

Release Notes: Junos[®] OS Release 17.1R3 for the ACX Series, EX Series, MX Series, PTX Series, QFX Series, and Junos Fusion

3 September 2020

Contents	Introduction 10
	Junos OS Release Notes for ACX Series 10
	New and Changed Features 11
	Release 17.1R3 New and Changed Features 11
	Release 17.1R2 New and Changed Features 11
	Release 17.1R1 New and Changed Features 11
	Changes in Behavior and Syntax 17
	General Routing 17
	Interfaces and Chassis 18
	MPLS 18
	Network Management and Monitoring 18
	Platform and Infrastructure 19
	Services Applications 20
	System Management 20
	User Interface and Configuration 20
	Known Behavior 21
	General Routing 21
	High Availability (HA) and Resiliency 22
	Known Issues 23
	General Routing 23
	Network Address Translation (NAT) and Stateful Firewall Services 24

Resolved Issues | 24

Resolved Issues: 17.1R3 | 25

Resolved Issues: 17.1R2 | 26

Resolved Issues: 17.1R1 | 26

Documentation Updates | 27

Protocol Independent Routing Properties | 27

Migration, Upgrade, and Downgrade Instructions | 27

Upgrade and Downgrade Support Policy for Junos OS Releases | 28

Product Compatibility | 29

Hardware Compatibility | 29

Junos OS Release Notes for EX Series Switches | 30

New and Changed Features | 30

Release 17.1R3 New and Changed Features | 32

Release 17.1R2 New and Changed Features | 32

Release 17.1R1 New and Changed Features | 32

Changes in Behavior and Syntax | 36

General Routing | 37

High Availability (HA) and Resiliency | 37

MPLS | 37

Network Management and Monitoring | 37

Services Applications | 39

System Management | 39

User Interface and Configuration | 39

Virtual Chassis | 39

Known Behavior | 40

General Routing | 40

High Availability (HA) and Resiliency | 41

Interfaces and Chassis | 41

Known Issues | 41

General Routing | 42

Infrastructure | 44

Junos Fusion Enterprise | 44

Layer 2 Features | 45

Multicast | 45

Network Management and Monitoring	45
Platform and Infrastructure	46
Spanning Tree Protocols	46
Virtual Chassis	46
Resolved Issues	47
Resolved Issues: 17.1R3	47
Resolved Issues: 17.1R2	55
Resolved Issues: 17.1R1	57
Documentation Updates	59
Documentation Updates	59
Migration, Upgrade, and Downgrade Instructions	60
Upgrade and Downgrade Support Policy for Junos OS Releases	60
Product Compatibility	61
Hardware Compatibility	61
Junos OS Release Notes for Junos Fusion Enterprise	62
New and Changed Features	63
Release 17.1R3 New and Changed Features	63
Release 17.1R2 New and Changed Features	63
Release 17.1R1 New and Changed Features	64
Changes in Behavior and Syntax	68
System Management	68
Known Behavior	68
Junos Fusion Enterprise	69
Known Issues	71
Junos Fusion Enterprise	72
Resolved Issues	73
Resolved Issues: 17.1R3	73
Resolved Issues: 17.1R2	74
Resolved Issues: 17.1R1	74
Documentation Updates	75
Migration, Upgrade, and Downgrade Instructions	75
Basic Procedure for Upgrading Junos OS on an Aggregation Device	76
Upgrading from Junos OS Release 16.1 to 17.1 in a JUNOS OS Fusion Enterprise System	78

Upgrading an Aggregation Device with Redundant Routing Engines	79
Preparing the Switch for Satellite Device Conversion	79
Converting a Satellite Device to a Standalone Switch	80
Upgrade and Downgrade Support Policy for Junos OS Releases	81
Downgrading from Release 17.1	81
Product Compatibility	82
Hardware and Software Compatibility	82
Hardware Compatibility Tool	82
Junos OS Release Notes for Junos Fusion Provider Edge	83
New and Changed Features	84
Release 17.1R3 New and Changed Features	84
Release 17.1R2 New and Changed Features	84
Release 17.1R1 New and Changed Features	84
Changes in Behavior and Syntax	85
System Management	85
Known Behavior	86
Known Issues	86
Resolved Issues	87
Resolved Issues: 17.1R3	87
Resolved Issues: 17.1R2	88
Resolved Issues: 17.1R1	88
Documentation Updates	88
Migration, Upgrade, and Downgrade Instructions	89
Basic Procedure for Upgrading an Aggregation Device	89
Upgrading an Aggregation Device with Redundant Routing Engines	92
Preparing the Switch for Satellite Device Conversion	92
Converting a Satellite Device to a Standalone Device	93
Upgrading an Aggregation Device	96
Upgrade and Downgrade Support Policy for Junos OS Releases	96
Downgrading from Release 17.1	96
Product Compatibility	97
Hardware Compatibility	97

Junos OS Release Notes for MX Series 5G Universal Routing Platforms | 98

New and Changed Features | 99

Release 17.1R3 New and Changed Features | 100

Release 17.1R2 New and Changed Features | 101

Release 17.1R1 New and Changed Features | 103

Changes in Behavior and Syntax | 125

Class of Service (CoS) | 127

General Routing | 127

Interfaces and Chassis | 127

Intrusion Detection and Prevention | 129

Junos OS XML API and Scripting | 129

Layer 2 VPN | 130

Management | 130

MPLS | 131

Network Management and Monitoring | 132

Operation, Administration, and Maintenance (OAM) | 134

Routing Protocols | 134

Security | 136

Services Applications | 136

Software Defined Networking | 137

Subscriber Management and Services | 137

System Management | 141

User Interface and Configuration | 141

VPNs | 141

Known Behavior | 142

Class of Service (CoS) | 143

EVPN | 143

Forwarding and Sampling | 144

General Routing | 145

High Availability (HA) and Resiliency | 146

Interfaces and Chassis | 146

MPLS | 147

Platform and Infrastructure | 147

Routing Protocols | 147

- Services Applications | 147
- Software Installation and Upgrade | 148
- Subscriber Management and Services | 149

Known Issues | 150

- Authentication and Access Control | 151
- Class of Service (CoS) | 151
- EVPN | 151
- Forwarding and Sampling | 152
- General Routing | 153
- High Availability (HA) and Resiliency | 159
- Infrastructure | 159
- Interfaces and Chassis | 159
- Layer 2 Ethernet Services | 160
- Layer 2 Features | 160
- MPLS | 161
- Network Management and Monitoring | 163
- Platform and Infrastructure | 163
- Routing Protocols | 165
- Services Applications | 167
- Subscriber Access Management | 167
- User Interface and Configuration | 168
- VPNs | 168

Resolved Issues | 169

- Resolved Issues: 17.1R3 | 169
- Resolved Issues: 17.1R2 | 215
- Resolved Issues: 17.1R1 | 224

Documentation Updates | 232

- Subscriber Management Access Network Guide | 232
- Subscriber Management Provisioning Guide | 232
- Subscriber Management VLANs Interfaces Guide | 233

Migration, Upgrade, and Downgrade Instructions | 234

- Basic Procedure for Upgrading to Release 17.1 | 235
- Procedure to Upgrade to FreeBSD 10.x based Junos OS | 235
- Procedure to Upgrade to FreeBSD 6.x based Junos OS | 237

Upgrade and Downgrade Support Policy for Junos OS Releases	239
Upgrading a Router with Redundant Routing Engines	240
Downgrading from Release 17.1	240
Product Compatibility	241
Hardware Compatibility	241
Junos OS Release Notes for PTX Series Packet Transport Routers	242
New and Changed Features	242
Release 17.1R3 New and Changed Features	243
Release 17.1R2 New and Changed Features	243
Release 17.1R1 New and Changed Features	243
Changes in Behavior and Syntax	252
General Routing	253
Interfaces and Chassis	253
Management	254
MPLS	254
Network Management and Monitoring	255
Routing Protocols	256
Services Applications	256
System Management	256
User Interface and Configuration	256
Known Behavior	257
General Routing	258
Interfaces and Chassis	258
High Availability (HA) and Resiliency	258
Known Issues	259
General Routing	259
Interfaces and Chassis	261
MPLS	261
Platform and Infrastructure	261
Routing Protocols	261
User Interface and Configuration	261
Resolved Issues	262
Resolved Issues: 17.1R3	262
Resolved Issues: 17.1R2	267

