard Interfaces (configured with IGMPv3), the KRT queue gets stuck while deleting a multicast next hop (MCNH) with the error **EPERM -- Jtree walk in progress**. [PR1342032](#)
- An SNMP walk might fail for LLDP-related OIDs. [PR1342741](#)
- The vFPC might get absent resulting in the total loss of traffic. [PR1343170](#)
- Support required for the **show system resource-monitor subscribers-limit chassis extensive** command on Summit. [PR1343853](#)
- An MX Series router is sending IPv6 RA and the DHCPv6 advertisements before IPCv6 ACK from CPE. [PR1344472](#)
- Unable to route over an RLT interface after upgrading from Junos OS Release 15.1 to Release 17.3. [PR1344503](#)
- The ancpd process might generate a core file when clearing ancp subscribers in a scaled scenario when enhanced-ip is configured. [PR1344805](#)
- The Framed-Route "0.0.0.0/0" will not be installed on an MX Series platform with Junos OS enhanced subscriber management releases. [PR1344988](#)
- The ARP packet uses the VRRP/virtual-gateway MAC address in the Ethernet header instead of the IRB MAC address. [PR1344990](#)
- A dot1x re-authentication issue. [PR1345365](#)
- The rpd process crash might be seen if **no-propagate-ttl** is set in a routing instance that has a specific route. [PR1345477](#)
- The MAC address of multiple interfaces are found to be duplicates. [PR1345882](#)
- The Routing Engine model changed from JNP10003-RE1 to RE-S-1600x8. [PR1346054](#)
- New PPPoE users might fail to log in. [PR1346226](#)
- The AC system error counter in the **show pppoe statistics** command is not working. [PR1346231](#)
- The VCCP-ADJDOWN detection is delayed on the Virtual Chassis backup router (VC-Bm) when deleting one VCP link on Virtual Chassis master router (VC-Mm). [PR1346328](#)
- Statistics daemon PFED might generate a core file on an upgrade between certain releases. [PR1346925](#)

- The twice-napt-44 sessions are not syncing to the backup SDG with stateful sync configured. [PR1347086](#)
- IPv6 MAC resolve will fail if the DHCPv6 client uses a non-EUI64 link-local address. [PR1347173](#)
- Remove libstdc++ dependency on the hypervisor to install the JDM rpm/deb package. [PR1347921](#)
- There is an issue with handling the community_action ("add") in a RPC call. [PR1348082](#)
- The FPC might crash due to a MIC error interrupt hogging. [PR1348107](#)
- Packet loop is detected when virtual routing and forwarding (VRF) multipath is enabled with **equal-external-internal** under an Layer 3 VPN instance and **install-nexthop** is enabled in a forwarding-table export policy regarding that Layer 3 VPN route. [PR1348175](#)
- A chassisd memory leak is observed on an MX10003 and an MX204 platform and it would eventually cause a Routing Engine switchover and crash. [PR1348753](#)
- The DHCPv6 solicit packet might be dropped on an MX Series Virtual Chassis with L2TP LNS when the packet is received over a VCP port and the anchor si- interfaces exist on the same Packet Forwarding Engine as the VCP port. [PR1348846](#)
- The **Major PEM 0 Input Failure** major alarm might be observed for a DC PEM. [PR1349179](#)
- The mspmand process might crash when executing the **show services nat deterministic-nat nat-port-block** command. [PR1349228](#)
- The mgd process generates a core file because of an issue in the nsindb infra. [PR1349288](#)
- The pccd might crash after a delegated LSP is removed in PCEP scenario. [PR1350240](#)
- The MTU value for subscriber's interface might be programmed incorrectly if **routing-services** or **protocol pim** is configured in dynamic-profile. [PR1350535](#)
- The subinfo process might crash when executing the **show subscribers address <> extensive** command for a DHCP IPv6 address. [PR1350883](#)
- The VCP port might not come back up after removing and adding it again. [PR1350845](#)
- The PPE Errors async xtxn error is observed when FPC is restarted or removed. [PR1350909](#)
- The pfd process might consume high CPU if subscriber or interface statistics are used at large scale. [PR1351203](#)
- A high CPU usage for the bbe-smgd process might be seen when L2BSA subscribers get stuck. [PR1351696](#)
- After GRES, the BGP neighbors at the master Routing Engine might reset and the BGP neighbors at the backup Routing Engine might take a long time to establish. [PR1351705](#)
- The bbe-smgd process might restart in a subscriber environment. [PR1352546](#)
- The DHCP relay-reply packets are dropped in the DHCPv6 relay scenario. [PR1352613](#)
- The offlining of MIC6-100G-CFP2 MIC through the CLI command might trigger the FPC card to crash. [PR1352921](#)
- The rpd process is permanently overusing CPU due to a logical system configuration commit. [PR1353548](#)

- Traffic interruption is observed after multiple Routing Engine switchover. [PR1354002](#)
- The dfw_bbe_filter_bind:1125 BBE filter bind type 0x84 index 167806251 returned 1. [PR1354435](#)
- The rpd might generate core files when adding an inter-region template in routing-instances. [PR1354629](#)
- Aggregated Ethernet operational state goes up even though some of the member interfaces configured under the Aggregated Ethernet are down. [PR1354686](#)
- The ifinfo process might crash in an MX BNG running an L2BSA service. [PR1354712](#)
- JSSCD static-subscribers do not properly update firewall information on the Packet Forwarding Engine when dynamic configuration changes are made to active subscribers. [PR1354774](#)
- A memory leak is found in agentd while running valgrind. [PR1354922](#)
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100Gbps). [PR1355168](#)
- Packets destined to Routing Engine might be dropped in the kernel when LACP is configured. [PR1355299](#)
- The fabric chip failure alarms are observed in a GRES scenario. [PR1355463](#)
- Syslog messages : **ui_client_connect_to_kmd_instance: KMD-SHOW connect to kmd-instance failed kmd-instance RE, fpc slot 0, pic slot 0.** [PR1355547](#)
- The flex-flow-sizing is not working on an MX204. [PR1356072](#)
- The rpd process will crash when issuing the **show dynamic-tunnels database terse** command for RSVP automatic mesh tunnels. [PR1356254](#)
- The L2C messages from PEM/PSM are reported if SNMP is enabled. [PR1356259](#)
- The **show pppoe underlying-interfaces** command in a scaled environment might cause a bbe-smgd memory leak. [PR1356428](#)
- The bbe-smgd generates core files in recursive loop between functions bbe_autoconf_if_l2_input and bbe_if_l3_input. [PR1356474](#)
- DHCP subscribers fail after a reconfiguration of the port from tagged to un-tagged mode. [PR1356980](#)
- Upgrading from Junos OS Release 15.1F2-S20 to Junos OS Release 15.1X12 using **validate** throws a Fabric Mixed Mode error. [PR1357423](#)
- A Routing Engine switchover during backup Routing Engine being not GRES ready might cause linecard restart, which causes the Routing Engine kernel to crash and multiple chassisd crashes. [PR1357427](#)
- Traffic might be sent to a wrong RLT member interface after RLT switchover. [PR1358320](#)
- An incorrect traffic load balance might be seen even if **locality-bias** is configured on MX Series Virtual Chassis. [PR1358635](#)
- FPC was offline with the **Disconnected after ISSU and before switchover** message during a unified ISSU from Junos OS Release 17.4 to Junos OS Release 18.2. [PR1359282](#)

- The **FRU-model-number** is not displayed for a few FRUs in the component sensor for an MX10008 and an MX10003. [PR1359300](#)
- The IPv6 subscriber might fail to access network. [PR1359520](#)
- The rpd cores at **Assertion failed rpd[10169]: file**
`"../../../../../../../../src/junos/usr/sbin/rpd/lib/rt/rt_attrib.c, line 3329: rt_template_get_rtn_ngw(nhp)`
`<= 1` on doing Routing Engine switchover with SRTE routes. [PR1360354](#)
- The rpd scheduler slip might be seen when frequently deleting, modifying, and adding groups which are applied on top level. [PR1361304](#)
- Spontaneous bbe-smgd core file might be seen on the backup Routing Engine. [PR1362188](#)
- The route stuck might be seen after BGP neighbor and route flapping. [PR1362560](#)
- Unexpected DCD_PARSE_ERROR_SCHEDULER messages are logged when MS-MPC/MS-MIC is brought offline or online. [PR1362734](#)
- A quick memory leak in bbe-smgd is observed if the dynamic profile variable name and the default associated value are configured to be the same. [PR1362810](#)
- The non-default routing-instance is not supported correctly for NTP packet in subscriber scenario. [PR1363034](#)
- Traffic destined to the MAC or IP address of VRRP VIP gets dropped on the platforms which have common TFEB terminals such as MX5/10/40/80/104. [PR1363492](#)
- A **pmbus_read_volt: sfb-07 - MAX20751-PF1-0.9v: pmbus** read failed for cmd 0x8b. [PR1363587](#)
- The xmlproxyd for internal interfaces is reporting uint32 instead of uint64. [PR1363766](#)
- The l2circuit on MPC7E/8E/9E with asynchronous-notification and ccc configured might keep flapping when the circuit is going up. [PR1363773](#)
- A traffic loop might occur even though that port is blocked by RSTP in a ring topology. [PR1364406](#)
- The traffic is still forwarded through the member link of an Aggregated Ethernet bundle interface even with **Link-Layer-Down** flag set. [PR1365263](#)
- Midplane attributes are not getting exported. [PR1365303](#)
- The next-hop of MPLS path might be stuck in hold state which might cause traffic loss. [PR1366562](#)
- Snmp mib walk for udp flood gives different output statistics than CLI. [PR1366768](#)
- The **show system resource-monitor fpc** might show non-existing Packet Forwarding Engine. [PR1367534](#)
- The **commit** or **commit check** might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)
- Subscriber filter not removed from the Packet Forwarding Engine when routing-services are enabled in the dynamic profile on an L2TP LNS. [PR1369968](#)
- Kernel crash might be seen after committing DEMUX related configuration. [PR1370015](#)

- The packet which size exceeds 8000 might be dropped by MS-MPC in ALG scenario. [PR1370582](#)
- FPC high CPU utilization or crash during hot-banking condition. [PR1372193](#)
- PCE initiated LSPs remain **Control status became local** after removing PCE configuration. [PR1374596](#)

High Availability (HA) and Resiliency

- After server links flap, the GNFs associated with the ports on the Control Board show the status message: **Switchover Status: Not Ready** message. [PR1306395](#)
- The ksyncd process might crash continuously on the new backup Routing Engine after performing GRES. [PR1329276](#)
- There is insufficient available space on the hard disk lead by the crashinfo files that are generated by the ksyncd process when GRES is configured in a large-scale configuration scenario. [PR1332791](#)
- VC-Bm cannot sync with VC-Mm when the the Virtual Chassis splits then reforms. [PR1361617](#)

Infrastructure

- The syscalltrace.sh might create a huge output file, which might cause the router to run out of storage space. [PR1306986](#)
- A cleanup at the thread exit is causing memory leaks. [PR1328273](#)
- On all Junos OS platforms, on a port configured with both dot1x static mac by-pass and normal authentication, the hosts configured for static mac by-pass may not be able to send traffic. [PR1335125](#)
- The kernel might crash and the system might reboot in an SNMP query reply scenario. [PR1351568](#)
- Junos OS is no longer going to database prompt at ~ +Ctrl+b. [PR1352217](#)

Interfaces and Chassis

- RL-dropped packets are not displayed by **show interfaces <ifl> detail/extensive** commands. [PR1249164](#)
- Out of sequence packets seen with LSQ interface. [PR1258258](#)
- L2TP subscribers might not be cleared if the access-internal routes fail to install. [PR1298160](#)
- Some CFM sessions do not come up after a DUT with MPC9 line cards is rebooted with scale configuration. [PR1300515](#)
- The MPC CPU might reach 100 percent when optical transport network (OTP) ultra forward error correction (UFEC) is configured. [PR1311154](#)
- Observing jpppd core **telemetry_start_timer,mosquitto_handle_connack,telemetry_mqtt_publisher** [PR1311396](#)
- The jpppd process generates a core file at **telemetry_start_timer,mosquitto_handle_connack,telemetry_mqtt_publisher**. [PR1311396](#)
- The ifinfo process might crash and generate core files when executing the **show interfaces name** command with a name greater than 128 characters. [PR1313827](#)

- The MX Series Virtual Chassis unified ISSU emits a benign error message if unsupported FRUs are present. [PR1316374](#)
- There is no route to an IP address from the directly connected route. [PR1318282](#)
- The **show interfaces interface-set** command is displaying wrong logical interface. [PR1319682](#)
- The IPv6 framed Interface ID field (from the **show subscribers extensive** command output) is not properly matching the negotiated one. [PR1321392](#)
- IPCP negotiation might fail for dual stack PPPoE subscribers. [PR1321513](#)
- Unexpected log messages might be seen if a BGP session flaps in a dynamic-tunnels GRE scenario. [PR1326983](#)
- Unexpected log messages might be seen on a router for a subscriber management scenario. [PR1328251](#)
- Traffic loss might be seen after deleting aggregated Ethernet bundle unit 1. [PR1329294](#)
- The cfmd process generates core files. [PR1329779](#)
- The interface might not work properly after the FPC restarts. [PR1329896](#)
- The dcd process might crash due to a memory leak and cause a commit failure. [PR1331185](#)
- The last IFL digit is sometimes truncated in jpppd trace logs. [PR1332483](#)
- The transportd process might crash when you run an snmp query on the jnxoptIfOChSinkCurrentExtTable with an unsupported interface index. [PR1335438](#)
- The MX Series router might occasionally drop the first LCP configure request packet when operating in PPPoE subscriber management configuration. [PR1338516](#)
- The 100G DWDM interface might be going down for 15 seconds after a loss of signal event. [PR1343535](#)
- When eth-oam is deactivated with a scale PM configuration (under hardware-assited-pm-mode), the FPC might become unstable and generate core files. [PR1347250](#)
- Suppressing cfmd logs : **jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0**. [PR1347650](#)
- The jpppd process generates core files spontaneously on the backup Routing Engine in a longevity test at **../src/junos/usr/sbin/jpppd/pppMain.cc:400**. [PR1350563](#)
- The VRRP VIP becomes unreachable after deleting one of the logical interfaces. [PR1352741](#)
- The FPC might be stuck at 100 percent for a long time when MC-AE with enhanced-convergence is configured with large-scale logical interfaces. [PR1353397](#)
- The FPC generates a core file related to cfmmman. [PR1358192](#)
- Clients might not get an IPv4 address in a PPPoE dual-stack scenario. [PR1360846](#)
- Approximately 50 percent of PPPoE subscribers (PTA and L2TP) and all ESSM sub lost after post unified ISSU during DT CST stress test. [PR1360870](#)
- On all Junos OS products, the CLI allows to configure more than 2048 sub-interfaces on LAG interface from 17.2R1. [PR1361689](#)

- The EOAM LTM messages might not get forwarded after system reboot in CFM scenario configured with CCC interface. [PR1369085](#)
- Subscriber cannot negotiate MLPPP session with MX LNS when dynamic-profile name contains more than 30 characters. [PR1370610](#)

Layer 2 Features

- The rpd process memory leak is observed upon any changes in a VPLS configuration such as deleting or re-adding VPLS interfaces. [PR1335914](#)
- The VPLS instance stays in NP state after the LDP session flaps. [PR1354784](#)
- The Routing Engine kernel might crash when OSPFv3 is configured with an IPsec key authentication over an IRB interface. [PR1357430](#)

Layer 2 Ethernet Services

- The MAC address might not be learnt due to spanning-tree state discarding in kernel table after a Routing Engine switchover. [PR1205373](#)
- The MX Series platforms might display a false positive CB alarm **PMBus Device Fail**. [PR1298612](#)
- DHCP IPv6 traffic might be dropped in a subscriber scenario. [PR1316274](#)
- The jdncpd process generates core files after making DHCP configuration changes. [PR1324800](#)
- The on-demand-address-allocation under dual-stack-group does not work for IPv6. [PR1327681](#)
- The snmpget for OID: dot3adInterfaceName might not work. [PR1329725](#)
- A memory leak might happen in l2cpd if the l2-learning process is disabled. [PR1336720](#)
- The DHCPv6 second Solicit message might not be processed when IA_NA and IA_PD are sent in a separate Solicit message. [PR1340614](#)
- DHCP client is not able to connect if VLAN is modified on the aggregate Ethernet interface associated with the IRB. [PR1347115](#)
- ZTP infra scripts are not included for MX PPC routers. [PR1349249](#)
- When DHCP subscribers are in an bound (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state if dhcp-service is restarted then the subscribers in this state are logged out. [PR1350710](#)
- The DHCP relay agent will discard a DHCP request message silently if the requested IP address has been allocated to the other client. [PR1353471](#)
- Restarting an FPC that hosts the micro-BFD link might cause LACP to generate a core file. [PR1353597](#)
- DHCPv6 relay ignores replies from server when renewing. [PR1354212](#)
- The DHCP lease query message is replied with incorrect source address. [PR1367485](#)
- DHCP Relay Binding state - rebinding state counter added to dhcpv4 and dhcpv6 binding sensors. [PR1368392](#)

MPLS

- When minimum-bandwidth and bandwidth commands are present in the configuration, the bandwidth selection of the lsp is inconsistent. [PR1142443](#)
- Ingress RSVP LSP fails to come up after issuing the **clear rsvp lsp all** command on the egress router. [PR1275563](#)
- The rpd might crash in an LDP Layer 2 circuit scenario. [PR1275766](#)
- LDP egress policy not advertising label for inet.3 BGP labeled-unicast route. [PR1289860](#)
- Traffic drop is observed during an NSR switchover for RSVP P2MP provider tunnels used by MVPN. [PR1293014](#)
- The traffic in P2MP tunnel might be lost when NG-MVPN uses RSVP-TE. [PR1299580](#)
- The rpd process might crash in rare conditions where **traffic-engineering** is configured. [PR1303239](#)
- The RSVP node-hello packet might not work correctly after the next hop for a remote destination is changed. [PR1306930](#)
- The kysncd process might crash after removing and inserting backup RE in analytics and "mpls sensor" scenario. [PR1303491](#)
- The RSVP node-hello packet might not work correctly after the next-hop for remote destination is changed. [PR1306930](#)
- The rpd process might crash if LDP updates the label for a BGP route. [PR1312117](#)
- The output of the **show mpls container-lsp** command is delayed. [PR1314960](#)
- An RSVP node-neighbor is found even when node-hello has been disabled. [PR1317241](#)
- The IPv4/IPv6 multicast traffic might get dropped in an MX Series Virtual Chassis scenario when the traffic comes in through an Layer 2 circuit and goes out through an aggregated Ethernet member interface across Virtual Chassis members. [PR1320742](#)
- The rpd might crash when LDP P2MP recursive is configured. [PR1321626](#)
- The rpd might crash due to a memory leak in an RSVP scenario. [PR1321952](#)
- Receipt of specially crafted UDP packets over MPLS may bypass stateless IP firewall rules. [PR1326402](#)
- SNMP OID counters for mplsLspInfoAggrOctets show constant value for some LSPs even though traffic is constantly increasing in **show mpls lsp statistics**. [PR1327350](#)
- In Junos OS Release 17.2X75-D40, a new feature related to "per AE member OAM" introduced additional processing on pfeman thread during link flaps. [PR1327988](#)
- Packet loss might be observed when **auto-bandwidth** is enabled for CCC connections. [PR1328129](#)
- The rpd might crash on the backup Routing Engine due to memory exhaustion. [PR1328974](#)
- Fate-sharing group cost does not re-set to the default value after a CLI change, removing explicit cost configuration. [PR1330161](#)

- After a MPLS LSP link flap and local repair, a new LSP instance is tried to be signaled but it may get stuck. [PR1338559](#)
- Whenever there is a decrease in the stats value across an LSP, the `mplsLspInfoAggrOctets` value takes two intervals to get updated. [PR1342486](#)
- An LDP label is generated for a serial interface subnet route unexpectedly. [PR1346541](#)
- The MPLS LSP does not come up after changing admin-group mapping. [PR1348208](#)
- The rpd crash might happen in an RSVP setup-protection scenario. [PR1349036](#)
- In a very rare scenario, the rpd might crash when LDP failed to allocate a self-ID for the P2MP FEC. [PR1349224](#)
- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. [PR1359956](#)
- Layer 2 Circuit might flap after an interface goes down even if the LDP session stays up when `l2-smart-policy` is configured. [PR1360255](#)
- The process rpd might crash during P2MP LSPs churn. [PR1363408](#)
- The rpd process might crash after RSVP is deactivated and then re-activated globally for multi times. [PR1366243](#)
- The rpd might crash in BGP LU and LDP scenario. [PR1366920](#)

Multicast

- DHCP6 Relay is not working unless DHCP is restarted. [PR1316210](#)
- Multicast traffic is not forwarded on the newly added P2MP branch or receiver. [PR1317542](#)
- Some IGMP groups might have wrong upstream interface due to discard route is installed in PIM. [PR1337591](#)

Network Management and Monitoring

- The syslog might generate duplicate entries of hostname and timestamp. [PR1304160](#)
- The mib2d might crash when SNMP polling occurs on interface mibs and while the FPC restarts or the interface flaps. [PR1318302](#)
- SNMP stops or becomes very slow after a very long period of time. [PR1328455](#)
- With interface-mib, the MX Series router is responding with **type : NoSuchInstance** for OIDs when multiple OIDs are polled in one SNMPGET request. [PR1329749](#)
- The eventd process fails to start up with the syslog configuration. [PR1353364](#)
- The `jnxDcuStatsEntry` and `jnxScuStatsEntry` OIDs are missing in a post interface configuration change. [PR1354060](#)
- The SNMP process crashes during polling the CFM stats. [PR1364001](#)

Platform and Infrastructure

- On MX Series routers, if a large number of routes are processed, then the Packet Forwarding Engine of the MS-MPC might crash. [PR1277264](#)
- Executing the **show services inline ip-reassembly statistics** command might cause a ukern sheaf memory leak. [PR1285833](#)
- The **apply-path** prefix is not inherited under policy after modifying the interface address. [PR1286987](#)
- The output values of command **show system resource-monitor** are not accurate. [PR1287592](#)
- The **interface-mac-limit** might fail for an aggregated Ethernet interface. [PR1303293](#)
- The source MACs might leak (or not learn) between different VPLS instances at the receiving end of VPLS PE devices. [PR1306293](#)
- An rpm probe with a probe interval of 1 second fails on MX Series routers. [PR1308952](#)
- Error messages are not observed during telnet with a username longer than an acceptable limit. [PR1312265](#)
- The mgd process might crash and a session gets terminated after the load override from netconf. [PR1313158](#)
- The issue addresses the ICMP error messages in the Packet Forwarding Engine and is forwarded to the correct pic in the AMS bundle. [PR1313668](#)
- VPLS instance fails to learn MAC addresses upon pseudowire switchover. [PR1316459](#)
- Rate-limit configured with a small temporal buffer size might cause packet loss. [PR1317385](#)
- Multicast traffic might get duplicated when MoFRR is configured. [PR1318129](#)
- The GNF FPC hangs at reboot during a unified ISSU. [PR1318394](#)
- The default severity of the correctable ECC errors on MX Series routers with MPC2E NG Q, MPC3E NG Q, or MPC5E has been changed from fatal to major. [PR1320585](#)
- Errors might be observed when the **fabric-header-crc-enable** feature is enabled. [PR1320874](#)
- The traffic with more than 2 VLAN tags might be incorrectly rewritten and sent out. [PR1321122](#)
- The RPM probes delegated to MS-MIC get stuck when any change is made to the BGP group statement. [PR1322097](#)
- The **no-propagate-ttl** option might not take effect if **chained-composite-next-hop ingress l3vpn extended-space** is configured. [PR1323160](#)
- The MAC might not be learned on MX Series routers with MPCs or MIC-based line cards due to the negative value of the bridge MAC table limit counter. [PR1327723](#)
- The packet might get dropped in an LSR if MPLS pseudowire payload does not have a control word and its destination MAC starts with '4'. [PR1327724](#)
- Traffic loss might be observed on the LT interface. [PR1328371](#)

- Directories and files under `/var/db/scripts` lose execution permission or directory 'jet' is missing under `/var/db/scripts` causing an **error: Invalid directory: No such file or directory** error during commit. [PR1328570](#)
- The tcpdump filter might not work in the egress direction on PS and LT logical interfaces. [PR1329665](#)
- The router hits the database prompt at `netisr_process_workstream_proto`. [PR1332153](#)
- RPM MIB's pingResultsMinRtt, pingResultsMaxRtt, and pingResultsAverageRtt response is "1" while target address is unreachable, it should be "0". [PR1333320](#)
- Traffic loss might be seen for some flows due to network churn. [PR1335302](#)
- Commit might fail with error reading from commit script handler **error: commit script failure**. [PR1335349](#)
- The MPC might crash after setting **max-queues** to a very large number. [PR1338845](#)
- Route corruption occurs in the Packet Forwarding Engine with CFM enabled on the aggregated Ethernet interface. [PR1338854](#)
- Configuring the same DHCP server in different routing instances is not supported in a DHCP relay scenario. [PR1342019](#)
- Commit error is observed when configuring the same VLAN ID on different logical interfaces of the same LT physical interface and the **ethernet-bridge** encapsulation is configured. [PR1342229](#)
- Route corruption in the Packet Forwarding Engine with connectivity-fault-management is enabled for l2ckt. [PR1342881](#)
- ZTP is not supported for vmhost images on next-generation Routing Engines on the MX Series platforms. [PR1343338](#)
- The IPv4 GPRS traffic over the aggregated Ethernet interface might be dropped if gtp-tunnel-endpoint-identifier is configured. [PR1347435](#)
- Output policing action does not work on IRB interfaces for VNIs. [PR1348089](#)
- FPC CPU utilization with LT interfaces is pegged continuously at 100 percent. [PR1348840](#)
- Running RSI through the console port might cause a system crash and reboot. [PR1349332](#)
- The ICMP error messages are not generated if 'don't fragment' packets exceed the MTU of the multiservice interface. [PR1349503](#)
- When viewing IPv6 addresses, **display rfc5952** does not work when combined with **display set**. [PR1349949](#)
- The chassisd process memory leak is observed. [PR1353111](#)
- The kernel crashes because the initialization of the logical Interface MAC filter function is missing for Packet Forwarding Engine extended port devices. [PR1353498](#)
- The FPC might crash due to the memory leak caused by the VTEP traffic. [PR1356279](#)

- Traffic is discarded silently along with `JPRDS_NH:jprds_nh_alloc()`, 651: JNH[0] failed to grab new region for NH messages. [PR1357707](#)
- When forwarding-class-accounting knob is enabled, on an interface, inside of a routing-instance of instance-type vrf, aggregate input forwarding-class statistics do not increment (egress statistics work fine). [PR1357965](#)
- Select CLI functions are not triggering properly (set security ssh-known-hosts load-key-file, set system master-password). [PR1363475](#)
- Same vlan-id not allowed on multiple IFLs of the same GR interface. [PR1365640](#)
- Subscribers over AE interface might have tail drops which will affect the fragmented packets due to QXCHIP buffer getting filled up. [PR1368414](#)
- The logical tunnel interface might be unable to send out control packets generated by RE. [PR1372738](#)

Routing Policy and Firewall Filters

- Condition based policy fails to take action even though condition is matched [PR1300989](#)
- The policy configuration might not be evaluated if the policy expression is changed. [PR1317132](#)
- Access-internal route might fail to be leaked between routing instances when **from instance** is configured in the policy. [PR1339689](#)
- The policy might not clean up after deleting configuration and cause the rpd to generate a core file. [PR1357724](#)

Routing Protocols

- The **show bgp summary** results are incorrect while assisting GR. [PR1045151](#)
- BGP extended communities with sub-type 4 erroneously displayed at LINK_BANDWIDTH. [PR1216696](#)
- The rpd generates core files in the ASBR when BGP is deactivated in the ASBR before all stale labels have been cleaned up. [PR1233893](#)
- The rpd might crash after deactivating or activating BGP. [PR1272202](#)
- After a bfdd restart, the issue is seen with a next-generation MVPN and Layer 2 VPN route exchange causing MVPN and VPLS traffic drop. [PR1278153](#)
- Routing loops might be seen after configuring BGP Prefix Independent Convergence (BGP PIC). [PR1282520](#)
- Few adj-sid details are not updated in an IS-IS database with a LAN + adjset scenario. [PR1288331](#)
- Multihop BFD sessions flap continuously. [PR1291340](#)
- The Impd crashes repeatedly when a logical system is configured on the same device. [PR1294166](#)
- The rpd process might crash because of the AS PATH check error that occurs when RIB groups are added first and later the routing instances are added. [PR1298262](#)

- MSDP sessions might flap when NSR or GRES is enabled. [PR1298609](#)
- While the device is booting up with the Junos OS Release 17.4R1 image, **error: channel 0: chan_shutdown_read: shutdown() failed for fd 10 [i0 o3]: Socket is not connected** messages might show up. [PR1300409](#)
- IBGP route damping is not taking effect on an IBGP inet-vpn address family. [PR1301519](#)
- Observed mcsnoopd core file at `__raise,abort,__task_quit,__task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal (enable_slip_detector=true, no_exit=true)` at `../src/junos/lib/libjtask/base/task_scheduler.c:275`. [PR1305239](#)
- BGP traceoption logs are still written when it is deactivated. [PR1307690](#)
- The rpd might generate a core file in `bgp_rt_send_message` at `../src/junos/usr/sbin/rpd/bgp/bgp_io.c:1460`. [PR1310751](#)
- The BGP session might flap when the connection between the master Routing Engine and the backup Routing Engine keeps flapping with NSR configured. [PR1311224](#)
- The rpd might crash when the neighbor IS-ISv6 router is restarted, causing a route churn. [PR1312325](#)
- Unexpected route age refresh might be observed if BGP PIC is configured. [PR1312538](#)
- The IS-IS SPF might be triggered by LSP updates containing changes only in reservable bandwidth in a TE extension. [PR1313147](#)
- The rpd might crash and generate a core file with distributed IGMP. [PR1314679](#)
- The rpd might constantly consume a high percentage of CPU in a BGP setup. [PR1315066](#)
- On a chassis with BMP configured, the rpd might crash when the rpd process is gracefully terminated. [PR1315798](#)
- The primary path of an MPLS LSP might switch to another address. [PR1316861](#)
- If a loop free alternative is configured, an Isdb entry cleanup might cause the rpd to crash. [PR1317023](#)
- The inactive route cannot be installed in a multipath next hop after disabling and enabling the next hop interface in an Layer 3 VPN scenario. [PR1317623](#)
- A BGP-LU update oscillates with a BGP-PIC. [PR1318093](#)
- IS-IS might choose a suboptimal path after the metric change in ECMP links. [PR1319338](#)
- Traffic might get discarded temporarily when BGP GR is triggered and the direct interface flaps. [PR1319631](#)
- There is an issue with tracing of the BGP Layer 2 VPN DF election community. [PR1323596](#)
- The rpd crash is seen when deactivating the static route if the next-hop interface is type P2P. [PR1323601](#)
- When the prefix limit is reached, increasing maximum-prefixes does not take effect. [PR1323765](#)

- The rpd process might crash continuously on both Routing Engines when **backup-spf-options remote-backup-calculation** is configured in the IS-IS protocol. [PR1326899](#)
- Multiple next hops might not be installed for an IBGP multipath route after an IGP route update. [PR1327904](#)
- With BGP/LDP/IS-IS configurations, deleted IS-IS routes might still be visible in the RIB. [PR1329013](#)
- The rpd might crash on the backup Routing Engine after BGP peer is deleted. [PR1329932](#)
- Manual GRES with an MX Series Virtual Chassis results in some packet loss on core facing interfaces. [PR1329986](#)
- The conditional route policy cannot withdraw all routes in a BGP add-path scenario. [PR1331615](#)
- LDP route in inet.3 is missing when both OSPF rLFA and LFA protections are available and rejected by the backup selection policy. [PR1333198](#)
- Discard next hop being installed when the primary LSP interface drops. When primary interface returns, discard next hop remains until BGP LU neighbor is cleared. This only impacts the cloned route (S=0). [PR1333570](#)
- For Junos OS Release 15.1 and later, IGMP joins are not processed with the **passive allow-receive** command configured on the IGMP interface. [PR1334913](#)
- BGP sessions get stuck in an active state after the remote end restarts the device. [PR1335319](#)
- The rpd crash might occur when receiving BGP updates. [PR1341336](#)
- Changes to the displayed value of AIGP in the **show route ... extensive** command. [PR1342139](#)
- Traffic black hole might be seen if a local device is receiving BFD-down. [PR1342328](#)
- The rpd might crash when BGP flaps. [PR1342481](#)
- The rpd generates a core file while running streaming telemetry test. [PR1347431](#)
- The rpd might crash if a route for RPF uses a qualified-next-hop. [PR1348550](#)
- The rpd might crash while restart routing or deactivate IS-IS. [PR1348607](#)
- The rpd might crash when the BGP route damping and the BGP multipath feature are configured. [PR1350941](#)
- Source-as community is not appended to the rendezvous point. The display issue is in the **show route** detail output. [PR1353210](#)
- Static Route flaps on commit when configured with resolve statement. [PR1366940](#)

Services Applications

- PCP mappings cannot be manually cleared when a NAT pool is shared between PCP and standard NAT. [PR1284261](#)
- The L2TP subscribers might get stuck in a terminating state during login. [PR1298175](#)

- LTS clients experience packet drop for large packets due to fragmentation in LTS. [PR1312691](#)
- AVP 145 is not present in IRQ when *ANCP DSL-type* = 0. [PR1313093](#)
- L2TP tunnel Tx and Rx byte count sometimes decrease when subscriber sessions are reduced within the tunnel. [PR1318133](#)
- SNMP MIBs are not yielding data related to sp-interfaces. [PR1318339](#)
- The MRU might be changed to 1492 instead of the default 1500 in an L2TP scenario. [PR1319252](#)
- IPCP active mode is not getting enabled for MLPPP on LNS. [PR1319580](#)
- Long route remains in forwarding table after subscriber session goes down. [PR1322197](#)
- The L2TP LTS might drop the first CHAP success packet from LNS due to delayed programming of /136 route on the Packet Forwarding Engine. [PR1325528](#)
- The jl2tpd might crash if the RADIUS server returns 32 tunnel-server-endpoints. [PR1328792](#)
- A few CSURQ messages might not respond when the number of sessions addressed in CSURQ is more than 107. [PR1330150](#)
- The l2tpd might crash when multiple l2tp related commands are executed together. [PR1337406](#)
- The **show services stateful-firewall flows count** command shows an incorrect flow count after a services configuration change. [PR1338704](#)
- Output of **show interfaces si-x/y/z.xxxxx extensive** CLI command shows an incorrect inet/inet6 MTU value for an MLPPP subscriber on MX Series L2TP LNSs. [PR1346049](#)
- The bbe-smgd process might crash if there are 65,535 L2TP sessions in a single L2TP tunnel. [PR1346715](#)
- Session limit per tunnel on LAC does not work as expected. [PR1348589](#)
- After performing an SNMP walk on the IKE SA that is deleted, IPsec tunnels might go down and an infinite loop scenario might be seen. [PR1348797](#)
- The UDP checksum inserted by an MS-DPC after a NAT64 is not valid when an incoming IPv4 packet has UDP checksum set to 0. [PR1350375](#)
- The **show services stateful-firewall flows counter** command shows high numbers. [PR1351295](#)
- The JI2tpd process might crash shortly after one of the L2TP destinations becomes unavailable. [PR1352716](#)
- L2TP tunnel-switch clients in subscriber session database reference the wrong routing instance. [PR1355396](#)
- In some corner cases, a few tunneled PPPoE subscribers might get stuck in a terminating state. [PR1363194](#)
- The L2TP subscribers might not be able to log in successfully due to the jl2tpd memory leak. [PR1364774](#)
- Actual Data Rate Downstream value not included in the L2TP ICRQ message from the LAC. [PR1370699](#)

Software Installation and Upgrade

- New versions of Junos OS do not have the tool for accessing an aux port - /usr/libexec/interposer. [PR1329843](#)
- Commit might fail in single-user mode [PR1368986](#)

Subscriber Access Management

- A memory leak might happen after clearing a subscriber either with a script or manually. [PR1312517](#)
- Service interim is missing for random users in a JSRC scenario. [PR1315207](#)
- The PPPoE subscribers might encounter a connection failure during login. [PR1317019](#)
- The unified ISSU is allowed to proceed when the account is suspended. [PR1320038](#)
- IP addresses are assigned discontinuously from the linked IP pools. [PR1323829](#)
- Authd considers RADIUS attribute *Framed-IPv6-Prefix* = ::/64 or *Delegated-IPv6-Prefix* = ::/56 as valid parameters. [PR1325576](#)
- An MX204 does not send a **RADIUS Accounting-Off** message. [PR1327822](#)
- Multiple RADIUS servers having a different dynamic-request-port is not supported. [PR1330802](#)
- Subscriber might get stuck in a terminated state when JSRC synchronization state is stuck in a FULL-SYNC in progress state. [PR1337729](#)
- In dual stack subscribers scenario with NDRA pool configured, the linked pools are not used when the first NDRA pool is exhausted. [PR1351765](#)
- When attempting to scale clients saw `sdbsts_lock_holder.bbe-smgd.pid10686.core` core files. [PR1358339](#)

User Interface and Configuration

- There is an increase in commit times. [PR1029477](#)
- The CLI session might die while issuing the **show configuration | compare rollback 1** command. [PR1331716](#)
- The **max-db-size** configuration might not work on some MX platforms. [PR1363048](#)

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the Layer 2 circuit flaps repeatedly. [PR1282875](#)
- Un-hide **set protocols pim mvpn family inet6 disable configuration** to allow users to disable inet6 on MVPN. [PR1317767](#)
- The rpd might crash after a unified ISSU in a large scale scenario with a PIM configuration. [PR1322530](#)
- Moving MC-LAG from LDP based pseudowire to BGP based pseudowire might cause the rpd to crash. [PR1325867](#)

- The multicast might be rejected when Junos OS PE devices received a C-Mcast route from other vendor PE devices. [PR1327439](#)
- MVPN sender-site configuration is not allowed with S-PMSI. [PR1328052](#)
- The rpd generates a core file on the backup Routing Engine with an next-generation MPVPN and NSR configuration. [PR1328246](#)
- The rpd might crash after committing interface related parameters (for example, MTU change, VRF RD or RT, QOS) on the PS interface with vlan-ccc encapsulation and no vlan-id. [PR1329880](#)
- The rpd might continuously crash on the backup Routing Engine and some protocols might flap on the master Routing Engine if hot-standby is configured for Layer 2 circuit or VPLS backup-neighbor. [PR1340474](#)
- The rpd might crash on the backup Routing Engine when changing the Layer 2 circuit **virtual-circuit-id** in an NSR scenario. [PR1345949](#)

Resolved Issues: 17.4R1

Class of Service (CoS)

- The Routing Engine level **scheduler-hierarchy** command misses a forwarding class when the "per-unit-scheduler" mode is configured. [PR1281523](#)

Forwarding and Sampling

- The Sampled process stops collecting data on Routing Engine based sampling supported platforms. [PR1270723](#)
- Firewall filter might not be matched when wildcard (*.*) is specified as the matching condition. [PR1274507](#)
- The sampled route reflector process (srrd) might crash in a large routes churn situation. [PR1284918](#)
- The mib2d process generated a core file @fw_counter_key2components. [PR1286448](#)
- The sampled process might crash and generate a core file if traceoptions are enabled. [PR1289530](#)
- Some accounting files might be missed if the remote archive site is unreachable. [PR1300764](#)
- There is memory leak on mib2d when polling firewall MIBs. [PR1302553](#)
- ACCT_FORK_LIMIT_EXCEEDED log level is ERROR even when backup-on-failure feature is enabled for accounting files. [PR1306846](#)
- The commit might fail if enabling nexthop-learning knob for J-Flow v9. [PR1316349](#)