Resolved Issues: 17.1R1	269
Documentation Updates	270
Migration, Upgrade, and Downgrade Instructions	270
Basic Procedure for Upgrading to Release 17.1	270
Upgrade and Downgrade Support Policy for Junos OS Releases	273
Upgrading a Router with Redundant Routing Engines	274
Product Compatibility	274
Hardware Compatibility	275
Junos OS Release Notes for the QFX Series	275
New and Changed Features	276
Release 17.1R3 New and Changed Features	277
Release 17.1R2 New and Changed Features	277
Release 17.1R1 New and Changed Features	277
Changes in Behavior and Syntax	296
Class of Service	297
General Routing	297
MPLS	297
Network Management and Monitoring	297
Security	300
Services Applications	300
Software Defined Networking	300
Software Installation and Upgrade	300
System Management	300
User Interface and Configuration	300
Virtual Chassis	301
VPNs	301
Known Behavior	302
EVPN	302
General Routing	303
High Availability and Resiliency	303
Layer 2 Features	304
MPLS	304
Routing Protocols	304

Known Issues | 304**EVPN | 305****General Routing | 306****MPLS | 307****Network Management and Monitoring | 307****Platform and Infrastructure | 307****Routing Protocols | 308****Resolved Issues | 309****Resolved Issues: 17.1R3 | 309****Resolved Issues: 17.1R2 | 321****Resolved Issues: 17.1R1 | 323****Documentation Updates | 326****Migration, Upgrade, and Downgrade Instructions | 326****Upgrading Software on QFX Series Switches | 327****Installing the Software on QFX10002 Switches | 329****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 | 329****Installing the Software on QFX10008 and QFX10016 Switches | 331****Performing a Unified ISSU | 335****Preparing the Switch for Software Installation | 336****Upgrading the Software Using Unified ISSU | 336****Product Compatibility | 339****Hardware Compatibility | 339****Upgrading Using ISSU | 340****Compliance Advisor | 340****Finding More Information | 340****Requesting Technical Support | 341****Self-Help Online Tools and Resources | 341****Opening a Case with JTAC | 342****Revision History | 342**

Introduction

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

These release notes accompany Junos OS Release 17.1R3 for the ACX Series, EX Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, PTX Series, and QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- New and Changed Features | 11
- Changes in Behavior and Syntax | 17
- Known Behavior | 21
- Known Issues | 23
- Resolved Issues | 24
- Documentation Updates | 27
- Migration, Upgrade, and Downgrade Instructions | 27
- Product Compatibility | 29

These release notes accompany Junos OS Release 17.1R3 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 17.1R3 New and Changed Features | 11](#)
- [Release 17.1R2 New and Changed Features | 11](#)
- [Release 17.1R1 New and Changed Features | 11](#)

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 17.1R3 New and Changed Features

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 17.1R3.

Release 17.1R2 New and Changed Features

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 17.1R2.

Release 17.1R1 New and Changed Features

This section describes the new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 17.1R1.

Application Level Gateways (ALGs)

- **Support for Application Level Gateways (ALGs) for NAT processing (ACX500)**—Starting with Junos OS Release 17.1R1, ACX500 routers support basic TCP, basic UDP, DNS, FTP, ICMP, TFTP, and UNIX Remote-Shell Services ALGs for NAT processing.

NOTE: The ALG for NAT is supported only on the ACX500 indoor routers.

[See [ALGs Available by Default for Junos OS Address Aware NAT on ACX500 Router.](#)]

Bridging

- **Support for DHCP option 82 over bridge domain (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX Universal Metro Routers supports configuring DHCP option 82 over bridge domain. ACX routers support option 82 type, length, and value (TLV) information for DHCP client messages over bridge domain.

[See [Using DHCP Relay Agent Option 82 Information.](#)]

Firewall

- **Support for stateful firewall (ACX500)**—Starting with Junos OS Release 17.1R1, ACX500 Universal Metro Routers supports configuring stateful firewall rules. Contrasted with a stateless firewall that inspects packets in isolation, a stateful firewall provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions for new communication attempts.

NOTE: The stateful firewall configuration is supported only on the ACX500 indoor routers.

[See [Junos Network Secure Overview.](#)]

Generic Routing

- **Support for generic routing encapsulation (GRE) (ACX Series)**—Starting with Junos OS Release 17.1R1, ACX Series Universal Metro Routers support configuring generic routing encapsulation (GRE). GRE provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets inside a transport protocol known as an IP encapsulation protocol.

[See [Understanding Generic Routing Encapsulation on ACX Series.](#)]

Interfaces and Chassis

- **Aggregated Ethernet load-balancing support for circuit cross-connect (CCC), VPLS, bridge domain, and Layer 3 VPN (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers support aggregated Ethernet (AE) operation over Layer 2 circuit, Layer 3 VPN, bridge domain, CCC, OAM, no-local-switching, and IGMP snooping. Also supported are AE class of service and firewall support for families such as bridge domain, VPLS, CCC, MPLS, IPv4, and IPv6. The firewall support extends the support for single-rate two-color policer and two-rate two-color policer.

[See [Understanding Ethernet Link Aggregation on ACX Series Routers.](#)]

Junos OS XML API and Scripting

- **Support for Python language for commit, event, op, and SNMP scripts (ACX500, ACX1000, ACX1100, ACX2000, ACX2100, ACX2200, and ACX4000)**—Starting in Junos OS Release 17.1R1, you can author commit, event, op, and SNMP scripts in Python on devices that include the Python extensions package in the software image. Creating automation scripts in Python enables you to take advantage of Python features and libraries as well as leverage Junos PyEZ APIs to perform operational and configuration tasks on devices running Junos OS. To enable execution of Python automation scripts, which must be owned by either root or a user in the Junos OS **super-user** login class, configure the **language python** statement at the **[edit system scripts]** hierarchy level, and configure the filename for the Python script under the hierarchy level appropriate to that script type. Supported Python versions include Python 2.7.x.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

Layer 2 Features

- **Support for pseudowire cross-connect (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers supports pseudowire cross-connect. The pseudowire cross-connect feature enables virtual circuit (VC) to terminate locally on a router and supports local switching of Layer 2 circuits. Layer 2 circuits allows the creation of point-to-point Layer 2 connections over an IP and MPLS-based network. Physical circuits with the same Layer 2 encapsulations can be connected together across such a network.

[See [Configuring Local Interface Switching in Layer 2 Circuits.](#)]

Mirroring

- **Support for port mirroring (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers supports port mirroring to mirror a copy of a packet to a configured destination, in addition to the normal processing and forwarding of the packet. Port mirroring is supported on both ingress and egress ports, using a protocol analyzer application that passes the input to mirror through a list of ports configured through the logical interface.

[See [Port, VLAN, and Flow Mirroring Overview.](#)]

MPLS

- **Support for the Path Computation Element Protocol (ACX Series)**—Starting with Junos OS Release 17.1R1, ACX Series Universal Metro Routers support the Path Computation Element Protocol (PCEP). A Path Computation Element (PCE) is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE. PCEP enables communications between a PCC and a PCE, or between two PCEs (defined in RFC 5440). PCEP is a TCP-based protocol defined by the IETF PCE Working Group, and defines a set of messages and objects used to manage PCEP sessions and to request and send paths for multidomain traffic engineered LSPs (TE LSPs). It provides a mechanism for a PCE to perform path computation for a PCC's external LSPs. The PCEP interactions include LSP status reports sent by the PCC to the PCE, and PCE updates for the external LSPs.

[See [PCEP Overview](#).]

Network Management and Monitoring

- **Support for hrProcessorTable object (ACX Series)**—Starting in Junos OS Release 17.1R1, support is provided for the **hrProcessorTable** object (object id: 1.3.6.1.2.1.25.3.3) described in the RFC2790, *Host Resources MIB*. The **hrProcessorTable** object provides the load statistics information per CPU for multi-core devices.

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

- **Support for RFC 2544 reflector (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers support the Layer 1 reflector functionality for performing RFC 2544 benchmarking tests. The device that is configured as a reflector reflects or sends back the packets as they are received on the pseudowire. This feature does not support any packet modification functionality. To enable your ACX5000 router to reflect the packets back to the initiator, you can configure any unused physical port on the router as the reflector port. Use the **reflector-port** statement at the **[edit services rpm rfc2544-benchmarking tests test-name]** hierarchy level to configure the reflector port.

[See [RFC 2544-Based Benchmarking Tests Overview](#).]

Operations, Administration, and Management (OAM)

- **SNMP support for Service OAM (SOAM) performance monitoring functions (ACX Series)**—Starting with Junos OS Release 17.1R1, ACX Series Universal Metro Routers support Service OAM (SOAM) performance monitoring functions that are defined in Technical Specification MEF 17, the Service OAM performance monitoring requirements specified in SOAM-PM, and the Service OAM management objects specified in Technical Specification MEF 7.1.

A new enterprise-specific MIB, SOAM PM MIB, that defines the management objects for Ethernet services operations, administration, and maintenance for performance monitoring, has been added and SNMP support is available for the MIB objects defined in Technical Specification MEF 36.

[See [Interpreting the Enterprise-Specific Service OAM MIB](#).]

Spanning-Tree Protocols

- **Support for bridge protocol data unit, loop protect, and root protect (ACX Series)**—Starting with Junos OS Release 17.1R1, ACX Series Universal Metro Routers support configuring bridge protocol data unit (BPDU), loop protect, and root protect on spanning-tree instance interface. You can configure BPDU protection on individual interfaces or on all the edge ports of the bridge.

[See [Understanding BPDU Protection for Spanning-Tree Instance Interfaces](#), [Understanding Loop Protection for Spanning-Tree Instance Interfaces](#), [Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network](#).]

Timing and Synchronization

- **Support for precision time protocol over integrated routing and bridging (ACX Series)**—Starting with Junos OS Release 17.1R1, ACX Series Universal Metro Routers support configuring precision time protocol (PTP) over integrated routing and bridging (IRB). You can configure a boundary clock node with PTP (IPv4) over IRB in a master-only mode across single or multiple IRB logical interfaces.

[See [Configuring Precision Time Protocol Over Integrated Routing and Bridging](#).]

- **Support for timing and synchronization (ACX Series)**—Starting with Junos OS Release 17.1R1, ACX Universal Metro Routers support external clock synchronization and automatic clock selection for Synchronous Ethernet, T1 or E1 line timing sources, and external inputs. The IEEE 1588v2 standard defines the Precision Time Protocol (PTP), which is used to synchronize clocks throughout a network. ACX Series routers support PTP ordinary clock and boundary clock features. ACX Series routers also support PTP over Ethernet.

[See [External Clock Synchronization Overview for ACX Series Routers](#), [Automatic Clock Selection Overview](#).]

- **Support for transparent clock (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers support the transparent clock functionality. Transparent clocks measure packet residence time for Precision Time Protocol (PTP) events. The packet delay variation experienced by PTP packets can be attributed to queuing and buffering delays inside the router. ACX5000 routers support only end-to-end transparent clock functionality as defined in the IEEE 1588 standard. The transparent clock functionality works for both PTP over IP (PTPoIP), and PTP over Ethernet (PTPoE).

To configure the transparent clock functionality, you must include the **e2e-transparent** statement at the **[edit protocol ptp]** hierarchy level.

Use the **show ptp global-information** command to check the status of the transparent clock functionality configured on the router.

[See [Understanding Transparent Clocks in Precision Time Protocol.](#)]

Tunneling

- **Support for remote loop-free alternate (LFA) over LDP tunnels in IS-IS and OSPF networks (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers support remote LFA over LDP tunnels in an IS-IS and OSPF network. Remote LFA increases the backup coverage for IS-IS and OSPF routes and provides protection especially for Layer 1 metro-rings. The IS-IS protocol creates a dynamic LDP tunnel to reach the remote LFA node from the point of local repair (PLR). The PLR uses this remote LFA backup path when the primary link fails.

[See [Configuring Remote LFA Backup over LDP Tunnels in an OSPF Network](#), [Configuring Remote LFA Backup over LDP Tunnels in an IS-IS Network](#).]

- **Support for automatic bandwidth allocation for label-switched paths (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers support automatic bandwidth allocation for label-switched paths (LSPs). Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

[See [Automatic Bandwidth Allocation for LSPs](#).]

VPLS

- **Mesh group support for VPLS routing (ACX5000)**—Starting with Junos OS Release 17.1R1, ACX5000 Universal Metro Routers support mesh group configuration for VPLS routing instances. A mesh group within the routing instance is a group of PE interface members with common forwarding attributes. The following are the default member attributes in a mesh group:
 - **no-local-switching**—Traffic will not switch between members of the same mesh group (known-unicast, multicast, broadcast, unknown-unicast).
 - **flood-to-all-other-mesh-group**—Traffic can flow from a member of one mesh group to any set of members of other mesh groups.

[See [Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS](#).]

SEE ALSO

Known Behavior 21
Known Issues 23
Resolved Issues 24
Documentation Updates 27
Migration, Upgrade, and Downgrade Instructions 27
Product Compatibility 29

Changes in Behavior and Syntax

IN THIS SECTION

- General Routing | 17
- Interfaces and Chassis | 18
- MPLS | 18
- Network Management and Monitoring | 18
- Platform and Infrastructure | 19
- Services Applications | 20
- System Management | 20
- User Interface and Configuration | 20

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 17.1R3 for the ACX Series.

General Routing

- For the **routing** command, starting in Junos 15.1F3, 15.1R2, 15.1R3, and 15.2R1, 64-bit mode is enabled by default on systems that support it and that have at least 16 GB of RAM.
- **Support for deletion of static routes when the BFD session goes down (ACX Series)**—Starting with Junos OS Release 17.1R3, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session-down message.

Interfaces and Chassis

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

MPLS

- **Representation for OSPF DR node**—Up until version -10 of the BGP-LS draft, the OSPF DR node representation was ambiguous. One could represent DR node as 'AdvertisingRouterId-InterfaceIpAdress' or 'InterfaceIpAdress-1'. Junos OS used to follow 'InterfaceIpAdress-1' format. Starting with version '-11' of the BGP-LS draft, the representation for OSPF DR node must be 'AdvertisingRouterId-InterfaceIpaddress'. Junos OS now follows the latest format.

Network Management and Monitoring

- **SNMP syslog messages changed (ACX Series)**—Starting in Junos OS Release 17.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD — **AgentX master agent failed to respond to ping. Attempting to re-register**
NEW — **AgentX master agent failed to respond to ping, triggering cleanup!**
 - OLD — **NET-SNMP version %s AgentX subagent connected**
NEW — **NET-SNMP version %s AgentX subagent Open-Sent!**

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

- **Update to SNMP support of apply-path statement (ACX Series)**—Starting in Junos OS Release 17.1R2, SNMP implementation for the **apply-path** configuration statement supports only two lists:
 - **apply-path "policy-options prefix-list <list-name> <*>"**
This configuration has been supported from day 1.
 - **apply-path "access radius-server <*>"**
This configuration is supported as of Junos OS Release 17.1R2.
- **Juniper MIBs loading errors fixed (ACX Series)**—Starting in Junos OS Release 17.1R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:
 - jnx-gen-set.mib
 - jnx-ifotn.mib
 - jnx-optics.mib

[See [MIB Explorer](#).]

- **Change in default log level setting (ACX Series)**—Starting in Junos OS Release 17.1R3, 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 and non-default logical system (ACX Series)**—Starting in Junos OS Release 17.1R1, 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](#).]

- **Need to reconfigure SNMPv3 configuration after upgrade (ACX5048 and ACX5096)**—Starting in Junos OS Release 17.1R2, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. This is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. It used to be that customers had to reconfigure SNMPv3 every time after a reboot. This problem was fixed. If you upgrade, you must still reconfigure SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3 again. To reconfigure SNMP v3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters.

[See [Configuring the Local Engine ID](#).]

Platform and Infrastructure

- **DMA recovery mechanism (ACX Series)**—Starting in Junos OS Release 17.1R3, 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 Packet Forwarding Engine reboot to recover from idle state.

The following recovery message is logged in the Routing Engine 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 Packet Forwarding Engine syslog message:

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

Services Applications

- **Device discovery with device-initiated connection (ACX Series)**—In Junos OS Release 17.1R1 and later releases, when you configure statements and options under the `[system services ssh]` hierarchy and commit the configuration, make sure that the system reaches a stable state before you commit any `outbound-ssh` configurations.