General Routing

- Enhanced IP/enhanced Ethernet and MS-DPC compatibility. [PR1035484](#)
- Ksyncd might crash due to transient replication errors between Routing Engines. [PR1161487](#)

- On MX240/480/960 platforms, due to a I2C bus hardware issue, error messages might appear. [PR1174001](#)
- SNMP trap sent for **PEM Input failure** alarm. [PR1189641](#)
- Stale VBF states occur without SDB sessions. [PR1204369](#)
- The rpd might crash on the backup Routing Engine after a Routing Engine switchover in MX Series subscriber environment. [PR1206804](#)
- The rpd might crash on platforms with 64-bit X86 RE if IPv6 is configured. [PR1224376](#)
- MPC2E-NG/MPC3E-NG generates a core file with specific MIC due to tight loop of PCI Express critical exceptions. [PR1231167](#)
- The MS-MPC card might crash when OSPFv3 IPv6 traffic goes through it. [PR1233459](#)
- FPCs on MX960 platform might be stuck in offline state with **FPC Incompatible with SCB** due to delayed PEM startup. [PR1235132](#)
- With vLNS (vBNG), a commit generates the message **warning: requires 'l2tp-inline-lns' license** even if a valid license is installed. [PR1235697](#)
- The "multicast-replication" setting cannot be reflected in the redundancy environment after rebooting both Routing Engines. [PR1240524](#)
- In a BGP/MPLS scenario, if the next-hop type of label route is indirect, disabling and enabling the "family mpls" of the next-hop interface might cause the route to go into a dead state. [PR1242589](#)
- XM chip-based line card might drop traffic under high temperature. [PR1244375](#)
- On MX2000 with MPC6E, EOAM LFM adjacency flaps when an unrelated MIC accommodated in the same MPC6E slot is brought online by configuring OAM pdu-interval 100 ms and pdu-threshold 3. [PR1253102](#)
- The "validation-state:unverified" routing entry might not be shown with proper location in show route output. [PR1254675](#)
- The rpd might crash during the next-hop change, if unicast reverse-path- forwarding (uRPF) is used. [PR1258472](#)
- Status LED for the ge-0/0/0 interface does not glow. [PR1259112](#)
- MPC might report a parity error with the **fast-lookup-filter** command configured. [PR1266879](#)
- When ISSU is performed under scaled scenarios where the Packet Forwarding Engine next-hop memory uses more than 4 Million Dwords, PPE traps and traffic loss might be observed during software-sync phase until the end of hardware-sync. [PR1267680](#)
- On MX Series routers, the **show chassis led** command should not be displayed in possible completions of the **show chassis** command. [PR1268848](#)
- A low memory condition putting the Service PIC into the red zone on the MS-MIC or MS-MPC card might cause the SIP ALG to generate a core file. [PR1268891](#)

- The FPC might go offline and the ABB fan might crash after enabling MACsec. [PR1270121](#)
- The mspmand log incorrectly generates messages about memory zone level. This occurs every 49.7 days and will recover by itself. This is a display issue and will not affect traffic. [PR1273901](#)
- CLI commands fail to execute for **show subscribers detail**, **show subscribers extensive**, **show subscribers count client-type <>** and other commands because the subscriber management database is unavailable. [PR1274464](#)
- Link stays down after a flap on MPC next-generation cards with QSFP+-40G direct attach copper (DAC) cable. [PR1275446](#)
- The Packet Forwarding Engine of service DPC might crash with large scale of routes for MX Virtual Chassis. [PR1277264](#)
- Layer 2 control BUS stuck causes SFP+ thread hogging and restarting of MPC. [PR1277467](#)
- Multicast traffic when using iflsets in universal call admission control policy mode does not flow as expected in certain use cases, and bbe-smgd might generate a core file. [PR1278543](#)
- VLAN out-of-band subscriber session fails in autoconfigured mode. The physical interface goes down even if it is physically up. [PR1279612](#)
- After a MS-MPC-PIC is turned offline or online or bounced (because of an AMS configuration change), sometimes the PIC can take approximately 400 seconds to come up. [PR1280336](#)
- **MIC Error code: 0x1b0001** alarm might not be cleared for MIC on MPC7/8/9 when the voltage has returned to normal. [PR1280558](#)
- Authenticated subscriber dynamic VLAN interface might get disconnected immediately after a successful connection. [PR1280990](#)
- jfirmware upgrade support is not available for Routing Engine BIOS. [PR1281050](#)
- The **ingress service-accounting-deferred** command is not providing the correct IP traffic statistics for L2BSA subscribers. [PR1281201](#)
- Establishment of IPsec SAs for link-type tunnels might fail under certain conditions. [PR1281223](#)
- Subscribers might not be able to connect to MX BNG in certain scenarios. [PR1281896](#)
- DHCP/PPPoE subscribers fail to bind after FPC restart and smgd restart with BBE_RTsock_GET_RTsock_IFL_FAIL_TERMINATED counter going up. [PR1281930](#)
- Inline J-Flow unrelated configuration changes related to a routing instance result in invalid or incomplete J-Flow data packets. The **commit full** command resumes proper functionality. [PR1282580](#)
- In a specific CE device environment in which **asynchronous-notification** is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. [PR1282875](#)
- Error messages related to "IFRT: 'IFL'", "IFRT: 'Aggregate interface'" and "IFRT: 'IFD'" are seen on configuration change. [PR1282938](#)
- VBF flows are not programmed correctly on aggregated Ethernet interfaces. [PR1282999](#)

- The MX: **show interfaces** command should display the cause for Intf down when the Packet Forwarding Engine disabled. [PR1283323](#)
- GRE OAM fails to come up when GRE tunnel source and family inet address are the same. [PR1283646](#)
- PPTP session could not be established on MS-MPC when both stateful firewall and NAT were enabled. Also, the address could not be translated. [PR1285207](#)
- The J-Flow data template sequence number is zero for MPLS flows. [PR1285975](#)
- With CoS-based forwarding, when the primary path of one of the next-hop LSPs flaps, traffic carried by the other next-hop LSP could get load-balanced across the primary and secondary path. [PR1285979](#)
- Internal latency increases the overtime for Packet Forwarding Engine sensors with streaming telemetry. [PR1286286](#)
- Unified ISSU is not supported from Junos OS Release 15.1 or later, because the source release includes one or more BBE features such as logical interface (IFL) options, CoS fragmentation map, MLPPP, advisory options, advanced services, and multicast distribution. [PR1286507](#)
- DDS culprit flows are not reported by CLI or logs during login to a MX Series router with a single Packet Forwarding Engine. [PR1286521](#)
- The routing protocol process (rpd) crashes during subscriber login or logout with multicast service enabled while performing GRES switchover. [PR1286653](#)
- Framed routes might get stuck in KRT queue. [PR1286849](#)
- A10NSP interface is not getting attached to the L2 routing instance after the routing instance name is renamed. [PR1287070](#)
- The rpd might generate a core file when the routing-options dynamic-tunnels configuration is changed. [PR1287109](#)
- **Host 0 RTC Battery failure** error messages are seen on PTX1000 and QFX10000-line after upgrading to Junos OS Release 16.1. [PR1287128](#)
- LTS functionality is not working on Junos OS 16.1R4-S2 if the **rewrite-rule** statement is applied to the dynamic profile. [PR1287788](#)
- SNMP query for IF-MIB::ifOutQLen reports **Wrong Type should be Gauge32 or Unsigned32** for a dynamic VLAN DEMUX0 interface. [PR1287852](#)
- The **services-oids-ev-policy.slax** and **services-oids.slax** files built in the Junos OS image are not the latest versions. [PR1287894](#)
- After offlining and 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 upgrade micro BFD flap is observed. [PR1288433](#)

- The smg-service process (daemon) might generate core files in the backup Routing Engine with a distributed IGMP configuration. [PR1288465](#)
- Performance issues can be seen when nontranslated traffic is introduced to a service-set using a large number of NAT terms. [PR1288510](#)
- After GRES **smid** was thrashing and was not restarted after a fatal SDB error. [PR1288871](#)
- Kernel "rtdata" memory leak is found on an MX Series Virtual Chassis with the **heartbeat** command enabled. [PR1289363](#)
- FPC memory leak might happen in a BBE subscriber environment. [PR1289365](#)
- The interfaces might got to a down state after performing GRES. [PR1289493](#)
- The **request system zeroize** command deletes the **/var/db/scripts** directory, which does not get re-created until the next USB/Netboot recovery. [PR1289692](#)
- The jnxContainersType MIB is not displayed for PIC and MIC as correctly as it is displayed on other Juniper platforms. [PR1289778](#)
- If the vmhost application is not running, then the alarm string will have "Application" name embedded in it. [PR1290150](#)
- NAT-T and DPD functionality do not work for aggressive mode. [PR1290689](#)
- Incorrect temperature is displayed for MPCP5/MPC7 in **show chassis fpc** output. [PR1290771](#)
- When IGMP protocol is enabled, there can be a leak of 56 bytes in the bbe-smgd process (daemon) during logout for every subscriber who had joined any multicast group during the session. [PR1290918](#)
- Rpd core file might be generated when restarting the process via CLI. [PR1291110](#)
- JDI-RCT-RPD: Device going to the DB prompt "db@jsr_jsm_send_ka_after_merge,send_proto_keepalive" was observed on master Routing Engine. [PR1291247](#)
- l2tp iccn fast retransmission occurs after tunnels go down. [PR1291557](#)
- The bbe-smgd process might crash and subscribers might get stuck when a large group of different types of subscribers login/logout. [PR1291969](#)
- The local preference cannot work correctly for EVPN type 5 route in multipath scenario. [PR1292234](#)
- An error in **vbf_filter_add_orphan_check** might be seen when the subscribers using filters log out or log in. [PR1292582](#)
- Error message might be seen while bringing up the subscriber in a subscriber management environment. [PR1293057](#)
- CPCDD might generate core files while using Routing Engine based http-redirect. [PR1293553](#)
- The **show extensible-subscriber-services sessions** command is displaying incorrect timestamp after a unified ISSU. [PR1293800](#)

- Loss of DHCP/PPPoE subscribers is observed during unified ISSU from Junos OS Release 16.1-20170718_161_r4_s5.0 to Release 16.1-20170718_161_r4_s5.0. [PR1294709](#)
- The krt queue might be stuck with the error of "RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0". [PR1295756](#)
- Unable to edit dynamic profiles after scaling up to 400 dynamic profiles. [PR1295446](#)
- The bbe-smgd process might generate a core file at bbe_mcast_ifl_vbf_encoder on service activation or deactivation along with smg-service process (daemon) restart. [PR1295938](#)
- The service-profile's CoS might be overrode by the client-profile's CoS when second family DHCP session added in dual-stack subscriber scenario. [PR1296002](#)
- TACACS remote user is unable to run JET applications because of a bad stored heap. [PR1296237](#)
- The mspmand process might crash if you use SCG services on MS-MPC/MS-MIC. [PR1296422](#)
- The continuous kernel might crash when a lot of terms are configured for firewall filters. [PR1296884](#)
- In ECMP fast reroute scenario, traffic might get silently dropped or discarded because of a next hop in "hold" state. [PR1297251](#)
- A memory leak is seen when **set protocols mld XXX** is changed and committed. [PR1297454](#)
- Multiple bbe-smgd core files are seen during a subscriber binding configuration with DT CST with as little as 200-300 subscribers and continual core files while scaling. Maximum scale cannot be achieved with multicast- enabled subscribers (related to IPTV profile). [PR1297612](#)
- During InFlight Daemon Kill test, rpd core files are seen with PPPoE and L2BSA flapping. [PR1298587](#)
- Commit error is thrown when trying to commit a configuration with apply groups. [PR1298649](#)
- The bbe-smgd process might crash when traceoption is enabled due to an invalid username character. [PR1298667](#)
- The bbe-smgd process constantly generates core files while ESSM+PPPoE stress test with concurrent GRES is running. [PR1298742](#)
- MX Series BNG does not respond to PADI after GRES on some ports/VLANs. [PR1298890](#)
- Junos Telemetry Interface: DREND errors are seen for components "mpcs-software-rev", "rom-software-rev", "software-rev", and "firmware-rev". [PR1299470](#)
- The "asynchronous-notification" feature cannot be implemented properly in a circuit that has MIC-3D-20GE-SFP-E/Tri Rate Copper SFP(740-013111). [PR1299574](#)
- Flat accounting files are not generated according to the configured timers. [PR1299597](#)
- Subscriber database is stuck in not-ready state after GRES. [PR1299940](#)
- After IS-IS-TE routes and BGP routes attribute change, traffic loss might be seen because BGP routes point to some stale labels. [PR1300425](#)

- Junos Telemetry Interface: The error **error: the SDN-Telemetry subsystem is not responding to management requests** is seen on issuing the CLI command **show agent sensors** if traceoptions is enabled for services analytics. [PR1300829](#)
- Configured logical interface might not be created correctly after commit. [PR1301823](#)
- The rpd might crash when toggling the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statement. [PR1302504](#)
- The log message **jam_cache_get.636 ERR:entity 0x997 not found, get cache failed** is continuously seen in jam_chassisd log file. [PR1302975](#)
- chassisd.core-tarball.0.tgz found during ISSU is aborted in FRU upgrade phase. [PR1303086](#)
- Incorrect MTU might be seen on PPP interfaces when PPP MTU is not defined in the dynamic profile. [PR1303175](#)
- The list of available routing instances is no longer provided for output of **show subscribers routing-instance ?command**. [PR1303199](#)
- Blocking PPPoE/DHCP to initiate VLAN auto-sensing if VLAN-OOB connected is in pending state. [PR1303338](#)
- MX Series MIB polling returns a value that has "sdg". Polling result should include "svc" generic value. [PR1303848](#)
- Truncated output appears for the **show pppoe lockout** CLI command. [PR1304016](#)
- Effective rate of E3 in framed mode is limited to 30 Mbps on certain channelized MICs. [PR1304344](#)
- RPF check strict mode is causing traffic drop in next-generation subscriber management release. [PR1304696](#)
- On MX2000 platform with MPC9E and SFB2 installed, certain high amount traffic volume might cause traffic drops with cell underflow messages. [PR1304801](#)
- Commit fails with error: **ffp_intf_ifd_hier_tagging_config_verify: Modified IFD "si-1/1/0" is in use by BBE subscriber, active L2TP LNS client**. [PR1304951](#)
- Inline J-Flow VMX: OIF field of VPLS data records sometimes reports the SNMP index value of the LSI interface instead of the egress physical interface. [PR1305411](#)
- MX Series router is sending immediate-interim for the services pushed by SRC. [PR1305425](#)
- Customers running 32-bit Junos OS might generate rpd core file when traceoptions are enabled. [PR1305440](#)
- Going forward, JET daemonize applications will not get respawned on a normal exit, which should be the ideal behavior of any App. [PR1305615](#)
- L2BSA subscriber connection attempts failed with vlan profile-request-error. [PR1305962](#)
- L2BSA subscribers came up, but no new ANCP session got established during the RADIUS disaster backup procedure. [PR1306872](#)

- Smihelperd generates core files when SNMP is polling for JUNIPER-SUBSCRIBER-MIB::jnxSubscriberGeneral.7.0. [PR1306966](#)
- Split horizon label is not allocated after switching a configuration of ESI from single-active to all-active. [PR1307056](#)
- The kmd process error UI_DBASE_OPEN_FAILED is seen because of too many open files. [PR1308380](#)
- License lost during Routing Engine switchover in scale-subscriber scenario. [PR1308620](#)
- CoS applied to a subscriber demux logical interface (IFL) is not working. [PR1308671](#)
- All the MICs on FPC, with ps interfaces configured, went offline during the restart of FPC in another slot. [PR1308995](#)
- Error message: **%PFE-3: fpc0 vbf_var_iflset_add:633: vbf container 11 not found in the msg for ifl .demux.6514** is often seen after MPC restart. [PR1309013](#)
- Incorrect values are found in the event-timestamp of RADIUS Accounting-Stop packets for L2BSA subscribers. [PR1309212](#)
- RPT BBE REGRESSIONS: DHCP client is stuck in selecting state while verifying untagged DHCP subscribers after modifying router configuration. [PR1309730](#)
- In next-generation subscriber-management release, bbe-smgd process memory leak is seen after deleting or adding the address pool. [PR1310038](#)
- The MS-MIC/MS-MPC memory utilization might stay at high level in the subscriber management scenario. [PR1310064](#)
- **SPD_CONN_OPEN_FAILURE** and **SPC_CONN_FAILURE** log messages are seen in the log for SI interfaces when running SNMP walk on Service PIC NAT OIDs. [PR1310081](#)
- The **krt_junos_sanity_check_ctrl_resp: rtsock** request finally succeeded after error 16' syslog message in the Junos OS Release 17.1R1.8. [PR1310678](#)
- After bsys reboot sometimes rpd is unresponsive on one or more GNFs. [PR1310765](#)
- In streaming telemetry, when a user logs in and logs out quickly from TACACS, the following message is displayed: **bad stored heap: heap-ptr=0x0 data-ptr=0x1481cbf8**. [PR1311482](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Counter at PPPoE session logical interface (IFL) incremented wrongly cause accounting packet contains wrong Acct-input-packets value and wrong Acct-input-octets value. [PR1312998](#)
- Rpd core is seen when any **show route inetcolor.0** command is executed from CLI. [PR1316078](#)
- **show auto-configuration out-of-band** CLI command with different configuration statements shows the same output. [PR1316661](#)
- After NSR to re1, switch back to RE0 has replication stuck for BGP and LDP. [PR1319784](#)
- Rpd core seen during configuration changes with BGP neighbors. [PR1320900](#)

- Commit operation gets stuck when commit check is performed with fast-synchronize option is enabled. [PR1322431](#)
- JDM Management is unreachable after flapping physical JDM and GNF/VNF management interfaces. [PR1323519](#)

High Availability (HA) and Resiliency

- Line Card reboots after GRES. [PR1286393](#)
- After flapping server CB ports GNFs shows "Switchover Status: Not Ready". [PR1306395](#)

Infrastructure

- "Last flapped " time stamp is not getting updated for fxp0 interface as it should be. [PR1244502](#)
- The **show system users** CLI command output displays users that are not using the router. [PR1247546](#)
- When **set system ports console log-out-on-disconnect** is enabled, system reboot or switchover can result in processes remaining in the wait state and failure of the syslog feature. [PR1253544](#)
- The device might fail to upgrade. [PR1298749](#)
- The syscalltrace.sh might create huge output file which could cause the router to run out of storage space. [PR1306986](#)

Interfaces and Chassis

- The output value is incorrect when querying the optical power of OTN interfaces in the router. [PR1216153](#)
- EX Series Packet Forwarding Engine and MX Series MPC7E/8E/9E PFE crash when fetching interface statistics with extended-statistics enabled (CVE-2017-10611). [PR1247026](#)
- At a high logical interface scale, an ifinfo process (daemon) generates a core file on executing the command **show interfaces extensive | no-more**. [PR1254189](#)
- The MRU of ae interface might reset to default value. [PR1261423](#)
- The MTU configuration option for vt interfaces should be removed because the MTU on this interface is already set to unlimited. [PR1277600](#)
- Monitor interface on aggregated Ethernet logical interfaces displays incorrect bps value compared to **show interface** output. [PR1283831](#)
- Interface flap while executing Routing Engine switchover if the member links of an ae interface are configured with framing settings. [PR1287547](#)
- No L2TP sessions come up on some si interfaces after an MPC restart followed by a Routing Engine switchover. [PR1290562](#)
- PPPoE/PPP subscriber might not be brought up with **reject-unauthorized-ipv6cp** configured. [PR1291181](#)
- Change in history records supported per EOAM performance-monitoring session. [PR1294123](#)
- Family inet shows as not-configured after adding or deleting the loopback address. [PR1294267](#)

- A VRRP track interface down does not trigger a mastership election immediately. [PR1294417](#)
- IRB interface shows incorrect bandwidth value. [PR1302202](#)
- AFEB might not come up if LFM is deactivated. [PR1306707](#)
- After executing the **request system reboot both** CLI command, the Juniper PPP daemon might become unresponsive. [PR1310909](#)
- The PPPoE subscriber might not login correctly after authentication failure in subscriber scenario. [PR1311113](#)
- MX Series Virtual Chassis unified ISSU emits benign error message if unsupported FRUs are present. [PR1316374](#)

Layer 2 Ethernet Services

- DHCP is not using the configured IRB MAC as the source MAC in DHCP offer unicast replies. [PR1272618](#)
- DHCPV6 client bound to IA_PD prefix on reception of DHCv6 Request for IA_NA, MX deletes the existing binding. [PR1286359](#)
- ARP requests not generated for IRB configured in VPLS over GRE tunnel. [PR1295519](#)
- PPPoE/DHCP clients cannot login to PPPoE/DHCP dual-stack subscriber scenario. [PR1298976](#)
- Multiple jdncpd core files are observed in jdncpd_update_groups at `../../../../src/junos/usr/sbin/jdncpd/jdncpd_config.c:2290`. [PR1311569](#)

Layer 2 Features

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

MPLS

- RSVP p2mp sub-LSPs having more than one sub-LSP in down state might not get re optimized after transit path goes down. [PR1174679](#)
- The rpd might crash when moving static LSP from one routing instance to another [PR1238698](#)
- Created time value in **show mpls lsp extensive** drifts by a second when the show command is issued multiple times. [PR1274612](#)
- Next generation MVPN 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 might not work if interface flaps. [PR1305228](#)
- Feature explicit-null might block host-bound traffic incoming from LSP. [PR1305523](#)
- The rpd process might crash during interface-down when UHP-based LSPs are configured. [PR1309397](#)

Network Management and Monitoring

- Command Esc-Q does not work when the syslog is disabled. The syslog message is still seen even if it is disabled by Esc-Q. [PR1269274](#)
- MIB2D-related syslog message **MIB2D_RTSLIB_READ_FAILURE: rtslib_iflm_snmp_pointchange** is seen when configurations are removed or restored. [PR1279488](#)
- MIB2D logs **RLIMIT curr 1048576000 max 1048576000** every time a commit is done. [PR1286025](#)
- The mib2d process might crash when polling the OID ifStackStatus.0 after a logical interface (IFL) of Io0 is deleted. [PR1286351](#)
- An alarm-mgmd core file is seen after upgrade due to an old version of the alarm.db file. [PR1296597](#)
- Implement prefix compression for subinterfaces from mib2d. [PR1297447](#)
- The **show arp no-resolve interface X** output for inexistent interface X is showing all unrelated static ARP entries. [PR1299619](#)
- After SNMP configuration activation the snmpd process started to consume a lot of CPU time. [PR1300016](#)

Platform and Infrastructure

- Traffic drop might occur under a large-scale firewall filter configuration. [PR1093275](#)
- The traffic might not be transmitted correctly from MPC/FPC in rare condition. [PR1170527](#)
- FPC crashes with the MAC accounting feature enabled. [PR1173530](#)
- The "forwarding-class-accounting enhanced" feature is not supported in combination with "forwarding-options hyper-mode". Using both features together results in traffic being silently discarded or dropped. [PR1198021](#)
- Packet Process Engine UCODE rebalancing getting enabled by default. [PR1207532](#)
- With a commit script configured, the mgd process might crash when configure anything in private configuration mode. [PR1244015](#)
- The RPM loss percentage values for "over all tests" via SNMP might be incorrect. [PR1272566](#)

- EVPN-VXLAN traffic gets dropped as **Incorrect vxlan fw path executed** due to a sampling configuration on the core interface. [PR1280539](#)
- The **request routing-engine login other-routing-engine** command might require password. [PR1283430](#)
- The traffic might be classified into the wrong queue when aggregated Ethernet interfaces with child legs are anchored on an MQ-based MPC without a queuing chip. [PR1284264](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- Administratively disabling an interface might cause high FPC CPU usage. [PR1285673](#)
- Transit traffic that has the second LSB set in the first octet of destination MAC will be punted to the Routing Engine when **mac-learn-enable** is configured. [PR1285874](#)
- Generate-event time-interval usage now triggers the event only on the actual expiry of the time interval. [PR1286803](#)
- Incorrect load-balancing on the aggregated Ethernet interface might occur if traffic goes from MS-DPC to MPC in enhanced-ip mode. [PR1287086](#)
- Packet Forwarding Engine heap memory leak is found in three routers with PPPoE subscribers. [PR1287870](#)
- mgd: error: **Couldn't open library: /usr/lib/render/libvccpd-render.tlv**. [PR1289158](#)
- Syslog error appears: not a proper library: **/usr/lib/render/libdcd-render.so: Cannot open "/usr/lib/render/libdcd-render.so"**. [PR1289974](#)
- The source MAC learned from Packet Forwarding Engines across ae interface might bounce between ae member Packet Forwarding Engines for a long time and might cause MLP-ADD storm. [PR1290516](#)
- Dynamic MAC learning might fail on GRE tunnel interface. [PR1291015](#)
- RMOPD might get stuck at sbwait upon receiving a specific response from the HTTP agent. [PR1292151](#)
- Transient flow control asserted by XLP MAC after upgrading the MX Series router to Junos OS Release 16.1. [PR1293232](#)
- The scale-subscriber license might leak on the backup Routing Engine during bulk subscriber logout. [PR1294104](#)
- The mgd process generates a core file after GRES in a subscriber environment. [PR1298205](#)
- **RMOPD_HW_TIMESTAMP_INVALID** is reported two to four times a day which raises an alarm when polled via jnxRpmResSumPercentLost MIB. [PR1300049](#)
- MPC might reset in firewall filter scenario during loading configuration on MX Series platform. [PR1300990](#)
- All traffic can be Tail/RED-dropped on some interfaces when **chassis fpc max-queues** is configured. [PR1301717](#)
- Classifier does not get applied on the aggregated Ethernet member links on DPC (I-chip) based platforms with CoS configured. [PR1301723](#)

- MX Series FPC wedges when creating more than 4000 logical tunnel interfaces per Packet Forwarding Engine. [PR1302075](#)
- When you execute the **mk destroy-all** command, it gives the error **Could not find jnx.wrlsb.mk**. [PR1302974](#)
- The interface-mac-limit might fail for aggregated Ethernet interface. [PR1303293](#)
- The Two-Way Active Measurement Protocol (TWAMP) Request-TW-Session message's Type-P Descriptor format is not RFC-compliant. [PR1305752](#)
- On MX Series routers with MPCs or MICs, the resource monitor (RSMON) thread might be stuck in a loop consuming 100 percent of FPC CPU. [PR1305994](#)

Routing Protocols

- No multicast forwarding in ASM mode occurs after unified ISSU. [PR1146621](#)
- RLFA computation might still consider a PQ-node not reachable via LDP, when LDP is deactivated. [PR1202392](#)
- The routing protocol process (rpd) on the backup Routing Engine might restart unexpectedly upon the addition of a new L2VPN routing instance. [PR1233514](#)
- When the **advertise-from-main-vpn-tables** configuration statement is used under BGP and the route reflector functionality is added, a refresh message is not sent, resulting in some missing routes. [PR1254066](#)
- MPLS over UDP tunnel creation fails in the absence of a VRF table. [PR1270955](#)
- A few BFD sessions are flapping while coming up after FPC restart/reboot. [PR1274941](#)
- Error messages might be seen when receiving BGP update messages with UNREACH NLRI. [PR1276758](#)
- After Routing Engine switchover (GRES+GR), default mdt failed to come up and core-facing interface flap was seen. [PR1279459](#)
- BGP updates might not be advertised to peers completely in certain condition. [PR1282531](#)
- The rpd process might crash due to a certain chain of events in a BGP-LU protection scenario. [PR1282672](#)
- The second multicast packet might be discarded on the rendezvous point router. [PR1282848](#)
- The rpd process might crash while deactivating the routing instance of pim static. [PR1284760](#)
- Some BGP-related traceoptions flag settings will not be effective immediately after the configuration commit, until the BGP sessions are flapped. [PR1285890](#)
- The rpd will run into a loop if bootstrap messages exceed the interface MTU size. [PR1287467](#)
- The rpd might crash if the dynamic rendezvous point goes down in ECMP topology and also PIM **join-load-balance automatic** is configured. [PR1288316](#)
- The rpd might crash after loading merge and rollback configuration with BGP traceoption. [PR1288558](#)
- Multicast flow reset might occur on OIF for RPT joined branch when PIM prune comes on another interface. [PR1293900](#)

- The rpd might crash if BGP flap happens. [PR1295062](#)
- ISSU might take more time to complete and the MPC card might go offline during ISSU reboot. [PR1298259](#)
- Inline BFD on IRB will be broken after GRES/NSR switchover, and the anchor FPC subsequent goes offline. [PR1298369](#)
- BGP might send an incorrect AS path when the alias is enabled and multiple peers are under the BGP group. [PR1300333](#)
- The rpd process might crash with a core file while deleting a multipath route. [PR1302395](#)
- Junos OS Release 16.2 and later releases might give the following error: **Request failed: OID not increasing: ospfIfIpAddress.0.0.0.0.0.** [PR1307753](#)
- Qualified next-hop resolution fails in some scenarios when there is a next-hop interface specified. [PR1308800](#)
- BGP labeled-unicast protection might break multicast Reverse Path Forwarding (RPF). [PR1310036](#)
- An rpd core file is observed while importing IS-IS routes. [PR1312325](#)
- BGP prefixes with three levels of recursion for resolution will get stuck with a stale next-hop at the first level after a link-down event. [PR1314882](#)

Services Applications

- Business service fails to get deactivated after Routing Engine switchover. [PR1280074](#)
- Backup Routing Engine goes to the database prompt with a vmcore if the configuration for the ASI interface that has gone down is deleted. [PR1281882](#)
- TLVs in ICRQ for actual-rate-downstream/actual-data-rate-upstream do not reflect PPPoE-IA value. [PR1286583](#)
- mspmand cored "@_arena_mALLOc" seen in Backup SDG's MS70. [PR1291664](#)
- L2TP subscribers are down after a GRES while verifying framed IPv6 route support for L2TP network server (LNS) at a higher scale with a maximum number of framed IPv6 routes. [PR1293783](#)
- Each subscriber session gets its own L2TP tunnel without "Tunnel-Client-Endpoint" from RADIUS. [PR1293927](#)
- The jl2tpd process might crash shortly after a GRES switchover. [PR1295248](#)
- [OC/ST] Continuous generation of *jl2tpd_era_Ins* log files occurs even though l2tp is not configured. [PR1302270](#)

Software Installation and Upgrade

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

Subscriber Access Management

- The DHCP subscriber might not get an IP address if the address pool utilization is tight. [PR1274870](#)
- Some RADIUS attributes might not be filtered out of the accounting-on/accounting-off message on an MX Series. platform. [PR1279533](#)
- IP assigned by RADIUS is incorrectly counted by the local pool after a Virtual Chassis switchover. [PR1286609](#)
- The authd process generates a core file at DynamicRequestEntry::addHistory authd_aaa_dyn_req. [PR1289215](#)
- Service interim for DHCP subscriber is not working in JSRC scenario. [PR1303553](#)
- The **show network-access aaa accounting** command might display additional entries. [PR1304594](#)
- Incorrect **Acct-Delay-Time in Radius Accounting-On** message is seen after rebooting the MX Series router acting as a BNG. [PR1308966](#)
- The delegated prefix from RADIUS is incorrectly parsed when the prefix is fewer than 20 bytes long. [PR1315557](#)

User Interface and Configuration

- Increasing commit times are seen. [PR1029477](#)
- The commitd process might generate a core file when removal of certain configuration is followed by a commit operation. [PR1267433](#)
- The commit might fail with the error of "Could not open configuration database" and "foreign file propagation (ffp) failed". [PR1287539](#)

VPNs

- Next generation MVPN SG entry and MVPN route persist after data stop. [PR1236733](#)
- Rpd memory leak is observed in a next generation MVPN environment. [PR1259579](#)
- Next generation MVPN IPv6 RP bootstrap type 3 S-PMSI AD route prefix ff02::d persist after BSR data stop. [PR1269234](#)
- L2circuits stitched via It peer interfaces might be stuck in "LD" (local site signaled down) status. [PR1305873](#)

SEE ALSO

Changes in Behavior and Syntax	149
Known Behavior	164
Known Issues	173
Documentation Updates	275
Migration, Upgrade, and Downgrade Instructions	276

Documentation Updates

IN THIS SECTION

- [Subscriber Management Access Network Guide | 275](#)
- [Subscriber Management Provisioning guide | 275](#)
- [Subscriber Management VLANs Interfaces Guide | 276](#)

This section lists the errata and changes in Junos OS Release 17.4R3 documentation for MX Series.

Subscriber Management Access Network Guide

- The guide failed to include a feature that enables you to override the information that the LAC sends to the LNS in L2TP Calling Number AVP 22 when the LAC is configured to use the Calling-Station-ID format. You can configure the access profile to override that value for AVP 22 with any combination of the agent circuit identifier and the agent remote identifier received by the LAC in the PADR packet.

[See [Override the Calling-Station-ID Format for the Calling Number AVP](#)].

- The guide incorrectly stated that the **linked-pool-aggregation** statement is located at the **[edit access address-assignment pool pool-name]** hierarchy level. In fact, this statement is located at the **[edit access]** hierarchy level.

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

Subscriber Management Provisioning guide

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

- Starting in Junos OS Release 15.1, the *Broadband Subscriber Sessions User Guide* and the [CLI Explorer](#) incorrectly included information about the **show extensible-subscriber-services accounting** command. This command is not present in the CLI. Instead, you can use accounting profiles to collect statistics from the Packet Forwarding Engine for Extensible Subscriber Services Manager (ESSM) subscribers. See [Flat-File Accounting Overview](#) for information about accounting for ESSM subscribers.

Subscriber Management VLANs Interfaces Guide

- The *Broadband Subscriber VLANs and Interfaces User Guide* did not clearly indicate that only demux0 is supported for demux interfaces. If you configure a different demux interface, such as demux1, the configuration commit fails.

SEE ALSO

[New and Changed Features | 115](#)

[Changes in Behavior and Syntax | 149](#)

[Known Behavior | 164](#)

[Known Issues | 173](#)

[Resolved Issues | 202](#)

[Migration, Upgrade, and Downgrade Instructions | 276](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.4 | 277](#)
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS | 277](#)
- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | 280](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 282](#)
- [Upgrading a Router with Redundant Routing Engines | 282](#)
- [Downgrading from Release 17.4 | 282](#)

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-x86-32-17.4R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-17.4R3.9-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-x86-32-17.4R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.4R3.9-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-17.4R3.9-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-17.4R3.9-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 | 115](#)

[Changes in Behavior and Syntax | 149](#)

[Known Behavior | 164](#)

[Known Issues | 173](#)

[Resolved Issues | 202](#)

[Documentation Updates | 275](#)

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [New and Changed Features | 284](#)
- [Changes in Behavior and Syntax | 285](#)
- [Known Behavior | 285](#)
- [Known Issues | 286](#)
- [Resolved Issues | 287](#)
- [Documentation Updates | 288](#)
- [Migration, Upgrade, and Downgrade Instructions | 288](#)

These release notes accompany Junos OS Release 17.4R3 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R3 New and Changed Features | 284](#)
- [Release 17.4R2 New and Changed Features | 284](#)
- [Release 17.4R1 New and Changed Features | 284](#)

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

Release 17.4R3 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 17.4R3.

Release 17.4R2 New and Changed Features

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

Release 17.4R1 New and Changed Features

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 17.4R1.

NOTE: vSRX version 15.1X49-D100 is compatible with the Junos OS Release 17.4R1 for NFX Series devices.

SEE ALSO

[Changes in Behavior and Syntax | 285](#)

[Known Behavior | 285](#)

[Known Issues | 286](#)

Resolved Issues	 287
Documentation Updates	 288
Migration, Upgrade, and Downgrade Instructions	 288

Changes in Behavior and Syntax

There are no changes in behavior and syntax for NFX Series in Junos OS Release 17.4R2.

SEE ALSO

New and Changed Features	 284
Known Behavior	 285
Known Issues	 286
Resolved Issues	 287
Documentation Updates	 288
Migration, Upgrade, and Downgrade Instructions	 288

Known Behavior

There are no known limitations in Junos OS Release 17.4R3 for NFX Series.

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

SEE ALSO

New and Changed Features	 284
Changes in Behavior and Syntax	 285
Known Issues	 286
Resolved Issues	 287
Documentation Updates	 288
Migration, Upgrade, and Downgrade Instructions	 288

Known Issues

IN THIS SECTION

- [Virtual Network Functions | 286](#)
- [Juniper Device Manager | 286](#)
- [Junos Control Plane | 286](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 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.

Virtual Network Functions

- 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](#)

Juniper Device Manager

- On NFX250 devices, logging in to JDM using TACACS credentials might result in permission denied messages on the shell. This does not have any functional impact. [PR1330419](#)

Junos Control Plane

- If the traffic in the out-of-band interface is more, the control plane connectivity might get blocked for sometime while the packets are processed. If this interruption persists, the connection between the PFE and control plane is cleared, which results in a PFE restart or shutdown. You must ensure that there is no heavy traffic flow in the management VLAN. [PR1270689](#)
- On NFX250 devices, ICMP ping does not work for packets with ICMP payload beyond 1480 bytes. As a workaround, execute the following procedure on both the peer devices:

1. Issue the **set chassis fpc 0 power off** command on the virtual control plane (VCP) to power off the FPC.
2. Access the hypervisor by issuing the **rsh ?JU __juniper_private4__ 192.168.1.1** command at the VCP BSD shell.
3. Set the mtu values on the hypervisor:
 - **ip link set hsxe0 mtu 9216**
 - **ip link set hsxe1 mtu 9216**
4. Verify that the mtu values are set to 9216 by issuing the **ip link show hsxe0 | grep mtu** and **ip link show hsxe1 | grep mtu** commands.
5. Exit from the hypervisor.
6. Access the Junos CLI and issue the **delete chassis fpc 0 power off** command.
7. Wait for the FPC to come up.

[PR1285852](#)

SEE ALSO

New and Changed Features 284
Changes in Behavior and Syntax 285
Known Behavior 285
Resolved Issues 287
Documentation Updates 288
Migration, Upgrade, and Downgrade Instructions 288

Resolved Issues

There are no fixed issues in Junos OS Release 17.4R3 for NFX Series.

SEE ALSO

New and Changed Features 284
Changes in Behavior and Syntax 285
Known Behavior 285

[Known Issues | 286](#)

[Documentation Updates | 288](#)

[Migration, Upgrade, and Downgrade Instructions | 288](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 documentation for NFX Series.

SEE ALSO

[New and Changed Features | 284](#)

[Changes in Behavior and Syntax | 285](#)

[Known Behavior | 285](#)

[Known Issues | 286](#)

[Resolved Issues | 287](#)

[Migration, Upgrade, and Downgrade Instructions | 288](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 288](#)
- [Basic Procedure for Upgrading to Release 17.4 | 289](#)

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 17.4R3:

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

SEE ALSO

[New and Changed Features | 284](#)

[Changes in Behavior and Syntax | 285](#)

[Known Behavior | 285](#)

[Known Issues | 286](#)

[Resolved Issues | 287](#)

[Documentation Updates | 288](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- New and Changed Features | 291
- Changes in Behavior and Syntax | 304
- Known Behavior | 312
- Known Issues | 314
- Resolved Issues | 319
- Documentation Updates | 329
- Migration, Upgrade, and Downgrade Instructions | 330

These release notes accompany Junos OS Release 17.4R3 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R3 New and Changed Features | 292
- Release 17.4R2 New and Changed Features | 292
- Release 17.4R1 New and Changed Features | 292

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R3 for the PTX Series.

Release 17.4R3 New and Changed Features

Interfaces and Chassis

- **LACP hold-up timer configuration support on LAG interfaces (PTX Series)**— You can configure a Link Aggregation Control Protocol (LACP) hold-up timer value for link aggregation group (LAG) interfaces.