You use the device discovery feature in the Devices workspace to add devices to Junos Space Network Management Platform. By default, Junos Space manages devices by initiating and maintaining a connection to the device.

[See [Device Discovery Overview](#).]

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the `peers` option at the `[edit system commit]` hierarchy level is not supported in batch configuration mode.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (ACX Series)**—Starting in Junos OS Release 17.1R1, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 17.1, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the show system schema module juniper-command output directory (ACX Series)**—Starting in Junos OS Release 17.1R1, when you issue the `show system schema module juniper-command` operational command in the Junos OS CLI, the device places the generated output files in the current working directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1R1, the generated output files are placed in the `/var/tmp` directory.

SEE ALSO

[New and Changed Features | 11](#)

[Known Behavior | 21](#)

[Known Issues | 23](#)[Resolved Issues | 24](#)[Documentation Updates | 27](#)[Migration, Upgrade, and Downgrade Instructions | 27](#)[Product Compatibility | 29](#)

Known Behavior

IN THIS SECTION

- [General Routing | 21](#)
- [High Availability \(HA\) and Resiliency | 22](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 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

- When you run the **restart clksyncd-service** CLI command, incorrect correction field values are seen when transparent clock is **INACTIVE**. This does not have any functional impact. [PR1067583](#)
- ACX Series hardware does not have support for unicast RPF statistics. The values currently shown in Junos OS CLI at logical interface level for unicast RPF statistics can be ignored. [PR1188020](#)
- ACX Series hardware supports unicast RPF mode at the physical interface level only, though Junos OS configurations are done at the logical interface level. To avoid confusion with respect to unicast RPF modes, only one mode (strict/loose) should be configured for all the logical interfaces within a physical interface. This applies to bridge-domain logical interface also, if IRB is present in the bridge domain and unicast RPF is enabled in the IRB logical interface. [PR1196908](#)
- In a VPWS termination case, when no explicit expiration or default expiration classifier is configured for classification, it is expected that the packets should be classified to be queued to default queue 0. Instead

they are getting queued to the assured forwarding queue. This issue is seen only with ACX4000.
[PR1201072](#)

- The following error logs will be seen in a scenario where the number of MPLS tunnels exceeds the limit for the platform: - `acx_nh_mpls_tunnel_install(),1076:acx_nh_mpls_tunnel_install: BCM TNL InitiatorSet failed for NH ##### (Table full)` . As a consequence of a previous error, the following logs could also be seen: - `fpc0 ACX_NH::acx_nh_l3_tag_hw_install(),##### :acx_nh_l3_tag_hw_install: Tunnel installed failed: NH ##### - fpc0 NH: Failed to find nh (3662) for deletion`. When the error persists for a while, you eventually get the following message:
`ACX_NH::acx_nh_mpls_tunnel_uninstall(),1171:acx_nh_mpls_tunnel_uninstall: BCM L3 Egress destroy object failed for (-10:Operation still running)`. Note: For the ACX5000 line of routers, the maximum number of MPLS tunnels supported is 1024. In scenarios with route flaps, new tunnels will be created and reach limits during addition and deletion. As a result, we recommend you avoid reaching over 1000 tunnels in a normal operation. > To check the number of tunnels use CLI command: request pfe execute the `show pfe-hw mpls-tunnel` command target fpc0. [PR1231621](#)

High Availability (HA) and Resiliency

- **Unified ISSU incompatibility with VPLS dynamic profiles (ACX Series)**—Using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 17.1R1 does not work if VPLS dynamic profiles are configured and enhanced subscriber management is not configured.
[See [ISSU System Requirements](#).]

SEE ALSO

New and Changed Features 11
Changes in Behavior and Syntax 17
Known Issues 23
Resolved Issues 24
Documentation Updates 27
Migration, Upgrade, and Downgrade Instructions 27
Product Compatibility 29

Known Issues

IN THIS SECTION

- [General Routing | 23](#)
- [Network Address Translation \(NAT\) and Stateful Firewall Services | 24](#)

This section lists the known issues in hardware and software in Junos OS Release 17.1R3 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

- In Junos OS Release 12.3X54-D10.6, the aggregated Ethernet interface on ACX Series routers is permanently down after reboot, when link speed is configured. [PR1022248](#)
- SNMP MIB walk on jnxDomCurrentTable and jnxDomNotifications is enabled. [PR1076943](#)
- In some cases it was observed that the fan is running but CLI command and syslog give fan failure alarms, and the other fan is spinning at high speed: **root@> show chassis alarms no-forwarding alarms currently active Alarm time Class Description 2010-01-01 00:12:04 UTC Minor Single FAN Failure root@> show chassis environment no-forwarding Class Item Status Measurement Fans Fan 1 Check Fan 2 OK Spinning at high speed.** [PR1127846](#)
- Load balancing does not work at LSR with multiple CCC/VPLS sessions. [PR1198435](#)
- ACX: does not forward DHCP-RELAY requests with IRB interface after upgrade. [PR1243687](#)
- On ACX Series routers with Dynamic Host Configuration Protocol (DHCP) relay configured, if rebooting scaling number of DHCP clients at the same time, the DHCP negotiations might fail and eventually cause outage. [PR1335957](#)
- The remote fault signaling is not supported for 1 Gigabit fiber SFP during auto-negotiation. Therefore in releases without the fix of this PR, we get a 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](#)
- On MX Series and ACX Series platforms, offlining and then bringing the MIC-3D-16CHE1-T1-CE-H card back online might cause the FPC to crash. [PR1402563](#)

Network Address Translation (NAT) and Stateful Firewall Services

- On the ACX500 routers, when service application logging is enabled at [**edit services service-set service-set-name syslog host *host-name* class**] hierarchy level and when packets containing errors are received at higher rate toward the service engine, the resource scale requirements at the service engine cannot be met and the service processor might reboot. As a workaround, you can disable the application logging. [PR1223500](#)
- On the ACX500 routers, when there is a fast ramp-up of scaled user applications, the resource requirements of the service engine cannot be met. As a workaround, you can disable the application logging. [PR1226153](#)

SEE ALSO

New and Changed Features 11
Changes in Behavior and Syntax 17
Known Behavior 21
Resolved Issues 24
Documentation Updates 27
Migration, Upgrade, and Downgrade Instructions 27
Product Compatibility 29

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.1R3 | 25](#)
- [Resolved Issues: 17.1R2 | 26](#)
- [Resolved Issues: 17.1R1 | 26](#)

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: 17.1R3

General Routing

- Upgrade or large configuration commit fails are observed on ACX5000 router. [PR1088017](#)
- ACX1100 with midplane part no 650-062965 might fail to initialize FPGA. [PR1134335](#)
- MPLS LSP are being affected because NH failed to be programmed. [PR1195419](#)
- The chassisd[1825]: pvidb_get_root_node: Error(2) retrieving rootnode value error messages might be seen. [PR1198817](#)
- Several error logs are seen on the ACX Series router when a link in the primary path of LSP flaps. [PR1204714](#)
- EVPN CLI is exposed on ACX1000, ACX2000, and ACX4000. [PR1208248](#)
- The tagged/untagged LLDP and LACP packets are dropped on the VPLS CE device facing the aggregated Ethernet interface. [PR1245242](#)
- The 100m/full-duplex setting does not work with SFP-T at rebooting. [PR1262752](#)
- Transit ARP packets are being punted to the Routing Engine. [PR1263012](#)
- ARP packets coming with a VLAN tag that not configured in the ACX Series router are hitting the default_arp_policer. [PR1271100](#)
- In some rare scenarios, ACX500 IPsec will have a license activation error with the **license not valid for this product** message. [PR1275149](#)
- SHEAF leak is seen on the ACX5000 line of routers. [PR1280492](#)
- ACX2x00-AC is reporting false PEM0 alarms periodically. [PR1310488](#)
- Error syslog on output/egress firewall filter occurs on ACX Series routers. [PR1316588](#)
- Network events might cause Layer 2 circuit traffic forwarding to fail with the **Table Full** message. [PR1319591](#)
- With auto-installation USB configured, interface-related commits might not take effect because of dcd error. [PR1327384](#)
- An IPv6 service outage might occur after executing the **clear ipv6 neighbor** command. [PR1330791](#)
- Unable to commit multiple Ethernet-ring instances on ACX Series routers. [PR1337497](#)
- The ARP reply packet might be dropped in a Layer 2 circuit secondary path when using IEEE-802.1 classifier. [PR1341126](#)
- Traffic destined for specific IP within a subnet gets discarded without notification. [PR1345098](#)
- NAT might not work and the spd might crash. [PR1346546](#)
- The fxpc might crash on Packet Forwarding engine command **show pfe context_vlan**. [PR1349721](#)
- ARP reply is dropped when temporal buffer-size is added on the NNI interface. [PR1363153](#)