You configure the hold-up timer to prevent excessive flapping of a child (member) link of a LAG interface due to transport layer issues. With transport layer issues, it is possible for a link to be physically up and still cause LACP state-machine flapping. LACP state-machine flapping can adversely affect traffic on the LAG interface. To prevent this, a hold-up timer value is configured. LACP monitors the PDUs received on the child link for the configured time value, but does not allow the member link to transition from the expired or defaulted state to current state. This configuration thus prevents excessive flapping of the member link.

To configure the LACP hold-up timer for LAG interfaces, use the **hold-time up timer-value** statement at the **[edit interfaces ae aeX aggregated-ether-options lacp]** hierarchy level.

See [hold-time up](#) and [Configuring LACP Hold-UP Timer to Prevent Link Flapping on LAG Interfaces](#).

Release 17.4R2 New and Changed Features

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

Release 17.4R1 New and Changed Features

Hardware

- **PTX10016 Packet Transport Router**—Starting in Junos OS Release 17.4R1, the PTX10016 Packet Transport Router provides 3.0 Tbps per slot forwarding capacity for the service providers and cloud operators. The router provides an opportunity for the cloud, telco, and data center operators for a smooth transition from 10-Gigabit Ethernet and 40-Gigabit networks to 100-Gigabit Ethernet high-performance networks. This high-performance, 21 rack unit (21RU) modular chassis provides 48 Tbps of throughput and 32 Bpps of forwarding capacity. The PTX10016 router has 16 slots for the line cards that can support a maximum of 2304 10-Gigabit Ethernet ports, 576 40-Gigabit Ethernet ports, or 480 100-Gigabit Ethernet ports.

You can deploy the PTX10016 router in the core of the network for the following functions:

- Label switching routing
- IP core routing
- Internet peering

PTX10016 Packet Transport Router supports two PTX10K line cards, LC1101 and LC1102. The LC1101 line card consists of thirty QSFP+ Pluggable Solution (QSFP28) cages that support 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The line card supports speed of either 40-Gbps or 100-Gbps.

It also supports 10-Gigabit Ethernet by channelizing the 40-Gigabit Ethernet ports. The default port speed is 100-Gbps. The default port speed is 100-Gbps. If the user plugs in 40Gigabit or 4x10Gigabit optic, the appropriate port speed has to be configured manually.

The LC1102 line card consists of 36 quad small form-factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. The QSFP+ ports support 40-Gigabit or 100-Gigabit Ethernet optical transceivers in selected ports. The default port speed on the LC1102 line card is channelized 10-Gbps. Out of these 36 ports, 12 ports are QSFP28 capable for supporting 100-Gigabit Ethernet. The line card supports 10-Gigabit Ethernet by channelizing the 40-Gigabit ports. Channelization is supported on fiber breakout cable using standard structured cabling techniques.

For more information, see [PTX10016 Packet Transport Router Hardware Guide](#) .

- **Support for the CFP2-DCO-T-WDM-1 transceiver on the P2-100GE-OTN PIC (PTX)**—Starting in Junos OS Release 17.4R1, you can install the CFP2-DCO-T-WDM-1 transceiver on the P2-100GE-OTN PIC. The CFP2-DCO-T-WDM-1 transceiver is a 100-Gigabit digital pluggable CFP2 digital coherent optical module.

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

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

[See [100-Gigabit Ethernet OTN PIC with CFP2 \(PTX Series\)](#) .]

High Availability (HA) and Resiliency

- **Resiliency Support for PTX10K-LC1101 and PTX10K-LC1102 (PTX10016)**—Starting with Junos OS Release 17.4R1, resiliency support is enabled for the following components:
 - PTX10K-LC1101 and PTX10K-LC1102
 - Routing and Control Boards
 - Switch Interface Boards

Interfaces and Chassis

- **Fabric Management Support (PTX100016)**—Starting in Junos OS Release 17.4R1, you can set up and manage the fabric connections between the Packet Forwarding Engines in the PTX100016 routers. Fabric management includes collecting fabric status and statistics, monitoring health of the hardware, and responding to CLI queries. It also tracks addition and removal of FRUs from the router and monitors faults in the data plane. It is enabled by default and can be monitored by using the following commands:

- **show chassis fabric summary**
- **show chassis fabric fpcs fpc fpc-slot**
- **show chassis fabric sibs**
- **show chassis fabric errors**
- **show chassis fabric reachability**

[See [Fabric Management Overview](#).]

- **Support for large-scale packet-forwarding features (PTX10000)**—Starting with Junos OS Release 17.4R1, PTX10000 router supports large scaling IPv4 and IPv6 forwarding information base (FIB). A maximum of 4 million routes are supported.
- **Support for pre-FEC BER monitoring when using the CFP2-DCO-T-WDM-1 transceiver (PTX Series)**—Starting in Junos OS Release 17.4R1, you can monitor the condition of an OTN link by using the pre-forward error correction (pre-FEC) bit error rate (BER) when using the CFP2-DCO-T-WDM-1 transceiver.

[See [Understanding Pre-FEC BER Monitoring and BER Thresholds](#).]

- **Support for a 16 Slot Chassis (PTX10016)**—Starting with Junos OS Release 17.4R1, the PTX10016 has 16 slots and supports core and edge profiles.

IPv6

- **Support for IPv6 statistics on PTX Series routers**—Starting in Junos OS Release 17.4R1, you can obtain the transit IPv6 statistics at both the physical interface and logical interface levels on third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008 by using both a CLI command and SNMP MIB counters. Use the **show interfaces statistics** command to display both physical interface and logical interface statistics. You can view only logical interface statistics if you use SNMP MIB counters. However, for aggregated Ethernet interfaces, the accounting is not done at the level of the child links and, thus, IPv6 statistics for child links are not displayed.

To start getting IPv6 statistics on third-generation FPCs, use the **route-accounting** statement at the **[edit forwarding-options family inet6]** hierarchy level. PTX Series routers with first-generation and second-generation FPCs do not display IPv6 statistics for physical interfaces or logical interfaces, and transit statistics on child links in aggregated Ethernet interfaces are also not taken into account.

NOTE: Egress accounting for IPV6 traffic is not performed for cases where MPLS packets arrives on TCC interface and egress out of the router as IPV6 packets.

[See [route-accounting](#) and [show interfaces extensive](#).]

Junos OS XML API and Scripting

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

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

Layer 2 Features

- **Support for Layer 2 protocols (PTX 10016)**—Starting in Junos OS Release 17.4R1, Layer 2 protocols are supported on PTX10016 routers that have third-generation FPCs installed. Layer 2 protocols include Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), VLAN Spanning Tree Protocol (VSTP), Link Layer Discovery Protocol (LLDP), and so on.

Layer 3 Features

- **Support for Layer 3 protocols (PTX 10016)**—Starting in Junos OS Release 17.4R1, Layer 3 protocols are supported on PTX10016 routers that have third-generation FPCs installed. Layer 3 protocols include the Multiprotocol Label Switching (MPLS), Layer 3 Virtual Private Network (L3VPN), Bidirectional Forwarding Detection (BFD), Layer 2 Virtual Private Network (L2VPN), Point-to-multipoint (P2MP), Fast ReRoute (FRR), Operations, Administration and Maintenance (OAM), Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP), Adaptive Load Balancing (ALB), and so on.

Management

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

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

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

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

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

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

- **Enhancements to LSP events sensor for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, telemetry data streamed through gRPC for LSP events and properties is reported separately for each routing instance. To export data for LSP events and properties, you must now include

`/network-instances/network-instance[name_'instance-name']` in front of all supported paths. For example, to export LSP events for RSVP Signaling protocol attributes, use the following path: `/network-instances/network-instance[name_'instance-name']/mpls/signaling-protocols/rsvp-te/`. Use the **telemetrySubscribe** RPC to specify telemetry parameters and provision the sensor. If your device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Junos Network Agent software package, which provides the interfaces to manage gRPC subscriptions.

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

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

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

- **Support for bypass LSP statistics for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export statistics for bypass label-switched paths (LSPs). Previously, only statistics for the primary LSP path were exported. The ability to export bypass LSP statistics helps to monitor the efficiency of global convergence when the bypass LSP is used to carry traffic during a link or node failure. Statistics are exported for the following:

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

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

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

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

The tables are:

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

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

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

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

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

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

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

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

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

- Traffic engineering global attributes
- Traffic engineering interface attributes

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

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

- **FPC1 and FPC2 support for CPU and NPU sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.4R1, you can export data for CPU memory and NPU memory and utilization for FPC1 and FPC2 on PTX Series routers through the Junos Telemetry Interface. Previously, only FPC3 was supported on these sensors. To provision the sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

[See [sensor \(Junos Telemetry Interface\)](#) and [Guidelines for gRPC sensors.](#)]

MPLS

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

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

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

- **Support for static adjacency segment identifier for IS-IS (PTX Series)**—Starting with Junos OS Release 17.4R1, you can configure static adjacency segment ID (SID) labels for an interface. You can configure two IPv4 adjacency SIDs (protected and unprotected), IPv6 adjacency SIDs (protected and unprotected) per level per interface. You can use the same adjacent SID for multiple interfaces by grouping a set of interfaces under an interface-group and configuring the adjacency-segment for that interface-group. For static adjacency SIDs, the labels are picked from either a static reserved label pool or from segment routing global block (SRGB).

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

- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (MX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage, which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP resignalling events. This is because the LSP is responsive to even minor increase or decrease in the utilization when current

reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of say 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP ressignalling events for small LSPs of say 10M reservation. Increasing the adjust threshold percentage by for example 40 percent minimizes LSP ressignaling for small LSPs. However, large LSPs do not react to bandwidth usage changes unless they are huge, for example, 400M. Starting in Junos OS Release 17.4R1, you can configure an absolute value based threshold along with the percentage based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations. Configure **adjust-threshold-absolute value** option at the **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]** hierarchy level.

- **Support for default time-out duration for self-ping on an LSP instance (PTX Series)**—Starting in Junos OS 17.4R1, the default time out duration for which the self-ping runs on an LSP instance is reduced from 65535 (runs until success) to 1800 seconds. You can also configure the self ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is enabled. The LSP types like CCC, P2MP, VLAN-based, and non-default instances do not support self-ping. You can configure **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.
- **Support for flap and MBB counter for LSP (PTX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on master routing engine only:
 - Flap counter-- Counts the number of times an LSP flaps down or up.
 - MBB counter— Counts the number of times an LSP incurs MBB.

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

- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (PTX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
- **Support for label history for MPLS protocol (PTX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entry per label. By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range.
The **clear mpls label history** command clears the label history details.

Routing Protocols

- **Support for importing IGP topology information into BGP-LS (PTX Series)**—Starting in Junos OS Release 17.4R1, you can import interior gateway protocol (IGP) topology information into BGP-Link State (BGP-LS) in addition to RSVP-traffic engineering (RSVP-TE) topology information through the `Isdist.0` routing table. This allows you to monitor both IGP and traffic engineering topology information.

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

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

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

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

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

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

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

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

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

- **Support for network instance-based BGP configuration (PTX Series)**—Starting in Junos OS Release 17.4R1, you can configure BGP in a specific network instance. After the network instance is configured, you will be prompted with options for BGP configuration such as global bgp, neighbor bgp, and so on.

[See [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

- **DDoS protection support (PTX3000, PTX-5000, PTX1000, PTX10000)**—Starting with Junos OS Release 17.4R1, protection from DDoS attack is provided on PTX3000, PTX 5000, PTX1000, and PTX10000 routers only if they have PE-based FPCs installed.

If the total amount of traffic that a Routing Engine can handle exceeds its limit, the Routing Engine becomes overloaded and is unable to handle the routing protocol messages and other important control plane packets. This results in an inconsistent control plane protocol state and that is termed as DDoS attack.

With the support for DDoS protection, the firewall filters and policers available in Junos OS are used to discard or rate-limit control plane traffic so that such malicious traffic does not overwhelm and bring down the Routing Engine. The Packet Forwarding Engine does not support rate-based policers; therefore, DDoS protection works based on bandwidth.

DDoS protection is supported with the following protocols:

- L3 protocols— IGMP v4/v6, OSPF-Hello, OSPF, LDP-Hello, LDP, PIM-Ctrl, PIM-Data, RSVP, RIP, BFD, MHOP BFD, MSDP, BGP, TELNET, FTP, SSH, SNMP, NTP, TACACS, DNS, GRE, ICMP, MLD, NDP, and EGPv6
- L2 protocols— STP, LACP, LLDP, OAM-CFM, OAM-LFM, ISIS, ISO-TCC, ETH-TCC, and PVST

Exceptions to DDoS protection support include the following:

- L3 protocols are per protocol level and not at packet type level.
- Unsupported L3 protocols— DHCP v4/v6, PTP, VRRP, DTCP, RADIUS-SERVER, RADIUS-ACCT, RADIUS-AUTH, DIAMETER, DIAMETER-TCP, DIAMETER-SCTP, L2TP, LMP, BFDv6, Martian-address, and PIM-REGISTER
- Unsupported L2 protocols— STP, DOT1X, GARP, FC, Bridge control, and PVST
- FPC1 and FPC2 on PTX5000 router are not supported.

For more information, see [Distributed Denial-of-Service \(DDoS\) Protection Overview](#).

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

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

- Program the EBGp route server.

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

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

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

- **Support for BGP advertising aggregate bandwidth across external BGP links for load balancing (MX Series)**—Starting in Junos OS Release 17.4R1, BGP uses a new link bandwidth extended community, **aggregate-bandwidth**, to advertise aggregated bandwidth of multipath routes across external links. BGP calculates the aggregate of multipaths that have unequal bandwidth allocation and advertises the aggregated bandwidth to external BGP peers. A threshold to the aggregate bandwidth can be configured to restrict the bandwidth usage of a BGP group. In earlier Junos OS releases, a BGP speaker receiving multipaths from its internal peers advertised the link bandwidth associated with the active route. To advertise aggregated bandwidth of multipath routes and to set a maximum threshold, configure a policy with **aggregate-bandwidth** and **limit bandwidth** actions at the [edit policy-options policy-statement *name* then] hierarchy level.

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

Security

- **Support for Layer 2 circuit pass-through (PTX Series)**—Starting in Junos OS Release 17.4R1, you can configure PTX Series routers to allow LACP, LLDP, OAM LFM, and OAM CFM packets to cross the Layer 2 circuit. To configure Layer 2 circuit pass-through, include the **l2circuit-control-passthrough** statement at the [set forwarding-options] hierarchy level.

NOTE: LACP can be configured only when the aggregated interface is configured with the ethernet-ccc encapsulation.

[See [l2circuit-control-passthrough](#).]

Services Applications

- **Reporting of true outgoing interface packets for inline flow monitoring (PTX Series)**—Starting in Junos OS Release 17.4R1, you can configure inline flow monitoring to report true packets for the outgoing interface. For ECMP, the actual outgoing interface used for a given flow is the true outgoing interface. To enable a true outgoing interface, include the **nexthop-learning enable** statement at the [set services flow-monitoring (version9 | version-ipfix) template *template-name*] hierarchy level.

[See [template \(Flow Monitoring IPFIX Version\)](#) or [version9 \(Flow Monitoring\)](#).]

- **Reporting of the true incoming interface for the sampled packets for inline flow monitoring (PTX Series)**—Starting in Junos OS Release 17.4R1, inline flow monitoring reports the true incoming interface

for the GRE-encapsulated packets entering the router for the configured inline flow monitoring filter criteria.

[See [Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers.](#)]

- **Support for inline JFlow version 9 flow templates (PTX 10016)**—Starting in Junos OS Release 17.4R1, you can use inline-J-Flow export capabilities with version 9 flow templates to define a flow record template suitable for IPv4 or IPv6 traffic.

Software Installation and Upgrade

- **Device serial number added to DHCP option 60 (PTX1000)**—Starting in Junos OS Release 17.4R1, DHCP option 60 (Vendor Class Identifier) includes the serial number of the device when you use zero touch provisioning to automate provisioning of the device configuration and software image. The serial number can uniquely identify the device in a broadcast network. The serial number appears in the format *Juniper-model-number*. For example, a PTX1000 router numbered DA000 appears as *Juniper-ptx1000-DA000*.

SEE ALSO

[Changes in Behavior and Syntax | 304](#)

[Known Behavior | 312](#)

[Known Issues | 314](#)

[Resolved Issues | 319](#)

[Documentation Updates | 329](#)

[Migration, Upgrade, and Downgrade Instructions | 330](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 305](#)
- [General Routing | 305](#)
- [Interfaces and Chassis | 305](#)
- [Management | 308](#)
- [MPLS | 308](#)
- [Multicast | 309](#)
- [Network Management and Monitoring | 309](#)

- Routing Policy and Firewall Filters | 310
- Security | 311
- Software Licensing | 311
- Subscriber Management and Services | 311

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 17.4R3 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](#).

General Routing

- **User confirmation prompt for configuring the sub-options of request vmhost commands (MX Series and PTX series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm a [yes,no] for the sub-options also.
 - **request vmhost reboot**
 - **request vmhost poweroff**
 - **request vmhost halt**

In previous releases, the confirmation prompt was available for only the main options.

Interfaces and Chassis

- **Secondary interface (em2) raises an alarm when the link is down (PTX1000)**—Starting in Junos OS Release 17.4R1, secondary interface (em2) raises alarm when the link goes down. Earlier, no alarm was raised when an em2 (secondary interface) went down. Currently, the behavior is changed and an alarm will be raised when the interface link goes down as shown below:

```

user@host# run show chassis alarms
3 alarms currently active
Alarm time           Class  Description
2017-09-12 23:41:20 PDT  Major  FPC Management2 Ethernet Link Down
2017-09-12 23:38:45 PDT  Major  FPC0: PEM 2 Not Powered
2017-09-12 23:38:45 PDT  Major  FPC0: PEM 0 Not Powered

```

- **Modified output of the request vmhost zeroize command**—Starting with Junos OS Release 17.2, the command **request vmhost zeroize**, upon execution, prompts the user for confirmation to proceed. The following line is displayed:

```

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

```

- **Power supply alarm is not raised when the input switch status is OFF or power not connected (PTX10008, PTX10016)**—Starting in Junos OS Release 17.4R2, the power supply alarm **A power supply input has failed** will not be raised if the INP1/INP2 switch status is off and the power is not connected. Earlier, an alarm was raised for the power entry module (PEM) that it was not powered on as **Not Powered** irrespective of the switch state. For the power supply status, execute the **show chassis power** or **show chassis power detail** CLI command. The **DC input** is the new output parameter that provides information about the status of the input feed.

Previous behavior:

user@host> show chassis power

```

PEM 0:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC output:  864 W (zone 0, 72 A at 12 V, 34% of capacity)

PEM 1:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC output:  864 W (zone 0, 72 A at 12 V, 34% of capacity)

System:
  Zone 0:
    Capacity:      7500 W (maximum 7500 W)
    Allocated power: 6525 W (975 W remaining)
    Actual usage:   2616 W

```

```
Total system capacity: 7500 W (maximum 7500 W)
Total remaining power: 975 W
```

```
...
```

Current behavior:

user@host> show chassis power

```
PEM 0:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC input:   OK (No feed expected, Both feed connected)
  DC output:  576 W (zone 0, 48 A at 12 V, 23% of capacity)
```

```
PEM 1:
  State:      Online
  Capacity:   2500 W (maximum 2500 W)
  DC input:   OK (No feed expected, Both feed connected)
  DC output:  576 W (zone 0, 48 A at 12 V, 23% of capacity)
```

```
...
```

[See [show chassis power](#).]

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (PTX Series)**—In Junos OS Release 17.4R3, the **show lacp interfaces | display xml** command displays a new XML tag element **<lacp-hold-up-state>**. The **<lacp-hold-up-state>** displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces were in a single **<lacp-hold-up-information>** XML tag. Now, for each interface it is displayed in a separate **<lacp-hold-up-information>** XML tag.

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](#).]

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

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

- **Change in Error Severity (PTX10016)**—Starting in Junos OS Release 17.4R3, on PTX10016 routers, the severity of the FPC error, shown in the syslog as **PE Chip::FATAL ERROR!! from PE2[2]: RT: Clear Fatal if it is detected LLMEM Error MEM:llmem, MEMTYPE: 1**, is changed from fatal to non-fatal (or minor). In case of this error, only a message is displayed for information purpose. To view the error details, you can use the show commands **show chassis fpc errors** and **show chassis errors active**.

[See [show chassis fpc errors](#)]

- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (PTX Series)**—Starting in Junos OS Release 17.4R3, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.
- **No chassis alarm when power consumption by an FPC exceeds 90% or 100% of the allocated power budget**—Starting in Junos OS Release 17.4R3, the PTX5000 routers do not raise a chassis alarm in the following events:
 - Power consumption by an FPC exceeds 90% of the allocated power budget.
 - Power consumption by an FPC exceeds 100% of the allocated power budget (in this case, a system log is registered).

Routing Policy and Firewall Filters

- **Error caused by firewall filters with syslog and accept action (PTX1000 or PTX series routers with type 3 FPCs)**—In this release of Junos OS, under rare circumstances, the host interface may stop sending packets and the connections to and from the peer might fail if an outbound firewall filter is configured with an action of **syslog** and **accept**. This condition applies to IPv4 and IPv6 traffic families. Juniper recommends that you do not use the **syslog** and **accept** action in the output filter for these systems.

An example configuration is provided (shows IPv4).

```
set interfaces interface name unit unit family inet filter output name
set firewall family inet filter name term 1 then syslog
set firewall family inet filter name term 1 then accept
```

[For more information, see [PR 1354580](#).]

Security

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

Software Licensing

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

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

Subscriber Management and Services

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

The new behavior is as follows:

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

2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

In earlier releases:

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

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

SEE ALSO

[New and Changed Features | 291](#)

[Known Behavior | 312](#)

[Known Issues | 314](#)

[Resolved Issues | 319](#)

[Documentation Updates | 329](#)

[Migration, Upgrade, and Downgrade Instructions | 330](#)

Known Behavior

IN THIS SECTION

● [General Routing | 313](#)

● [Interfaces and Chassis | 313](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 for PTX Series.

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

General Routing

- For CFP2-DCO-T-WDM-1 pluggable, Rx payload type is shown incorrectly (shown 0 vs 7). [PR1300423](#)
- When CFP2-DCO-T-WDM-1 plugged in PTX PIC, when backward fr is enabled on far end, the convergence time is higher as extra delay (average 500 msec) incurred in triggering FRR, due to SW based polling. [PR1303820](#)
- Forwarding-option filter (FTF) with DSCP action is not supported on PTX1000 and other Provider Edge chipset platforms. [PR1310747](#)
- In the specific case of semi-graceful RCB reboot initiated by the internal shell command **vhclient init 0**, GRES takes longer to complete i.e 3 minutes as opposed to 21 seconds. The regular **request vmhost reboot** CLI command (graceful) and a jack-out-jack-in of the Routing Engine (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](#)
- Micro-BFD configuration with interface addresses is not yet supported on PTX on FPC3. [PR1341513](#)
- Newer B0 DCO modules(740-087314) HGFECC implementation is different and standardized vs. A0 (740-072229) which has different implementation causing link not to come up for interop between B0 and A0. [PR1394134](#)

Interfaces and Chassis

- On PTX10008 and PTX10016 routers, if you remove the redundant Switch Interface Board (SIB) after upgrading Junos OS from Release 17.4R1 or Release 17.2X75-D90 to a later release, then an alarm is not generated. This is a known behavior and has no impact on the performance of the router.

SEE ALSO

[New and Changed Features | 291](#)

[Changes in Behavior and Syntax | 304](#)

[Known Issues | 314](#)

[Resolved Issues | 319](#)

[Documentation Updates | 329](#)

[Migration, Upgrade, and Downgrade Instructions | 330](#)

Known Issues

IN THIS SECTION

- General Routing | 314
- Infrastructure | 318
- Interfaces and Chassis | 318
- Layer 2 Features | 318
- MPLS | 318
- Routing Protocols | 319

This section lists the known issues in hardware and software in Junos OS Release 17.4R3 for the PTX Series.

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

General Routing

- While upgrading from Junos OS Release 15.1F based images to Junos OS 16.x and later releases or downgrading from Junos OS Release 16.x to Junos OS Release 15.1F images, if the **validate** option is enabled, chassisd might crash and upgrade or downgrade will fail. This issue should not be seen if both base and target images are from Junos OS Release 15.1F or Junos OS Release 16.x and later. [PR1171652](#)
- The management process (daemon) might crash if the Openconfig package is installed immediately or within minutes of Network Agent package installation. This is a transient issue and will not impact any functionality. There is no action needed from the user's side in response to the crash. As a workaround, install Openconfig before installing Network Agent. [PR1265815](#)
- When an FPC goes offline or restarts, the source FPC sends traffic to destination FPC. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error**[PR1268678](#)
- On PTX5000 with FPC type 3 in rare condition FPC might crash during lo0.0 inet6 input filter. [PR1268875](#)

- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- On a PTX Series PIC with the CFP2-DCO-T-WDM transceiver installed, after repeated configuration rollback, the link sometimes takes a long time to come up. [PR1301462](#)
- When CFP2-DCO-T-WDM-1 plugged in PTX PIC, after FPC restart sometimes Carrier frequency offset tca is raised even when tca not enabled. [PR1301471](#)
- iLatency (calculated by differing producer timestamp and gRPC server timestamp) can sometimes be negative for Packet Forwarding Engine related telemetry packets due to drift in Routing Engine and Packet Forwarding Engine NTP servers. [PR1303376](#)
- When the **set forwarding-options l2circuit-control-passthrough** command 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](#)
- On a PTX Series router with a third-generation FPC, the error message is displayed when the FPC goes online or offline. [PR1322491](#)
- PTX3000 reports CCL (Chip to Chip Link) CRC errors while FPC3-SFF-PTX-1X is turned offline through CLI command or press offline button. The syslog error is generated by an FPC just before it goes offline, so there is no detectable traffic loss. The following messages are seen : Apr 2 08:43:00 fpc4 CMSNGFM: cmsngfpc_fm_send_spry_ctrl_ack: ev_id:11 fm_st:ALL fm_type:FPC_OFF fm_op:DEL Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 0 detected link error for S00F0_0(11,0,11)->FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: Logging statistics for FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: SOT:0x00000037649c2c43e Apr 2 08:43:00 fpc2 CCL: FrameCnt:0x000000000000419dc Apr 2 08:43:00 fpc2 CCL: LastCRCErrCnt:0x00000003 Apr 2 08:43:00 fpc2 CCL: AggrCRCErrCnt:0x0000000000000003 Apr 2 08:43:00 fpc2 CCL: AggrBERCnt:0x0000000000000001 Apr 2 08:43:00 fpc2 CCL: pe0-Avg-28nm-link-10-18 CRC error history (last 5 polls): Apr 2 08:43:00 fpc2 CCL: 0x0 0x0 0x0 0x0 0x3 Apr 2 08:43:00 fpc2 CCL: FEC Uncorrectable FEC Correctable Apr 2 08:43:00 fpc2 CCL: 00000004, 00000000 Apr 2 08:43:00 fpc2 CCL: 00000000, 00000000 Apr 2 08:43:00 fpc2 BEGIN Rx serdes info for asic pe0-0 serdes 18 Apr 2 08:43:00 fpc2 Signal & port condition for serdes_num 18 Apr 2 08:43:00 fpc2 Rx Signal : Signal Not OK Apr 2 08:43:00 fpc2 Rx Electrical Idle : High Apr 2 08:43:00 fpc2 Rx Frequency Lock: Set Apr 2 08:43:00 fpc2 Rx Port : Ready Apr 2 08:43:00 fpc2 DFE TAPs : -- snip -- Apr 2 08:43:00 fpc2 CCL: FrameCnt:0x00000000000041a0d Apr 2 08:43:00 fpc2 CCL: LastCRCErrCnt:0x00000003 Apr 2 08:43:00 fpc2 CCL: AggrCRCErrCnt:0x0000000000000003 Apr 2 08:43:00 fpc2 CCL: AggrBERCnt:0x0000000000000001 Apr 2 08:43:00 fpc2 CCL: pe0-Avg-28nm-link-14-22 CRC error history (last 5 polls): Apr 2 08:43:00 fpc2 CCL: 0x0 0x0 0x0 0x0 0x3 Apr 2 08:43:00 fpc2 CCL: FEC Uncorrectable FEC Correctable Apr 2 08:43:00 fpc2 CCL: 00000004, 00000000 Apr 2 08:43:00 fpc2 CCL: 00000000, 00000000 Apr 2 08:43:00 fpc2 BEGIN Rx serdes info for asic pe0-0 serdes 22 Apr 2 08:43:00 fpc2 Signal & port condition for serdes_num 22 Apr 2 08:43:00 fpc2 Rx Signal : Signal Not OK Apr 2 08:43:00 fpc2 Rx Electrical Idle : High Apr 2 08:43:00 fpc2 Rx Frequency Lock: Set Apr 2 08:43:00 fpc2 Rx Port : Ready Apr 2 08:43:00 fpc2 DFE TAPs : -- snip -- Apr 2 08:43:00 fpc2 CCL: Logging errors for FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: BER Err Apr 2 08:43:00 fpc2 CCL: Frame Lock Loss Apr 2 08:43:00 fpc2 CCL: Align Loss Apr 2 08:43:00 fpc2 CCL: Header Comparison Error Apr

2 08:43:00 fpc2 CCL: Header Preamble Error Apr 2 08:43:00 fpc2 CMSNGFM:
 cmsngfpc_platform_fm_periodic: PFE 0 detected link error for S00F1_0(14,0,14)->FPC02FE0(1,00) Apr
 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 1 detected link error for
 S00F0_0(11,0,11)->FPC02FE1(0,00) Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic:
 PFE 1 detected link error for S00F1_0(14,0,14)->FPC02FE1(1,00) User@PTX3000> show chassis
 hardware detail Hardware inventory: FPC 0 REV 43 750-057064 ACPV7514 FPC3-SFF-PTX-1XCPU
 BUILTIN BUILTIN SMPC PMB FPC 2 REV 40 750-057064 ACPJ9145 FPC3-SFF-PTX-1XCPU BUILTIN
 BUILTIN SMPC PMB FPC 4 REV 43 750-057064 ACPR8506 FPC3-SFF-PTX-1XCPU BUILTIN BUILTIN
 SMPC PMB SIB 0 REV 10 750-057067 ACPJ8829 SIB3-SFF-PTX SIB 1 REV 10 750-057067 ACPJ8683
 SIB3-SFF-PTX SIB 2 REV 10 750-057067 ACPJ8843 SIB3-SFF-PTX SIB 3 REV 10 750-057067 ACPJ8920
 SIB3-SFF-PTX [PR1348733](#)

- On next-generation Routing Engine (NG-RE), a failure of the Hardware Random Number Generator (HWRNG) leaves the system in a state where is not enough entropy available to operate. [PR1349373](#)
- The log of "SMART ATA Error Log Structure error: invalid SMART checksum." might be seen on FPC with WINTEC mSata SSD. [PR1354070](#)
- Unsuccessful connection attempts will not be logged on the backup SPMB. [PR1369731](#)
- When an Routing Engine reboots and comes up again it sends Gratuitous ARP packets to the internal interfaces in order to advertise its MAC address. These packets get in to the UKERN running on the FPC, which drops these packets. The messages seen here are printed just before dropping these packets. These error messages are harmless and do not disrupt working of any feature. [PR1374372](#)
- In certain scenarios where flows are sampled through aggregate bundles when jflow sampling is enabled, the following harmless error logs can be seen: [Tue Oct 30 18:17:40.648 LOG: Info] **expr_get_local_pfe_child_ifl: cannot find child ifl of agg ifl 74 for this fpc** [Tue Oct 30 18:17:40.648 LOG: Info] **flowtb_get_cpu_header_fields: Failed to find local child ifl for 74** [Tue Oct 30 18:17:40.648 LOG: Info] **fpc0 cannot find stream on [hostname]** [PR1379227](#)
- Output of Jflow sensor is changed for a FPC that is not configured. [PR1379770](#)
- Due to transient Hardware condition single-bit error (SBE) event are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. [PR1384435](#)
- On all PTX Series routers or with CoS deployed, all the physical member interfaces of aggregated Ethernet (AE) might drop the packets in lower priority queues when micro-bursts are received. These micro-burst are typically due to the speed differential between ingress interface (for example, 100G) and egress interface (for example, 10G). Typically it occurs when a large burst of high priority traffic and lower priority traffic arrive simultaneously. [PR1385454](#)
- Junos may falsely report warning messages at **Module voltage** and **Module voltage low warning** fields in **show interfaces diagnostics optics** output. [PR1394266](#)
- Sometimes the SFP+ read does not work on one or more ports of the "LC1103 - 2C / 6Q / 60X". If this happens, the corresponding SFP+ module will not get detected and will not be displayed at the Routing

Engine CLI under the output of **show chassis hardware**. As a workaround, re-seat the SFP+ Module.
[PR1412897](#)

- In PTX with FPC3-PTX PIC, one of the interfaces on them might not come up after an interface of peer device flapping in short intervals and then restart the local FPC. Due to the BCM8238x chip of Broadcom with a wrong re-timer leading to the local interface remain in "down" state. [PR1428307](#)
- The timestamp reported for packet arrival in NetFlow records will report inaccurate time due to the synchronization issue with NTP. [PR1431498](#)
- Due to a bug in software, it may fail to clean up old entries during filter change operation. The filter manager is associating filters to transit packets where the ASIC unable to locate the program as it has been deleted earlier. Hence all transit packets are hitting `flt.Dispatcher.flt_err` and got dropped on Packet Forwarding Engine. [PR1433648](#)
- JDI-L3VPN-REGRESSIONS: More packet loss seen after ISSU with InterAS L3VPN OptionB configuration [PR1435578](#)
- On PTX1000 platform, in case of a jlock hog lasts for more than 5 seconds, FPC reboot might be seen. [PR1439929](#)
- On PTX Series, if particular 100G port is used, CPU might hang or interface might be stuck down on the 100G port. This issue may cause traffic disruption in the network. [PR1440526](#)
- On the platforms that do not support Router Advertisement Guard (RA Guard), such as PTX Series, after issuing the **show access-security router-advertisement-guard** command, the process `jdhcpd` might crash. [PR1446034](#)
- Currently ISIS is sending system host-name instead of system-id in OC paths in `lsdb` or Adjacency `xpaths` in periodic streaming and on-change notification. [PR1449837](#)
- On PTX5000 Router, the 100G interface might not come up after flapping due to optic reliability issues. [PR1453217](#)
- In PTX 3000/5000 platforms with CB2-PTX (Control Board), there is an existence of an errata on a clock signal component manufactured by a third-party supplier, which might cause the Switch Processor Mezzanine Board (SPMB) and Switch Interface Boards (SIBs) failure, eventually, traffic will be interrupted. [PR1460992](#)
- On PTX10K platforms, FPC might restart if there is some corruption in BCM (Broadcom) switch (a small internal ethernet switch, instead of PFE engine) inside the FPC. It is a timing issue. The reason is that the PCIe speed configuration for BCM switch is not correct. And this issue is resolved in some FPC U-boot versions. [PR1464119](#)
- When tunnel-services are configured on a PIC, the optics measurements that subscribed via gRPC might not be streamed. [PR1468435](#)

Infrastructure

- If you pulled out a USB from the system while files are being copied, the kernel will panic and the system will restart. [PR1425608](#)
- Junos packages may have incorrectly registered as "unsupported". [PR1427344](#)

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after upgrade. This is because of the old version of `/var/db/cfm.db`. [PR1281073](#)

Layer 2 Features

- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case the subscriber might need more time to get IP address assigned. The subscriber might remain in this state until it's lease expires if it has previously bound with the address in the option 50. [PR1435039](#)

MPLS

- LDP to BGP stitching with eBGP indirect next hop having an implicit null label does not work. It does work when BGP indirect next hop has a real label.

As a workaround:

- Ensure the peer advertises a real label by adding another router between the egress and ingress PE devices.
- Use IBGP that gets resolved over LDP or RSVP-TE LSPs. This will ensure that the BGP indirect next hop has a real label. [PR1254702](#)
- The optimization timer is being updated in an incorrect manner in the code path. Due to this a particular check fails when the exponential increase function is called. [PR1416948](#)
- In a corner case on Junos platform, where the family ccc is configured along with any other existing family within the same interface, like inet, inet6, etc. (basically, Junos never allows to do so, but somehow a customer did it). And if the family ccc is deleted from the interface, which might cause kernel crash and the device reboot automatically, so all the traffic will be interrupted. [PR1478806](#)

Routing Protocols

- In BGP segment routing traffic engineering (SRTE) scenario, process rpd might crash when configuration statement "extended-nexthop-color" is added or removed from the BGP configuration. [PR1442952](#)

SEE ALSO

New and Changed Features	291
Changes in Behavior and Syntax	304
Known Behavior	312
Resolved Issues	319
Documentation Updates	329
Migration, Upgrade, and Downgrade Instructions	330

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.4R3 | 320
- Resolved Issues: 17.4R2 | 323
- Resolved Issues: 17.4R1 | 327

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: 17.4R3

Forwarding and Sampling

- The flow label statistics are retrieved periodically by pfd for PTX or TVP platforms, if the statistics reply becomes very big number, the pfd might crash hence affecting traffic. [PR1452363](#)

General Routing

- The mgd might crash when Ephemeral DB is used. [PR1305424](#)
- Unnecessary interface-related log messages might be seen after configuration commit on PTX series routers. [PR1309575](#)
- System might hang during initial configuration when doing a USB install or request vmhost zeroize/request system zeroize. [PR1311553](#)
- OCST /junos/system/linecard/interface/./interfaces/:Packet Forwarding Engine packet drop is seen on PTX5000 when there is 100ms RTT delay between DUT and collector. [PR1316429](#)
- PTX1000:flabel Mem alloc failure is followed by FPC core files.[PR1318595](#)
- Interface down due to **PFE Marked Disabled** on PECHIP, causing traffic loss. [PR1320413](#)
- Invalid programming of interfaces during Packet Forwarding Engine initialization may lead to traffic black hole on PTX platform. [PR1331299](#)
- Linux based FPC should close chassis TCP control connection immediately when J-UKERN is crashing. [PR1347536](#)
- PTX Series: FPC process crashes after J-Flow processes a malformed packet (CVE-2019-0014). [PR1348417](#)
- Traffic drop might be seen after GRES if uRPF is configured. [PR1354285](#)
- Adding a warning when commit that the host interface may stop sending packets on PTX1000, PTX5000 and PTX10000 when using outbound firewall filter with syslog option. [PR1354580](#)
- Extended traffic loss when performing ISSU/GRES with aggregated Ethernet interface configured with LACP. [PR1365316](#)
- The 'Normal discards' pfe statistics traffic counter might increase at a higher rate when Inline-Jflow or sFlow is enabled. [PR1368208](#)
- Unexpected incrementing of counters on the interface. [PR1370062](#)
- Traffic might be dropped on third-generation FPCs on PTX Series routers. [PR1378392](#)
- Layer 3 VPN traffic might be dropped due to one core-facing interface down. [PR1380783](#)
- FPC might crash on PTX Series after lo0 filter change. [PR1380917](#)
- BFD flaps are seen on PTX or QFX10000 platforms with inline BFD. [PR1389569](#)
- Forwarding issue on mixed link-speed aggregated Ethernet interface after FPC reloads. [PR1390417](#)

- The lcmd core file and FPC restart might be seen. [PR1391443](#)
- The **show chassis fpc** command on PTX1000 and PTX10000 series routers shows incorrect buffer memory utilization. [PR1397612](#)
- When issuing "request vmhost reboot routing-engine both", no user confirmation is asked. [PR1397912](#)
- CPU overuse may be observed on PTX/QFX10000 Series platform. [PR1399369](#)
- Unexpected alarm might be shown on NG-RE [PR1399654](#)
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- The TCP connection between ppmdd and ppmann might be dropped due to a kernel issue. [PR1401507](#)
- The log message **JAM HW data base open failed for ptx5kpic_3x400ge-cfp8** during commit. [PR1403071](#)
- Incorrect mem stat message is seen in FPC logs of PTX Type 1 FPC. [PR1404088](#)
- PTX3000: FPCs are not able to come online for tens of minutes after a reboot of the chassis. [PR1404611](#)
- 100G SR4 Optics with part number 740-061405 should be displayed as **QSFP-100G-SR4-T2**. [PR1405399](#)
- No chassis alarm is raised on PTX1000 when PEM is removed or power lost to PEM. [PR1405430](#)
- Layer 2 VPN might flap repeatedly after the link up between PE and CE. [PR1407345](#)
- The Packet Forwarding Engine might get disabled unexpectedly due to a auto correctable non-fatal hardware error on PTX Series router. [PR1408012](#)
- The link flaps occur when a 100g QSFP is inserted into PTX Series which LFM (Link-Fault Management) is configured. [PR1408204](#)
- The port at FPC (for example, JNP10K-LC1101) might fail to come up. [PR1409585](#)
- Indirect-next-hop pointing to unknown unilist stuck with weight 65535 may occur after a link flap. [PR1409632](#)
- Hostname does not update at FPC shell after system configuration change on CLI. [PR1412318](#)
- Junos PCC may reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- The Layer 2circuit egress PE might drop the traffic in FAT+CW enabled L2circuit scenario when another FAT+CW enabled L2circuit PW flaps. [PR1415614](#)
- Traffic loss could be seen for duration of hold-time down timer when flapping an interface with hold-time down timer configured. [PR1418425](#)
- RX alarms are not set as according to the threshold value configured for the DCO Tunable Optics. [PR1419204](#)
- An interface may go to downstate on QFX10000/PTX10000 platform. [PR1421075](#)
- Virtual Chassis may become unstable and FXPC core files when there are a lot of configured filter entries. [PR1422132](#)

- 4x10G interfaces on PTX3000/PTX5000 FPC type 3 might not come up after frequently flap for a large amount of time. [PR1422535](#)
- While committing huge configuration, the customer is seeing the error **error: mustd trace init failed**. [PR1423229](#)
- Traffic is dropped after FPC reboot with aggregated Ethernet member links deactivated by a remote device. [PR1423707](#)
- Specific interface on P3-15-U-QSFP28 PIC card remains down until another interface comes up. [PR1427733](#)
- On PTX10000 platforms certain interfaces might go to down state. [PR1427883](#)
- The jumbo frame size packets are dropped when max MTU is configured. [PR1428094](#)
- The l2cpd process might crash and generate a core dump when interfaces are flapping. [PR1431355](#)
- Traffic loss might be seen on PTX10000 platforms using line card LC1105. [PR1433300](#)
- IPv6 neighbor solicitation packets getting dropped on PTX. [PR1434567](#)
- Interfaces on PTX Series might not come up after FPC restart or port flap. [PR1442159](#)
- Gladiator: BCM FW needs to be upgraded to DE2E. [PR1445473](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)

Infrastructure

- The FPC might go down on some vmhost based PTX platforms. [PR1367477](#)
- The error **jlaunchd: disk-monitoring is thrashing, not restarted** might be seen. [PR1380032](#)
- The command **request system recover oam-volume** might fail on PTX Series routers. [PR1425003](#)

Interfaces and Chassis

- PE Chip:pe0[0]: IPW: **oversize_drop error** causes Major error on FPC. [PR1375030](#)
- The syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** upon LFM related configuration commit on aggregated Ethernet interfaces. [PR1423586](#)
- Some ports on PTX Series routers might remain down after rebooting the FPC/device at remote side. [PR1429315](#)
- On PTX3000 or PTX5000 routers with FPC3 installed, some 100G ports might remain down after rebooting the FPC or device at remote side. [PR1429315](#)

MPLS

- The rpd might crash when executing Routing Engine switch-over under BGP environment and route churn occurs. [PR1373313](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition. [PR1374102](#)

- A RSVP-signaled LSP might stay in down state after a link in the path flaps. [PR1384929](#)
- The rpd might crash when LDP route with indirect next-hop is deleted. [PR1398876](#)
- A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based-stats are used. [PR1401152](#)
- LDP routes might flap if committing any configuration changes. [PR1416032](#)
- Traffic loss might be seen after LDP session flaps rapidly. [PR1436119](#)
- The transit packets might be dropped if an LSP is added or changed on PTX Series device. [PR1447170](#)

Platform and Infrastructure

- The dcd Micro BFD seems to be failing in dcd_commit_check log file even when BFD is not configured. [PR1300796](#)
- The **commit synchronize** command might fail due to several internal connections stuck. [PR1394370](#)
- Some files are missing during log archiving. [PR1405903](#)

Routing Protocols

- Routing Engine-based micro BFD packets do not go out with configured source IP when the interface is in logical-system. [PR1370463](#)
- The rpd process might crash after executing commit the configuration related to mapping-server-entry. [PR1379558](#)
- RPD core files on backup routing-engine during neighbor-ship flap when using **authentication-key** with size larger than 20 characters. [PR1394082](#)
- Syslog message is seen whenever prefix sid coincides with the node sid. [PR1403729](#)
- Dynamic routing protocol flapping with vmhost Routing Engine switchover on NG-RE. [PR1415077](#)
- Route churn might be seen after changing maximum-prefixes configuration from value A to vlaue B. [PR1423647](#)

VPNs

- Memory leak might happen if PIM messages received over an MDT (mt- interface) in Draft-Rosen MVPN scenario. [PR1442054](#)

Resolved Issues: 17.4R2

General Routing

- The commands **restart na-grpc-server** and **restart na-mqtt** for MTRE does not work. [PR1284121](#)
- PTX1000s routers are periodically exporting IPFIX flow packets with high octet values. [PR1286427](#)

- On a PTX1000, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3-20170721 fails frequently on enabling sampling. [PR1296533](#)
- Interfaces might go down when Packet Forwarding Engine encounters **TOE::FATAL ERROR**. [PR1300716](#)
- On a PTX3000 platform, powering on a FPC (PTX-IPLC-B-32) card might cause the other FPC cards to reboot. [PR1302304](#)
- Internal latency is high during the initial subscription of sensors. [PR1303393](#)
- Repeated log message **%PFE-3 fpcX expr_nh_index_tree_ifl_get and expr_nh_index_tree_ipaddr_get** is observed when the sampling packet is discarded with a log(or syslog) configuration statement under the firewall filter. [PR1304022](#)
- The "interface hold-time down" timer does not take effect on a PTX5000 with an optical interface. [PR1307302](#)
- Packet Forwarding Engine error messages are flooding as **expr_sensor_update_cntr_to_sid_tree** after a deletion and rollback as **protocols isis source-packet-routing node-segment** . [PR1309288](#)
- On a PTX10000, a suppress chassis alarm for switched off PEM is seen. [PR1311574](#)
- The SIB LED on the FPD comes to GREEN-STEADY state even before an SIB comes online. [PR1311632](#)
- PTX10000 router does not bounce FPC without warning or alarm for different port speed settings. [PR1311875](#)
- The rpd process generates a core file after multiple session flaps on a scale setup. [PR1312169](#)
- Many logs are generated after executing many Vhclient related commands. [PR1315128](#)
- Memory leak in chassisd daemon is noticed while streaming active telemetry subscriptions. [PR1315672](#)
- The Packet Forwarding Engine on a PTX FPC3 or PTX10000 line card might be disabled if the interface connecting to the Packet Forwarding Engine goes down. [PR1315823](#)
- The physical interfaces might generate framing errors when ports are connected to an odd interface. [PR1317827](#)
- On MX0016 router, after Jack-out/Jack-in, FPCs shows up as "No-Power" for some time; FPC, however, comes up. [PR1319156](#)
- There is no traffic flowing with IPv6 payload prefixes on PTX Series platform. [PR1319273](#)
- PTX10000 routers for 100G LR4 optics with part number 740-061409 changes the display of the **show chassis hardware** command to QSFP-100G-LR4-T2. [PR1322082](#)
- The rpd process might crash when the OpenConfig package is upgraded with JTI streaming data in the background. [PR1322553](#)
- On Junos OS MPC7, MPC8, and MPC9, PTX-FPC3 (FPC-P1, FPC-P2), PTX3000-FPC3, and PTX1000 line cards might crash upon receipt of specific MPLS packet (CVE-2018-0030). [PR1323069](#)
- On a PTX1000, the local time on an FPC might be different from the local time on a Junos-VM or VM-host. [PR1325048](#)

- The GRE traffic is not decapsulated by the firewall filter. [PR1325104](#)
- PTX Series routers MKA sessions are not coming up after changing CA parameters such as transmit-interval, and key-server-priority. [PR1325392](#)
- MPLS traceroute fails across the PTX Series platform. [PR1327609](#)
- On a PTX5000 with FPC3 line cards, on a PTX10000, and on PTX1000 platforms, output firewall filters that are configured with "syslog" and "discard" actions do not perform the "syslog" action. [PR1328426](#)
- PTX10000 line card might reboot continuously after upgrading to Junos OS Release 17.2R1 or later if HMC BIST fails. [PR1330618](#)
- There is a link instability after a link-down event on PTX Series routers. [PR1330708](#)
- A PTX5000 FPC might reboot in certain rare scenarios when **interface-specific policer** is configured. [PR1335161](#)
- The directory where the init and configuration files are stored are removed when **request system zeroize** is executed. Hence, the telemetry process does not start after the zeroize is done. [PR1336004](#)
- A member of an IPv4 unicast next hop might get stuck in "Replaced" state after an interface flap. [PR1336201](#)
- Disabling a breakout 10G port on interface et-0/0/5 will unexpectedly disable another breakout 10G port on interface et-0/0/5. [PR1337975](#)
- FPC/FPC2/FPC E on a PTX Series does not forward traffic. [PR1339524](#)
- The link goes down on a PTX3000. The PTX5000 with an FPC3 is inserted after the router reboots or the link flaps. [PR1340612](#)
- The interface might flap continuously after reboot on PTX5000 and PTX3000 platform with P3-24-U-QSFP28 PIC. [PR1342681](#)
- On a PTX1008, the 30-port coherent line card (DWDM-IC) does not come up. [PR1344732](#)
- The FPC was rebooted a few minutes after loading the configuration. [PR1346467](#)
- Sensors are not getting cleared up after doing Routing Engine switchover. [PR1347779](#)
- MPLS traceroute for P2MP LSPs configured with link-protection causes FPC crash. [PR1348314](#)
- The interface of 15 100G ports PIC might delay 60 seconds to come up. [PR1357410](#)
- P2MP LSP replication traffic loss on an aggregated Ethernet bundle after a member link is down on PTX Series routers. [PR1359974](#)
- The route stuck might be seen after BGP neighbor and route flapping. [PR1362560](#)
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with **Link-Layer-Down** flag set. [PR1365263](#)

- On a PTX Series IPLC (OPT3-SFF-PTX FPC), a first J-UKERN crash triggers multiple secondary J-UKERN crashes. [PR1365791](#)
- The **commit** or **commit check** might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. [PR1368992](#)

Infrastructure

- The ixlv interface statistics are not accounting properly. [PR1313364](#)
- PTX Series routers might get to abnormal state because of the malfunction of the protection mechanism for the F-Label. [PR1336207](#)

MPLS

- Traffic drops during NSR switchover for RSVP P2MP provider tunnels used by MVPN. [PR1293014](#)
- Traffic loss for static LSP configured with the **stitch** configuration statement. [PR1307938](#)
- The rpd process might crash on the backup Routing Engine due to memory exhaustion. [PR1328974](#)
- The rpd might crash with MPLS traceoption configured. [PR1329459](#)
- MPLS LSP statistics are not shown in the CLI command. **show mpls lsp ingress statistics**. [PR1344039](#)
- On PTX1000, PTX10000, or QFX10000 platform, when hybrid memory cube (HMC) error occurs, label switched paths (LSPs) might go down because of the incorrect bandwidth requested for auto-bandwidth adjustment. [PR1374102](#)

Platform and Infrastructure

- Continuous log messages occur. For example, you see **tftpd[23724]: Timeout #35593 on DATA block 85**. [PR1315682](#)
- Traffic black hole seen along with **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH** messages. [PR1349332](#)
- Traffic black hole seen along with **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH** messages. [PR1357707](#)

Routing Protocols

- With BGP LU FRR in an inter-AS scenario, a very high FRR time is visible once the link is up. [PR1307258](#)
- The rpd process might constantly consume high CPU in a BGP setup. [PR1315066](#)
- The primary path of an MPLS LSP might switch to another address. [PR1316861](#)
- The rpd process might crash after deactivating the passive interface under IS-IS. [PR1318180](#)
- The rpd process might crash continuously on both Routing Engines when **backup-spf-options remote-backup-calculation** is configured in an IS-IS protocol. [PR1326899](#)
- The rpd process generates a core file at **ispfc_incrementally_mend_one_postf_sp_tree (postf_spf_res=<optimized-out>, topo=<optimized-out>)** at

```
.././.././.././../src/junos/usr.sbin/rpd/lib/igp-spf-compute/igp_spf_compute_ti_lfa.c:3364" PTX1K.  
PR1339296
```

- A protocol churn might lead the rpd process to crash. [PR1341466](#)
- The rpd process generates a core file while running a streaming telemetry test. [PR1347431](#)

VPNs

- In a specific CE device environment in which asynchronous notification is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly.[PR1282875](#)

Resolved Issues: 17.4R1

General Routing

- PTX1000 : `ch_get_product_attribute.324`: Cannot find chassisd error message appears while loading image. [PR1217505](#)
- The rpd might crash on platforms with 64-bit X86 RE if IPv6 is configured. [PR1224376](#)
- On PTX Series platforms, chassisd thread is not getting CPU resources for 200 seconds and multiple chassisd core files are continuously generated. [PR1226992](#)
- The "validation-state:unverified" routing entry might not be displayed with proper location when using the show route command. [PR1254675](#)
- The routing protocol process (rpd) might crash after flapping BGP sessions and routes. [PR1269327](#)
- 100Base-ER4 (740-045420) is displayed as "UNKNOWN" in **show chassis hardware** in Junos OS Release 15.1R5.5. [PR1280089](#)
- FPC cards might go offline due to fabric healing in PTX3000 with SIB-SFF-PTX-240-S platform. [PR1282983](#)
- "Host 0 RTC Battery failure" error messages are seen on PTX1000, QFX10000-series after upgrade to Junos version 16.1. [PR1287128](#)
- The MPLS TTL might be reset to 255 on third-generation PTX Series FPCs if **mpls no-propagate-ttl** protocols configuration statement is configured. [PR1287473](#)
- LSP traffic gets silently dropped or discarded after link goes down in bypass path. [PR1291036](#)
- The rpd core file might be generated when restarting the process through CLI. [PR1291110](#)
- Incorrect SNMP OID values are sent in SNMP traps for removal or insertion of front panel display on PTX Series routers. [PR1294741](#)
- LINK LED is "RED" when the port is disabled on PTX Series routers. [PR1294871](#)
- The rpd core might be generated after interface or BGP flapping. [PR1294957](#)
- The chassisd process might run out of memory and restart on PTX1000 platform. [PR1295691](#)

- PTX5K/SyncE (ESMC): clock is not getting locked if the source interface is a member link of an ae bundle. [PR1296015](#)
- CoS escalation: Alarms and syslog errors are seen with priority strict-high on AF4 queue, on the oversubscription cases (1X100G egress to 1X10G egress setup). [PR1297343](#)
- **PE Chip: FATAL ERROR!! from pe0[0]: HMCIF:** might trigger FPC crash or slow route/next-hop installation processing. [PR1300180](#)
- PTX Series FPC3 will drop MPLS packets if its oif has inet MTU that is less than the MPLS packet size. [PR1302256](#)
- Heap memory leak might be observed on PTX Series FPCs during a multicast route installation into the Packet Forwarding Engine. [PR1302303](#)
- The third-generation FPC (FPC3-SFF-PTX) is not booting up on Control Board/Routing Engine systems. [PR1303295](#)
- On PTX3000 and PTX5000 platforms, the 100G interfaces might not come up. [PR1303324](#)
- If MPLS LSP self-ping is enabled (self-ping is enabled by default), kernel might panic with the error message **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)
- PTX3000 with RCB-PTX Routing Engine might not be online or recognize IPLCs. [PR1304124](#)
- The 10g interface might flap if it is set to 100g speed. [PR1315079](#)
- The physical interfaces might generate framing errors when ports are connecting odd interfaces. [PR1317827](#)
- The physical interfaces might generate framing errors when ports are connecting to odd interface. [PR1317827](#)
- No traffic is flowing with IPV6 payload prefixes on PTX platform. [PR1319273](#)
- PFT : RCB restarts continuously after executing **request system reboot**. [PR1320977](#)

Infrastructure

- The **show system users** CLI command output displays a larger number of users than that are actually using the router. [PR1247546](#)

Interfaces and Chassis

- The interface might flap when performing Routing Engine switchover if the member link of an ae interface is configured with framing settings. [PR1287547](#)
- 100-Gigabit Ethernet interfaces might not come up if **otn-options laser-enable** is configured on PTX Series platforms. [PR1297164](#)
- LFM discovery state appears as Fault for aggregate Ethernet interface after GRES. [PR1299534](#)

Multiprotocol Label Switching (MPLS)

- Stale RSVP LSP entry after NSR switchover and session is not refreshed. [PR1292526](#)

- The rpd might crash if the MPLS LSP path change occurs. [PR1295817](#)

Platform and Infrastructure

- Mgd generates core file when downgrading from Junos OS Release 17.3-20170721 to 16.1X65D40.2. The mgd core is also overwritten if attempted multiple times. [PR1296504](#)

Routing Protocols

- A few BFD sessions are flapping while coming up after FPC restart/reboot. [PR1274941](#)
- The rpd generated core files multiple times when it received an “OPEN” message from an existing BGP peer. [PR1299054](#)
- With BGP LU FRR in Inter-As scenario, a very high FRR time is seen once link is up. [PR1307258](#)
- Assignment of SUB-TLV values for Segment routing TE policy SUB-TLVs. [PR1315486](#)

SEE ALSO

New and Changed Features 291
Changes in Behavior and Syntax 304
Known Behavior 312
Known Issues 314
Documentation Updates 329
Migration, Upgrade, and Downgrade Instructions 330

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 documentation for PTX Series.

SEE ALSO

New and Changed Features 291
Changes in Behavior and Syntax 304
Known Behavior 312
Known Issues 314
Resolved Issues 319
Migration, Upgrade, and Downgrade Instructions 330

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 330
- Upgrading a Router with Redundant Routing Engines | 330
- Basic Procedure for Upgrading to Release 17.4 | 331

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

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

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

Basic Procedure for Upgrading to Release 17.4

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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

To download and install Junos OS Release 17.4R3:

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 17.4R3 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-17.4R3.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-17.4R3.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 | 291](#)

[Changes in Behavior and Syntax | 304](#)

[Known Behavior | 312](#)

[Known Issues | 314](#)

[Resolved Issues | 319](#)

[Documentation Updates | 329](#)

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- [New and Changed Features | 335](#)
- [Changes in Behavior and Syntax | 347](#)
- [Known Behavior | 354](#)
- [Known Issues | 358](#)
- [Resolved Issues | 369](#)
- [Documentation Updates | 391](#)
- [Migration, Upgrade, and Downgrade Instructions | 391](#)

These release notes accompany Junos OS Release 17.4R3 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- [Release 17.4R3 New and Changed Features | 336](#)
- [Release 17.4R2 New and Changed Features | 336](#)
- [Release 17.4R1 New and Changed Features | 336](#)

This section describes the new features for the QFX Series switches in Junos OS Release 17.4R3.

NOTE: The following QFX Series platforms are supported in Release 17.4R3: QFX5100, QFX5110, QFX5200, QFX10002, QFX10008, and QFX10016.

Release 17.4R3 New and Changed Features

- There are no new features or enhancements to existing features for QFX Series Switches in Junos OS Release 17.4R3.

Release 17.4R2 New and Changed Features

Restoration Procedures and Failure Handling

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

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

Release 17.4R1 New and Changed Features

Hardware

- **QFX10000-30C-M line card (QFX10008 and QFX10016 switches)**--Starting with Junos OS Release 17.4R-S2, the QFXF10000-30C-M line cards provides 30 ports of either 100-gigabit or 40-gigabit QSFP28 with MACsec features.

Class of Service (CoS)

- **Priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic (QFX5110 and QFX5200 switches)**—Starting in Junos OS Release 17.4R1, to support lossless traffic across Layer 3 connections to Layer 2 subnetworks on QFX5110 and QFX5200 switches, you can configure priority-based flow control (PFC) to operate using 6-bit DSCP values from Layer 3 headers of untagged VLAN traffic, rather than IEEE 802.1P priority values in Layer 2 VLAN-tagged packet headers. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).

To enable DSCP-based PFC, map a forwarding class to a PFC priority using the **pfc-priority** statement, define a congestion notification profile to enable PFC on traffic specified by a 6-bit DSCP value, and set up a classifier for the DSCP value and the PFC-mapped forwarding class.

[See [Understanding PFC Using DSCP at Layer 3 for Untagged Traffic](#).]

EVPNs

- **Support for LACP in EVPN active-active multihoming (QFX5100, QFX5100 Virtual Chassis, QFX5110, and QFX5200 switches)**—Starting with Junos OS Release 17.4R1, an extra level of redundancy can be

achieved in an Ethernet VPN (EVPN) active-active multihoming network by configuring the Link Aggregation Control Protocol (LACP) on both the endpoints of the link between the multihomed customer edge (CE) and provider edge (PE) devices. The link aggregation group (LAG) interface of the multihomed CE-PE link can either be in the active or in the standby state. The interface state is monitored and operated by LACP to ensure fast convergence on isolation of a multihomed PE device from the core. When there is a core failure, a traffic black hole can occur at the isolated PE device. With the support for LACP on the CE-PE link, at the time of core isolation, the CE-facing interface of the multihomed PE device is set to the standby state, thereby blocking data traffic transmission from and toward the multihomed CE device. After the core recovers from the failure, the interface state is switched back from standby to active.

To configure LACP in EVPN active-active multihoming network:

- On the multihomed CE device include the `lacp active` statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy.
- On the multihomed PE device include the `lacp active` statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy, and include the `service-id` number statement at the `[edit switch-options]` hierarchy.

[See [Understanding LACP for EVPN Active-Active Multihoming](#).]

- **EVPN pure type-5 route support (QFX5110 switches)**—Starting with Junos OS Release 17.4R1, you can configure pure type-5 routing in an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which is used to carry the MAC address of the sending switch and to provide next-hop reachability for the prefix. To configure pure type-5 routing include the `ip-prefix-routes advertise direct-nexthop` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level. To enable two-level equal-cost multipath (ECMP) next hops in an EVPN-VXLAN overlay network, you must also include the `overlay-ecmp` statement at the `[edit forwarding-options vxlan-routing]` hierarchy level.

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

- **SPRING support for EVPN (QFX10000 switches)**—Starting in Junos OS Release 17.4R1, Junos OS supports using Source Packet Routing in Networking (SPRING) as the underlay transport in EVPN. SPRING tunnels enable routers to steer a packet through a specific set of nodes and links in the network.

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

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

- **Support for duplicate MAC address detection and suppression (QFX10000 switches)**— When a MAC address relocates, PE devices can converge on the latest location by using sequence numbers in the extended community field. Misconfigurations in the network can lead to duplicate MAC addresses. Starting in Junos OS Release 17.4R1, Juniper supports duplicate MAC address detection and suppression.

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

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

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

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

General Routing

- **Enhancement to show chassis forwarding-options command (QFX5200 Virtual Chassis)**—Starting in Junos OS Release 17.4R1, the show **chassis forwarding-options** command displays information about memory banks for QFX5200 Virtual Chassis only for the master. This information is not displayed for all the other members. Memory banks can be partitioned among different types of forwarding table entries through the Unified Forwarding Table feature. Values remain the same across all members. All configuration changes for the Unified Forwarding Table are made through the Master.

[See [show chassis forwarding-options](#).]

Interfaces and Chassis

- **Support for resilient hashing for LAGs and ECMP (QFX10000)**—Starting with Junos OS Release 17.4R1 on QFX10000 switches, you can prevent the reordering of flows to active paths in link aggregation groups (LAGs) or ECMP when one or more paths fail. Only flows that are on inactive paths are redirected. It overrides the default behavior of disrupting all existing, including active, TCP connections when an active path fails. You can optionally set a specific value for the resilient-hash seed that differs from the hash-seed value that will be used by the other hash functions on the switch. A resilient hashing configuration on ECMP is applied through use of a route policy.

[See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping](#).]

- **Enterprise profile for Precision Time Protocol (PTP) (QFX10002 switches)**—Starting with Junos OS Release 17.4.1, 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 timeout duration for self-ping on an LSP instance (QFX Series)**—Starting in Junos OS 17.4R1, the default timeout duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also manually configure the self-ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level. By default, self-ping is enabled. The LSP types such as CCC, P2MP, VLAN-based , and non-default instances do not support self-ping . You can configure the **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.
- **Support for label history for MPLS protocol (QFX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entry per label . By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range.
The **clear mpls label history** command clears the label history details.
- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (QFX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage that

is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP 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 resignalling 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 | 347](#)

[Known Behavior | 354](#)

[Known Issues | 358](#)

[Resolved Issues | 369](#)

[Documentation Updates | 391](#)

[Migration, Upgrade, and Downgrade Instructions | 391](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 348](#)
- [EVPNs | 348](#)
- [General Routing | 348](#)
- [Interfaces and Chassis | 348](#)
- [Management | 349](#)
- [MPLS | 349](#)
- [Network Management and Monitoring | 350](#)
- [Routing Policy and Firewall Filters | 352](#)
- [Security | 352](#)
- [Software Licensing | 353](#)
- [Virtual Chassis | 353](#)

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

Class of Service (CoS)

- When you configure a **transmit-rate**, you must also configure a **guaranteed-rate** under **traffic-control-profiles**. If you commit a configuration of a **transmit-rate** without a **guaranteed-rate**, a warning message is displayed and the default scheduler map is applied.

EVPNs

- **Change to the show vlans evpn command (QFX5100 switches)**—Starting with Junos OS Release 17.4R2, the **show vlans evpn** command is replaced by the **show ethernet-switching evpn** command.

General Routing

- **Change in default value for port ID TLV for QFX5200 switches**—In Junos OS Release 17.4R1, for QFX5200 switches, the default value used for port ID TLV in LLDP messages is interface name, not SNMP index.

Interfaces and Chassis

- **Commit Error thrown when GRE interface and Tunnel source interface configured in different routing instances (QFX Series)**—In Junos OS Releases 17.3R4 and 17.4R3, QFX series switches does not support configuring GRE interface and the underlying tunnel source interface in two different routing instances. If you try this configuration, it will result in a commit error with the following error message:

error: GRE interface (gr-0/0/0.0) and its underlying tunnel source interface are in different routing-instances

error: configuration check-out failed

[See [Understanding Generic Routing Encapsulation](#) .]

- **Packets with MTU size greater than the default value are dropped (QFX5110)**—In Junos OS Releases 17.3R3 and 17.4R2, on QFX5110 switches, setting maximum transmission unit (MTU) on the L3 interface does not take effect and packets with MTU size greater than the default value are dropped.

[See [mtu](#).]

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (QFX Series)**—In Junos OS Release 17.4R3, the **show lacp interfaces | display xml** command displays a new XML tag element **<lacp-hold-up-state>**. The **<lacp-hold-up-state>** displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information

for all interfaces were in a single `<lacp-hold-up-information>` XML tag. Now, for each interface it is displayed in a separate `<lacp-hold-up-information>` XML tag.

- **The resilient-hash statement is no longer available under aggregated-ether-options (QFX5200 switches)**—Starting in Junos OS Release 17.4R3, the **resilient-hash** statement is no longer available in the `[edit interfaces aex aggregated-ether-options]` hierarchy level. Resilient hashing is not supported on LAGs on QFX5200.

[See [aggregated-ether-options](#).]

Management

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

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

MPLS

- **Support for Flap and MBB counter for LSP (QFX Series)**—Starting in Junos OS Release 17.4R1, the **show mpls lsp extensive** command introduces the following two counters for LSP on the master routing engine (RE) only:
 - Flap counter-- Counts the number of times a LSP flaps down or up.
 - MBB counter— Counts the number of times a LSP incurs MBB.

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

- **Display of labels in received record route for unprotected LSPs by show mpls lsp extensive command (QFX Series)**—The **show mpls lsp extensive** command displays the labels in received record route (RRO) for protected LSPs. Starting in Junos OS Release 17.4R1, the command also displays the labels associated with the hops in RRO for unprotected LSPs as well. The label recording in RRO is enabled by default.
- **Support for default timeout duration for self-ping on an LSP instance (QFX Series)**—Starting in Junos OS 17.4R1, the default timeout duration for which the self-ping runs on an LSP instance is reduced from 65,535 (runs until success) to 1800 seconds. You can also manually configure the self-ping duration value between 1 to 65,535 (runs until success) seconds using the **self-ping-duration value** command at the `[edit protocols mpls label-switched-path label-switched-path]` hierarchy level. By default, self-ping is enabled. The LSP types such as CCC, P2MP, VLAN-based , and non-default instances do not support

self-ping . You can configure the **no-self-ping** command at the **[edit protocols mpls label-switched-path label-switched-path]** hierarchy level to override the behavior of self-ping running by default.

- **Support for label history for MPLS protocol (QFX Series)**—Starting in Junos OS Release 17.4R1, configure **max-entries number** option at the **[edit protocols mpls label-history]** hierarchy level to display label allocation, release history, and associated information such as RSVP session that helps debug label related error such as stale label route and deleted label route. You can configure the limit for the maximum number of MPLS history entries per label . By default, label history is off and there is no maximum limit for the number of entries for each label. The **show mpls label history label-value** command displays the label history for a given label value and the **show mpls label history label-range start-label end-label** command displays the history of labels between the given label range. The **clear mpls label history** command clears the label history details.
- **Support for adjusting the threshold of autobandwidth based on the absolute value for LSP (QFX Series)**—Current autobandwidth threshold adjustment is done based on the configured percentage which is hard to tune to work well for both small and large bandwidth reservations. For a given threshold percentage, when the bandwidth reservation is small there can be multiple LSP ressignaling events. This is because the LSP is responsive to even minor increases or decreases in the utilization when current reservation is small. For example, a small threshold adjustment of 5 percent allows large LSPs of around 1G to respond to changes in bandwidth of the order of 50M. However, that same threshold adjustment results in too many LSP 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.
- When the **no-propagate-ttl** statement is configured on a QFX5200 switch in an MPLS network, the TTL value is not is not copied and decremented on the transit devices during a swap operation. When the switch acts as an ingress device for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the switch acts as the penultimate provider switch, it pops the MPLS header without writing the MPLS TTL into the IP packet. PR1368417

Network Management and Monitoring

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

Before this change:

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

After this change:

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

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

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

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

- **SNMP syslog messages changed (QFX Series)**—In Junos OS Release 17.4R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD — AgentX master agent failed to respond to ping. Attempting to re-register
NEW — AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD — NET-SNMP version %s AgentX subagent connected
NEW — NET-SNMP version %s AgentX subagent Open-Sent!

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

- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (QFX Series)**—Starting in Junos OS Release 17.4R3, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

Routing Policy and Firewall Filters

- **Support for configuring the GTP-TEID field for GTP traffic (QFX5000 line of switches)**—Starting in Junos OS Release 17.3R3 and 17.4R2, the **gtp-tunnel-endpoint-identifier** statement is supported to configure the hash calculation of IPv4 or IPv6 packets that are included in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations. The **gtp-tunnel-endpoint-identifier** configuration statement is configured at the **[edit forwarding-options enhanced-hash-key family inet]** hierarchy level.

In most of the cases, configuring **gtp-tunnel-endpoint-identifier** statement is sufficient for enabling GTP hashing. After enabling, if GTP hashing does not work, it is recommended to capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In this cases, use **gtp-header-offset** statement to set a proper offset value. Once the header offset value is resolved, run **gtp-tunnel-endpoint-identifier** command for enabling GTP hashing successfully.

[See [gtp-tunnel-endpoint-identifier](#) and [gtp-header-offset](#).]

Security

- **Support to log the SSH key changes**—Starting with Junos OS 17.4R1, the configuration statement **log-key-changes** is introduced at the **[edit system services ssh]** hierarchy level. When the **log-key-changes** configuration statement is enabled and committed (with the **commit** command in configuration mode), Junos OS logs the changes to the set of authorized SSH keys for each user (including the keys that were added or removed). Junos OS logs the differences since the last time the **log-key-changes** configuration statement was enabled. If the **log-key-changes** configuration statement was never enabled, then Junos OS logs all the authorized SSH keys.
- **Syslog or log action on firewall drops packets (QFX5000 switches)**—Starting in 17.4R3, if you configure a syslog or log action on an ingress firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.

Software Licensing

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

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

Virtual Chassis

- **Adaptive load balancing (ALB) feature (Virtual Chassis Fabric)**—Starting in Junos OS Release 17.4R1, the adaptive load balancing (ALB) feature for Virtual Chassis Fabric (VCF) is being deprecated to avoid potential VCF instability. The **fabric-load-balance** configuration statement in the **[edit forwarding-options enhanced-hash-key]** hierarchy is no longer available to enable and configure ALB in a VCF. When upgrading a VCF to a Junos OS release where ALB is deprecated, if the configuration has ALB enabled, you should delete the **fabric-load-balance** configuration item before initiating the upgrade.

[See [Understanding Traffic Flow Through a Virtual Chassis Fabric](#) and [fabric-load-balance](#).]

- **New configuration option to disable automatic Virtual Chassis port conversion (QFX5100 Virtual Chassis)**—Starting in Junos OS Release 17.4R2, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in a QFX5100 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:
 - LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
 - The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
 - The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

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

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

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

SEE ALSO

[New and Changed Features | 335](#)[Known Behavior | 354](#)[Known Issues | 358](#)[Resolved Issues | 369](#)[Documentation Updates | 391](#)[Migration, Upgrade, and Downgrade Instructions | 391](#)

Known Behavior

IN THIS SECTION

- [Class of Service \(CoS\) | 355](#)
- [EVPN | 355](#)
- [General Routing | 355](#)
- [Interfaces and Chassis | 356](#)
- [Junos Fusion Satellite Software | 356](#)
- [Layer 2 Features | 356](#)
- [MPLS | 356](#)
- [Routing Protocols | 357](#)
- [Platform and Infrastructure | 357](#)
- [Virtual Chassis | 358](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 for the QFX Series.

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

Class of Service (CoS)

- With pechip version 1.1, if dot1p rewrites are configured on an interface, then packets that are not matching to a rewrite rule will not retain their previous value. Set the rewrite rule value to 0. This functionality is fixed in pechip version 2.0 [PR1294471](#)

EVPN

- A provider edge (PE) device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE device. The IGP instance running in the VRF on the PE might be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- 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](#)
- Even though an ARP route is learned locally, the **show arp** command output on the provider edge (PE) device on which the route was learned might display the route as **permanent remote**. In Junos OS releases earlier than Junos OS Release 17.4R1, *permanent remote* means that the ARP route was learned from a remote PE device such as an EVPN Type 2 route (MAC+IP route).

This issue might occur under the following conditions:

- A customer edge (CE) device is multihomed to QFX10000 switches in an EVPN-VXLAN topology with a two-layer IP fabric or collapsed IP fabric.
- The QFX switches function as Layer 3 only, or Layer 2 and Layer 3 PE devices.
- The QFX switches run Junos OS Release 17.4R1 or later.

To work around this issue, you can view locally learned ARP routes by entering the **show evpn database origin local** command on the PE devices. [PR1324824](#)

General Routing

- L3 multicast traffic does not converge to 100 percent and continuous drops are observed after bringing down/up the downstream interface or while an FPC comes online after FPC restart. This happens with multicast replication for 1000 VLAN/IRB's. [PR1161485](#)
- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED stays unlit. [PR1317750](#)
- On v44 stage3 the mib2d crashes at **mib2d_write_snmpidx** after loading the configuration from SNP to PDT configuration. [PR1300892](#)
- L2 and L3 are not supported together. You cannot have encapsulation (**inet**, **flexible-ethernet-services**, and **vlan-bridge**) on the same interface. [PR1358200](#)

- Having the network control protocols such as BGP, PIM, IGP, LDP, etc not to starve while node management activities taking place on the device, FTP SSH is being rate-limited on the WAN (IRB) interfaces. [PR1371509](#)
- On QFX10K/QFX5K switches, packet drops can occur for the traffic that has to use an EVPN type-5 overlay tunnel if the first FPC(FPC0) is down on the other end of the tunnel. In this case, the destination switch which has the FPC0 down receives the packet and drops it. [PR1423928](#)

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 Satellite Software

- PEM alarm behavior is same for Junos OS (standalone) also for SD (SNOS). The **PEM # Not Present** present alarm is triggered only if that PEM FRU is removed from the box at runtime. This alarm will be cleared, once the PEM is inserted back or board is rebooted. [PR1287856](#)

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging has been enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. [PR1075230](#)

MPLS

- Layer 2 circuits on aggregated Ethernet interfaces are not supported on QFX5100, QFX5110, and QFX5200 switches. [PR1333730](#)
- On QFX5100, QFX5110, QFX5200 switches with Layer 2 circuit configured on the PE switches, enabling VLAN bridge encapsulation on a CE interface drops packets if flexible Ethernet services and VLAN CCC encapsulation are configured on the same logical interface. You can configure only one encapsulation type, either **set interfaces xe-0/0/18 encapsulation flexible-ethernet-services** or **set interfaces xe-0/0/18 encapsulation vlan-ccc**. [PR1329451](#)

Routing Protocols

- During a graceful Routing Engine switchover (GRES) on QFX10000 switches, some IPv6 groups might experience momentary traffic loss. This issue occurs when IPv6 traffic is running with multiple paths to the source, and the **join-load-balance** statement for PIM is also configured. [PR1208583](#)
- For the QFX10002 and QFX10008 switches, you might observe an increase in the convergence time of OSPF routes when compared to Junos OS Release 17.3. An average increase of 1.5 seconds is seen for 100,000 OSPFv3 routes. [PR1297541](#)
- A QFX10000 switch running Junos OS Release 17.3Rx or 17.4Rx software might experience a small and continuous traffic loss under the following conditions: 1) The switch is configured as a Layer 2, Layer 3 or both VXLAN gateway in an EVPN-VXLAN topology with either a two-layer or collapsed IP fabric. 2) The switch has default ARP and MAC aging timer values. Under these conditions, the following types of traffic flows might be impacted: 1) Bidirectional Layer 3 traffic in a multihomed topology, and 2) Unidirectional Layer 3 traffic in a single-homed topology. Note that this issue does not impact bidirectional Layer 3 traffic in a single-homed topology. [PR1309444](#)

Platform and Infrastructure

- On a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On a QFX5110-32C switch, if a splitter cable is connected to a peer end device capable of 10G CV/MX card, ports will not come up due to varied pre-empt settings for the splitter and DAC cables. There is a hardware limitation where we have no way in EEPROM to differentiate between splitter and DAC cable to apply different settings. As a workaround, use manual channelisation on the QFX5110-32C side. [PR1280593](#)
- ERPS convergence takes time after a GRES switchover and hence traffic loss is observed for a brief period. [PR1290161](#)
- On QFX Series, the logical interface (IFD) and the physical interface (IFL) go down when traffic exceeds the ratelimit. Storm control is supported only on interfaces configured in family Ethernet-switching. Moreover, in this family, only one IFL is supported per IFD. Thus, bringing down the IFD is acceptable. Flexible VLAN tagging is not supported on the interfaces enabled for storm control. [PR1295523](#)
- Traffic drop occurs when sending Layer 3 traffic across an MPLS LSP. [PR1311977](#)
- Traffic drop occurs when sending traffic over "et" interfaces due to CRC errors. [PR1313977](#)
- On Junos OS Automation Enhancement images there is a way to use the Python interpreter in interactive mode. When Python interpreter is used in an interactive mode on a shell, the prompt does not seem to return immediately. This is an example of a session: -- % python Python 2.7.8 (default, Nov 10 2017, 01:45:13) [GCC 4.2.1 (for JUNOS)] on junos Type "help", "copyright", "credits" or "license" for more information. >>> >>> print "hello" >>> hello -----> waiting here, hit 'enter' here to return the python prompt >>> quit() >>> % -- The regular script is not impacted. [PR1324124](#)

Virtual Chassis

- Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on an QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur and its considered to be known behavior. Release note reference:
https://www.juniper.net/documentation/en_US/junos/information-products/topi-c-collections/release-notes/17.2/topic-118735.html [PR1347902](#)

SEE ALSO

[New and Changed Features | 335](#)

[Changes in Behavior and Syntax | 347](#)

[Known Issues | 358](#)

[Resolved Issues | 369](#)

[Documentation Updates | 391](#)

[Migration, Upgrade, and Downgrade Instructions | 391](#)

Known Issues

IN THIS SECTION

- [EVPN | 359](#)
- [Forwarding and Sampling | 360](#)
- [General Routing | 360](#)
- [Interfaces and Chassis | 365](#)
- [Layer 2 Ethernet Services | 365](#)
- [Layer 2 Features | 366](#)
- [MPLS | 366](#)
- [Network Management and Monitoring | 367](#)
- [Platform and Infrastructure | 367](#)
- [Routing Protocols | 367](#)
- [Virtual Chassis | 368](#)

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.4R3.