- IPsec SA as OSPFv3 authentication is not working in Junos OS Releases 16.2R2 and 17.3R2. [PR1363487](#)
- The fxpc might crash after an interface is changed on ACX5000 routers. [PR1378155](#)
- Certain builds of Junos OS do not allow you to upgrade or commit configuration changes when the SI service interface is used. [PR1393729](#)

Layer 2 Ethernet Services

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

Platform and Infrastructure

- On devices running Junos OS, the next-hop index allocation fails and private index space gets exhausted through incoming ARP requests to the management interface (CVE-2018-0063). [PR1360039](#)

Resolved Issues: 17.1R2

General Routing

- The SNMP MIB walk for jnxOperatingState incorrectly shows the CB status as down. [PR1191995](#)
- 10-Gigabit Ethernet interface fault detection behavior changed. [PR1223457](#)

Resolved Issues: 17.1R1

There are no fixed issues in Junos OS 17.1R1 for ACX Series.

SEE ALSO

New and Changed Features	 11
Changes in Behavior and Syntax	 17
Known Behavior	 21
Known Issues	 23
Documentation Updates	 27
Migration, Upgrade, and Downgrade Instructions	 27
Product Compatibility	 29

Documentation Updates

IN THIS SECTION

- [Protocol Independent Routing Properties | 27](#)

This section lists the errata and changes in Junos OS Release 17.1R3 for the ACX Series documentation.

Protocol Independent Routing Properties

- **Support for deletion of static routes when the BFD session goes down (ACX Series)**—Starting with Junos OS Release 17.1R3, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when BFD receives a session-down message.

SEE ALSO

[New and Changed Features | 11](#)

[Changes in Behavior and Syntax | 17](#)

[Known Behavior | 21](#)

[Known Issues | 23](#)

[Resolved Issues | 24](#)

[Migration, Upgrade, and Downgrade Instructions | 27](#)

[Product Compatibility | 29](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 28](#)

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 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 15.1, 16.1 and 16.2 are EEOL releases. You can upgrade from Junos OS Release 15.1 to Release 16.1 or even from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

SEE ALSO

New and Changed Features	 11
Changes in Behavior and Syntax	 17
Known Behavior	 21
Known Issues	 23
Resolved Issues	 24
Documentation Updates	 27
Product Compatibility	 29

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 29](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 11
Changes in Behavior and Syntax 17
Known Behavior 21
Known Issues 23
Resolved Issues 24
Documentation Updates 27
Migration, Upgrade, and Downgrade Instructions 27

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- New and Changed Features | 30
- Changes in Behavior and Syntax | 36
- Known Behavior | 40
- Known Issues | 41
- Resolved Issues | 47
- Documentation Updates | 59
- Migration, Upgrade, and Downgrade Instructions | 60
- Product Compatibility | 61

These release notes accompany Junos OS Release 17.1R3 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 17.1R3 New and Changed Features | 32
- Release 17.1R2 New and Changed Features | 32
- Release 17.1R1 New and Changed Features | 32

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

NOTE: The following EX Series switches are supported in Release 17.1R3: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.1R3, 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 17.1A1 for EX4300 and EX4600 Switches](#).

Release 17.1R3 New and Changed Features

Restoration Procedure Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (EX Series)**—Starting in Junos OS Release 17.1R3, 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 the 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](#).]

Release 17.1R2 New and Changed Features

- There are no new features or enhancements to existing features for EX Series in Junos OS Release 17.1R2.

Release 17.1R1 New and Changed Features

Hardware

- **New Routing Engine for EX9200 Switches**—Starting with Junos OS Release 17.1R1, EX9200 switches support the new Routing Engine EX9200-RE2.

[See [Routing Engine Module in an EX9200 Switch](#).]

- **New Configurations for EX9200 Switches**—Starting with Junos OS Release 17.1R1, EX9200 switches are available in the following configurations:

- EX9204-AC-BND2
- EX9204-RED3B-AC
- EX9204-RED3B-DC
- EX9204-BASE3B-AC
- EX9208-BASE3B-AC
- EX9208-RED3B-AC
- EX9208-RED3B-DC
- EX9214-BASE3B-AC
- EX9214-RED3B-AC
- EX9214-RED3B-DC

See

- [EX9204 Switch Configurations](#)
- [EX9208 Switch Configurations](#)
- [EX9214 Switch Configurations](#)

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

- **Access control and authentication (EX4300-EX4600 mixed VC)**—Starting with Junos OS Release 17.1R1, EX4600 switches operating in a mixed Virtual Chassis with EX4300 switches support controlling access to your network using 802.1X authentication and MAC RADIUS authentication.

802.1X authentication provides port-based network access control (PNAC) as defined in the IEEE 802.1X standard. Supported features include guest VLAN, private VLAN, server fail fallback, dynamic changes to a user session, RADIUS accounting, and configuration of port-filtering attributes on the RADIUS server using VSAs.

MAC RADIUS authentication is used to authenticate end devices independently of whether they are enabled for 802.1X authentication. You can permit end devices that are not 802.1X-enabled to access the LAN by configuring MAC RADIUS authentication on the switch interfaces to which the end devices are connected.

Access control features in a mixed EX4300-EX4600 Virtual Chassis are supported only on EX4300 ports.

This feature was previously supported in an “X” release of Junos OS.

[See [Access Control on a Mixed EX4300-EX4600 Virtual Chassis](#).]

Class of Service (CoS)

- **Support for classification of multdestination traffic (EX4300)**—Multidestination traffic includes BUM (broadcast, unknown unicast, and multicast) traffic and Layer 3 multicast traffic. By default on EX4300 Series switches, all multidestination traffic is classified to the **Mcast-BE** traffic class mapped to queue 8. Beginning with Junos OS Release 17.1R1, you can classify multidestination traffic to four different queues, queues 8-11, based on either the IEEE 802.1p bits or the DSCP IPv4/v6 bits. You can classify multidestination traffic by including the **multi-destination** statement at the **[edit class-of-service]** (to apply globally) or to an individual interface at the **[edit class-of-service interfaces interfaces-name]** hierarchy. Classification at an individual interface takes precedence over global classification.

[See [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers](#).]

- **Firewall filter with policer action as forwarding-class and loss priority (PLP) (EX4300 switches)**—Starting with Junos OS Release 14.1X53-D35 and Junos OS Release 17.1R1, on EX4300 switches you can configure the firewall with policer action as forwarding-class and loss priority (PLP). When the traffic hits the policer, PLP changes as per the action rule. The supported PLP designations are low, medium-low, medium-high, and high. You configure policer actions at the **[edit firewall]** hierarchy level.

[See [then \(Policer Action\)](#).]

High Availability (HA) and Resiliency

- **New options for the show vrrp track command (EX Series)**—Starting in 17.1R1, the **show vrrp track routes** command gives you the option to view all tracked routes. Another new option for the **show vrrp track** command, **all**, is equivalent to the already existing command **show vrrp track**.

[See [show vrrp track](#).]

Interfaces and Chassis

- **LLDP-MED power negotiation (EX4300 Switches)** —Starting with Junos OS Release 17.1R1, EX4300 switches support Link Layer Discovery Protocol Media Endpoint Discovery (LLDP-MED) power negotiation with high power (802.3at) devices. LLDP-MED power negotiation enables the PoE controller to dynamically allocate power to an interface based on the power required by the connected powered device.

[See [Power over Ethernet \(PoE\) User Guide for EX4300 Switches](#).]

- **Half-duplex link support (EX4300 switches)**—Starting with Junos OS 17.1R1, half-duplex communication is supported on all built-in network copper ports on EX4300 switches. *Half-duplex* is bidirectional communication, but signals can flow in only one direction at a time. *Full-duplex* communication means that both ends of the communication can send and receive signals at the same time. Half-duplex is configured by default on EX4300 switches. This feature was previously supported in an “X” release of Junos OS.

[See [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#).]

Junos OS XML API and Scripting

- **Support for Python language for commit, event, op, and SNMP scripts (EX Series)**—Starting in Junos OS Release 17.1R1, you can author commit, event, op, and SNMP scripts in Python on devices that include the Python extensions package in the software image. Creating automation scripts in Python enables you to take advantage of Python features and libraries as well as leverage Junos PyEZ APIs supported in Junos PyEZ Release 1.3.1 and earlier releases to perform operational and configuration tasks on devices running Junos OS. To enable execution of Python automation scripts, which must be owned by either root or a user in the Junos OS **super-user** login class, configure the **language python** statement at the **[edit system scripts]** hierarchy level, and configure the filename for the Python script under the hierarchy level appropriate to that script type. Supported Python versions include Python 2.7.x.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

Management

- **Support for adding non-native YANG modules to the Junos OS schema (EX Series)**—Starting in Junos OS Release 17.1R1, you can load custom YANG models on devices running Junos OS to add data models that are not natively supported by Junos OS but can be supported by translation. Doing this enables you to extend the configuration hierarchies and operational commands with data models that are customized for your operations. The ability to add data models to a device is also beneficial when you want to create device-agnostic and vendor-neutral data models that enable the same configuration or RPC to be used on different devices from one or more vendors. You can load custom YANG modules by using the **request system yang add** operational command.

[See [Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS.](#)]

OpenFlow

- **Support for OpenFlow v1.0 and v1.3.1 (EX4600 switches)**—Starting with Junos OS Release 17.1R1, EX4600 switches support OpenFlow v1.0 and v1.3.1. OpenFlow v1.0 enables you to control traffic in a network by adding, deleting, and modifying flows in the switch. You can configure one OpenFlow virtual switch and one active OpenFlow controller at the **[edit protocols openflow]** hierarchy level on each EX4600 switch in the network.

Also, OpenFlow v1.3.1 allows the action specified in one or more flow entries to direct packets to a base action called a group. The group action further processes these packets and assigns a more specific forwarding action to them. You can view groups that were added, modified, or deleted from the group table by using the **show openflow groups** command. You can view group statistics by using the **show openflow statistics groups** command.

[See [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS.](#)]

Software Installation and Upgrade

- **Support for unified in-service software upgrade (ISSU) (EX9200-6QS)**—Starting with Junos OS Release 17.1R1, you can perform a unified ISSU on the EX9200-6QS line card. 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 [Unified ISSU System Requirements](#).]

SEE ALSO

Changes in Behavior and Syntax	 36
Known Behavior	 40
Known Issues	 41
Resolved Issues	 47
Documentation Updates	 59
Migration, Upgrade, and Downgrade Instructions	 60
Product Compatibility	 61

Changes in Behavior and Syntax

IN THIS SECTION

- [General Routing](#) | [37](#)
- [High Availability \(HA\) and Resiliency](#) | [37](#)
- [MPLS](#) | [37](#)
- [Network Management and Monitoring](#) | [37](#)
- [Services Applications](#) | [39](#)
- [System Management](#) | [39](#)
- [User Interface and Configuration](#) | [39](#)
- [Virtual Chassis](#) | [39](#)

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 17.1R3 for the EX Series.

General Routing

- **Support for deletion of static routes when the BFD session goes down (EX Series)**—Starting with Junos OS Release 17.1R3, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session-down message.

[See [Enabling BFD on Qualified Next Hops in Static Routes for Route Selection](#).]

High Availability (HA) and Resiliency

- **In-service software upgrade (EX4600 switches)**—Starting with Junos OS Release 17.1R1, you cannot perform an ISSU from a Junos OS Release earlier than 17.1R1 to later Junos OS releases.

MPLS

- **Representation for OSPF DR node**—Up until version -10 of the BGP-LS draft, the OSPF DR node representation was ambiguous. One could represent DR node as 'AdvertisingRouterId-InterfaceIp Address' or 'InterfaceIp Address-1'. Junos OS used to follow 'InterfaceIp Address-1' format. Starting with version '-11' of the BGP-LS draft, the representation for OSPF DR node must be 'AdvertisingRouterId-InterfaceIp address'. Junos OS now follows the latest format.

Network Management and Monitoring

- **SNMP syslog messages changed (EX Series)**—Starting in Junos OS Release 17.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - Old message—**AgentX master agent failed to respond to ping. Attempting to re-register**
New message—**AgentX master agent failed to respond to ping, triggering cleanup!**
 - Old message—**NET-SNMP version %s AgentX subagent connected**
New message—**NET-SNMP version %s AgentX subagent Open-Sent!**

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

- **Update to SNMP support of apply-path statement (EX Series)**—Starting in Junos OS Release 17.1R2, SNMP implementation for the **apply-path** configuration statement supports only two lists:
 - **apply-path "policy-options prefix-list <list-name> <*>"**
This configuration has been supported from day 1.
 - **apply-path "access radius-server <*>"**
This configuration is supported as of Junos OS 17.1R2 release.

- **MIB loading errors fixed (EX Series)**—Starting in Junos OS Release 17.1R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:

- jnx-gen-set.mib
- jnx-ifotn.mib
- jnx-optics.mib

[See [MIB Explorer](#).]

- **Change in default log level setting (EX Series)**—In Junos OS Release, 17.1R3, 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 and logical interfaces.

After this change:

- IFD LinkUp -> LOG_NOTICE (because this is an important message but occurs less frequently)
- 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 and a non-default logical system (EX Series)**—Starting in Junos OS Release 17.1, 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](#).]

- **Need to reconfigure SNMPv3 configuration after upgrade (EX4600)**—Starting in Junos OS Release 17.1R2, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. This is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. It used to be that customers had to reconfigure SNMPv3 every time after a reboot. This problem was fixed. If you upgrade, you must still reconfigure SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3 again. To reconfigure SNMP v3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters.

[See [Configuring the Local Engine ID](#).]

Services Applications

- **Device discovery with device-initiated connection (EX Series)**—In Junos OS Release 17.1R1 and later releases, when you configure statements and options under the **[system services ssh]** hierarchy and commit the configuration, make sure that the system reaches a stable state before you commit any **outbound-ssh** configurations.

You use the device discovery feature in the Devices workspace to add devices to Junos Space Network Management Platform. By default, Junos Space manages devices by initiating and maintaining a connection to the device.

[See [Device Discovery Overview](#).]

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the **peers** option at the **[edit system commit]** hierarchy level is not supported in batch configuration mode.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (EX Series)**—Starting in Junos OS Release 17.1R1, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 17.1R1, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the show system schema module juniper-command output directory (EX Series)**—Starting in Junos OS Release 17.1, when you issue the **show system schema module juniper-command** operational command in the Junos OS CLI, the device places the generated output files in the current working directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1, the generated output files are placed in the **/var/tmp** directory.

Virtual Chassis

- Starting with Junos OS Release 17.1R1, EX9200 Virtual Chassis is no longer supported. You should not upgrade an existing EX9200 Virtual Chassis to Junos OS Release 17.1R1 or later. For deployments with EX9200 switches, we recommend planning or moving to MC-LAG or Junos Fusion Enterprise architectures instead of using a Virtual Chassis.

SEE ALSO

[New and Changed Features](#) | 30

Known Behavior	40
Known Issues	41
Resolved Issues	47
Documentation Updates	59
Migration, Upgrade, and Downgrade Instructions	60
Product Compatibility	61

Known Behavior

IN THIS SECTION

- General Routing | 40
- High Availability (HA) and Resiliency | 41
- Interfaces and Chassis | 41

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 for the EX Series.

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

General Routing

- On EX4600 switches, the Zero Touch Provisioning might take more than normal time to complete because TFTP might take a long time to fetch the required data. [PR980530](#)
- On an EX4300 or a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- 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 following events: 1. Process (for example, the pfx and dot1x process) restart or system restart 2. Link flaps [PR1294526](#)

High Availability (HA) and Resiliency

- **Unified ISSU incompatibility with VPLS dynamic profiles (EX Series)**—Using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 17.1R1 does not work if VPLS dynamic profiles are configured and enhanced subscriber management is not configured.