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

EVPN

- Mac-move-shutdown stops working if a “physical loop” is introduced continuously in quick succession of 10 minutes. The issue is not seen every time but can occur only if physical loop is introduced at least four times. If the loops span a long period, the issue is not seen. A test is performed to check the overall impact on basic features. There is no issue seen on basic learning or major impact on any protocol. This is a negative scenario, but it is unlikely to occur in a customer network where the multiple loops occur in a short time span. [PR1284315](#)
- 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](#)
- The chained-composite-next-hop (CNH) is a must for EVPN pure type 5 with VXLAN encapsulation. Without this Packet Forwarding Engine might not program the tunnel next hop. This should be explicitly set it on QFX5110. **set routing-options forwarding-table chained-composite-next-hop ingress evpn.** QFX10000 it is applied as part of default configuration. **user@host> show configuration routing-options forwarding-table | display inheritance defaults.** [PR1303246](#)
- In a EVPN collapsed L2/L3 multi-homed gateways topology, when traffic is sent from IP fabric towards EVPN, some traffic loss is seen. If the number of hosts behind EVPN gateways is increased, the traffic loss becomes higher. This issue is seen with QFX10000. [PR1311773](#)
- ARP gets deleted and relearned during the first ARP refresh with EVPN-VXLAN multihomed CE, so traffic drops and recovers for first ARP refresh. [PR1327062](#)
- In an EVPN environment, proxy ARP and ARP suppression is enabled on the PE device by default for reducing the flooding of ARP packets. However, in the case of ARP probe packets used in the process of Duplicate Address Detection (DAD), the client might treat the IP address that it is in use as duplicated address after receiving the proxied packets from PE device. [PR1427109](#)

Forwarding and Sampling

- Commit failure with error might be seen and the dfwd crashes when applying a firewall filter with action "then traffic-class" or "then dscp" to an interface. [PR1452435](#)

General Routing

- L3 multicast traffic does not converge to 100 percent and continuous drops are observed after bringing down/up the downstream interface or while an FPC comes online after FPC restart. This happens with multicast replication for 1000 VLAN/IRB's. [PR1161485](#)
- When per-packet load balancing is removed or deleted, next hop index might change. [PR1198092](#)
- Single-bit and multiple-bit ECC errors are not logged on QFX5110 switches. [PR1251917](#)
- On the QFX10000-12C-DWDM coherent line card, it is possible that sometimes the link flaps when MACsec is enabled on Ethernet interfaces. [PR1253703](#)
- The management process (daemon) might crash if the Openconfig package is installed immediately or within minutes of network agent package installation. This is a transient issue and will not impact any functionality. There is no action needed from the user side in response to the crash. As a workaround, install Openconfig before installing network agent. [PR1265815](#)
- On QFX5100-VC, the buffer is corrupted on port 0 (*/*/0) and error message **MACDRAINTIMEOUT** and **dcbcm_check_stuck_buffers** are observed, which could eventually lead to port 0 (*/*/0) flapping. [PR1284590](#)
- On QFX5100 switches, LACP link protection switchover/revert is not working when LACP link-protection is configured with backup-state "down". **set interfaces ae0 aggregated-ether-options link-protection backup-state down**. When configuring LACP, the state of the backup link should not be configured manually as down. This is not supported if LACP is configured. [PR1286471](#)
- When link protection with the backup port state "down" and LACP are configured, if backup state "down" is removed from the configuration, both ports should be up and the primary port should pass all egress traffic. In some instances, however, traffic might pass through the backup port instead of the primary port. [PR1297597](#)
- Traffic drop occurs on sending traffic over "et" interfaces due to CRC errors. [PR1313977](#)
- Family Ethernet-switching cannot be used when **flexible-vlan-tagging** is configured. It is unsupported. The behavior is non-deterministic with this configuration and there is a possibility of seeing dcpfe core file. [PR1316236](#)
- Port LEDs on QFX5100 do not work. If a device connects to a port on QFX5100, the port LED stays unlit. [PR1317750](#)
- On a QFX Series devices with a third-generation FPC, the error message is displayed when the FPC goes online or offline. [PR1322491](#)

- The management process (mgd) might panic after modifying aggregated Ethernet interface members under **ethernet-switching vlan** stanza. After mgd panic, your remote session is terminated as a result. [PR1325736](#)
- On QFX52xx standalone devices with VXLAN configured, user configured ingress ACL scale limit is 256 terms. [PR1331730](#)
- The mib2d core file might be generated in mib2d_write_snmpidx at snmpidx_sync.c on both ADs. [PR1354452](#)
- On QFX5110, the FEC for 100G optics is not being displayed when expected behavior is for FEC to be shown as NONE. On QFX10002-36Q, the FEC for 40G optics is being displayed as NONE when expected behavior is for FEC not to be displayed. On QFX10008, the FEC for 40G optics is being displayed as NONE when expected behavior is for FEC not to be displayed. [PR1360948](#)
- When MC-LAG is configured with force-up enabled on MCLAG nodes, the LACP admin key should not match the key of the access or CE device. [PR1362346](#)
- Currently, other than QFX5100-24Q and EX4600, PIC1 is not supported on any other platforms inline with QFX5100. The command below cannot be used on PIC1 **set chassis fpc 0 pic 1 port <x> channel-speed disable-auto-speed-detection**. This will result in a commit error **[edit chassis fpc 0 pic 1 port 2 channel-speed] channel-speed disable-auto-speed-detection**. PIC1 is not valid for auto-speed disable mode error **configuration check-out failed**. So, if you want to disable auto-channelisation on PIC1, you have to disable auto-speed-detection for whole FPC **set chassis fpc 0 auto-speed-detection disable**. [PR1362647](#)
- On QFX52100, filter with the routing instance applied to family inet logical interfaces causes traffic to be discarded on unrelated interfaces. [PR1364020](#)
- From Junos OS Release 17.3R1, on QFX10002 platform, in a rare condition, the IPFIX flow statistics (packet/byte counters) are incorrect in the exported record. Since the statistics are not collected properly, the flow might timeout and get deleted because of the inactive timeout, causing the number of exported records to be sent out unexpected. Traffic spikes generated by IPFIX might be seen. [PR1365864](#)
- On the QFX10000 line of switches, with EVPN-VXLAN, the following error is seen **expr_nh_fwd_get_egress_install_mask:nh type Indirect of nh_id: # is invalid**. [PR1367121](#)
- The L2 bridge domain might fail to create on Packet Forwarding Engine after changing VLAN configuration. For example, there are 3 VLANs V1001, V1002 and V1003. V1001 is deleted and V1002's VLAN ID and VNI is changed to that of V1001 and a new VLAN V1200 is added with the VLAN ID and VNI of VLAN V1002. After the above changes, V1200 is not created in Packet Forwarding Engine and the other 2 VLANs are functioning as expected. The reason for the new VLAN not created is due to out of order messages. This is a timing issue. [PR1371611](#)
- On QFX10000 platforms, when a same filter is applied on both input and output directions at same time, packets might be dropped after removing that filter. [PR1372957](#)
- Starting in Junos OS Release 17.1R1, the MAC address of interfaces on the QFX10002-36Q and QFX10002-72Q will change. On the QFX10002-36Q, after you upgrade to Junos OS Release 17.x, the

last octet of interface MAC addresses will increase by 3. On the QFX10002-72Q, after the upgrade to Junos OS Release 17.x, the last octet of interface MAC addresses will increase by 6. [PR1375349](#)

- In certain scenario's where flows are sampled through aggregate bundles when jflow sampling is enabled, the following harmless error logs can be seen [Tue Oct 30 18:17:40.648 LOG: Info] **expr_get_local_pfe_child_ifl: cannot find child ifl of agg ifl 74 for this fpc** [Tue Oct 30 18:17:40.648 LOG: Info] **flowtb_get_cpu_header_fields: Failed to find local child ifl for 74** [Tue Oct 30 18:17:40.648 LOG: Info] **fpc0 cannot find stream on [hostname]**. [PR1379227](#)
- Due to transient hardware condition single-bit error (SBE) event are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. [PR1384435](#)
- On QFX10008 and QFX10016 platforms, traffic loss might be observed because of switch modular failure on the Control Board (CB). This failure further causes all SIBs to be marked as faulty and causes FPCs to restart until Routing Engine switchover occurs. [PR1384870](#)
- With **MLD-snooping** enabled and when we have two receivers in the same VLAN interested in the same group address but from a different source, traffic will be received on only one receiver which sent the latest MLD report. This is because we do not install S, G routes in hardware when **MLD-snooping** is enabled. [PR1386440](#)
- When show command is taking a long time to display results, the STP might change states as BPDUs are no longer processed and cause lots of outages. [PR1390330](#)
- On QFX10000 switches, the major alarm **FPC Management Ethernet Link Down** might be displayed for management Ethernet (em0 or em1) interface that is administratively down. The alarm message has no service impact and can be ignored. [PR1391949](#)
- On QFX5100, traffic initiated from a server connected to an interface will be dropped at the interface on the switch if the interface is configured with **family ethernet-switching** with VXLAN and the configuration is changed to family inet. [PR1399733](#)
- When **mac-table-aging-time** is configured, bridge domain sequence get incremented unnecessarily. As a result, all MACs get flushed when the change message is received by I2-learning daemon with new sequence number. [PR1403358](#)
- On QFX10000 platforms, in EVPN-VXLAN scenarios, ping between Spine to Spine loopback over TYPE 5 tunnel might not work. [PR1405786](#)
- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. [PR1411015](#)
- On a QFX5120 system Transition from VXLAN/EVPN collapsed to non-collapsed L2/L3 gateway and vice versa needs switch reload due to stale source vtep IP. [PR1405956](#)
- During normal operation on QFX10002 platforms, if the chassis-control process restarts, the hardware might not get properly programmed. This causes packets to be dropped on the output interface. [PR1414434](#)

- On QFX10000 platforms with EVPN scenario, if an EVPN instance is created using the statement **set protocols evpn encapsulation mpls**, then the MAC learning might not happen on the CE-facing interface if the interface is configured with trunk-mode, because the solution of EVPN/MPLS is not currently supported on QFX10000 Series devices. [PR1416987](#)
- On QFX5110 and QFX5120 platforms, uRPF check in strict mode will not work properly. [PR1417546](#)
- ERSPAN traffic is not tagged when the output interface is a trunk port. [PR1418162](#)
- On QFX10002, QFX10008, and QFX10016 platforms, there is an aggregated Ethernet interface which has atleast 2 child links, which are located on different Packet Forwarding Engine chips, and this aggregated Ethernet interface is added to a VXLAN VLAN with IRB as an access interface, if aggregated Ethernet membership changes, for example, removing one child link from the aggregated Ethernet, traffic loss might be seen on the aggregated Ethernet interface. [PR1418396](#)
- **show interface** indicates "Media type: Fiber" on QFX5100-48T running QFX 5e Series image. This is a display issue. **Physical interface: xe-0/0/0, Enabled, Physical link is Down Interface index: 650, SNMP ifIndex: 515 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Auto, Speed: Auto, BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Auto-negotiation: Disabled, Remote fault: Offline, Media type: Fiber <<<<< Here!! Should be "Copper" Device flags : Present Running Down Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000 Link flags : None.** [PR1419732](#)
- Multiple EX Series platforms might be unable to commit baseline configuration after zeroize.

```
{master:0}[edit] root# commit check Mar 26 05:50:48 mustd: UI_FILE_OPERATION_FAILED: File /var/run/db/enable-process.data doesn't exist Mar 26 05:50:48 mgd[1938]: UI_FILE_OPERATION_FAILED: Failed to open /var/run/db/enable-process.data+ file error: Failed to open /var/run/db/enable-process.data+ file error: configuration check-out failed: daemon file propagation failed.
```

[PR1426341](#)
- CRC errors can be seen when other manufacturer device is connected to QFX10000 on a 100G link with QSFP-100GBASE-LR4-T2. Other manufacturer device report CRC errors and input errors on those 100G links. The QFX10000 interfaces do not show any errors. It might cause packet loss. [PR1427093](#)
- On QFX10002, QFX10008, and QFX10016 Series platforms with enhanced MC-LAG scenario, the dcpfe process might crash and restart if the ARP/NDP next hop is changed. [PR1427994](#)
- On QFX5110, QFX5120, and QFX5210, optical interface like 1G/10G SFP/SFP+ might take almost 3 minutes to reduce the tx power to "0" on the other end of the interface, after issuing **request system reboot at now** command. [PR1431900](#)
- On QFX10002, QFX10008, and QFX10016 Series platforms with enhanced convergence is configured in an MC-LAG scenario, if a line-card that has MC-LAG links is rebooted, the MC-LAG might not function correctly after the line-card comes back up. The impact is that it might not block the BUM traffic received on the interchassis link (ICL) and might cause the MAC movement and packet loss on the downstream devices. [PR1444100](#)

- In QFX10000 Series platforms, if a firewall filter with multiple match conditions is configured on interfaces which are Up and the firewall filter is modified (either a new action is added or the condition is added/removed etc.), the FPC might crash and restart. It might affect the service/traffic. [PR1432116](#)
- A firewall configuration change operation might not be done correctly within the Packet Forwarding Engine causing transit packets drops. [PR1433648](#)
- On QFX10008/QFX10016 platform, xSTP recognizes 1G SFP-T optic interface as LAN type link even if it is in full-duplex mode. This might cause the xSTP to converge slowly. As a workaround, configure the xSTP link type from LAN to Pt-Pt (point to point) using command **set protocols <vstp> <vlan-X> interface <interface-name> mode point-to-point**. [PR1439095](#)
- There is a sequence issue when Virtual Chassis member rebooted in aggregated interface. After rebooting VC member, Routing Engine kernel inject MAC entry to FPC that rebooted. Because of the sequence issue, Routing Engine added MAC entry, originally source MAC entry, to FPC as remote MAC entry. And MAC entry is never be aged out because it is remote entry. [PR1440574](#)
- QFX 10002, QFX10008, and QFX10016 when upgrading these systems from Junos OS Release 18.1 or previous to 18.2 or later releases , a minor error is reported. But upgrade/downgrade goes through fine. One side effect of this error is that, if upgrade or downgrade is happening as part of ZTP, then ZTP fails. ZTP keeps on upgrading (or downgrading) forever and ZTP never completes. [PR1446540](#)
- In QFX5100 Virtual Chassis scenario, Cyclic Redundancy Check (CRC) error might be seen on the Virtual Chassis Port (VCPs) when the VCPs are "BCM84328 PHY" ports. The CRC error indicates there is data corrupt, the issue might reduce the system performance. The issue can be avoided by using non-"BCM84328 PHY" ports as VCPs to build the VC. [PR1449406](#)
- The sFlow sample packets might stop on one aggregated ethernet member link if ingress sFlow is configured on the member link. This might cause inaccurate monitoring on the network traffic. [PR1449568](#)
- On QFX10K platforms, under the scale scenario more than 500 AE IFLs, if the classifier configuration frequent churns or link flaps, the CoS classification will not work on the impacted interfaces. [PR1450265](#)
- In EVPN-VXLAN with service provider style config, if VLAN name associated with access ports is changed then virtual bridge domain may not be created. This is because Bridge domain add notification for the new VLAN comes before Bridge domain delete for the old vlan. Due to this, virtual Bridge domain will not be created and MAC's will not be learnt. [PR1454095](#)
- On QFX51/EX4300/EX4600 VC/VCF scenario with Vxlan used, when configuring a firewall filter and commit, the firewall filter might not be able to be applied in a particular VC/VCF member for TCAM space running out. [PR1455177](#)
- On QFX Series platforms with Link Aggregation Group (LAG) interface, if periodic "SFP diagnostic" is configured with short interval (e.g. test sfp periodic diagnostic-interval 3), the LAG interfaces might have intermittent flaps and therefore bring service impact due to this issue. [PR1458363](#)
- On QFX5100 and EX4600 platforms, the fxpc (packet forwarding engine manager) process might crash when multiple BGP IPV6 sessions (for instance around 500) are flapped and then restored at the same time. [PR1459759](#)

- On QFX10K platforms, FPC might restart if there is some corruption in BCM switch (a small internal ethernet switch, instead of PFE engine) inside the FPC. It is a timing issue. The reason is that the PCIe speed configuration for BCM switch is not correct. And this issue is resolved in some FPC U-boot versions. [PR1464119](#)
- On QFX5100-48T, the 10G interface might not come up or negotiate at the speed of 1G with Broadcom 10G 57800-T daughter card. In the issue state, speed will be set to 1G which might make the interface down and result in traffic impact. [PR1465196](#)
- EPR iCRC errors in QFX10000 series platforms might cause protocols down. FPC will be in wedged state and will not pass traffic on that PFE if hitting this issue. EPR iCRC errors are normal and caused by transient hardware conditions. EPR iCRC errors are not expected to impact the protocols, and only one CRC failed packet will be dropped. But due to incorrect handling of this error, it affects protocols and causes FPC wedge. [PR1466810](#)
- When tunnel-services are configured on a PIC, the optics measurements that subscribed via gRPC might not be streamed. [PR1468435](#)
- On QFX5K platform, when MPLS node-link-protection is configured on all nodes (PE and P device), the IP routed packets might be looped on the MPLS PHP node (P device) if continuous interface flaps at ingress/egress of PE devices. [PR1469998](#)

Interfaces and Chassis

- Traffic drop is observed when trying to configure aggregated Ethernet interface description. [PR1305794](#)
- A QFX switch may send out ARP reply unicast packets as a result of an ARP request sent for the device's VRRP MAC address. [PR1454764](#)
- When dynamic DHCP sessions are existing in the device, if multiple commits in parallel are performed, the commit might hang up. [PR1470622](#)

Layer 2 Ethernet Services

- In MC-LAG with force-up scenario, the LACP PDU loop might be seen when both MC-LAG nodes and access device using same admin key. [PR1379022](#)
- On QFX5000 Series or EX4300/EX4600/ platforms with Spine-Leaf scenario, when some (two or more than two) underlay interfaces with ECMP are brought down on leaf devices, the multi-hop BFD overlay sessions between spines and leafs might flap. And if BFD flaps, the protocols depending on BFD (typically, IBGP protocols) might also flap, which leads to traffic impact. [PR1416941](#)

Layer 2 Features

- On QFX10016, after delete and readding of 1000 lag interfaces, traffic drops might be seen until ARP are refreshed even though all lag interfaces comes up. [PR1289546](#)
- On QFX Series platforms, if vlan-id-lists are configured under a single physical interface, QinQ might be malfunctioning for certain vlan-id-list(s). [PR1395312](#)
- On QFX5000 platforms, the fxpc might continuously crash when a firewall filter is applied on a logical unit of a dsc interface. It has a traffic impact. [PR1428350](#)
- On QFX5100/QFX5110/EX4600 platforms, if copper base SFP-T is used, it might not get up on physical layer and the MAC/ARP learning might not work if it gets up. The PR fixes both layer-1 and layer-2 issues in this scenario. [PR1437577](#)
- On EX/QFX platforms with STP disabled, the LLDP function might fail when a Juniper device connects to a non-Juniper one. In this scenario, the LLDP PDU with destination MAC 01:80:c2:00:00:00, which is one of the three reserved MAC addresses for LLDP in IEEE 802.1AB, will be ignored by Juniper LLDP process, and this causes the LLDP function failure. This issue has service impact. [PR1462171](#)

MPLS

- LDP to BGP stitching with eBGP indirect next hop having an implicit null label does not work. As a workaround, ensure the peer advertises a real label by adding another router between the egress and ingress PE devices. Use the IBGP that gets resolved over LDP or RSVP-TE LSPs. This will ensure that the BGP indirect next hop has a real label. [PR1254702](#)
- Statistics of transit traffic does not increment LSP statistics signaled by RSVP-TE. [PR1362936](#)

Network Management and Monitoring

- On all platforms, after the AGENTX session timeout between master(snmppd) and sub-agent, the sub-agent might crash and restart. [PR1396967](#)

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)

Routing Protocols

- OSPF stuck at Exchange state for lag interfaces in a QFX5100-EX4300 mixed Virtual Chassis setup whose child members belong to EX4300 after Rebooting Virtual Chassis. [PR1459329](#)
- For single-hop eBGP session, upon interface down event, do not do GR helper logic.
- On QFX10000 line platforms, during route next hop churn or earliest deadline first (EDF) job priority changes, memory corruption might occur, leading to processing issues and constant packet drop. [PR1243724](#)
- BGP as protocol strongly recommends configuration of local-address for each multihop iBGP/eBGP peer configuration. As a recommendation local-address should be route-able lo0 address. Using loopback address reduces dependency with interfaces. Note: Multihop is by default enabled for iBGP Peers. [PR1323557](#)
- When cleaning up routes as the peer goes down, we observe a 30 percent degradation in time taken in Junos OS Release 17.2X75-D91 as compared to Junos OS Release 17.2. [PR1329921](#)
- On a scaled setup, when the host table is full and the host entries are installed in LPM table, OSPF sessions might take more time to come up. [PR1358289](#)
- In some specific scenarios, the configuration of bpdu-block-on-edge might cause both QFX5110 to claim as VRRP masters. [PR1367439](#)
- L3 traffic travelling through QFX5000 do not get converged, after various triggers. [PR1379418](#)
- On QFX Series switches except for QFX10000, if host destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log/syslog action (for example, **filter <> term <> then log/syslog**), such packets should not be dropped and reach the Routing Engine. [PR1379718](#)
- On QFX10002\QFX10008\QFX10016 Series platforms with EVPN/VXLAN deployment scenario, the transit statistics of Integrated Routing and Bridging (IRB) interface might fail to be counted for the EVPN/VXLAN traffic, but it works for the regular IRB interface. [PR1383680](#)
- There is no functionality impact due to this error message. [PR1407175](#)

- On QFX5200 and QFX5110 platform or Junos OS on White Box (AS7816), interface flap might cause FPC watchdog timeout which then further triggers the FPC/dcpfe crash. As a result, traffic impact might be observed at that time. [PR1408428](#)
- On QFX5110 and QFX5200 platforms, the dcpfe might crash if any interface flaps. [PR1415297](#)
- By default BGP multipath is for load balance with BGP neighbors in same AS. For load balance with BGP neighbors in different AS, the statement **multiple-as** is further needed. However, if the statement **multiple-as** is only configured in some BGP groups but not in all BGP groups, the expected load balance will not work. [PR1430899](#)
- When Precision Time Protocol (PTP) transparent clock is enabled, PTP adds the residence time to the Correction Field of the PTP packets as they pass through the device. On QFX5K platforms with PTP transparent clock enabled, the IPv4 fragmented packets of UDP datagram might be broken by PTP in some rare scenario, and the corrupted packets will be discarded by system. This issue has traffic impact. [PR1437943](#)
- On QFX5K/EX4600 with service provider (SP) style VLAN configuration (in this method, each VLAN-ID is locally significant to a physical interface), if interface-mac-limit/mac-table-size is configured (that is, software MAC learning is enabled) and the scale of MAC addresses on the box is more than 2000, traffic might be dropped after Q-in-Q enabled interface is flapped or a change is made to the vlan-id-list. [PR1441402](#)
- OSPF stuck at Exchange state for lag interfaces in a QFX5100-EX4300 mixed VC setup whose child members belong to EX4300 after Rebooting VC. [PR1459329](#)
- Multicast statistics related errors like **brcm_ipmc_route_counter_delete:3900Multicast stat destroy failed (-10:Operation still running)** will be observed during ISSU and these messages are harmless and does not affect multicast functionality. [PR1460791](#)

Virtual Chassis

- ACX5000 reports false parity error messages like **soc_mem_array_sbusdma_read**. The ACX5000 SDK can raise false alarms for parity error messages like **soc_mem_array_sbusdma_read**. This is a false positive error message. [PR1276970](#)

SEE ALSO

[New and Changed Features | 335](#)

[Changes in Behavior and Syntax | 347](#)

[Known Behavior | 354](#)

[Resolved Issues | 369](#)

[Documentation Updates | 391](#)

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.4R3 | [369](#)
- Resolved Issues: 17.4R2 | [380](#)
- Resolved Issues: 17.4R1 | [387](#)

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: 17.4R3

Authentication and Access Control

- Without dot1x configuration, the syslog `dot1xd[2192]: task_connect: task PNACAUTH./var/run/authd_control addr /var/run/authd_control: Connection refused` is generated repeatedly. [PR1406965](#)

Class of Service (CoS)

- CoS is incorrectly applied on Packet Forwarding Engine, leading to egress traffic drop. [PR1329141](#)

EVPN

- QFX10000 import default IPv6 route to VRF causes infinite entries to get created in 'evpn ip-prefix-database' and become unstable. [PR1369166](#)
- Some MACs are in EVPN database, but not in ethernet-switching table after clearing entry in ethernet table. [PR1377496](#)
- ARP refresh functionality might fail in an EVPN scenario. [PR1399873](#)
- A few minutes of traffic loss might be observed during recovery from link failure [PR1396597](#)
- In the non-collapsed (centralized) topology, when one of the 2 spines deactivates the underlay protocol (ospf), the leaf still points the virtual-gw-mac's next hop to the down spine. [PR1403524](#)

- ARP entry is still pointing to failed VTEP after PE-CE link fails for multihomed remote ESI. [PR1420294](#)
- The device may proxy the ARP Probe packets in an EVPN environment [PR1427109](#)
- ARP request/NS might be sent back to the local segment by DF router [PR1459830](#)

Forwarding and Sampling

- The kernel crash might be observed when there is a firewall filter modification. [PR1365265](#)
- Firewall filter terms named "internal-1" and "internal-2" are ignored. [PR1394922](#)
- The l2ald process might observe memory leak on Junos OS platforms. [PR1455034](#)

General Routing

- The 1G copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms. [PR1286709](#)
- Syslogs contain messages with %PFE-3: fpc0 ifd null, port 28 dc-pfe: %USER-3: ifd null, port 28 : %PFE-3: fpc0 ifd null, port 29 dc-pfe: %USER-3: ifd null, port 29. [PR1295711](#)
- Oinker and TCP connection drop is seen during large file SCP or FTP to the system (high intr{ virtio_p} seen). [PR1295774](#)
- Port 0 does not come up in QFX5100-48T member in a mixed VCF. [PR1323323](#)
- MACsec causes dot1xd JTASK_SCHED_SLIP or FPC disconnect. [PR1322302](#)
- AI-script does not get auto upgrade unless it is manually done after a Junos OS upgrade. [PR1337028](#)
- QFX5000 platforms might display fpc0 error **requesting CMTFPC SET INTEGER, illegal setting 37 observed after upgrade**. [PR1340897](#)
- On QFX10000, FPC process crashes after J-Flow processes a malformed packet. [PR1348417](#)
- QFX5100 40G port has an interoperability issue with some other vendors. [PR1349664](#)
- When VOIP VLAN is set as NATIVE VLAN on the port, the interface still shows up as a tagged interface and drops all untagged traffic. [PR1349712](#)
- Bogus DDOS counter values and syslog messages might be seen after clearing DDOS statistics for a specific protocol on QFX10000 Series switches. [PR1351212](#)
- Unable to create QFX5200 VC w/100G DACs. [PR1360721](#)
- VME interface might be unreachable after link flap of em0 on master FPC. [PR1362437](#)
- The following log messages are seen: **kernel: tcp_timer_keep: Dropping socket connection**. [PR1363186](#)
- Extended traffic loss when performing a unified ISSU or GRES with an aggregated Ethernet interface configured with LACP. [PR1365316](#)
- SFP-T might not work on QFX5100 and QFX5110 devices. [PR1366218](#)
- In certain routing topologies with sFlow configured, sampled packets might be duplicated and sFlow records are not sent to the collector. [PR1370464](#)

- Packet Forwarding Engine is in a bad state after performing optics insertion or removal on a port. [PR1372041](#)
- The IPv6 routed packet might be transmitted through interface whose VRRP state is in non-master. [PR1372163](#)
- The backup member switch might fail to become the master switch after switchover on QFX5100, QFX5200, and EX4600 Virtual Chassis platform. [PR1372521](#)
- MAC refresh packet might not be sent out from the new primary link after RTG failover. [PR1372999](#)
- TPI-50840 BUM traffic received on QFX5110 is not flooded to all remote VTEPs. [PR1373093](#)
- The QSFP-100G-LR4-T2 has been incorrectly identified as QSFP-100G-LR4. Hence, the operational parameters might not be programmed correctly. [PR1373758](#)
- LLDP might stop working fully between QFX10000 and non-Juniper device. [PR1374321](#)
- On QFX5110, ethernet-switching flood group shows incorrect information. [PR1374436](#)
- The rpd process might crash when route flap and LSP flap occur with CBF enabled. [PR1374558](#)
- RIPV2 update packets might not send with IGMP-snooping enabled. [PR1375332](#)
- Packet Forwarding Engine wedge might be observed if there are interfaces going to down state. [PR1376366](#)
- The auto-negotiation interface might go down if the opposite device supports only 10/100M auto-negotiation. [PR1377298](#)
- Debug log message, **expr_nh_flanel_check_overwrite: Caller nh_id params**, classified as error log when it should be LOG_INFO. [PR1377447](#)
- Deleting an IRB interface might affect other IRB interface if the same custom MAC address is configured. [PR1379002](#)
- The overlay-ecmp might not work as expected on QFX5110 in an EVPN-VXLAN environment. [PR1380084](#)
- There is an inconsistency in applying scheduler map with excess-rate on the physical interface and aggregated Ethernet interface. [PR1380294](#)
- L3 VPN traffic might be dropped because of one core-facing interface down. [PR1380783](#)
- Packet Forwarding Engine on QFX5000 might show DISCARD next-hop for overlay-bgp-lo0-ip in the VXLAN scenario. [PR1380795](#)
- IRB interface does not turn down when master of Virtual Chassis is rebooted or halted. [PR1381272](#)
- Traffic is silently dropped when FPC goes offline in an MC-LAG scenario. [PR1381446](#)
- The 40G-SR4 transceiver might not be recognized after upgrading to QFX100e OS. [PR1381545](#)
- LACP might get stuck in detached state on QFX5000 platforms in VXLAN scenario. [PR1382209](#)
- New CLI statement to enable copying of Open vSwitch Database (OVSDb) to RAM on Virtual Chassis backup Routing Engine instead of SSD. [PR1382522](#)

- The Packet Forwarding Engine might crash if the GRE destination IP is resolved over another GRE tunnel. [PR1382727](#)
- Static default route with next-table inet.0 does not work. [PR1383419](#)
- The log messages **RPD_KRT_Q_RETRIES: list nexthop ADD: No such file or directory** might be shown continuously after the rpd process restarts. [PR1383426](#)
- The DMA failure errors might be seen when the cache flush or the cache is full. [PR1383608](#)
- DHCP packets might be dropped on a Junos Fusion Data Center scenario. (QFX10000 series) [PR1383623](#)
- The Virtual Chassis might not come up after upgrading to QFX5E platforms (TVP-based platforms for QFX5100 or QFX5200 switches). [PR1383876](#)
- BFD sessions might flap consistently. [PR1384601](#)
- VM core file might be seen on the Junos OS Release 18.1R3. [PR1384750](#)
- All 1G SFP copper and 1G fiber optic links remain UP on QFX10008 after all SIBs or FPCs are offline. [PR1385062](#)
- ARP/ethernet-table is pointing to down the aggregated Ethernet interface if MTU is changed. [PR1385199](#)
- The IPv6 packet might not be routed when IPv6 packet is encapsulated over IPv4 GRE tunnel on QFX10000. [PR1385723](#)
- The spine EVPN routes might get stuck in a hidden state with next hop as unusable after FPC is offline in the spine. [PR1386147](#)
- The QFX10000-12C DWDM line card might crash when booting up. [PR1386400](#)
- The rpd process might get stuck with KRT queue in VRF scenario. [PR1386475](#)
- DDOS statistics and logging is not working for internal queues such as Q42 and Q4. [PR1387508](#)
- Traffic drop might be seen on QFX10000 platform with EVPN VXLAN configured. [PR1387593](#)
- QFX5100, QFX5110, QFX5200, and QFX5210 Virtual chassis could not be formed normally. [PR1387730](#)
- Certain log messages might be observed on QFX Series platforms. [PR1388479](#)
- MAC learning might stop working on some LAG interfaces. [PR1389411](#)
- Link problems might occur with 100G-AOC on QFX Series platforms. [PR1389478](#)
- FPC might crash on QFX5100 platforms in a large-scale scenario. [PR1389872](#)
- The input rate statistics might not increase if there are non-standard packets flow. [PR1389908](#)
- The vmcore file might be seen when routing changes are made on the peer spine in an EVPN VXLAN scenario. [PR1390573](#)
- An incorrect error message might be seen when J-Flow sensors are configured with reporting rate less than 30 seconds. [PR1390740](#)
- sdk-vmmd might consistently write to the memory. [PR1393044](#)

- 10-Gigabit Ethernet copper link flapping might happen during TISSU operation of QFX5100-48T switches. [PR1393628](#)
- IPv6 next-hop programming issue might be observed on QFX10000/PTX1000/PTX10000 devices. [PR1393937](#)
- The dhcp-security binding table might not be updated because of the renew request with '0.0.0.0' value in 'ciaddr'. [PR1394341](#)
- L2ALD core file is seen when I2-learning traceoptions are enabled. [PR1394380](#)
- On QFX5110 VC, Fan tray output is not displayed for backup Routing Engine. [PR1394655](#)
- DRAM and buffer utilization fields are not correct for QFX10000 platforms. [PR1394978](#)
- Unable to install licenses automatically on QFX Series platforms. [PR1395534](#)
- The subscriber bindings might not be successful on QFX and EX Series platforms. [PR1396470](#)
- On QFX5110, Fan LED turns Amber randomly. [PR1398349](#)
- The DHCPv6 relay packets are dropped when both the UDP source and destination ports are 547. [PR1399067](#)
- CPU hog might be observed on QFX10000 Series platform. [PR1399369](#)
- SFP-LX10 does not work on QFX5110. [PR1399878](#)
- Only one Packet Forwarding Engine can be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- The dcpfe crashes after adding or deleting a large number of LSPs several times. [PR1400868](#)
- The authd might stop when issuing **show network-access requests pending** command during the authd restarting. [PR1401249](#)
- The MTU might change to a Jumbo default size on Packet Forwarding Engine side after deleting and re-adding the interface. [PR1402588](#)
- File permissions are changed for **/var/db/scripts** files after reboot. [PR1402852](#)
- The STP does not work when aggregated interfaces number is "ae1000" or above in QFX5000 and "ae480" or above in other QFX or EX Series platforms. [PR1403338](#)
- Storm-control profile does not take effect after reboot of device/member of a Virtual Chassis if traffic flow increases beyond the threshold post reboot. [PR1403424](#)
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface. [PR1403528](#)
- The VRRP VIP might not work when it is configured on the LAG interface. [PR1404822](#)
- Executing command **request system configuration rescue save** might fail with error messages. [PR1405189](#)
- DHCP is not working for some clients in dual AD fusion setup on EP ports. [PR1405495](#)
- On QFX10002, SNMP trap for PSU removal or insertion is not generated. [PR1405877](#)

- The Packet Forwarding Engine might get disabled unexpectedly due to a auto correctable non-fatal hardware error on PTX or QFX10002/QFX10008/QFX10016. [PR1408012](#)
- The DHCP discover packets might be dropped over VXLAN tunnel if DHCP relay is enabled for other VXLAN/VLANs. [PR1408161](#)
- Fan failure alarms might be seen on QFX5100-96S after upgrading to Junos OS Release 17.3R1. [PR1408380](#)
- Restarting line card on QFX10008 and QFX10016 with MC-LAG enhanced-convergence might cause intra-vlan traffic to get silently dropped and discarded. [PR1409631](#)
- The FPC might crash and might not come up if interface-num or next-hop is set to maximum value under vxlan-routing on QFX Series platforms. [PR1409949](#)
- LLDP memory leak when ieee_dcbx packet is received in auto-negotiation mode followed by another dcbx packet with none of ieee_dcbx TLVs present. [PR1410239](#)
- Storm control is not shutting down mc-ae interface. [PR1411338](#)
- PEM alarm for backup FPC remains on master FPC though backup FPC that is detached from Virtual Chassis. [PR1412429](#)
- Junos PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- Virtual Chassis ports using DAC might not establish link on QFX5200. [PR1414492](#)
- Rebooting QFX5200-48Y using **request system reboot** does not take physical links offline immediately. [PR1419465](#)
- An interface might go to downstate on QFX10000/PTX10000 platform. [PR1421075](#)
- BFD might get stuck in slow mode on QFX10002/QFX10008/QFX100016 platform. [PR1422789](#)
- QFX5100-48T 10G interface might be auto-negotiated at 1G speed instead of 10G. [PR1422958](#)
- The interface cannot get up when the remote-connected interface only supports 100M in QFX5100 Virtual Chassis setup. [PR1423171](#)
- Traffic is dropped after FPC reboot with aggregated Ethernet member links deactivated by the remote device. [PR1423707](#)
- The J-Flow export might fail when channelization is configured on FPC QFX10000-30C. [PR1423761](#)
- All interfaces creation fails after NSSU. [PR1425716](#)
- Heap memory leak might be seen on QFX10000 platforms [PR1427090](#)
- Rebooting or halting Virtual Chassis member might cause 30 seconds down on RTG link. [PR1427500](#)
- Licenses used flag for ovsdb on **show system license** might not be flagged even though ovsdb is configured and working. [PR1428207](#)
- On EVPN-VXLAN L2ALD core files are generated when number of VXLAN hardware IFBDS exceeds the maximum limit of 16382. [PR1428936](#)

- On QFX10008 after Routing Engine switchover, LED status is not set for missing fan tray. [PR1429309](#)
- When **forward-only** is set within dhcp-reply, dhcp declines are not forwarded to the server. [PR1429456](#)
- DHCP-relay might not work in an EVPN-VXLAN scenario. [PR1429536](#)
- Interface on QFX Series does not come up after the transceiver is replaced with one having a different speed. [PR1430115](#)
- On QFX10000, **hold-down timer** configured interface are processing incoming packets leading to packet forwarding through the ASIC. [PR1430722](#)
- On QFX switch, **Validation of meta data files failed** message is observed. [PR1431111](#)
- The dcpfe might crash on all line cards on QFX10000 in scaled setup. [PR1431735](#)
- All ingress traffic might be dropped on 100m fixed speed port with no-auto-negotiation enabled [PR1431885](#)
- Outer VLAN tag might not be pushed in the egress VXLAN traffic towards the host for QinQ scenario. [PR1432703](#)
- On QFX10000 platforms, SIB and FPC minor Link Error alarms might happen on QFX10000 switches due to a single CRC. [PR1435705](#)
- LASER TX remained enabled while interface is disabled using the Routing Engine CLI configuration. [PR1436286](#)
- DHCP discover packets sent to IP addresses in the same subnet as IRB interface cause the QFX5110 to send bogus traffic out of **dhcp-snooping** enabled interfaces. [PR1436436](#)
- Unknown SNMP trap (1.3.6.1.4.1.2636.3.69.1.0.0.1) sent on QFX5110 restart. [PR1436968](#)
- On QFX5110, QFX5200, QFX5210 line of switches, there is no jnxFruOK SNMP trap message when only the power cable is disconnected and connected back. [PR1437709](#)
- The DHCP Snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it. [PR1438351](#)
- PSU status keeps "Check" when power supply is disconnected. [PR1441920](#)
- Flow control does not work as expected on 100G interface of QFX5110. [PR1442522](#)
- Chassis alarm message **Management Ethernet Link Down** will be displayed on QFX 10000 Series switches. [PR1391949](#)
- The TCP connection for external or internal might be dropped due to a kernel issue [PR1401507](#)
- QFX5k : Transit traffic loss when one of LAG child interfaces deleted or deactivated [PR1408178](#)
- The PTX1000/PTX10002/QFX10002 may stop forwarding packets after the "chassis-control" process restarts [PR1414434](#)
- Traffic loss might be seen on the ae interface on QFX10000 platforms [PR1418396](#)
- Traffic loss might be seen after NSSU operation [PR1418889](#)

- CRC errors can be seen when other manufacturer device is connected to QFX10K with QSFP-100GBASE-LR4-T2 optics [PR1427093](#)
- QFX5100-VCF - 'rollback' for uncommitted config takes 1 hour [PR1427632](#)
- On QFX10k/PTX10k platforms certain interfaces might go to down state [PR1427883](#)
- The jumbo frame size packets are dropped when max MTU is configured [PR1428094](#)
- The l2cpd process might crash and generate a core dump when interfaces are flapping [PR1431355](#)
- The et interfaces might not come up on QFX10000-60S-6Q [PR1431743](#)
- Traffic loss might be seen on QFX10k/PTX10k platforms using line card LC1105 [PR1433300](#)
- VC Mezz temp and QIC sensor get failure on QFX [PR1433525](#)
- The mc-ae interface may get stuck in waiting state in dual mc-ae scenario [PR1435874](#)
- The FPC might crash if both the AE bundle flapping on local device and the configuration change on peer device occur at the same time [PR1437295](#)
- BGP neighbourship might not come up if the MACsec feature is configured [PR1438143](#)
- Interfaces configured with flexible-vlan-tagging might loss connectivity [PR1439073](#)
- The EX4600/QFX5100 VC might not come up after replacing VC port fiber connection with DAC cable [PR1440062](#)
- When a line-card is rebooted, the MC-LAG may not get programmed after the line-card comes back online [PR1444100](#)
- On QFX10008 traffic impact might be seen when the JSRV interface is used [PR1445939](#)
- On QFX10K platforms removing EVPN-VXLAN L3 Gateway on the IRB interface from spine switches might cause black holing of traffic [PR1446291](#)
- Qfx10008: FPC0 cored after running the pfe command "show cos sched-usage" [PR1449645](#)
- "show cos scheds-per-pfe" and "show cos pfe-scheduler-ifds" pfe commands will restart forwarding planes on QFX10008 switches [PR1452013](#)
- Vgd core might happen when tunnel getting deleted twice [PR1452149](#)
- Config change in VLAN all option might affect the per-VLAN configuration [PR1453505](#)
- Slow packet drops might be seen on QFX5000 platforms [PR1466770](#)
- Ingress drops to be included at CLI from interface statistics and added to InDiscards [PR1468033](#)

Infrastructure

- Packets with the DEI/CFI bit set to 1 in the L2 header might not be forwarded. [PR1326855](#)
- Traffic gets silently dropped and discarded with indirect next hop and load balancing. [PR1376057](#)

Interfaces and Chassis

- Constant dcpfe process crash might be seen when using an unsupported GRE interface configuration. [PR1369757](#)
- On QFX5200 MCLAG, `parse_remove_ifl_from_routing_inst()` ERROR : No route inst on et-0/0/16.16386 errors are seen after l2cpd daemon is restarted. [PR1373927](#)
- Changing the value of `mac-table-size` to default might lead all FPC to reboot. [PR1386768](#)
- The logical interfaces in EVPN routing instances might flap after committing configurations. [PR1425339](#)
- The traffic might be forwarded to wrong interfaces in MC-LAG scenario. [PR1465077](#)

Junos Fusion Satellite Software

- Extended Port (EP) LAG might go down on the Satellite Devices (SDs) if the related Cascade Port (CP) links to an Aggregation Device (AD) goes down. [PR1397992](#)
- ARP Request packet might be dropped at egress SD when ingress and egress ECID is same. [PR1458930](#)

Layer 2 Ethernet Services

- Junos core file `jdhcpd.core.0` found in `dhcpx6_packet_handle` is seen. [PR1329390](#)
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. [PR1373807](#)
- The malfunction of core isolation feature in EVPN-VXLAN scenarios might cause traffic drop. [PR1417729](#)
- After GRES switchover, LACP will be down on peer device and never been recovered automatically [PR1395943](#)

Layer 2 Features

- Storm control configuration might be disabled for the interface. [PR1354889](#)
- LACP packets might be dropped with `native-vlan-id` configured after reboot. [PR1361054](#)
- When `native-vlan-id` is configured for aggregated Ethernet LACP session to multihomed server goes down. [PR1369424](#)
- On QFX5000 Series switches, a stack buffer overflow vulnerability in Packet Forwarding Engine manager (FXPC) process is seen. [PR1371400](#)
- DHCP discover packets might be dropped if there is VXLAN configured. [PR1377521](#)
- Packets might be dropped on AD in Junos Fusion Data Center environment. [PR1377841](#)
- The dcpfe process might crash while changing MTU of physical ports for GRE. [PR1384517](#)
- The LACP might be detached state when deleting `native-vlan-id` on aggregated Ethernet interface with `flexible-vlan-tagging` configured. [PR1385409](#)
- The dcpfe core might be observed when doing "restart routing" or BGP neighbors flaps when EVPN-TYPE 5 routes are present. [PR1387360](#)

- RTG MAC refresh packets will be sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. [PR1389695](#)
- On QFX Series platforms, error message **Failed with error (-7) while deleting the trunk 1 on the device 0** is observed. [PR1393276](#)
- On QFX5000 platforms symmetric hashing can be done, though it can not be enabled and stored in the Junos OS configuration. [PR1397229](#)
- DCPFE is restarted at the `_bcm_field_td_counter_last_hw_val_update` routine after upgrading spine with latest image. [PR1398251](#)
- QFX5110 Virtual Chassis generates DDOS messages of different protocols on inserting a 1G/10G SFP or forming VCP connection. [PR1410649](#)
- The traffic with triple or more 802.1Q tags might fail to forward. [PR1415769](#)
- Stale entries might be observed in a layer 3 VXLAN gateway scenario. [PR1423368](#)
- Transit DHCPv6 packets might be dropped on QFX5100 and QFX5200 platforms. [PR1436415](#)
- dcpfe core file is generated in QFX5200. [PR1362557](#)
- Unequal LAG hashing might happen on QFX Series devices. [PR1455161](#)

MPLS

- Traffic loss might be observed after changing configuration under "protocols mpls" in ldp-tunneling scenario [PR1428081](#)
- The LSP might remain UP even if no path is acceptable due to CSPF failure. [PR1365653](#)
- The rpd might crash when executing Routing Engine switchover under BGP environment and route churn occurs. [PR1373313](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition. [PR1374102](#)
- LSP "statistics" and "auto-bandwidth" functionality might not take effect with single hop LSPs. [PR1390445](#)
- The l2circuit traffic might silently get dropped at **EVPN SPINE/MPLS LSP TRANSIT** device if VXLAN access interface flaps on remote PE node(QFX5110). [PR1435504](#)

Network Management and Monitoring

- The AGENTX session timeout between master (snmpd) and subagent triggers some daemon crash [PR1396967](#)

Platform and Infrastructure

- The **Platform failed to bind rewrite** message might be seen when chassis control restart is done with the CoS rewrite rule configured on an aggregated Ethernet interface. [PR1315437](#)

Routing Protocols

- vrf-fallback on QFX5000 is not supported in ALPM mode. [PR1345501](#)
- Some storm control error logs might be seen on QFX Series platforms. [PR1355607](#)
- The pfe process might crash and all interfaces might flap. [PR1369011](#)
- The rpd process might crash after committing the configuration related to **mapping-server-entry**. [PR1379558](#)
- BUM packets might get looped if EVPN multihoming interface flap. [PR1387063](#)
- It might fail to update next hop in hardware for existing ECMP route when **ecmp-resilient-hash** is configured. [PR1387713](#)
- On EVPN-VXLAN NON-COLLAPSED autonegotiation errors and flush operation failed errors are seen after power cycle of the device. [PR1394866](#)
- The rpd soft core and inappropriate route selection might be seen when L2VPN is used. [PR1398685](#)
- ICMPv6 RA packets generated by Routing Engine might be dropped on the backup member of Virtual Chassis if **igmp-snooping** is configured. [PR1413543](#)
- The same traffic flow might be forwarded to different ECMP next hops on QFX5000 platforms. [PR1422324](#)
- The rpd process might generate a core file because of the improper handling of Graceful Restart stale routes. [PR1427987](#)
- On QFX5110 devices, we might not be able to ping the IRB address when being received as a type 5 route. [PR1433918](#)
- DDOS violation for protocols with shared host-path queue even when PPS rate is below the configured bandwidth value. [PR1440847](#)
- The rpd process might crash in inter-AS option B L3VPN scenario if CNHs is used. [PR1442291](#)
- When VRF fallback is enabled, running "show pfe route ip hw lpm" may crash the switch. [PR1367584](#)
- The IRB transit traffic might not be counted for EVPN/VXLAN traffic. [PR1383680](#)
- Junos OS: vSRX, SRX1500, SRX4K, ACX5K, EX4600, QFX5100, QFX5110, QFX5200, QFX10K and NFX Series: console management port device authentication credentials are logged in clear text. [PR1408195](#)

- Loopback address exported into other VRF instance might not work on EX/QFX/ACX platforms. [PR1449410](#)
- MPLS LDP may still use stale MAC of the neighbor even the LDP neighbor's MAC changes. [PR1451217](#)
- The egress interface in PFE for some end-hosts may not be correct on the layer 3 gateway switch after it is rebooted. [PR1460688](#)

Spanning Tree Protocols

- The l2cpd might crash if the VSTP traceoptions and VSTP VLAN all commands are configured. [PR1407469](#)

User Interface and Configuration

- Switch may unable to commit baseline configuration after zeroize. [PR1426341](#)

Resolved Issues: 17.4R2

Class of Service (CoS)

- You cannot filter packets with DstIP as 224/4 and DST MAC = QFX_intf_mac on a loopback interface using a single match condition for source address 224.0.0.0/4. [PR1354377](#)

EVPN

- Next hop installation error messages are seen on QFX10000 line switches. [PR1258930](#)
- EVPN-VXLAN QFX10000: `jprds_dlu_alpha_add : 222 JPRDS_DLU_ALPHA KHT` addition failed. [PR1258933](#)
- VXLAN-EVPN: IPv6 packet loss after a normal traffic run rate. [PR1267830](#)
- Subinterfaces from the same physical port do not work if configured under the same VXLAN VLAN. [PR1278761](#)
- For a VLAN with an IRB interface as the routing interface, set the vlan-id parameter to "none" to ensure proper traffic routing. [PR1287557](#)
- QFX10000 VXLAN with MPLS underlay traffic loss is seen at the RSVP egress. [PR1289666](#)
- VXLAN traffic loss is observed after deleting and adding VLANs. [PR1318045](#)
- A core link flap might result in an inconsistent global MAC count. [PR1328956](#)
- The partial multicast traffic might be dropped in an EVPN-VXLAN multi homing scenario with non-default **virtual-switch/evpn routing-instance** configured. [PR1334408](#)
- The MAC movement between remote VTEP and local VTEP might cause traffic to be transmitted incorrectly in an EVPN-VXLAN scenario. [PR1335431](#)
- Configuring **encapsulate-inner-vlan** on the partial VXLANs might cause traffic impact. [PR1337953](#)
- In an EVPN-VXLAN environment, BFD flaps cause VTEP flaps and cause the Packet Forwarding Engine to crash. [PR1339084](#)

- Rpd has unreproducible cored with scaling EVPN-VXLAN configuration on QFX10K platform. [PR1339979](#)
- The rpd core might be seen if deleting the default switch in an EVPN-VXLAN environment. [PR1342351](#)
- In an EVPN-VXLAN scenario, the traffic might get dropped as the core-facing interfaces goes down. [PR1343515](#)
- Traffic might be lost on a Layer 2 and Layer 3 spine node in a multihome EVPN scenario. [PR1355165](#)
- The QFX10000 might drop transited traffic coming from MPLS network to VXLAN/EVPN. [PR1360159](#)
- Increased risk of a routing crash with temporary impact on traffic on QFX10000 or QFX5100 nodes with certain configuration changes or clearing L2 or L3 learning information in a high-scale EVPN-VXLAN configuration environment. [PR1365257](#)
- Proxy ARP may not work as expected in an EVPN environment. [PR1368911](#)
- QFX10k / Import default ipv6 route to VRF causes infinite entries to get created in 'evpn ip-prefix-database' and become unstable. [PR1369166](#)

High Availability (HA) and Resiliency

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

Infrastructure

- QFX5100: Enabling mac-move-limit stops ping on **flexible-vlan-tagging** enabled interface. [PR1357742](#)

Interfaces and Chassis

- Multicast data packets are looping in MC-LAG. [PR1281646](#)
- Upgrading might encounter a commit failure if **redundancy-group-id-list** is not configured under ICCP. [PR1311009](#)
- CVLANs range is 16, which might not pass traffic in a Q-in-Q scenario. [PR1345994](#)
- MC-LAG peer doesn't send ARP request to the host. [PR1360216](#)

Layer 2 Ethernet Services

- A jdhcpd core file is generated after making DHCP configuration changes. [PR1324800](#)

Layer 2 Features

- Device transmits packets that exceed the interface MTU. [PR1306724](#)
- NLB heartbeat packets might be dropped on a QFX10000. [PR1322183](#)
- ARP entry might be learned on STP blocking ports. [PR1324245](#)
- The DHCP discover packets might be looped in an MC-LAG and a DHCP-relay scenario. [PR1325425](#)
- QFX5100: With multiple logical units configured on an interface, **input-vlan-map POP** is not removing outer VLAN-tag when Q-in-Q and VXLAN are involved. [PR1331722](#)

- The operation of pushing a VLAN tag does not work for VXLAN local switching tunneled Q-in-Q traffic. [PR1332346](#)
- Interface with **flexible-vlan-tagging** and **family ethernet-switching** does not work on a QFX10000. [PR1337311](#)

MPLS

- QFX5100: ISSU is not supported with an MPLS configuration. [PR1264786](#)
- Traffic drop during a NSR switchover for RSVP P2MP provider tunnels used by MVPN . [PR1293014](#)
- MPLS forwarding might not happen properly for some LSPs. [PR1319379](#)
- The rpd process might crash on backup Routing Engine due to memory exhaustion. [PR1328974](#)
- The hot standby for the L2 circuit does not work on a QFX5000. [PR1329720](#)
- RSVP sessions go down for ingress LSPs with no-cspf enabled. [PR1339916](#)
- LSP is not received by QFX5110. [PR1351055](#)
- NO-propagate-TTL acts on MPLS Swap operation. [PR1366804](#)
- LSP with auto-bandwidth enabled goes down during HMC error condition. [PR1374102](#)

Platform and Infrastructure

- After upgrading the QFX5100 to Junos OS Release 16.1 or later from Junos OS Release 15.1, the commit warning **/boot/ffp.cookie+** might be seen. [PR1283917](#)
- SFP management Ethernet port C0 might not come up. [PR1298876](#)
- Run-time pps statistics value might show zero for a subinterface of the aggregated Ethernet interface. [PR1309485](#)
- Traffic loss might be seen if traffic is sent through the 40G interface. [PR1309613](#)
- Some log messages are seen on the QFX5110 platform when plugging in an SFP-SX. [PR1311279](#)
- One aggregated Ethernet member cannot send out sFlow sample packets. [PR1311559](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- A memory leak is seen for dot1xd. [PR1313578](#)
- Some certain IGMP join packets cannot be processed correctly at a high rate. [PR1314382](#)
- Transit traffic over a GRE tunnel might hit the CPU and trigger a DDoS violation on the L3 next hop. [PR1315773](#)
- On an L2 next-generation switch platform (QFX5100/QFX10000), l2cpd might drop core files repeatedly if an interface is connected to a VoIP product with LLDP and LLDP-MED enabled. [PR1317114](#)
- Packets such as TDLS without an IP header are looped between virtual gateways. [PR1318382](#)

- The optic interface transmits power even after it has been administratively shutdown. [PR1318997](#)
- The packet might be dropped between 4-60 seconds when the master Routing Engine is rebooted in a virtual chassis. [PR1319146](#)
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX-VC ISSU. [PR1320370](#)
- The MAC address is stuck with "DR" flag on spine node even though packets are received on the interface from the source MAC. [PR1320724](#)
- FPCs go offline in some situations. [PR1321198](#)
- On the QFX10016 EVPN-VXLAN scaled testbed, it takes up to 3 minutes for traffic to converge when configured. [PR1323042](#)
- The openflow session cannot be established correctly with controller and interface options configured on QFX5100 switches. [PR1323273](#)
- Update new firmware versions for jfirmware package for 100G-PSM4 and 100G-AOC issues. [PR1323321](#)
- EVPN Type 5: Unicast traffic is getting dropped on the backup forwarder. [PR1323907](#)
- The next hop of _all_ces__ flood details might go missing. [PR1324739](#)
- The GRE traffic is not decapsulated by the firewall filter. [PR1325104](#)
- VLAN or VLAN bridge might not be added or deleted if there is an IFBD HW token limit exhaustion. [PR1325217](#)
- ARP request packets might not be flooded on a QFX5110. [PR1326022](#)
- The major alarm about 'Fan & PSU Airflow direction mismatch' might be seen by removing the management cable. [PR1327561](#)
- Deleting one VXLAN might cause a traffic loop on another VXLAN in a multi homing EVPN-VXLAN scenario with a service provider style interface. [PR1327978](#)
- QFX10002: Major alarm should be cleared once the chassis has more PEM units installed than the **minimum PEM** configuration. [PR1327999](#)
- Directories and files under **/var/db/scripts** lose execution permission or directory 'jet' is missing under **/var/db/scripts** causing **error: Invalid directory: No such file or directory** error during commit. [PR1328570](#)
- FAN tray removal or insertion trap is not generated for a backup FPC. [PR1329031](#)
- The **etherStatsCRCAlignErrors** counters might disappear in the SNMP tree. [PR1329713](#)
- After commit, members of Virtual Chassis or VCF are split and some members might get disconnected. [PR1330132](#)
- An rpd process core file generated on a new backup Routing Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler** after disabling NSR+GRES. [PR1330750](#)
- The **out of HMC range** and **HMC READ faild** error messages are seen. [PR1332251](#)

- Traffic does not pass through VCP ports after rebooting the Virtual Chassis members. [PR1332515](#)
- EVPN-VXLAN: DF drops multicast traffic. [PR1333069](#)
- On QFX10K8/QFX10K16 platforms, SIB LEDs on the fan tray are off after the replacement of the Fan Tray Controllers (FTC). [PR1334006](#)
- The DHCPv6 SOLICIT message is dropped. [PR1334680](#)
- AI-script does not get auto re-install upon a JUNOS upgrade on Next Generation-Routing Engine. [PR1337028](#)
- The DF of an EVPN instance might flood all the ARP request back to the Ethernet Segment. [PR1337275](#)
- On QFX5100 platforms, LR4 QSFP can take up to 15 min to come up after Virtual Chassis reboot. [PR1337340](#)
- SNMP jnxBoxDescr OID returns different value when upgrading to Junos OS Release 17.2. [PR1337798](#)
- On the QFX10000 platforms, VRRP function does not work well when it is configured on sub-interfaces. [PR1338256](#)
- The traffic coming from the remote VTEP PE might be dropped. [PR1338532](#)
- The analyzer status might show as down when port mirroring is configured to mirror packets from an aggregated Ethernet member. [PR1338564](#)
- The VXLAN traffic might not be transmitted correctly with an IRB interface as the underlay interface of the VTEP tunnel. [PR1338586](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- Multicast traffic drop is seen if downstream IRB interfaces have snooping enabled. [PR1340003](#)
- On the QFX5200: there is an inconsistent result after using **deactivate xxx** command on 'pfc-priority' and 'no-loss' context. [PR1340012](#)
- L3 traffic is not getting converged properly upon disabling the ECMP link between spine and leaf with EVPN-VXLAN configurations. [PR1343172](#)
- BPDU packets might get dropped and bpdu-block-on-edge might not work. [PR1343330](#)
- Broadcast frames might be modified with the ethertype 0x8850. [PR1343575](#)
- EVPN-VXLAN: VLAN with **flexible-tag** mode , the xe statistics appears to not be updated for ingress. [PR1343746](#)
- LACP packets are getting dropped with native-vlan-id configured after reboot. [PR1361054](#)
- QFX5000 Virtual-Chassis acting as EVPN-VxLAN ARP Proxy might cause ARP resolution to fail. [PR1365699](#)
- Hashing does not work for the IPv6 packet encapsulated in VxLAN scenario. [PR1368258](#)
- When native-vlan-id is configured for aggregated Ethernet LACP session to multihomed server goes down. [PR1369424](#)

- A port might still work if it's deleted from an aggregated Ethernet interface. [PR1372577](#)
- Implement the **edit interfaces interface-name ether-options] configured-flow-control** option for the QFX Series. [PR1343917](#)
- For EVPN-VXLAN, the ARP packet uses VRRP/virtual-gateway MAC in an Ethernet header instead of an IRB MAC address. [PR1344990](#)
- In the QFX5100, fan RPM fluctuates when temperature sensor reaches its threshold. [PR1345181](#)
- FXPC process might generate a core file when removing VXLAN configuration. [PR1345231](#)
- Backup Routing Engine might experience a crash, causing vmcore to be generated on master Routing Engine, master Routing Engine performance will not be affected. [PR1346218](#)
- CPU and memory statistics not populating for the backup switch in a QFX5110 Virtual Chassis. [PR1346268](#)
- An incorrect inner VLAN tag is sent from the QFX10000 platform with Q-in-Q configured on the Layer 3 sub interface. [PR1346371](#)
- Statistics daemon pfed might generate core files on an upgrade between certain releases. [PR1346925](#)
- On QFX5110 switches, a DCPFE core file might be generated after removing Type-5 tunnel in an EVPN-VXLAN configuration. [PR1346980](#)
- A QFX5100-48T 10G interface might be auto negotiated at 100M speed instead of 10G. [PR1347144](#)
- On QFX5110-48S-4C platforms, part numbers and serial numbers are not displayed for any of the 10G optics/DAC connected. [PR1347634](#)
- The ARP might not update and packets might get dropped at the Routing Engine. [PR1348029](#)
- On a QFX5100, a BGP session flaps when changes are made on the extended-vni-list under the EVPN hierarchy and if the BGP neighborship is through an IRB. [PR1349600](#)
- QFX5100 40G port has an interoperability issue with some other vendors. [PR1349664](#)
- Blackholing traffic with destination MAC matching the virtual gateway MAC might be seen. [PR1348659](#)
- The pfed process might consume high CPU if subscriber or interface statistics are used at large scale. [PR1351203](#)
- A DCPFE process might crash on QFX10000 switches. [PR1351503](#)
- The GTP traffic might not be hashed correctly for an aggregated Ethernet interface. [PR1351518](#)
- Telemetry traffic does not leave the local box when telemetry server is reachable via a VR routing-instance. [PR1352593](#)
- QFX5100 arp fail after change interface MAC address. [PR1353241](#)
- RPC output not showing failure when running **request system software add** with software already staged. [PR1353466](#)
- SFP-LX10 on QFX5110 might fail to connect with another device. [PR1353677](#)

- The alarm errors might be seen during the bootup on a QFX10000. [PR1354582](#)
- Untagged packets might not be forwarded through the trunk port. [PR1355338](#)
- Commit error observed if box is downgraded from 18.2/18.3 release to 17.3R3. [PR1355542](#)
- On QFX5110 platforms, LX10 SFP needs to be reinserted after autonegotiation is enabled or disabled. [PR1355746](#)
- EVPN-VXLAN: the VXLAN traffic might be lost in EVPN type 2 and type 5 scenario. [PR1355773](#)
- "Load averages" output under **show chassis routing-engine** shows "nan" periodically. [PR1356676](#)
- The IGMP membership report packets might not be forwarded over an interface on a QFX10000. [PR1360137](#)
- On QFX10k, virtual-gateway-address should be only configured on a irb interface associated with a vxlan VLAN. [PR1360646](#)
- Unable to create QFX5200 VC w/100G DACs. [PR1360721](#)
- The GTP traffic might not be hashed correctly on aggregated Ethernet interface. [PR1361379](#)
- The **clear services accounting statistics inline-jflow fpc-slot 0** command should be supported in QFX Series. [PR1362396](#)
- QFX5100VC: Unable to connect management address through vme interface. [PR1362437](#)
- On QFX10008, QFX10016, PTX1000, PTX5000, PTX10008, PTX10016 platforms, MPLS exp rewrite might not work for IPV6 and IPV4 traffic. [PR1364391](#)
- Root password recovery process doesn't work. [PR1365740](#)
- On QFX5100/QFX5110/QFX5200 platforms, ISIS adjacency goes down when mtu 9192 is configured. [PR1368913](#)
- On QFX10000 platforms, before the 17.3R3 code, the maximum number of ESI IFLs was 4000 in the Packet Forwarding Engine. [PR1371414](#)
- TPI-50840 BUM traffic received on 5110 is not flooded to all remote vteps. [PR1373093](#)

Routing Protocols

- Observed mcsnoopd core file at `__raise,abort,__task_quit__,task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal(enable_slip_detector=true,no_exit=true)` at `../src/junos/lib/libjtask/base/task_scheduler.c:275` [PR1305239](#)
- Packet drop is seen when programming for GRE traffic. [PR1308438](#)
- Diffserv bits/ToS bits are not getting copied from Inner IP header to GRE header. [PR1313311](#)
- Some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update. [PR1320723](#)
- On the QFX5100, consistent hashing is not getting programmed. [PR1322299](#)

- IS-IS Layer 2 hello packets are dropped when they come from another vendor's device. [PR1325436](#)
- The loopbacked IRB interface is not accessible to a remote network. [PR1333019](#)
- The dcpe process crash is seen in a route leak scenario on the QFX10000. [PR1334714](#)
- The rpf-check-policy does not work as expected. [PR1336909](#)
- Ping fails if MTU is different on the interfaces. DF is not working as expected. [PR1345495](#)
- vrf-fallback on QFX5K is not supported in ALPM mode. [PR1345501](#)
- On QFX10000 platforms, Netconf SSH TCP port 830 traffic hitting host path/unclassified queue. [PR1345744](#)
- On QFX5100 platforms, parity errors in L3 IPv4 table in the Packet Forwarding Engine memory might cause traffic black holing. [PR1364657](#)

Software Installation and Upgrade

- Commit may fail in single-user mode. [PR1368986](#)

Virtual Chassis

- QFX-Virtual Chassis: Sometimes, the multicast packets are received 2x 3x times than expected. [PR1306239](#)

Resolved Issues: 17.4R1

Class of Service (CoS)

- On QFX5100 switches, traffic might be dropped when there is more than one forwarding class under **forwarding-class-sets**. [PR1255077](#)
- The transmit rate applied with **forwarding-class-set** does not work properly. [PR1277497](#)

EVPNs

- On QFX5100 switches with EVPN-VXLAN deployed, broadcast and multicast traffic might not be sent to other switches through VTEP interfaces. [PR1293163](#)
- On QFX10000 switches with EVPN deployed, packet corruption is seen with Packet Forward Engine trap code (129) egp.v4_chksum when sending L3 inter-VNI traffic with the underlay vlan-tagging inet interface. [PR1295491](#)
- The dynamic routing protocols might not work correctly over the IRB interface in an EVPN-VXLAN scenario with ECMP. [PR1301521](#)
- QFX5110-48S: L3 VPN traffic is dropped for some instances when EVPN-VXLAN configuration is removed and reapplied. [PR1307590](#)

Hardware

- FEC is disabled by default on 100G-LR optics for QFX5200 switches. [PR1286389](#)
- The 1G copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms. [PR1286709](#)
- ULC-60S-6Q LC on QFX10008: The port becomes unusable after inserting a third-party SFP-T optic. [PR1294394](#)
- Update new firmware versions for jfirmware package for 100G-PSM4 and 100G-AOC issues. [PR1323321](#)

High Availability (HA) and Resiliency

- Normal VRRP MAC is triggering a MAC move, and logical interfaces on the BD are getting shut down. [PR1285749](#)

Infrastructure

- Create new command: "enable-tcp-nodelay" and allow flash sub-jobs to run for max quantum. [PR1136167](#)
- Disabled 10-Gigabit Ethernet interfaces might stay up on QFX10000 line switches. [PR1300775](#)
- The 40-Gigabit Ethernet connection between two QFX5100-24Qs might not come up sometimes. [PR1178799](#)
- QFX10002 and QFX10008: BFD sessions over IRB interfaces with Junos OS Releases 17.1R1, 17.1R2, 17.2R1, and 17.3R1 are centralized. [PR1284743](#)

Interfaces and Chassis

- Random interfaces do not come up after a line card is rebooted. [PR1262839](#)
- Copper ports flap on QFX5100-48T when short-reach-mode is enabled. [PR1248611](#)
- The 40-Gigabit Ethernet interface might flap between QFX5100 and other products. [PR1273861](#)
- QFX10000-12C-DWDM: an ot- interface link flap is observed whenever an optics TCA alarm is raised; however, there is no LOS and no traffic loss is observed. [PR1279351](#)
- On QFX5100 switches, an AE interface might flap upon commit if an explicit speed is configured on an AE member interface [PR1284495](#)
- On QFX10000 line switches, the input and output rates for 10-Gigabit, 40-Gigabit, or 100-Gigabit Ethernet interfaces are not 0 if the interface is down. [PR1291412](#)
- Traffic might not be received on a 1-Gigabit Ethernet interface if autonegotiation is disabled and speed/duplex is configured on both the QFX Series switch and the peer host. [PR1292275](#)
- High heap memory utilization might be seen if multiple SFP-T optics are inserted or **set interface <> link-mode full-duplex** is enabled. [PR1294208](#)
- The 40-Gigabit Ethernet interface might not come up if a specific vendor's DAC cable is used. [PR1296011](#)
- QFX10008/10016: Commit error is seen when configured with mixed speed. [PR1301923](#)

Junos Fusion Satellite Software

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

Layer 2 Features

- To set up PTP BC forwarding on a QFX10002, configure routing on the interface or add a static ARP entry on the remote PTP device. [PR1275327](#)
- Feature swap-swap might not work as expected in a Q-in-Q scenario. [PR1297772](#)
- QFX5100 crashes and the fxcp process generates a core file. [PR1306768](#)

MPLS

- QFX10008 is dropping egress MPLS traffic, if the egress interface is an IRB with access L2 AE interface. [PR1279827](#)

Network Management and Monitoring

- UFT for non-local member is not shown in the CLI. [PR1243758](#)
- LAG interface input bytes counter continuously decreases when no packets come in. [PR1266062](#)
- SNMP process is not running on QFX Series switches with incorrect source addresses. [PR1285198](#)
- On QFX5100, an incorrect alarm type might be displayed. [PR1291622](#)
- Previous learned MAC address from remote ESI cannot be changed to local. [PR1303202](#)
- The sflow records are missing "extendedType ROUTER" fields as well as an outbound interface for traffic that is using BGP multipath. [PR1303236](#)
- QFX5110-48S: digital optical monitoring statistics cannot be received through the CLI in Junos OS Releases 15.1X53 through 17.x. [PR1305506](#)

Platform and Infrastructure

- A hostname synchronization issue occurs between the Junos OS VM instance and the Linux host in TVP platforms. [PR1283710](#)
- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- The dcpfe process might crash and restart on MC-LAG active and standby nodes when there is ARP/NDP next-hop change. [PR1299112](#)
- OSPFv3 authentication using IPsec SA does not work if you are using IPsec to authenticate OSPFv3 neighbors on some QFX Series platforms. [PR1301428](#)

Port Security

- On QFX10000 switches, MACsec sessions are not coming up on a Layer 3 logical interface. [PR1282995](#)
- Proxy-ARP and ARP suppression are not yet supported for the QFX10000 line. [PR1293707](#)

Routing Protocols

- When the static link protection mode configured backup state is down, the primary port goes to down state instead of the secondary port, and the secondary remains in up state. [PR1276156](#)
- Analytics JSON data format is reporting a incorrect value for 'rxbps' counter. [PR1285434](#)
- On QFX5100 switches, if a term with the policer action is configured, **dc-pfe: list_destroy()** messages might be displayed on commit. [PR1286209](#)
- GRE tunnel traffic does not switch over to the alternate path if the primary path to the tunnel destination changes. [PR1287249](#)
- UDP traffic with destination port 520 and 521 is discarded on QFX5110 switches after a Junos OS upgrade. [PR1287271](#)
- 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

- 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

Changes in Behavior and Syntax	347
Known Behavior	354
Known Issues	358
Documentation Updates	391
Migration, Upgrade, and Downgrade Instructions	391

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.4R3.

SEE ALSO

New and Changed Features	335
Changes in Behavior and Syntax	347
Known Behavior	354
Known Issues	358
Resolved Issues	369
Migration, Upgrade, and Downgrade Instructions	391

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 392
- Installing the Software on QFX10002 Switches | 394
- 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 | 394
- Installing the Software on QFX10008 and QFX10016 Switches | 396
- Performing a Unified ISSU | 400
- Preparing the Switch for Software Installation | 401
- Upgrading the Software Using Unified ISSU | 401
- Upgrade and Downgrade Support Policy for Junos OS Releases | 404

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  
-R3.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
-R3.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
-R3.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
-R3.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-10-17.4R3.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 401](#)
- [Upgrading the Software Using Unified ISSU on page 401](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

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

To verify that nonstop active routing is enabled:

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

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

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

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.

Upgrading the Software Using Unified ISSU

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

To upgrade the switch using unified ISSU:

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

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

where *package-name.tgz* is, for example, *jinstall-host-qfx-5-17.3R1-signed.tgz*.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-17.4
-R3.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-17.4
-R3.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 | 335](#)

[Changes in Behavior and Syntax | 347](#)

[Known Behavior | 354](#)

[Known Issues | 358](#)

[Resolved Issues | 369](#)

[Documentation Updates | 391](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- New and Changed Features | 405
- Changes in Behavior and Syntax | 417
- Known Behavior | 422
- Known Issues | 424
- Resolved Issues | 427
- Documentation Updates | 449
- Migration, Upgrade, and Downgrade Instructions | 449

These release notes accompany Junos OS Release 17.4R3 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

- Release 17.4R3 New and Changed Features | 406
- Release 17.4R2 New and Changed Features | 406
- Release 17.4R1-S1 New and Changed Features | 406
- Release 17.4R1 New and Changed Features | 408

This section describes the new features and enhancements to existing features in Junos OS Release 17.4R3 for the SRX Series devices.

Release 17.4R3 New and Changed Features

There are no new features in Junos OS Release 17.4R3 for the SRX Series devices.

Release 17.4R2 New and Changed Features

There are no new features in Junos OS Release 17.4R2 for the SRX Series devices.

Release 17.4R1-S1 New and Changed Features

Junos OS Release 17.4R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/SRX4200, SRX5400, SRX5600, and SRX5800.

Junos OS Release 17.4R1-S1 supports SRX4600 device.

Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 15.1X49-D80 through 15.1X49-D100. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D100 are not available in 17.4 releases.

NOTE: Junos OS for SRX Series Software documentation includes information about SRX4600 Services Gateway.

New features for security platforms in Junos OS Release 17.4R1 and Junos OS Release 17.4R1-S1 include:

Chassis Cluster

- **Media Access Control Security (MACsec) (SRX4600)**—Starting in Junos OS Release 17.4R1-S1, Media Access Control Security (MACsec) is supported on HA control and fabric ports of SRX4600 devices in chassis cluster mode to secure point-to-point Ethernet links between two nodes in a cluster.

In the SRX chassis cluster implementation, the control and fabric link carry secure traffic between two nodes in clear text format. Because of this, it is important to encrypt the data between the two nodes. MACsec is an industry-standard security technology that provides secure communication and identifies and prevents most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks. MACsec can be used in combination with other security protocols to provide end-to-end network security.

See [Understanding Media Access Control Security \(MACsec\)](#).

Hardware

- **SRX4600 Services Gateway**—Starting with Junos OS Release 17.4R1-S1, SRX4600 Services Gateways are available as the next-generation, high-performance, and scalable security services devices. The services gateway supports 75-Gbps Internet mix (IMIX) throughput, is suited for large enterprises and small to medium data centers. The SRX4600 Services Gateway provides industry-leading next-generation firewall capabilities (AppID, UserFW, IPS, UTM, and so on) and advanced threat detection and mitigation capabilities features such as SecIntel and SkyATP. The Services Gateway features two high-performance Intel Xeon processors with 14 cores per processor.

Platforms and Infrastructure

- **Software support for SRX4600 devices**—Starting in Junos OS Release 17.4R1-S1, Junos OS supports the SRX4600 Services Gateway. The SRX4600 device is a high-end dynamic services gateway that consolidates security functionality, networking services, and uncompromised performance for medium to large enterprises. With advanced security and threat mitigation capabilities, SRX4600 device can be used for campus edge integrated firewall, data center edge firewall, data center core firewall, LTE security gateway, and Gi/SGi firewall.

SRX4600 device supports Juniper's Software-Defined Secure Network (SDSN) framework, including Sky Advanced Threat Prevention (Sky ATP), which is built around automated and actionable intelligence that can be shared quickly to recognize and mitigate threats.

The SRX4600 device supports the following software features:

- Stateful firewall
- Application security suite
- UTM (Sophos AV, Web filtering, content filtering, and antispam)
- IDP
- Advanced anti-malware
- High availability (Chassis cluster)

- Dual HA control ports (10G)
- MACsec support for HA ports
- Ethernet interfaces through QSFP28 (100G modes), QSFP+ (40G/4x10G modes) and SFP+ (10G mode)
- IPsec VPN, including AutoVPN and Group VPNv2
- QoS and network services
- J-Web
- Routing policies with multicast

The SRX4600 implements use of an individual thread for each session that is dedicated to management of that session and its flow. As a result, out-of-order packet problems that can occur with concurrent processing are eliminated.

Installation packages available for SRX4600 devices are, Preboot Execution Environment (PXE), USB install media package, and CLI upgrade.

You can use the **show chassis hardware** command to display the part number and the model number of the SRX4600 device.

You can use the **show security ipsec tunnel-distribution** command to display the number of VPN tunnels anchored in each thread ID.

[See [Understanding Flow Processing on the SRX4600 Device](#).]

Security

- **Secure Boot (SRX4600)**—Starting in Junos OS Release 17.4R1-S1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms.

[See [Feature Explorer](#) and enter **Secure Boot**.]

Release 17.4R1 New and Changed Features

Junos OS Release 17.4R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/SRX4200, SRX5400, SRX5600, and SRX5800.

Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 15.1X49-D80 through 15.1X49-D100. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D100 are not available in 17.4 releases.

ALG

- **H.323 gateway-to-gateway support (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, the gateway-to-gateway call feature is supported on the H.323 ALG. This feature introduces

one-to-many mapping between an H.225 control session and H.323 calls as multiple H.323 calls go through a single control session.

[See [Understanding H.323 ALG.](#)]

- **NAT64 support for H.323 ALG (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, the H.323 ALG supports NAT64 rules in an IPv6 network.