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

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

Interfaces and Chassis

- The same IP address can be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface is assigned with the identical address after commit. There are no warning messages seen during the commit; only syslog messages are seen indicating incorrect configuration. [PR1221993](#)

SEE ALSO

[New and Changed Features | 30](#)

[Changes in Behavior and Syntax | 36](#)

[Known Issues | 41](#)

[Resolved Issues | 47](#)

[Documentation Updates | 59](#)

[Migration, Upgrade, and Downgrade Instructions | 60](#)

[Product Compatibility | 61](#)

Known Issues

IN THIS SECTION

- [General Routing | 42](#)
- [Infrastructure | 44](#)
- [Junos Fusion Enterprise | 44](#)
- [Layer 2 Features | 45](#)

- Multicast | 45
- Network Management and Monitoring | 45
- Platform and Infrastructure | 46
- Spanning Tree Protocols | 46
- Virtual Chassis | 46

This section lists the known issues in hardware and software in Junos OS Release 17.1R3 for the EX Series.

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

General Routing

- On EX9200 switches, the analyzer configurations with analyzer input and output statements, containing members of the same VLAN or the VLAN itself, are not supported. With such configurations, packets can mirror in a loop, resulting in LU chip errors. As a workaround, use the **mirror-once** option if the input is for ingress mirroring. If it is for ingress and egress mirroring, configure the output interface as an access interface. [PR1068405](#)
- On EX9200 and EX4300 switches, 802.1X supplicants might not be reauthenticated by server fail fallback authentication after the server becomes reachable. [PR1157032](#)
- On EX4300 switches, when 802.1X single-supplicant authentication is initiated, multiple "EAP Request Id Frame Sent" packets might be sent. [PR1163966](#)
- 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 the connection is incorrectly configured as 40-Gigabit Ethernet. [PR1175918](#)
- On an EX9200-40XS line card, if you toggle the **MACsec encryption** option multiple times, encryption and protected MACsec statistics might be updated incorrectly. As a workaround, restart the line card. [PR1185659](#)
- On an EX9200-6QS line card, storm control might not work for multicast traffic. [PR1191611](#)
- On all platforms running Junos OS, the **file copy** CLI command uses `/var/home/<user>` as a temporary staging directory for a nonroot user, and uses `/var/tmp` for the root user. When you issue the **file copy user@x.x.x.x:/dir/ /var/tmp/** CLI command to copy a file to the device, and if the file you are trying to transfer is larger than the temporary staging directory size, the copy operation might fail. [PR1195599](#)
- On EX Series Virtual Chassis that support PoE, when the master Routing Engine member is rebooted, PoE devices connected to the master might not come back online after the reboot. As a workaround, when configuring PoE interfaces, use the **set poe interface all configuration** command instead of

configuring specific interfaces individually. To recover connections after observing this issue, disable and re-enable the ports affected by the issue. [PR1203880](#)

- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, fpc7 KERNEL/PFE APP=NH OUT OF SYNC: **error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none.** No service impact is seen in MPC2 and MPC3 type cards. [PR1205593](#)
- When a configuration that moves the Packet Forwarding Engine offline and another configuration that brings the Packet Forwarding Engine back online are committed in quick succession, the **out-of-synchronization** syslog errors might occur. Most of the time these are benign errors, but sometimes these errors might crash the Packet Forwarding Engine. [PR1232178](#)
- On EX Series switches except EX4300, EX4600, and EX9200, the switch cannot send DHCP option 2 when extended DHCP local server is configured. The switch sends DHCP option 2 incorrectly when a traditional DHCP server is configured. [PR1252437](#)
- On EX Series switches (except EX4300, EX4600, or EX9200), in a Virtual Chassis scenario, a LAG interface with **bpdud-block** disabled might go into a down state after the master Flexible PIC Concentrator (FPC) switch is rebooted. [PR1262703](#)
- When the em0 interface is unplugged, **Management Ethernet Links Down Alarms** might flap. [PR1271325](#)
- On EX Series switches (excluding EX4300, EX4600, and EX9200) that are in a DHCP relay with option 82 scenario, the jdhcpd memory might leak. The process will stop working with the following logged messages **/kernel: Process (3126, jdhcpd) attempted to exceed RLIMIT_DATA: attempted 131076 KB Max 131072 KB.** [PR1277433](#)
- Configuration statements that were allowed in Junos OS Release 12.3 are now invalid in Junos OS Releases 14.1X53 and 15.1. As a result, when you upgrade an EX Series switch from Junos OS Release 12.3 to Release 14.1X53 or Release 15.1R1, the switch might lose its configuration and run in a line-card mode or go into amnesiac mode. [PR1281947](#)
- The error in TQ-chip MACsec software a MACsec session might not reestablish after a physical link flap. Additionally, an FXPC core file might be generated because of this error. [PR1283314](#)
- When the EX4300-32F's 1/10 Gigabit Ethernet ports are reset, MACsec sessions might stay down and will not be able to reestablished. [PR1299484](#)
- Every load override and rollback operation increases the refcount by 1 and after it reaches the maximum value (65,535), an mgd crash will be observed and the session will get killed. When mgd crashes, the active lock might remain, preventing any further commits. [PR1313158](#)
- Some configurations that are valid for Junos OS Release 12.3 are not valid for Junos OS Release 15.1. When you try to upgrade from Junos OS Release 12.3 to Junos OS Release 15.1 with such configurations, the post-upgrade device goes into amnesiac(brick) mode. [PR1313501](#)

- EX4300 Virtual Chassis system might fail to register some jnxOperating SNMP OIDs related to the Routing Engine. This behavior is more likely if Virtual Chassis members 0 and 1 (FPC0 and FPC1) are not selected as Routing Engine. [PR1368845](#)
- On EX9200 Series platforms, if a packet-length keyword under the firewall filter is applied on the interface egress, the configuration is not committed, because of the commit-check failure. [PR1378901](#)
- In an aggregated interfaces and STP scenario, the STP does not work when the aggregated interface number is "ae1000" or above in QFX5000 and "ae480" or above in other QFX Series or EX Series platforms. Such interfaces remain in incorrect STP discarding state and might not forward packets. [PR1403338](#)
- On EX9200 device with MCLAG configuration and other features enabled, there is a loss of 20 seconds during the restart of routing daemon. This traffic loss varies with the configuration that is done. [PR1409773](#)

Infrastructure

- 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. For a permanent fix, upgrade the chassis to Junos OS Release 15.1R8 or later. [PR1360962](#)

Junos Fusion Enterprise

- On a Junos Fusion, when using LLDP, the power through MDI and extended power through MDI TLVs are not transmitted. [PR1105217](#)
- On a Junos Fusion Enterprise, Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) fast start does not work. [PR1171899](#)
- On a Junos Fusion Enterprise, when the satellite devices of a cluster are rebooted, the output of the CLI command **show chassis satellite** shows the port state of the cascade ports as **Present**. [PR1175834](#)
- On a Junos Fusion Enterprise, a loss of connectivity of the link connecting the standalone switch might cause conversion of the switch from Junos OS to SNOS to fail. [PR1232798](#)
- On a Junos Fusion Enterprise, the satellite device might not come online when the system is converted from cluster to non-cluster mode without accompanying topology changes. [PR1251790](#)
- When on the aggregation device Junos Fusion Enterprise is powered OFF or powered ON, it might take 6 to 30 seconds for the traffic to converge. [PR1257057](#)
- During Routing Engine switchover on a Junos Fusion Enterprise, the BUM traffic is duplicated to indirectly connected satellite devices. This occurs because there is no current support to notify the GRES event to indirectly connected satellite devices. [PR1298434](#)

- The ppm-lite process might generate a core file on the Fusion satellite devices. It is unexpectedly treating IEEE PORT VLAN ID TLV on LLDP packets as a DCBXv1.01 TLV. [PR1364265](#)
- Power over Ethernet (PoE) over Link Layer Discovery Protocol (LLDP) negotiation is not supported in Junos Fusion Enterprise (JFE) setup. The issue results in powering up failure when a device makes PoE over LLDP negotiation with the JFE. [PR1366106](#)