[See [Understanding H.323 ALG.](#)]

Application Security

- **Advanced policy-based routing (APBR) with midstream support (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, SRX Series Services Gateways support advanced policy-based routing (APBR) with an additional enhancement to apply the APBR in the middle of a session (midstream support). With this enhancement, you can apply APBR for a non-cacheable application and also for the first session of the cacheable application.

You can fine-tune the outbound traffic with APBR configuration (for example, limiting route changes and terminating sessions) to avoid issues such as excessive transitions due to frequent route changes.

The enhancement provides more flexible traffic-handling capabilities that offer granular control for forwarding packets.

[See [Understanding Advanced Policy-Based Routing.](#)]

- **Application tracking enhancements to support category and subcategory (SRX Series, vSRX instances)**—Starting from Junos OS Release 17.4R1, AppTrack session create, session close, and volume update logs include new fields **category** and **subcategory**. AppTrack syslog message provide general information about the application type, and including category and subcategory of the application in the message, helps in categorizing the applications.

[[Understanding AppTrack.](#)]

Authentication and Access

- **User firewall support for IPv6 (SRX Series, vSRX instances)**—Starting in Junos OS Release 17.4R1, SRX Series devices support IPv6 addresses for user firewall (UserFW) authentication. This feature allows IPv6 traffic to match any security policy configured for source identity. Previously, if a security policy was configured for source identity and “any” was specified for its IP address, the UserFW module ignored the IPv6 traffic. IPv6 addresses are supported for the following authentication sources:
 - Active directory authentication table
 - Device identity with active directory authentication
 - Local authentication table
 - Firewall authentication table

[See [Overview of Integrated User Firewall](#).]

Chassis Cluster

- **Preemptive delay timer (SRX Series)**—Starting with Junos OS Release 17.4R1, a failover delay timer is introduced on SRX Series devices in a chassis cluster to limit the flapping of redundancy group state between the secondary and the primary nodes in a preemptive failover.

Back-to-back failovers of a redundancy group in a short interval can cause the cluster to exhibit unpredictable behavior because of flapping of the active and backup systems.

To prevent this, a delay timer can be configured to delay the immediate failover for a configured period of time—between 1 and 21,600 seconds. In addition, you can configure the preemptive limit to restrict the number of failovers (1 to 50) in a given time period (1 to 1440 seconds) when preemption is enabled for a redundancy group.

This enhancement enables the administrator to introduce a failover delay, which can reduce the number of failovers and result in a more stable network state due to the reduction in active / backup flapping within the redundancy group.

[[Understanding Chassis Cluster Redundancy Group Failover](#).]

Class of Service (CoS)

- **Support for CoS on dlo 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 dlo interface for 4G wireless modems: behavior aggregate classifiers, multifield classifiers, policers, shapers, schedulers, and rewrite rules. The dialer interface, dlo, is a logical interface for configuring properties for modem connections.

[See [LTE Mini-PIM Overview](#).]

- **Support CoS on Logical Tunnel Interface in a Chassis Cluster on SRX300, SRX320, SRX340, SRX345, and SRX550M devices**— Starting with Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, queuing is supported on logical tunnel (lt) interfaces to allow CoS configuration.

[See [CoS Queuing for Tunnels Overview](#).]

- **Support for port-based egress traffic shaping and policing on SRX Series devices**— Starting with Junos OS Release 15.1X49-D100 and Junos OS Release 17.4R1, you can configure egress traffic shaping and policing at the physical port level, which limits the egress traffic rate of all logical interfaces on the port.

[See [shaping-rate \(CoS Interfaces\)](#).]

Flow-based and Packet-based Processing

- **Hash-based session distribution (SRX5400, SRX5600, SRX5800)**— Starting with Junos OS Release 17.4R1, traffic is hashed and distributed to different SPUs by the IOC, based on a hash-based session distribution algorithm. This enhancement provides an even hash distribution among all SPUs by using a larger fixed-length hash table. In earlier Junos OS releases, the traffic distribution was uneven among all SPUs in some cases due to a smaller fixed-length hash table.

[See [Understanding Load Distribution in SRX5800, SRX5600, and SRX5400 Devices and vSRX](#).]

GPRS

- **Support for GTP handover group (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 devices and vSRX instances)**—Starting with Junos OS Release 17.4R1, GTP handover group configuration is supported on GTP profiles. An administrator can configure a GTP profile and associate a GTP handover group to a GTP profile.

A GTP handover group is a set of SGSNs or serving gateway (SGW) with a common address-book library. When a GTP handover group name is referenced by a GTP profile, the device checks to see if the current SGSN/SGW address and the proposed SGSN/SGW address are contained within the same GTP handover group. If both the current and proposed SGSN/SGW addresses are contained within the same GTP handover group, then the handover is allowed. If both the current and proposed SGSN/SGW addresses are not within the same GTP handover group, then the profile for the default handover group is used.

This feature enables the administrator to define policies that determine whether handover can happen between individual SGSNs/SGW and/or groups of SGSNs/SGW for roaming.

[See [GTP Handover Group Overview](#).]

Hardware

- **SRX345 Services Gateway (DC power supply model)**—The SRX345 Services Gateway now includes a DC model. The DC model has a single internal power supply, which is not field-replaceable. The DC model supports the same features as those supported on the existing SRX345 Services Gateways. The minimum Junos OS release supported on the DC model is 17.4R1. The services gateway can be managed using the CLI, Junos Space, and J-Web.

[See [SRX345 Services Gateway Description](#).]

Interface and Chassis

- **MACsec support (SRX300, SRX320, SRX340 and SRX345)**—Starting in Junos OS Release 17.4R1, Media Access Control Security (MACsec) is supported on all MACsec-capable ports of SRX300, SRX320, SRX340 and SRX345 devices.

On SRX300 line devices MACsec is supported on the following ports:

- SRX300 and SRX320: 2 ports (on two fixed SFP interfaces.)
- SRX340 and SRX345: 16 ports (on eight fixed SFP interfaces + eight fixed Ethernet ports)

[See [Understanding Media Access Control Security \(MACsec\)](#).]

- **PPPoE support on SRX Series and vSRX devices**—Starting in Junos OS Release 17.4R1, SRX series devices and vSRX support Point-to-Point Protocol over Ethernet (PPPoE). You can connect multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. The hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

[See [Understanding PPPoE Interfaces.](#)]

- **RFC 4638 support for SRX300, SRX320, SRX340, SRX345, and SRX550M devices**— Starting in Junos OS Release 17.4R1, you can use the PPP-Max-Payload option to override the default behavior of the PPPoE client by providing a maximum size that the PPP payload can support in both sending and receiving directions. The PPPoE server might allow the negotiation of an MRU larger than 1492 and the use of an MTU larger than 1492.

[See [Understanding MTU and MRU Configuration for PPP Subscribers.](#)]

Installation and Upgrade

- **Upgraded FreeBSD support (SRX1500, SRX4100, SRX4200, and vSRX instances)**—Starting with Junos OS Release 17.4R1, the Junos Control Plane (JCP) virtual machine (VM) in the SRX Series devices is upgraded to support FreeBSD 11. Two virtual CPUs (VCPU) are allocated for JCP VM in the Linux host to improve Routing Engine performance for SRX4100 and SRX4200 devices and vSRX instances. For vSRX, additional vCPU will be allocated if you allocate more CPUs than the minimum required. For SRX1500 devices, no additional CPUs are available to allocate for JCP VM.

[See [Understanding Junos OS with Upgraded FreeBSD for SRX Series Devices.](#)]

Logical System

- **Logical system (LSYS) support (SRX1500)**—Starting in Junos OS Release 17.4R1, the logical system feature is supported on SRX1500 devices in addition to the existing support on SRX Series devices such as SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800. A logical system provides virtualization on a device that is partitioned into multiple logical administrative segments. Each segment can have its own security, routing, and bridging attributes.

[See [Understanding Logical Systems for SRX Series Services Gateways.](#)]

Management

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

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

NAT

- **Source NAT resource allocation improved (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 17.4R1, source NAT resources handled by the central point architecture have been offloaded to the SPUs when the SPC number is more than four, resulting in more efficient resource allocation.

[See [Understanding Central Point Architecture Enhancements for NAT.](#)]

Routing Policy and Firewall Filters

- **Maximum number of addresses per security policy increased (SRX550M)**—Starting in Junos OS Release 17.4R1, the maximum number of addresses per policy has been increased from 1024 to 2048 for SRX550M. SRX300, SRX320, SRX340 and SRX345 devices already support 2048 source and 2048 destination addresses per policy.

Routing Protocols

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

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

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

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

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

System Logging

- **Support for log warning messages on throughput overuse (SRX4100)**—Starting with Junos OS Release 17.4R1, when Internet mix (IMIX) throughput exceeds the limitation for an SRX4100 device, new log warning messages are logged. These log warning messages remind you that there is throughput overuse.

[See [Log File Sample Content](#).]

- **On-box reporting enhancements (SRX Series, vSRX instances)**—Starting in Junos OS Release 17.4R1, SRX4600 devices support the on-box reporting feature, which is already supported on SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200 devices and vSRX instances. Also, the on-box reports are now enhanced to provide comprehensive and detailed reports.

The on-box reporting feature now provides the following enhancements:

- AppTrack API gets information on application category, subcategory, and risk level. An RTLOG module uses this API to get and send information to the local log management process (daemon).
- Reports for applications, categories, subcategories, risk levels, and botnet threats are now by count and volume.
- Application information is generated in UTM log reports.
- Logs can now be listed from latest to oldest. Previously, logs were sorted only from oldest to latest.
- SRX4600 devices now have a hard disk partition available to save traffic logs.

[See [Understanding On-Box Logging and Reporting](#).]

Screens

- **UDP flood screen whitelist (SRX300, SRX320, SRX340, SRX345, SRX1400, SRX4100, and SRX4200 devices, and vSRX instances)**—Starting with Junos OS Release 17.4, UDP flood whitelist mechanism is implemented on SRX300, SRX320, SRX340, SRX345, SRX1400, SRX4100, and SRX4200 devices, and vSRX instances.

When UDP is enabled in a zone, all the UDP traffic performs UDP flood attack detection. The UDP packets that are above the threshold level will be dropped. To avoid these packet drops and instead allow these packets to bypass UDP flood detection, the UDP flood screen whitelist is implemented. To support UDP flood whitelist, the traffic from addresses in the whitelist groups will bypass UDP flood check. Both IPv4 and IPv6 whitelists are supported and can be configured using a single address or a subnet address. UDP flood whitelist supports a maximum of 32 whitelist groups and each group has 32 or fewer IPv4 or IPv6 addresses.

See [Understanding Whitelists for UDP Flood Screens](#).

UTM

- **Custom URL category support for SSL forward proxy (SRX Series)**—Starting with Junos OS Release 17.4R1, the whitelisting feature is extended to include custom URL categories supported by UTM in the whitelist configuration of SSL forward proxy. In this implementation, the Server Name Indication (SNI) field is extracted by the UTM module from client hello messages to determine the URL category. SNI is an extension of the SSL/TLS protocol. Each URL category has a unique ID. The list of URL categories in the whitelist is parsed and the corresponding category IDs are pushed to the Packet Forwarding Engine for each SSL forward proxy profile. The SSL forward proxy then determines through APIs whether to accept the proxy or to ignore the session.

[See [SSL Proxy Overview](#)]

- **Enhanced Web Filtering (EWF) reputation and categorization behavior support for EWF category (SRX Series)**—Starting from Junos OS Release 17.4R1, predefined base filters, defined in a category file, are supported for individual EWF categories. Each EWF category has a default action in a base filter, which is attached to the user profile to act as a backup filter. If the categories are not configured in the user profile, then the base filter takes the action. Online upgradation of base filters is also supported. Further, users can apply global reputation values, provided by the Websense ThreatSeeker Cloud (TSC). For the non-category URLs, the global reputation value is used to perform filtering, and from this release onward, the reputation base scores are configurable.

[See [Understanding Enhanced Web Filtering Process](#).]

- **Local Web filtering enhancement to support custom category configuration (SRX Series)**—Starting from Junos OS Release 17.4R1, support for custom category configuration is available for EWF, local, and Websense redirect profiles. The **custom-message** option is also supported in a category for local Web filtering and Websense redirect profiles. You can create multiple URL lists (custom categories) and apply them to a UTM Web filtering profile with actions such as permit, permit and log, block, and quarantine.

To create a global whitelist or blacklist, apply a local Web filtering profile to a UTM policy and attach it to a global rule.

[See [Understanding Local Web Filtering](#).]

- **Support for new Websense EWF categories (SRX Series)**—Starting from Junos OS Release 17.4R1, you can download and dynamically load new Enhanced Web Filtering (EWF) categories. The downloading and dynamic loading of the new EWF categories do not require a software upgrade. Websense occasionally releases new EWF categories. EWF classifies websites into categories according to host, URL, or IP address and performs filtering based on the categories.

[See [Understanding Redirect Web Filtering](#).]

VPN

- **Increased number of IKE security associations supported (SRX5600, SRX5800)**—Starting from JunosOS Release 17.4R1, SRX5600 with 5 SPC2 cards, and SRX5800 with 10 SPC2 cards can support up to 50,000 IKE security associations (SAs) (each SPC2 card supports upto 20,000 IKE SAs (5,000 IKE SAs / SPU)) for AutoVPN networks in point-to-point secure tunnel mode with multiple traffic selectors. There are no changes in configuration.

[See [Understanding AutoVPN](#).]

- **IPv6 address support for point-to-point AutoVPN networks that use traffic selectors (SRX Series, vSRX instances)**—Starting with Junos OS Release 17.4R1, AutoVPN networks that use secure tunnel interfaces in point-to-point mode support IPv6 addresses for traffic selectors and for IKE peers.

NOTE: IPv6 addresses are not supported for AutoVPN networks in point-to-multipoint secure tunnel mode.

[See [Understanding AutoVPN](#) and [Understanding AutoVPN with Traffic Selectors](#).]

- **IPsec VPN performance optimization (SRX5400, SRX5600, SRX5800)**—Starting with Junos OS Release 17.4R1, IPsec VPN performance is optimized when the VPN session affinity and performance acceleration features are enabled. Session affinity is enabled with the **set security flow load-distribution session-affinity ipsec** command, while performance acceleration is enabled with the **set security flow ipsec-performance-acceleration** command.

[See [Accelerating the IPsec VPN Traffic Performance](#) and [Understanding VPN Session Affinity](#).]

SEE ALSO

[Changes in Behavior and Syntax](#) | 417

[Known Behavior](#) | 422

[Known Issues](#) | 424

[Resolved Issues](#) | 427

Changes in Behavior and Syntax

IN THIS SECTION

- [Release 17.4R3 Changes in Behavior and Syntax | 417](#)
- [Release 17.4R2 Changes in Behavior and Syntax | 419](#)

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 17.4R3.

Release 17.4R3 Changes in Behavior and Syntax

Application Security

- Starting in Junos OS Release 17.4R3, you can set up automatic update of the application signature package in new format. Now you can use the YYYY-MM-DD.hh:mm format to configure the time for automatic download for application signatures. For example, following statement sets the start time as 10 AM on June 30, 2019:

```
user@host# set services application-identification download automatic start-time 2019-06-30.10:00:00
```

You can configure the automatic updates using the new format once you upgrade your previous Junos OS version to any of the above supported Junos OS version.

Authentication and Access Control

- **Enhanced output for show security firewall-authentication jims statistics (SRX Series)**—Starting in Junos OS Release 17.4R3, the output for the **show security firewall-authentication jims statistics** operational command is enhanced to display the statistics of both the primary and secondary JIMS servers. For example, the **show security firewall-authentication jims statistics** operational command displays the following sample output:

```
root@user> show security firewall-authentication jims statistics
```

```

Primary server:
  Push success counter: 1
  Push failure counter: 0

Secondary server:
  Push success counter: 1
  Push failure counter: 0

```

[See [show security firewall-authentication jims statistics](#).]

Chassis Clustering

- **MACsec on Chassis cluster (SRX4600)**—Starting in Junos OS Release 17.4R3, any new MACsec chassis cluster port configurations or modifications to existing MACsec chassis cluster port configurations will require the chassis cluster to be disabled and displays a warning message **Modifying cluster control port CA will break chassis cluster**. Once disabled, you can apply the preceding configurations and enable the chassis cluster.

[See [Configuration Considerations When Configuring MACsec on Chassis Cluster Setup](#).]

- **Chassis cluster with SPC card (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 17.4R3, when a SPC is the control plane as well as hosting the control port, this creates a single point of failure. If the SPC goes down on the primary node, the node is automatically rebooted to avoid split brain.

[Connecting SRX Series Devices to Create a Chassis Cluster](#)

Network Management and Monitoring

- **NSD Restart Failure Alarm (SRX Series)**—Starting in Junos OS Release 17.4R3, a system alarm is triggered when the Network Security Process (NSD) is unable to restart due to the failure of one or more NSD subcomponents. The alarm logs about the NSD are saved in the messages log. The alarm is automatically cleared when NSD restarts successfully.

The **show chassis alarms** and **show system alarms** commands are updated to display the following output when NSD is unable to restart - **NSD fails to restart because subcomponents fail**.

[See [Alarm Overview](#).]

- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (SRX Series)**—Starting in Junos OS Release 17.4R3, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

User Access and Authentication

- **SSH protocol version v1 option deprecated from CLI (SRX Series)**—Starting in Junos OS Release 17.4R3, the nonsecure SSH protocol version **v1** option is not available at the `[edit system services ssh protocol-version]` hierarchy level. The SSH protocol version **v2** is the default option to remotely manage systems and applications. The SSH protocol version **v1** deprecation enables Junos OS to be compatible with OpenSSH 7.4 and later versions.

Junos OS Release 17.4R2 and earlier releases supported the SSH protocol version **v1** option to remotely manage systems and applications.

[See [protocol-version](#).]

- **Enabling and disabling SSH login password or challenge-response authentication (SRX Series)**—Starting in Junos OS Release 17.4R3, you can disable either the SSH login password or challenge-response authentication at the `[set system services ssh]` hierarchy level.

In Junos OS releases earlier than Release 17.4R3, you can enable or disable both SSH login password and challenge-response authentication simultaneously at the `[set system services ssh]` hierarchy level.

[See [Configuring SSH Service for Remote Access to the Router or Switch](#).]

Release 17.4R2 Changes in Behavior and Syntax

Chassis Cluster

- **IP Monitoring**—Starting with Junos OS Release 17.4R2, on all SRX Series devices, if the reth interface is in bundled state, IP monitoring for redundant groups is not supported on the secondary node. This is because the secondary node sends reply using the lowest port in the bundle which is having a different physical MAC address. The reply is not received on the same physical port from which the request is sent. If the reply comes on the other interface of the bundle, then the internal switch drops it.
- **Power Entry Module**—Starting with Junos OS Release 17.4R2, when you use DC PEM on SRX Series devices operating in chassis cluster mode, the output of `show chassis power` command shows **DC input: 48.0 V input (57000 mV)**. The value **48.0 V input** is a fixed string and can be interpreted as a measured input voltage. The acceptable range of DC input voltage accepted by the DC PEM is 40 to 72 V. The **(57500 mV)** is a measured value, but is not related with the input. It is the actual output value of the PEM and the value is variable. The **DC input:** from `show chassis power` and **Voltage:** information from `show chassis environment pem` command output are removed for each PEM.
- SRX5400, SRX5600, and SRX5800 devices operating in a chassis cluster might encounter the em0 or em1 interface link failure on either of the nodes, which results in split-brain condition. That is, both devices are unable to detect each other. If the failure occurs on the secondary node, the secondary node is moved to the disabled state.

This solution does not cover the following cases:

- em0 or em1 failure on primary node
- HA process restart

- Preempt conditions
- Control link recovery

IDP

- Custom Attack (SRX Series)—Starting with Junos OS Release 17.4R2, the maximum number of characters allowed for a custom attack object name is 60. You can validate the statement using the CLI **set security idp custom-attack** command.

Forwarding and Sampling

- Support for Address Resolution Protocol (ARP) throttle and ARP detect [SRX5400, SRX5600, and SRX5800]—Starting in Junos OS Release 17.4R2, an ARP throttling mechanism is introduced for SRX Series devices.

Excessive ARP processing results in high utilization of Routing Engine CPU resources, resulting in deprivation of CPU resources to other Routing Engine processes. To provide protection against excessive ARP processing, you can now use the following configuration statements:

- **edit forwarding-options next-hop arp-throttle *seconds***
- **edit forwarding-options next-hop arp-detect *milliseconds***



CAUTION: We recommend that only advanced Junos OS users attempt to configure the ARP throttle and ARP detect feature. An improper configuration could result in high CPU utilization of the Routing Engine, which could affect other processes on your device.

[See [arp-throttle](#) and [arp-detect](#)].

System Logging

- **System log host support (SRX300, SRX320, SRX340, SRX345 Series devices)**— Starting in Junos OS Release 17.4R2, when the device is configured in stream mode, you can configure maximum of eight system log hosts.

In Junos OS Release 17.4R1 and earlier releases, you can configure only three system log hosts in the stream mode. If you configure more than three system log hosts, then the following error message is displayed **error: configuration check-out failed**.

User Interface and Configuration

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

VPNs

- **Certificate revocation list (SRX Series)**—Local certificates are being validated against certificate revocation list (CRL) even when CRL check is disabled. Starting in Junos OS Release 17.4R3, this can be stopped by disabling the CRL check through the Public Key Infrastructure (PKI) configuration. When CRL check is disabled, PKI will not validate local certificate against CRL.

[See [revocation-check \(Security PKI\)](#) and [Understanding Online Certificate Status Protocol and Certificate Revocation Lists](#).]

SEE ALSO

[New and Changed Features | 405](#)

[Known Behavior | 422](#)

[Known Issues | 424](#)

[Resolved Issues | 427](#)

[Documentation Updates | 449](#)

[Migration, Upgrade, and Downgrade Instructions | 449](#)

Known Behavior

IN THIS SECTION

- Authentication and Access | 422
- Chassis Clustering | 422
- J-Web | 423
- Layer 2 Ethernet Services | 423
- Platform and Infrastructure | 423
- User Interface and Configuration | 424
- VPNs | 424

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.4R3 for the SRX Series.

Authentication and Access

- On SRX Series devices with 256K user firewall authentication entries, in case of a failover or when Packet Forwarding Engine restart occurs, the **show services user-identification** command will generate response timeout. This timeout will last for at least 10 minutes. [PR1302269](#)

Chassis Clustering

- On all SRX Series devices, if you enable IP monitoring for redundancy groups, the feature might not work properly on the secondary node if the reth interface has more than one physical interface configured. This is because the backup node will send traffic using the MAC address of the lowest port in the bundle. If the reply does not come back on the same physical port, then the internal switch will drop it. [PR1344173](#)
- SRX5400, SRX5600, and SRX5800 devices operating in a chassis cluster might encounter the em0 or em1 interface link failure on either of the nodes, which results in split-brain condition. That is, both devices are unable to detect each other. If the failure occurs on the secondary node, the secondary node is moved to the disabled state.

The following cases are not covered:

- em0 or em1 failure on primary node
- HA process restart

- Preempt conditions
- Control link recovery

J-Web

- On SRX Series devices, you cannot view the custom log files created for event logging in J-Web. [PR1280857](#)
- 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 a 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](#)
- Generation of reports will work in Internet Explorer and Chrome browsers. To generate a report in Firefox, delete the existing **ff** profile and relaunch Firefox with new profile. [PR1303722](#)
- Uploading a certificate using the Browse button stores the certificate in the device at the `/jail/var/tmp/uploads/` location. The certificate is deleted when you execute the **request system storage cleanup** command. [PR1312529](#)
- The values of address and address-range are not displayed in the Inline address-set creation pop-up window of Juniper Identity Management Service (JIMS). [PR1312900](#)

Layer 2 Ethernet Services

- On SRX Series devices, the PPPoE and DHCPv6 cannot work together on the `pp0` interface. [PR1229836](#)

Platform and Infrastructure

- On SRX4600 devices, the USB disk is not made available to Junos OS. However, the USB disk is available for host OS (Linux) with full access. USB is still used in the booting process (install and recovery functions). [PR1283618](#)
- On SRX Series devices, when you perform a downgrade from a Junos OS release with upgraded FreeBSD (Junos OS Release 17.3+ for SRX5000 Series devices, Junos OS Release 17.4+ for SRX1500 and SRX4100 or 4200) to Junos OS Release 15.1X49, use the **force** option with the **request system software add** command. If you do not use the force option, an error message opens indicating that you have not used the option. Note that in such a downgrade, the configuration and other files might be lost from the device due to file system repartitioning. [PR1350558](#)

User Interface and Configuration

- In a few SRX Series setups, committing a configuration with a considerable number of logical system configuration can take a little more time than usual. The reason can be taking backup of previous configurations might take a little longer to finish. [PR1339862](#)

VPNs

- On SRX5400, SRX5600, and SRX5800 devices, when CoS is enabled on st0 interface and the incoming traffic rate destined for the st0 interface is higher than 3,00,000 packets per second (pps) per SPU, the device might drop some of the high priority packets internally and shaping of outgoing traffic might be impacted. We recommend that you configure an appropriate policer on the ingress interface to limit the traffic below 3,00,000 pps per SPU. [PR1239021](#)

SEE ALSO

[New and Changed Features | 405](#)

[Changes in Behavior and Syntax | 417](#)

[Known Issues | 424](#)

[Resolved Issues | 427](#)

[Documentation Updates | 449](#)

[Migration, Upgrade, and Downgrade Instructions | 449](#)

Known Issues

IN THIS SECTION

- [Chassis Clustering | 425](#)
- [Flow-Based and Packet-Based Processing | 425](#)
- [Intrusion Detection and Prevention \(IDP\) | 425](#)
- [J-Web | 426](#)
- [Platform and Infrastructure | 426](#)
- [VPNs | 426](#)

This section lists the known issues in hardware and software in Junos OS Release 17.4R3.

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 SRX340 and SRX345 devices, the Media Access Control Security (MACsec) on a physical port might not initialize correctly when a new node is joined to chassis cluster, which might lead the chassis cluster to be in a split-brain condition. [PR1396020](#)

Flow-Based and Packet-Based Processing

- On an SRX4600 device, the output of **show route forwarding-table** displays the next hop IP address twice if the next hop is st0 interface. The routing functionality is not impacted. [PR1290725](#)
- When both chassis cluster and PIM sparse mode are configured in SRX Series devices, multicast session leaks on secondary CP when PIM register messages are received in SRX Series devices. [PR1360373](#)
- On SRX Series devices, when the flow traceoptions with the packet filter are enabled, the traces of other sessions that are not configured in the packet filter might be captured in the logs. However, when the packet filters are removed, the traces are dumped in to the log file for some time less than 30 seconds. [PR1367124](#)
- In a multi-threaded environment, the service offload counters sometimes show incorrect values. It is a multi-threaded issue, and more than one thread can decrease the counter sometimes. [PR1381312](#)
- On SRX Series devices, in a chassis cluster with Z mode traffic and local (non-reth) interfaces are configured, when using ECMP routing between multiple interfaces residing on both node0 and node1, if a session is initiated through one node and the return traffic comes in through the other node, packets might be dropped due to reroute failure. [PR1410233](#)
- On an SRX4600 device, input and output bytes or bps statistic values might not be identical for the same size of packets. [PR1415117](#)
- On SRX Series devices, syslog severity level of “msg subtype is end of policy” is set to “error” although this message can be ignored. The severity level of this message is changed to “Info” as an ignorable message. [PR1435233](#)

Intrusion Detection and Prevention (IDP)

- On SRX Series devices in chassis cluster mode, IDP signature update might sometimes fail on one node because the AppID process gets stuck while unzipping of the downloaded signature file. [PR1336145](#)
- When IDP signature automatic update is scheduled, the secondary node might fail to upgrade the signature pack. [PR1358489](#)

J-Web

- On SRX Series devices, the root password configured at first J-Web access (Skip to J-Web feature) does not work if the password length is shorter than 8 characters. [PR1371353](#)

Platform and Infrastructure

- On SRX300, SRX320, SRX340, and SRX345 devices, if there is power outage too many times in a short period of time, the device might end up getting stuck in the loader prompt. This is resolved by upgrading the boot loader from Junos OS Release 15.1X49-D110 or later image using the following command **request system firmware upgrade re bios** and then rebooting the device. [PR1292962](#)
- On SRX1500, SRX4100, SRX4200, and vSRX platforms with **packet-capture** configured, packet capture does not work after you change, delete, or add the **maximum capture size**. [PR1304723](#)
- In Junos OS Release 17.4R1, the 1-Gigabit interface is not supported for an SRX4600 platform. [PR1315073](#)
- On SRX5600 and SRX5800 devices in chassis cluster, when a second Routing Engine is installed to enable dual control links, the **show chassis hardware** operational command might show the same serial number for both the second Routing Engines on both the nodes. [PR1321502](#)
- On SRX5400, SRX5600, and SRX5800 devices, the EM interface is an internal interface. If the EM interface is down, the control link is lost and the device cluster is in an abnormal status. [PR1342362](#)
- On an SRX1500 device, the activity LED (right LED) for the 1-Gigabit Ethernet and 10-Gigabit Ethernet ports (xe-0/0/16 through xe-0/0/19) does not light up when the interface is up and passing traffic correctly. [PR1380928](#)
- In a chassis cluster redundancy group failover scenario, on SRX5600 and SRX5800 devices, if the failover is caused by interface monitoring failure, the failover on the Packet Forwarding Engine side (that is data plane) might be slow (for example, impact on BFD session up to several seconds). This issue might result in protocol and traffic outage. [PR1385521](#)
- On SRX550M devices, when **encapsulation flexible-ethernet-services** is configured together with LACP protocol on aggregated Ethernet (AE) interfaces, the interface does not come up. [PR1448161](#)

VPNs

- IPsec uses ESP as the default protocol in IPsec proposal, if the administrator does not explicitly configure the protocol. However, this is not indicated as such in the schema for Security Director. [PR1061838](#)
- When an SRX Series device is an IPsec VPN initiator behind a NAT device, disabling NAT on the NAT device causes the next IKE negotiation to fail because UDP port 4500 is still in use. Use the CLI command **clear security ike security-associations** to recover and successfully establish a new IKE SA on UDP port 500. [PR1273213](#)

- If multiple traffic selectors are configured for a peer with Internet Key Exchange version 2 (IKEv2) reauthentication, only one traffic selector is rekeyed at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors will be cleared without immediate rekeying. A new negotiation of these traffic selectors is triggered through other mechanisms (for example, by traffic or by a peer). [PR1287168](#)
- On SRX Series devices, with NCP as client, sometimes IKE SA might not be displayed in the CLI output after RG1 failover. [PR1352457](#)
- VPN tunnels might flap when adding or deleting configuration group on SRX devices that are part of a Chassis Cluster. [PR1390831](#)
- On SRX Series devices, if IPsec VPN is configured with Network Address Translation-Transversal (NAT-T) and the size of an IP packet going into the tunnel is larger than 1400 bytes, IPsec Encapsulating Security Payload (ESP) packets might be received with fragments, indicating that post-fragmentation occurred. However, only pre-fragmentation is expected for the large packets, not post-fragmentation. [PR1424937](#)

SEE ALSO

[New and Changed Features | 405](#)

[Changes in Behavior and Syntax | 417](#)

[Known Behavior | 422](#)

[Resolved Issues | 427](#)

[Documentation Updates | 449](#)

[Migration, Upgrade, and Downgrade Instructions | 449](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.4R3 | 428](#)
- [Resolved Issues: 17.4R2 | 436](#)
- [Resolved Issues: 17.4R1 | 445](#)

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

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

Resolved Issues: 17.4R3

Application Layer Gateways (ALGs)

- The 5060 SIP active session will be affected after the status of SIP ALG is changed to disabled. [PR1373420](#)
- DNS requests with EDNS options might be dropped by DNS ALG. [PR1379433](#)
- The SUN RPC data traffic might be dropped after interface-related configuration is changed. [PR1387895](#)
- On SRX Series devices, SIP/FTP ALG does not work when SIP traffic with source NAT goes through the SRX Series device. [PR1398377](#)
- H.323 voice packets might be dropped on SRX devices. [PR1400630](#)
- The TCP reset packet is dropped when any TCP proxy-based feature and **rst-invalidate-session** are enabled simultaneously. [PR1430685](#)
- When H.323 packet passes the same SRX twice through different VRs instead of LSYS, the H.323 ALG Source NAT connection might fail but the Static NAT connections works fine. It also causes the H.323 connection not to be established successfully. [PR1436449](#)

Application Security

- Application firewall block message and redirect URL do not work for HTTPS websites. [PR1356483](#)
- Future group membership updates are not recognized by integrated user firewall after a user's sAMAccountName is changed while the distinguished name (DN) remains the same. [PR1394049](#)
- Packet loss might occur on unrelated traffic when AppQoS rate limiter is applied on SRX4600 and SRX5000 devices using SPC3. [PR1394085](#)
- Fail to match permit rule in application firewall ruleset. [PR1404161](#)
- Juniper Sky ATP does not escape the \ inside the username before the metadata is sent to the cloud. [PR1416093](#)
- The ipfd process might crash if security intelligence feature is configured. [PR1425366](#)
- Automatic application identification download stops after going over the year and rebooting. [PR1436265](#)
- The flowd or srpxfe process might crash when advanced anti-malware service is used. [PR1437270](#)
- Security logs cannot be sent to external syslog server through TCP. [PR1438834](#)

Chassis Clustering

- On SRX550 and SRX550M platforms, SFP-T based interfaces may fail to move to Link Up state after a chassis reboot. [PR1347874](#)
- Multiple flowd process files are seen on node1 after an RG0 failover. [PR1372761](#)

- Traffic loss occurs when the primary node is rebooting. [PR1372862](#)
- On SRX Series devices in chassis cluster, if reroute occurs on the IPv4 wings of a NAT64 or NAT46 session, the active node sends RTO message to the backup session to update the rerouted interface. [PR1379305](#)
- The packets might be dropped in an SRX Series devices in a chassis cluster environment if sampling or packet capture is configured. [PR1379734](#)
- The flowd process might stop if doing an ISSU upgrade. [PR1386522](#)
- On SRX4600 platform, in chassis cluster scenario, if configuring 4 100G interfaces on PIC 0, after reboot for the changed configuration to take effect, all the 4 interfaces might be in the 'down' state [PR1387701](#)
- The cluster IDs larger than 10 will cause FPCs to remain in offline on SRX4600 device in a chassis cluster. [PR1390202](#)
- GTPv2 Modify Bearer Request packets not containing F-TEID IE in bearer context are dropped during GTP inspection. [PR1399658](#)
- In SRX cluster environment, when a huge number of domain names are configured inside security policies and these policies are being used by some flow sessions, after RGO failover, traffic with domain name address might fail for 3-5 minutes. [PR1401925](#)
- The flowd process stops when updating or deleting a GTP tunnel. [PR1404317](#)
- The SRX Series devices might be potentially overwritten with an incorrect buffer address when detailed logging is configured under the GTPv2 profile. [PR1413718](#)
- PDP context response messages containing more than 6 Packet Data Network Gateway (PGW) connections will cause response packet to be dropped. [PR1422877](#)
- RGO failover sometimes causes FPC offline/present status. [PR1428312](#)

Class of Service (CoS)

- Configuring **host-outbound-traffic** under class of service might cause certain devices to stop. [PR1359767](#)

Flow-based and Packet-based Processing

- The half-duplex mode is not supported on SRX340 and SRX345 devices. [PR1149904](#)
- On all SRX platforms, the IPSec "replay error" might be seen on the IKE peer in cluster Z mode, because the IPSec traffic is encapsulated on the wrong node that leads to disorder of esp packet and the VPN peer gets replay error. As a result, the IPSec packet might be dropped on the peers. [PR1349724](#)
- IPsec VPN traffic loss is observed in an active-active HA mode. [PR1373161](#)
- False log message **/kernel: check_configured_tpid: : default tpid (0x8100) not configured. pic allows maximum of 0 tpid** is seen on SRX Series branch devices. [PR1373668](#)
- PIM register message might be dropped on SRX Series devices. [PR1378295](#)

- During SRX1500, SRX4100, SRX4200, SRX4600 and vSRX platforms reboot, users are not able to enter boot menu to select option to recover password. [PR1381653](#)
- On SRX1500 device, the IPv4 multicast packets might not be broadcast from the IRB interface. [PR1385934](#)
- Large file downloads slow down for many seconds. [PR1386122](#)
- Traffic might be processed by the VRRP backup when multiple VRRP groups are configured. [PR1386292](#)
- VDSL is not stable if there are sudden noises after configuring VDSL SOS feature. [PR1387133](#)
- Display issue in **show usp memory segment shm data module** and **show jsf shm module vty fwdd** commands on SRX Series branch devices. [PR1387711](#)
- On SRX4600 platform, when the same UDP source port is used for multiple flows, the traffic using the same source port might be stopped after session created. [PR1388735](#)
- The SRX Series devices do not send the messages **frag needed** and **DF set** back to the source host during path MTU discovery. [PR1389428](#)
- SRX Series devices might not strip VLAN added by native VLAN ID command. [PR1397443](#)
- SRX Series devices connection to JIMS keeps flapping causes failover to secondary JIMS. [PR1398140](#)
- On SRX4600 and SRX5000 line of devices, BGP packets might be dropped under high CPU usage. [PR1398407](#)
- VLAN push might not work on SRX1500. [PR1398877](#)
- Maximum feed number of IPFilter and GeoIP increased to 256. [PR1399314](#)
- The authd process might stop when issuing the **show network-access requests pending** command during the authd restart. [PR1401249](#)
- SRX Series devices cannot obtain IPv6 address through DHCPv6 when using PPPoE interface with logical unit number greater than zero. [PR1402066](#)
- Unable to access to SRX Series devices if messages **kern.maxfiles limit exceeded by uid 65534, please see tuning(7)** are seen. [PR1402242](#)
- Downloads might stall or completely fail when utilizing services that are reliant on TCP proxy. [PR1403412](#)
- Throughput or latency performance of TCP traffic drops when TCP traffic is passing through from one logical system to another logical system. [PR1403727](#)
- Split brain condition is experienced if the SPC2 or SPC3 card goes offline in the primary node. [PR1403872](#)
- The flowd process stops and all cards are brought offline. [PR1406210](#)
- The flowd process might stop if **enable-session-cache** option is configured under the SSL termination profile. [PR1407330](#)
- On SRX1500 devices, traffic is blocked on all interfaces after configuring **interface-mac-limit** on one interface. [PR1409018](#)

- On all SRX platforms, high memory utilization is observed if Advanced Anti-Malware service (AAMW) is enabled. This issue might use up all the memory and cause traffic slow. [PR1409606](#)
- Traffic might be dropped if SOF is enabled in a chassis cluster in active/active mode. [PR1415761](#)
- The command **show security firewall-authentication jims statistics** shows output statistics of both primary JIMS server and secondary JIMS server. [PR1415987](#)
- Traffic logging shows **service-name junos-dhcp-server** for UDP destination port 68. [PR1417423](#)
- Traffic might be lost on the SRX Series device if IPsec session affinity is configured with the **ipsec-performance-acceleration** command. [PR1418135](#)
- The 4G network connection might not be established if LTE mPIM card is in use. [PR1421418](#)
- The **show security flow session session-identifier < sessID>** is not working if the session ID is bigger than 10M on SRX4600 platform. [PR1423818](#)
- On the SRX300 series and SRX550 series platforms, when an unconnected port is added to the LAG, traffic over updated LAG might be assigned to the unconnected port causing packet loss. [PR1423989](#)
- SRX340 and SRX345 devices reporting high temperature alarms when operating within expected temperatures. [PR1425807](#)
- PIM neighbors might not come up on SRX Series devices in a chassis cluster. [PR1425884](#)
- When configuring GRE tunnel (GRE over IPsec tunnel) or IPsec tunnel on SRX Series devices, the MTU of the tunnel interface is calculated incorrectly. [PR1426607](#)
- The flowd process might stop on SRX5000 line of devices. [PR1430804](#)
- The second IPsec ESP tunnel might not be able to establish between two IPv6 IKE peers. [PR1435687](#)
- The ipfd process might stop when Security Intelligence (SecIntel) is used. [PR1436455](#)
- SRX4600 devices may encounter a SPMC version mismatch error causing SPM to remain in Present/Offline state after using USB Install method. [PR1437065](#)
- On SRX1500 devices, the member of dynamically created VLANs information is not displaying on show VLAN. [PR1438153](#)
- The flowd process stops on SRX550 or SRX300 platforms, when SFP module is plugged in. [PR1440194](#)
- SRX Series branch devices with RPM **probe-server hardware timestamp** configured does not respond with correct timestamp to the RPM client. [PR1441743](#)

Interfaces and Chassis

- The virtual IP of the VRRP on an SRX4600 device might not respond to host-inbound traffic. [PR1371516](#)
- SRX Series: Crafted packets destined to fxp0 management interface on SRX340 and SRX345 devices can lead to DoS (CVE-2019-0038). [PR1377152](#)
- The pkid process might stop after RGO failover. [PR1379348](#)
- SRX4600 10-Gigabit Ethernet interface optics diagnostic access issue. [PR1395806](#)

- On SRX4600 platform, the 40-Gigabit Ethernet interface might flap continuously due to MAC local fault. [PR1397012](#)
- Midplane FRU model number is not displayed. [PR1422185](#)

Interfaces and Routing

- Incorrect (double) ingress pps rate on MPLS-enabled interface. [PR1328161](#)
- Control traffic loss might be seen on SRX4600 platform. [PR1357591](#)
- Switching interface mode between family ethernet-switching and family inet/inet6 might cause traffic loss. [PR1394850](#)

Intrusion Detection and Prevention (IDP)

- Unable to deploy IDP because the IDP configuration cannot be committed. [PR1374079](#)
- IDP might crash with the custom IDP signature. [PR1390205](#)

J-Web

- The chassis cluster image is not displayed on the J-Web dashboard. [PR1382219](#)
- On the SRX300 and SRX4000 Series devices, the J-Web shows that the CPU is overheating. [PR1389981](#)
- The next-hop IP address is not displayed in the routing table in J-Web. [PR1398650](#)
- On all SRX Series devices, the special character without a quotation mark used in preshared key (**security ike policy**) is removed silently after a commit operation on J-Web. This will cause the VPN connection to be down due to the preshared key mismatch between the peers. [PR1399363](#)
- Configuring using the CLI editor in J-Web generates an mgd core dump and commit does not work. [PR1404946](#)
- The httpd-gk process stops, leading to dynamic VPN failures and high Routing Engine CPU utilization (100 percent). [PR1414642](#)
- J-Web configuration change for an address set using the search function results in a commit error. [PR1426321](#)
- J-Web shows incorrect port-mode under **Configure>Interfaces>Link Aggregation**. [PR1430414](#)
- IRB interface is not available in zone option of J-Web. [PR1431428](#)
- In SRX chassis cluster setup, node1 is not available for monitoring using J-Web. [PR1443819](#)

Layer 2 Ethernet Services

- DHCPv6 clients might fail to get addresses on SRX Series platforms. [PR1392723](#)
- IPv6 default route might not be installed from the received **Router Advertisement** message. [PR1411921](#)

Multiprotocol Label Switching (MPLS)

- The rpd might restart unexpectedly when **no-cspf** is configured and lo0 is not included under RSVP protocol. [PR1366575](#)

Network Address Translation (NAT)

- The SRX Series devices might send the **noSuchInstance** value to SNMP server in get response during commit. [PR1357840](#)
- NAT64 and traceroute do not work correctly on SRX Series devices. [PR1376890](#)
- The nsd process stops and causes the Web filter to stop working. [PR1406248](#)
- On SRX Series devices, when using NAT64 translation, RTSP uses a wrong string to rewrite the message payload, which might result in the message being dropped in a remote device. [PR1443222](#)

Network Management and Monitoring

- The **set system no-redirects** setting does not take effect for the reth interface. [PR894194](#)
- On SRX Series devices, after the AGENTX session timeout between master (snmpd) and subagent, the subagent might crash and restart. [PR1396967](#)
- MIB OID **dot3StatsDuplexStatus** shows wrong status. [PR1409979](#)
- SNMPD might generate core files after restarting NSD process through **restart network-security gracefully**. [PR1443675](#)

Platform and Infrastructure

- Junos OS: Login credentials are vulnerable to brute force attacks through the REST API (CVE-2019-0039). [PR1289313](#)
- High httpd utilization after reboot failover. [PR1352133](#)
- SRX4200 devices using **show chassis** commands may not display any outputs. [PR1363645](#)
- SRX1500 continues to raise alarm on fan **Fan Tray 0 Fan 0 Spinning Degraded**. [PR1367334](#)
- Packet capture feature does not work after removing the sampling configuration. [PR1370779](#)
- IP monitoring failure results in multiple interfaces disappearing from the forwarding table. [PR1371500](#)
- Slowness in cold sync when there are many Packet Forwarding Engines installed in the SRX Series devices in a chassis cluster. [PR1376172](#)
- Useridd CPU is higher than 100 percent for more than 1 hour. [PR1377684](#)
- On SRX5000 or SRX3000 platforms, some uspic failed messages might be seen while running **show interface extensive** command from CLI or Junos Space. [PR1380439](#)
- On SRX Series devices, when doing an ISSU upgrade, the reth interface might flap and cause traffic loss in rare occasions. [PR1381475](#)
- Traffic loss seen in Layer 2 VPN with GRE tunnel. [PR1381740](#)

- Junos upgrade might fail with validate option after the `/cf/var/sw` directory is accidentally deleted. [PR1384319](#)
- Login class with allowed days and specific access-start and access-end does not work as expected. [PR1389633](#)
- Memory leak might occur on the data plane during composite next-hop installation failure. [PR1391074](#)
- The flowd process stops if it goes into a dead loop. [PR1403276](#)
- Complete device outage might be seen when an SPU VM core file is generated. [PR1417252](#)
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10(based system). [PR1417321](#)
- On SRX Series devices, flowd process stops might be seen. [PR1417658](#)
- Routing Engine CPU utilization is high and eventd process is consuming a lot of resources. [PR1418444](#)
- REST API does not work on SRX550HM. [PR1430187](#)
- On SRX Series devices with any licensed feature enabled, a false license alarm might be generated. This issue is not feature or license specific; it is random and can happen for any licensed feature. This issue only generates a false alarm, and has no functionality impact. [PR1431609](#)
- On SRX platforms with lots of IPsec VPN tunnels configured (for example - 6700 IPsec VPN tunnels configured on SRX5400), after system bootup , the kmd and iked processes repeatedly generate `ipc_pipe_write:353 num_sent=-1 errno=35 Resource temporarily unavailable` with IPsec VPN tunnels temporarily being down. [PR1434137](#)
- On SRX4100 and SRX4200 devices, when LACP is configured on the reth interface, the interface flaps when Routing Engine is busy. [PR1435955](#)

Routing Policy and Firewall Filters

- The timeout value of `junos-http` is incorrect. [PR1371041](#)
- Application firewall action is not correct in the output of CLI `show security policies application-firewall`. [PR1378993](#)
- The nsd process stops and generates a core file. [PR1388719](#)
- Memory leak in nsd prevents change from taking effect. [PR1414319](#)
- The flowd process (responsible for traffic forwarding in SRX Series devices) stops on SRX Series devices while deleting a lot of policies from Junos Space. [PR1419704](#)
- If restarting NSD fails, there is no indication or symptom, and users are not notified. So a new alarm is added to indicate this failure. [PR1422738](#)
- DNS cache entry does not time out from device even after TTL=0. [PR1426186](#)
- Packet Forwarding Engine stops might be seen on SRX1500 platform. [PR1431380](#)

- SRX550M running Junos OS Release 18.4R1 shows PEM 1 output failure message, whereas with Junos OS Release 15.1X49 or Junos OS Release 18.1R3.3 it does not show any alarms. [PR1433577](#)
- SRX1500 device only allows a maximum of 256 policies with counting enabled. [PR1435231](#)

Unified Threat Management (UTM)

- EWF server status shows UP when 443 is specified as server port. [PR1383695](#)
- Whitelist/blacklist does not work for HTTPS traffic going through Web proxy. [PR1401996](#)
- SRX Series: srpxfe process crash while JSF/UTM module parses specific HTTP packets (CVE-2019-0052). [PR1406403](#)
- The device might not look up the blacklist first in a local Web filtering environment. [PR1417330](#)

VPNs

- The kmd process might stop when configuring IPsec VPN and BGP on SRX1500 platform. [PR1336235](#)
- Dot (.) usage in CA profile name causes issues when the pkid process is restarted. [PR1351727](#)
- Tunnel flap is seen after doing RG0 failover. [PR1357402](#)
- IPsec tunnel might flap when there are concurrent IKEv2 Phase 1 SA rekeys. [PR1360968](#)
- The IPsec traffic might be blocked by SRX5000 line of devices if they are acting as IPsec transit devices. [PR1372232](#)
- In a rare situation, VPN tunnels might not be configured successfully and the VPN tunnels might not come up. [PR1376134](#)
- Packet loss is seen in IPsec Z-mode scenario. [PR1377266](#)
- The kmd process might stop and cause VPN traffic outage after executing **show security ipsec next-hop-tunnels**. [PR1381868](#)
- Adding or deleting site-to-site manual NHTB VPN tunnels to an existing st0 unit causes existing manual NHTB VPN tunnels under the same st0 unit to flap. [PR1382694](#)
- A few VPN tunnels do not forward traffic after RG1 failover. [PR1394427](#)
- The kmd process might stop when SNMP polls for the IKE SA. [PR1397897](#)
- VPN tunnels flap after adding or deleting a configuration group in edit private mode on a clustered setup. [PR1400712](#)
- The kmd process stops and generates a core file after running the **show security ipsec traffic-selector** command. [PR1428029](#)

Resolved Issues: 17.4R2

Application Layer Gateways (ALGs)

- On SRX1400 device, the NFS traffic to port 2049 might drop. [PR1307763](#)
- The configure download URL displays warning message **requires appid-sig license**. [PR1324858](#)
- On SRX Series devices with SIP ALG enabled, the SIP ALG might drop SIP packets which have a **referred-by** or **referred-to header** field containing multiple header parameters. [PR1328266](#)
- SIP calls drop, when the limit per SPU crosses 10,000 calls. [PR1337549](#)

Authentication and Access Control

- On SRX Series devices, PFE might crash and huge number of core files might be generated within a short period of time. [PR1326677](#)
- On SRX Series devices, incomplete Request Support Information (RSI) might be seen. [PR1329967](#)
- On SRX Series devices, the sessions might close because of the idle **Timeout junos-fwauth-adapter** logs. [PR1330926](#)
- The uacd process is unstable after upgrading to Junos OS Release 12.3X48 and later releases. [PR1336356](#)
- On SRX Series devices, the **show version detail** command returns an error message: **Unrecognized command (user-ad-authentication)** while configuring the useridd settings. [PR1337740](#)
- A new configuration is available to configure the web-authentication timeout. [PR1339627](#)

Chassis Clustering

- The route information might not be synchronized between node0 and node1 when configuring the firewall filter or APBR to use the non-default routing-instance. [PR1292235](#)
- Flowd process core files are generated after adding 65536 VPN tunnels using traffic selector with the same remote IP. [PR1301928](#)
- On devices enabled with chassis cluster, the ISSU upgrade might fail and display an error message **ISSU aborted and exiting ISSU window**. [PR1306194](#)
- On SRX1500, SRX4100 and SRX4200 devices, ISSU might fail if LACP and interface monitoring are configured. [PR1305471](#)
- File Descriptor might leak on SRX Series chassis clusters with Sky ATP enabled. [PR1306218](#)
- When services offloading feature is enabled, the device changes TCP checksum value to 0x0000. [PR1317650](#)
- When ISSU is performed from a Junos OS Release prior to 15.1X49-D60 to a Junos OS Release 15.1X49-D60 or later, flowd process generates core files. [PR1320030](#)
- The device might stop forwarding traffic after RG1 failover from node0 to node1. [PR1323024](#)

- When RGO failover or primary node reboot happens, some of the logical interfaces might not be synchronized to the other node if the system has around 2,000 logical interfaces and 40,000 security policies. [PR1331070](#)
- After the primary node or the secondary node restarts, the FPC module goes offline on the secondary node. [PR1340116](#)
- In and active/active cluster, route change timeout does not work as expected. [PR1314162](#)

Class of Service (CoS)

- Packets go out of order on SPC2 cards with IOC1 or FIOC cards. [PR1339551](#)

Flow-Based and Packet-Based Processing

General Routing

- SRX1500 devices might power off unexpectedly because of incorrect device temperature readings which reported a too high temperature, leading to an immediate pro-active power-off of the device to protect the device from overheating. However in these cases the temperature was not actually too high and a power-off would not be required. When this occurs, the following log message is shown in file `/var/log/hostlogs/lcmd.log`: Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor. [PR1241061](#)
- On SRX4100 and SRX4200 devices, packet loss is observed when the value of packet per second (pps) through the device is very high. This occurs because of the update of the **application interval statistics** statement, which has a default timer value of 1 minute. You can avoid this issue by setting the interval to maximum using the **set services application-identification statistics interval 1440** command. [PR1290945](#)
- The **show host server name-server host** CLI command fails when the source address is specified under the name-server configuration. [PR1307128](#)
- A memory leak might occur in the appidd process while updating an application signature package. [PR1308863](#)
- On SRX4600 devices, when you run the **clear security flow session** command, time taken to clear the session depends on the total session number. For example, the clear session takes nine minutes to clear 57M session. [PR1308901](#)
- On SRX Series devices, if destination NAT and session affinity are configured with multiple traffic selectors in IPsec VPN, the traffic selector match might fail. [PR1309565](#)
- The flowd process might stop and generate a core file during failover between node 0 and node 1. [PR1311412](#)
- On SRX Series devices, the IPsec tunnel might fail to be established if datapath debug configuration include the options **preserve-trace-order**, **record-pic-history**, or both. [PR1311454](#)
- The SRX Series device drops packets citing the reason "Drop pak on auth policy, not authed". [PR1312676](#)
- The flowd process might stop if the SSL-FP profile is configured with whitelist. [PR1313451](#)

- If IDP and SSL forward proxy whitelist are configured together, the device might generate a core file. [PR1314282](#)
- On SRX550M devices, phone-home.core is generated after the zeroization procedure. [PR1315367](#)
- If the Sky ATP cloud feed updates, the Packet Forwarding Engine might stop causing intermittent traffic loss. [PR1315642](#)
- On SRX Series devices, the IPsec VPN tunnel with traffic-selector is configured and the packets TTL is set to 1, the flowd process stops and generates a core file on both the nodes. [PR1316134](#)
- Periodic PIM register loop is observed during switch failure. [PR1316428](#)
- On SRX Series devices, the **fin-invalidate-session** command does not work when the Express Path feature is enabled on the device. [PR1316833](#)
- Return traffic through the routing instance might drop intermittently after changing the zone and routing-instance configuration on the st0.x interface. [PR1316839](#)
- SRX300 devices DHCP client cannot obtain IP addresses. [PR1317197](#)
- Default route is lost after system zero. [PR1317630](#)
- SSL firewall proxy does not work if root-ca has fewer than four characters. [PR1319755](#)
- The OSPF peers are unable to establish neighbors between the LT interfaces of the logical systems. [PR1319859](#)
- On SRX Series devices, after logical system is configured, about 10 logical systems are not working. [PR1323839](#)
- The flowd process generates core files on both nodes causing an outage. [PR1324476](#)
- The MPC cards might drop traffic in the event of high temperatures. [PR1325271](#)
- Software next-hop table is full with log messages RT_PFE: NH IPC op 1 (ADD NEXTHOP) failed, err 6 (No Memory) peer_class 0, peer_index 0 peer_type 10. [PR1326475](#)
- If the serial number of the certificate for the SSL proxy has two consecutive zeros, the certificate authentication fails. [PR1328253](#)
- When you use CFLOW, the source address for flow packets is not displayed. [PR1328565](#)
- On SRX Series devices, the one-way jitter traps are not generated when the TWAMP is configured. [PR1328708](#)
- The FPC is dropped or hangs in the present state when the intermittent control link heartbeat is observed. [PR1329745](#)
- On SRX Series devices with stream logging configured, high CPU load is observed. [PR1331011](#)
- The IPv6 traffic does not work as expected on IOC3 with the services offloading (npcache) feature. [PR1331401](#)
- NTP synchronization fails and switches to a local clock. [PR1331444](#)

- Inaccurate Jflow records might be seen for output interface and next hop. [PR1332666](#)
- The whitelist function in syn-flood does not work. [PR1332902](#)
- The **show vlans detail no-forwarding** command in the RSI does not display any information, because the **no-forwarding** option is not supported. [PR1336267](#)
- Two-way active measurement protocol (TWAMP) client, when configured in a routing instance, does not work after a reboot. [PR1336647](#)
- On the front panel LED, the red alarm goes on after an RG0 failover is triggered when the flowd process stops. [PR1338396](#)
- The unfiltered traffic is captured after traceoptions are deactivated. [PR1339213](#)
- SSH to the loopback interface of SRX Series devices does not work properly when AppTrack is configured. [PR1343736](#)
- The flowd process might stop when SYN-proxy function is used. [PR1343920](#)
- SNMP MIB walk provides incorrect data counters for total current flow sessions. [PR1344352](#)
- SRX1500 devices might encounter a failure while accessing the SSD drive. [PR1345275](#)
- On SRX Series devices, when you upgrade to a Junos OS Release with "no-validate" option and if there are unsupported configurations with the new version, then configuration push fails and the ksyncd process stops. [PR1345397](#)
- The REST API is not working on the SRX320-POE device. [PR1347539](#)
- File download stops over a period of time when TCP proxy is activated through Antivirus or Sky ATP. [PR1349351](#)
- When a J-Flow related configuration is deleted, the forwarding plane begins to drop packets. [PR1351102](#)
- If the Trusted Platform Module (TPM) is enabled, the configuration integrity failure occurs when there is a power loss for few seconds after the commit. [PR1351256](#)
- On SRX1500 device, after the SSL forward proxy is configured, the system stops and generates a core file. [PR1352171](#)
- The flowd process generates a core file when the SIP ALG is enabled. [PR1352416](#)
- When the routing instance is configured, the **UTM Anti-Spam:DUT** process do not send the DNS query. [PR1352906](#)
- On SRX Series devices, if the memory buffer is accessed without checking the mbuf and the associated external storage, the flowd process might stop. [PR1353184](#)
- On SRX Series devices in a chassis cluster, if an IPv6 session is being closed and at the same time the related data-plane Redundancy Group (RG1+) failover occurs, this IPv6 session on the backup node might hang and cannot be cleared. [PR1354448](#)
- The PIM register might stop the message from the source First Hop Router (FHR). [PR1356241](#)

- On SRX300, SRX320, SRX340, and SRX345 devices, with LTE mini-PIM the DHCP relay packets are not forwarded. [PR1357137](#)
- On SRX5000 series devices, when the IPsec performance acceleration feature is enabled, packets going in or out of a VPN tunnel are dropped. [PR1357616](#)
- On SRX5400, SRX5600, and SRX5800 devices, the MIB walk tool is not working when screens are applied to the security zones. [PR1364210](#)

Interfaces and Chassis

- Unable to add IRB and aggregated Ethernet interfaces. [PR1310791](#)
- On SRX1500 devices, pp0.0 interface link status is not up. [PR1315416](#)
- An error is not seen at each commit or commit check if autonegotiation is disabled but the speed and duplex configurations are not configured on the interface. [PR1316965](#)
- If an interface is configured with the Ethernet switching family, we recommend that you do not configure **vlan-tagging**. [PR1317021](#)
- The interface might be brought down by IP monitoring at the time of committing a configuration because of incorrect interface status computing. [PR1328363](#)

Interfaces and Routing

- JIMS server stops responding to requests from SRX Series devices. [PR1311446](#)
- On SRX Series devices in a chassis cluster, the IRB interface does not send an ARP request after clearing the ARP entries. [PR1338445](#)
- Packet reorder occurs on the traffic received on the PPP interface. [PR1340417](#)
- On SRX Series devices, when the VPLS interface receives a broadcast frame, the device sends this frame back to the sender. [PR1350857](#)
- On the SRX1500, when the LACP is configured with interfaces ae0 and ae1, the mac address is displayed as 00:00:00:00:00:00 and 00:00:00:00:00:01 for interfaces ae0 and ae1 respectively. [PR1352908](#)
- The **set protocols rstp interface all** command does not enable RSTP on all interfaces. [PR1355586](#)

Intrusion Detection and Prevention (IDP)

- The control plane CPU usage is high when using IDP. [PR1283379](#)
- IDP signatures might not get pushed to the Packet Forwarding Engine if there is a policy in logical systems. [PR1298530](#)
- The IDP PCAP feature has been improved. [PR1297876](#)
- The output of **show security idp status** command does not accurately reflect the number of decrypted SSL or TLS sessions being inspected by the IDP. [PR1304666](#)
- The file descriptor might leak during a security package auto update. [PR1318727](#)

- On SRX4600 devices, the maximum SSLRP session count is observed to be approaching 100,000. In the CLI, configuring a maximum of 100,000 sessions are allowed, whereas in SSLFP, 600,000 sessions are allowed. Thus, the **set security idp sensor-configuration ssl-inspection sessions** command is now modified to allow a maximum of 600,000 sessions. However, for other devices the original session limit value of 100,000 is retained. [PR1329827](#)
- Loading IDP policy fails because of less available heap memory. [PR1347821](#)

J-Web

- J-Web system snapshot throws error. [PR1204587](#)
- In J-Web when you click the SKIP TO JWEB OPTIONS, the Google Chrome browser automatically redirects. [PR1284341](#)
- J-Web does not display all global address book entries. [PR1302307](#)
- On SRX300, SRX320, SRX340, and SRX345 devices, CPU usage is high when generating on-box reporting on the J-Web. [PR1310288](#)
- J-Web authentication fails when a password includes the backslash. [PR1316915](#)
- J-Web dashboard displays wrong last updated time. [PR1318006](#)
- J-Web display problems for security policies are observed. [PR1318118](#)
- J-Web displays the red alarm for temperature value within the threshold. [PR1318821](#)
- J-Web does not display wizards on the dashboard. [PR1330283](#)
- Unable to delete the dynamic VPN user configuration. [PR1348705](#)
- When the J-Web fails to get resource information, the Routing Engine CPU usage is displayed as 100 percent. [PR1351416](#)
- Security policies search button on the J-Web does not work with Internet Explorer version 11. [PR1352910](#)

Layer 2 Ethernet Services

- In DHCP relay configuration, the option **VPN** has been renamed to **source-ip-change**. [PR1318487](#)
- On SRX1500 devices, VLAN popping and pushing does not work over Layer 2 circuits. [PR1324893](#)
- DHCP rebind and renew packets is not calculated in BOOTREQUEST. [PR1325872](#)
- The default gateway route might be lost after the failover of RG0 in a chassis cluster. [PR1334016](#)
- The subnet mask address is not sent as a reply to the **DHCPINFORM** request. [PR1357291](#)

Network Address Translation (NAT)

- The default-gateway route received by DHCP when some interface in the chassis cluster has been configured as a DHCP client is lost in about 3 minutes after RG0 failover. [PR1321480](#)
- On SRX Series devices, the Sky ATP connection leak causes the service plane to be disconnected from the Sky ATP cloud. [PR1329238](#)

- Arena utilization on a FPC spikes and then resumes to a normal value. [PR1336228](#)

Network Management and Monitoring

- SRX300 device is unresponsive as a result of cf/var: filesystem full error. [PR1289489](#)
- CLI options are available to manage the packet forwarding engine handling the ARP throttling for NHDB resolutions. [PR1302384](#)

Platform and Infrastructure

- When you perform commits with apply-groups, VPN might flap. [PR1242757](#)
- The packet captured by datapath-debug on an IOC2 card might be truncated. [PR1300351](#)
- Inconsistent flow-control status on reth interface is observed. [PR1302293](#)
- On SRX5400, SRX5600, and SRX5800 devices, DC PEM is used on the box, the output of **show chassis environment pem** and **show chassis power** commands do not show DC input value correctly. [PR1323256](#)
- On SRX5400, SRX5600, and SRX5800 devices, SPC2 XLP stops processing packets in the ingress direction after repeated RSI collections. [PR1326584](#)
- When SecIntel is configured, IPFD CPU utilization might be higher than expected. [PR1326644](#)
- The log messages file contains **node*.fpc*.pic* Status:1000 from if_np for ifl_copnfig op:2 for ifl :104** message. [PR1333380](#)
- Log message **No Port is enabled for FPC# on node0** is generated every 5 seconds. [PR1335486](#)
- In RSI, a mandatory argument is missing for the **request pfe execute** and the **show usp policy counters** commands. [PR1341042](#)
- On SRX Series devices in a chassis cluster, configuration commit might succeed even though the external logical interface configuration (reth) associated with the Internet Key Exchange (IKE) VPN gateway configuration is deleted. This might lead to configuration load failure during the next device boot-up. [PR1352559](#)
- On SRX4100 devices, interfaces are shown as half-duplex, but there is no impact on the traffic. [PR1358066](#)

Routing Policy and Firewall Filters

- The firewall authentication does not list the correct polices when the NSD process is busy. [PR1312697](#)
- The number of address objects per policy for SRX5400, SRX5600, SRX5800 devices is increased from 4,096 to 16,000. [PR1315625](#)
- The flowd process stops when AppQoS is configured on the device. [PR1319051](#)
- Flowd process stops after configuring a huge number of custom applications. [PR1347822](#)

- On SRX Series devices, with a large number of firewall authentication entries, the flowd process might stop. [PR1349191](#)
- On SRX Series devices, a large scale commit, for example, 70,000 lines security policy might stop the NSD process on the Packet Forwarding Engine (PFE). [PR1354576](#)

Routing Protocols

- On SRX1500 devices, the IS-IS adjacency remains down when using an IRB interface. [PR1300743](#)
- Dedicated BFD does not work on SRX Series devices. [PR1312298](#)
- On a chassis with BMP configured, if the rpd termination timeout is happening while the BMP main task has failed to terminate and delete itself (seen when rpd is gracefully terminated), the rpd might stop. [PR1315798](#)
- When BGP traceoptions are configured and enabled, the traces specific to messages sent to the BGP peer (BGP SEND traces)are not logged The traces specific to received messages (BGP RECV traces) are logged correctly. [PR1318830](#)
- OpenSSL Security Advisory [07 Dec 2017]. Refer to <https://kb.juniper.net/JSA10851> for more information. [PR1328891](#)
- The ppmmd process might stop, after one node is upgraded and failover completes. [PR1347277](#)
- On SRX Series devices, dedicated BFD does not work. [PR1347662](#)

Software Installation and Upgrade

- The **request system reboot node in/at** command results in an immediate reboot instead of rebooting at the allotted time. [PR1303686](#)
- On SRX1500 devices, the fan speed often fluctuates. [PR1335523](#)

System Logs

- A warning syslog message is displayed when the number of security screens installed exceed the IOC capacity. [PR1209565](#)
- The following log messages are displayed on the device: **L2ALM Trying peer/master connection, status 26.** [PR1317011](#)

User Firewall and Authentication

- User firewall has a command to fetch the user-group mapping from the active directory server. [PR1327633](#)

Unified Threat Management (UTM)

- The ISSU upgrade might fail because of the Packet Forwarding Engine generating a core file. [PR1328665](#)

Upgrade and Downgrade

- The command **show system firmware** displays the old firmware image. [PR1345314](#)

VLAN Infrastructure

- On SRX Series devices in transparent mode, the flowd process might stop when matching the destination MAC. [PR1355381](#)

VPNs

- The IRB interface does not support VPN. [PR1166714](#)
- Next hop tunnel binding (NHTB) is not installed occasionally during rekey for VPN using IKEv1. [PR1281833](#)
- IPsec traffic statistic counters return 32-bit values. [PR1301688](#)
- Auto Discovery VPN (ADVPN) tunnels might flap with the spoke error no response ready yet, leading to IKEv2 timeout. [PR1305451](#)
- On SRX Series devices, core files are observed under certain conditions with VPN and when NAT-T is enabled. [PR1308072](#)
- PKID syslog for key-pair deletion is required for conformance. [PR1308364](#)
- On SRX Series devices, ESP packet drops in IPsec VPN tunnels with NULL encryption algorithm configuration are observed. [PR1329368](#)
- SNMP for jnxIpSecTunMonVpnName does not work. [PR1330365](#)
- The kmd process might generate a core file when all the VPNs are down. [PR1336368](#)
- On SRX5400, SRX5600, and SRX5800 devices, the chassis cluster control link encryption does not work. [PR1347380](#)
- The kmd process might stop if multiple IKE gateways uses the same IKE policy. [PR1337903](#)
- All IPsec tunnels are in both active and inactive state. [PR1348767](#)
- S2S tunnels are not redistributed after IKE or IPsec are reactivated in a configuration. [PR1354440](#)

Resolved Issues: 17.4R1

Application Layer Gateways (ALGs)

- On SRX Series devices SIP packet might drop when SIP traffic performs destination NAT. [PR1268767](#)
- The pfed process stops and generates core files. [PR1292992](#)
- H323 ALG decode Q931 packet error was observed even after disabling H323 ALG. [PR1305598](#)
- HTTP ALG is listed within **show security match-policies**, when the HTTP ALG does not exist. [PR1308717](#)

Chassis Cluster

- Node 0 is going into db prompt after applying Layer 2 switching configuration and rebooting. [PR1228473](#)
- HA configuration synchronization monitoring does not work if **encrypt-configuration-files** is enabled. [PR1235628](#)
- The ISSU or ICU operation might fail if upgrade is initiated from Junos Space on multiple SRX clusters. [PR1279916](#)
- ALG traffic and other traffic with tcp-proxy gets stuck after back-to-back RG1 failover when using PPPoE on the reth interface. [PR1286547](#)
- Warning messages are incorrectly tagged as errors in the RPC response from the SRX Series device when you configure a change through NETCONF. [PR1286903](#)
- After software upgrade, the cluster goes into a brief split-brain state when rebooting RG0 on the secondary node. [PR1288819](#)
- In an SRX1500 cluster, if control-link-recovery is configured, ISSU might not complete successfully and the cluster will end up with different software releases. [PR1303948](#)
- IP monitoring on the secondary node shows unknown status after rebooting. [PR1307749](#)
- On SRX Series devices, the traffic logging impact issue after ISSU is fixed. [PR1284783](#)

Class of Service (CoS)

- on SRX devices, self-generated TCP session from RE destined to an lt-0/0/0.x nexthop is not established. [PR1286866](#)

Flow-Based and Packet-Based Processing

- The software-NH value increases and and causes a traffic outage. [PR1190301](#)
- SRX1500 devices might power-off unexpectedly because of incorrect device temperature readings which reportedly is a too high temperature, leading to an immediate proactive power-off of the device to protect the device from overheating. When this condition occurs, the following log message is shown in file `/var/log/hostlogs/lcmd.log`: `Jan 25 13:09:44 localhost lcmd[3561]: srx_shutdown:214: called with FRU TmpSensor`. [PR1241061](#)
- Duplicate hops or a higher than expected hop count is seen in L2 traceroute. [PR1243213](#)

- Configuring dpd results in timeouts for TCP encapsulation sessions. [PR1254875](#)
- A down interface in the **mirror-filter** command might cause a core file in certain situations. [PR1270724](#)
- Core files are seen on SRX1500 when J-Flow is enabled. [PR1271466](#)
- SRX320 with MPIM: IPv6 static route on dl0.0 is not active, so it cannot work for dial-on-demand. [PR1273532](#)
- Multicast traffic sent to the downstream interface in the destination MAC address is set to all zeros. [PR1276043](#)
- Output hangs while checking pki ca-certificate ca-profile-group details. [PR1276619](#)
- SRX1500 randomly stops forwarding traffic. [PR1277435](#)
- When using integrated user firewall, the useridd process might consume high CPU. [PR1280783](#)
- When executing operational commands for creating rescue configuration, some errors will be reported but the rescue configuration will still be created. [PR1280976](#)
- User firewall users are not assigned their roles. [PR1282744](#)
- Certain SCTP packets are dropped. [PR1285089](#)
- The pfed process stop and core files are generated by committing traceoptions configure. [PR1289972](#)
- More CPU threshold warnings are seen than in the previous releases. [PR1291506](#)
- CoS scheduler and shaping does not work on IRB interface. [PR1292187](#)
- Cryptographic weakness is seen on SRX300 line devices TPM Firmware (CVE-2017-10606) [PR1293114](#)
- The APN profile password is displayed in cleartext. [PR1295274](#)
- On SRX Series devices running the user firewall feature, under some conditions, flowd or useridd might generate core files. The Packet Forwarding Engine might get restarted, and RG1+ failover occurs. [PR1299494](#)
- SRX Series device fail to upgrade the Junos image when you use the unlink and partition options at the same time. [PR1299859](#)
- When you run the **show interfaces queue rethx** command, the output displays ingress queue information. [PR1309226](#)
- On SRX Series devices, the Stream Control Transmission Protocol (SCTP) packet has an incorrect SCTP checksum after the payload is translated by the device. [PR1310141](#)

Interfaces and Chassis

- On SRX1500 devices with SFP+-10G-CU3M DAC, 10-Gigabit Ethernet interface does not work. [PR1246725](#)
- On SRX1500, 10-Gigabit Ethernet interface might not come up between the SRX Series device and another type of device when using SFP+-10G-CU3M DAC. [PR1279182](#)

- Ping to VRRP (VIP) address failed when VRRP on vlan-tagging. This only affected IOC2 and IOC3 cards in SRX5000 line devices. SRX1500, SRX4100, and SRX4200 devices are not impacted. [PR1293808](#)
- RPM packets do not go through the LT interface under certain configurations. [PR1303445](#)

J-Web

- SRX Series devices cannot be upgraded with Junos image using J-Web. [PR1297362](#)
- Configuration upload using J-Web does not work. [PR1300766](#)
- In J-Web, when logical system adds a custom application, the applications 'any' are not present in **Logical System Configure > Security > Security Policy > Add Policy**. [PR1303260](#)
- J-Web removes the backslash character on the source identity object when the commit changes. [PR1304608](#)

Layer 2 Ethernet Services

- ARP issues are seen when using Layer 2 switching with the IRB interface. [PR1266450](#)
- On SRX1500 devices in an Ethernet switching mode, an IRB interface located in a custom routing instance is not reachable. [PR1234000](#)
- The **change no-dns-propagation** command should be changed to **no-dns-install**. [PR1284852](#)
- DHCPv6 prefix delegation does not start with the first available subnet [PR1295178](#)

Network Address Translation (NAT)

- On SRX Series devices, the periodic execution of the **show security zones detail** command causes the NSD process to fail in releasing unused memory, causing memory leak. [PR1269525](#)
- The proxy-arp does not work intermittently after RGO failover. [PR1289614](#)
- Commit check might allow a Source NAT pool without addresses to be committed, leading to flowd core file generation when the misconfigured pool is utilized by traffic. [PR1300019](#)
- Active source NAT causes an NSD error and the session closes. [PR1313144](#)

Network Management and Monitoring

- On the SRX340 device, one Routing Engine does not reply for the SNMP request after power-on or RGO failover in a cluster. [PR1240178](#)
- On SRX Series devices, when J-Flow is enabled for multicast traffic **extern nexthop** is installed during the multicast composite next hop. However, when you uninstall the composite next hop, it does not free the **extern nexthop**, which results in the jtree memory leak. [PR1276133](#)
- The mib2d process might crash when polling the OID ifStackStatus.0 after a logical interface of lo0 is deleted. [PR1286351](#)
- The **show arp no-resolve interface X** command for nonexistent interface X is showing all unrelated static ARP entries. [PR1299619](#)

Platform and Infrastructure

- SRX300 line devices reboot when Juniper RE-USB-4G-S (yellow or orange) USB is inserted. [PR1214125](#)
- The flowd process might crash during route update. [PR1249254](#)
- Unexpected behavior with IP monitoring is seen. [PR1263078](#)
- The TTL (Time To Live) of some Z-mode packets is reduced to zero incorrectly, if IOC2 or IOC3 interface is configured as HA fabric port. [PR1270770](#)
- DNS cache does not get populated in multiple virtual router (VR) environments. [PR1275792](#)
- Memory leak occurs on SRX Series devices chassis cluster when em0 or em1 interface is down. [PR1277136](#)
- On SRX5000 line devices, under a heavy flood of IPv6 Neighbor Discovery Protocol (NDP) packets, some incoming IPv6 neighbor advertisements (NA) might be dropped because of a queue being full. This issue has been resolved by using a different queue for IPv6 NA packets. [PR1293673](#)
- XLP lost heartbeat (SPU hang) is not detected in a timely manner by hardware monitoring. [PR1300804](#)

Routing Policy and Firewall Filters

- Secured e-mail application is not available. [PR1273725](#)
- On SRX Series devices, the DNS configured in the address-book fails to resolve the IP address, if the case (uppercase or lowercase) in the DNS query and the DNS response do not match. [PR1304706](#)
- The NSD process might crash when replacing the name of a logical-system. [PR1307876](#)

System Logging

- The logs from syslog **RT_FLOW: FLOW_REASSEMBLE_SUCCEED: Packet merged** might cause high CPU usage on the Routing Engine. [PR1278333](#)

Unified Threat Management (UTM)

- The Packet Forwarding Engine CPU utilization is high when using the UTM antivirus feature. [PR1282719](#)

VPNs

- The st0 global counter statistics do not increment. [PR1171958](#)
- The second client is disconnected when the assigned IP address is changed in the access profile for the first client. [PR1246131](#)
- IPsec traffic through tunnel fails without configuring the authentication algorithm under the IPsec proposal on the SRX1500; however, it works on the SRX5600. [PR1285284](#)

SEE ALSO

New and Changed Features 405
Changes in Behavior and Syntax 417
Known Behavior 422
Known Issues 424
Documentation Updates 449
Migration, Upgrade, and Downgrade Instructions 449

Documentation Updates

There are no errata or changes in Junos OS Release 17.4R3 for the SRX Series documentation.

SEE ALSO

New and Changed Features 405
Changes in Behavior and Syntax 417
Known Behavior 422
Known Issues 424
Resolved Issues 427
Migration, Upgrade, and Downgrade Instructions 449

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Scripts for Address Book Configuration | 450](#)

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 | 450](#)
- [Running Upgrade and Downgrade Scripts | 451](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 452](#)

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 451](#)).

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 | 405](#)

[Changes in Behavior and Syntax | 417](#)

Known Behavior | 422

Known Issues | 424

Resolved Issues | 427

Documentation Updates | 449

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the [High Availability User Guide](#).

For additional information about using ISSU on security devices, see the [Chassis Cluster User Guide for SRX Series Devices](#).

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

Compliance Advisor

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

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. <https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home


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

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. apps.juniper.net/compliance/.

Documentation Feedback

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

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



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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

13 January 2021—Revision 5, Junos OS Release 17.4R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 4, Junos OS Release 17.4R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 April 2020—Revision 3, Junos OS Release 17.4R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 February 2020—Revision 2, Junos OS Release 17.4R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 January 2020—Revision 1, Junos OS Release 17.4R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 July 2019—Revision 14, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 June 2019—Revision 13, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 May 2019—Revision 12, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 March 2019—Revision 11, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 March 2019—Revision 10, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 February 2019—Revision 9, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 February 2019—Revision 8, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 January 2019—Revision 7, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 December 2018—Revision 6, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 November 2018—Revision 5, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 October 2018—Revision 4, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 October 2018—Revision 3, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 September 2018—Revision 2, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2018—Revision 2, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 August 2018—Revision 1, Junos OS Release 17.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 February 2018—Revision 5, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

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

4 January 2018—Revision 3, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 December 2017—Revision 2, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 December 2017—Revision 1, Junos OS Release 17.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

Copyright © 2020 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.