Layer 2 Features

- The eswd process might crash after doing a Routing Engine switchover in EX Series Virtual Chassis scenario. The crash happens because of disordered processing of a VLAN or its member by eswd and L2PT modules. As the order of processing does not remain the same every time, the crash is random across switchovers. [PR1275468](#)
- On EX Series platforms, except for EX4300/EX4600/EX9200, the multiple spanning tree protocol (MSTP) might not be able to detect the topology changes after a nonstop software upgrade (NSSU) process, which might lead to a packet loop. The topology change count is shown as 0 after that. **user@switch> show spanning-tree bridge msti 2 ++++++Output Snipped+++++ STP bridge parameters for MSTI 2 MSTI regional root : 8194.78:fe:3d:b1:e4:01 Hello time : 2 seconds Maximum age : 20 seconds Forward delay : 15 seconds Number of topology changes : 0 >>>>> showing 0 Topology change last recvd. from : 88:a2:5e:35:70:04 Local parameters Bridge ID : 8194.78:fe:3d:b1:e4:01 Extended system ID : 0 Internal instance ID : 2.** [PR1284415](#)
- The ERP route update fails during the addition of a new member to the ERP-configured VLAN. [PR1301595](#)
- The following syslog messages occur during ERPS PDU in ERPS setup every few minutes on ERPS owner: **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.** This message can be ignored. [PR1372422](#)

Multicast

- IGMP query packets might be duplicated between Layer 2 interfaces with IGMP snooping enabled. [PR1391753](#)

Network Management and Monitoring

- On EX Series switches except EX4300, EX4600, EX9200, when redundant trunk group (RTG) switchover, then the **/var/log/shadow.log** or **/var/log/shadow_debug.log** is rotated. And it might cause Packet Forwarding Engine process to crash. [PR1233050](#)
- The default syslog level is **LOG_NOTICE** in the default configuration. **SNMP_TRAP_LINK_UP** for the physical interface is logged as **LOG_INFO**. To help debug physical link up issues, **SNMP_TRAP_LINK_UP** events will be logged by default. [PR1287244](#)

Platform and Infrastructure

- On EX4300, EX4600, and QFX5100 switches, if a remote analyzer has an output IP address that is reachable through a route learned by BGP, the analyzer might be in a down state. [PR1007963](#)
- On EX4300-VC, if you configure a Q-in-Q S-VLAN interface with MC-LAG, when the backup EX4300 is acting as master, you might lose connection to the management IP address through the interface. As a result, management traffic might be dropped. [PR1131755](#)
- In case SP style configuration is used, deactivated or deleted one of the logical interfaces on LAG might cause traffic failure passing through the same LAG interface. [PR1422920](#)

Spanning Tree Protocols

- On EX Series switches except for EX4300, EX4600, and EX9200, the VoIP interfaces might be blocked by Rapid Spanning Tree Protocol (RSTP) if voice VLAN is running VLAN Spanning Tree Protocol (VSTP) and data VLAN is running RSTP respectively. [PR1306699](#)

Virtual Chassis

- If the linecard role FPC is removed from and rejoined to the Virtual Chassis, then the LAG interface on the master or backup switch is not reprogrammed in the rejoined FPC. [PR1255302](#)
- On EX Series switches except for EX4300/EX4600/EX9200, the packet drop might be seen during the failover or switchover from the master switch to backup switch in a Virtual Chassis. This is because of the delay in ARP update during the failover or switchover of the master Routing Engine (RE) . [PR1278214](#)

SEE ALSO

New and Changed Features	 30
Changes in Behavior and Syntax	 36
Known Behavior	 40
Resolved Issues	 47
Documentation Updates	 59
Migration, Upgrade, and Downgrade Instructions	 60
Product Compatibility	 61

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.1R3 | 47](#)
- [Resolved Issues: 17.1R2 | 55](#)
- [Resolved Issues: 17.1R1 | 57](#)

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: 17.1R3

Authentication and Access Control

- The LLDP-MED protocol cannot forward the correct POE class. [PR1296547](#)
- The dot1x process might stop authenticating if continuous reauthentication requests from dot1x clients do not get processed. [PR1300050](#)
- Dot1xd core file might be generated when dot1x interface is configured with EAP-PEAP as an authentication protocol. [PR1322891](#)

Class of Service (CoS)

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

Forwarding and Sampling

- Unexpected messages might be seen in logs. [PR1270686](#)

General Routing

- LACP does not work when MACsec is enabled. [PR1093295](#)
- The **storm control action-shutdown** configuration does not work as expected. [PR1130099](#)
- After an access is rejected, the dot1x process might crash because of a memory leak. [PR1160059](#)
- On an 802.1X-enabled interfaces, clients might not be able to access the network when they are connected or disconnected for a short period of time. [PR1230073](#)
- An LCD corruption issue is observed when an EX Series switch boots up. [PR1233580](#)

- The EOAM LFM adjacency on an MX Series MPC or an EX9200 might flap when an unrelated MIC, which 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](#)
- The management daemon (MGD) might crash after invoke a specific RPC, SSH/console need to be reconnected. [PR1271024](#)
- FPC connections might drop with the following syslog messages: **CHASSISD_MAIN_THREAD_STALLED: main chassis-control thread stalled for XXX sec -- exiting**. [PR1276605](#)
- The l2ald memory might leak for every IPv6 ND (Neighbor Discovery) message it receives from a peer MC-LAG and it is not freeing the allocated memory. [PR1277203](#)
- Interfaces configured with 100 Mbps speed might go down after reboot. [PR1283531](#)
- The **show security macsec statistics** command does not display the expected results. This is a MACsec issue. [PR1283544](#)
- The VLAN association does not get updated in the Ethernet switching table when the device is configured in single-supplicant mode. [PR1283880](#)
- The jhdcpd process might generate core files if **dhcpx6-security** is configured. [PR1287074](#)
- The dot1x process might crash on EX4300 switches when traffic is flooded and if a VLAN configuration commit is in progress. [PR1293011](#)
- On executing the **load replace terminal** command and attempting to replace the interface stanza might terminate the current CLI session and leave the user session hanging. [PR1293587](#)
- A warning message **Interface matching is supported only in a stand-alone** might be seen when had a commit operation with "from interface" condition in firewall filter on single device on a single device. [PR1296767](#)
- Network analytics does not transmit data. [PR1297535](#)
- The eswd process might generate a core file if **apply-groups** is configured under **interface-range**. [PR1300709](#)
- An l2ald crash might occur with no apparent trigger. [PR1302344](#)
- The **show snmp mib walk** command used for jnxMIMstMstiPortState does not display any output on an EX4600 running Junos OS Release 17.1R2. [PR1305281](#)
- Traffic loss is observed while performing NSSU. [PR1311977](#)
- PEM alarms and layer 2 circuit failures are observed on EX9200 Series switches. [PR1312336](#)
- The **dhcp-security** binding table might not get updated. [PR1312670](#)
- A memory leak is seen for the dot1xd process. [PR1313578](#)
- A vmcore file might be displayed and the device might reboot after the ICL is changed from an aggregated Ethernet interface to a physical interface. [PR1318929](#)

- The EX Series switches do not send RADIUS requests after the **interface-range** configuration is modified. [PR1326442](#)
- Traffic going through an aggregated Ethernet interface might be dropped because of a mastership change. [PR1327578](#)
- The rpd might crash on new backup Routing Engine when chassis switchover is triggered without GRES. [PR1330750](#)
- On an EX9200 switch, when the anchor FPC has no active child, BPDUs are not sent out on the other active child [VSTP/MSTP]. [PR1333872](#)
- On EX9200 switches, the MQSS error with error code 0x2203cb is observed. [PR1334928](#)
- The l2cpd crash might be seen in vstp scenario during Routing Engine switchover. [PR1341246](#)
- The statistics pfd process might generate a core file on an upgrade between certain releases. [PR1346925](#)
- The EX4600 switch detects a **LATENCY OVER-THRESHOLD** event with a wrong value. [PR1348749](#)
- After an FPC becomes online, the other FPC's CPU usage might go up to 100 percent and have a traffic loss for around 30 seconds. [PR1346949](#)
- The latency over-threshold event might be detected with an incorrect value. [PR1348749](#)
- The 40G interfaces might not forward traffic. [PR1349675](#)
- PPE errors async xtxn error when FPC is restarted or removed. [PR1350909](#)
- Commit error is observed when box is downgraded from Junos OS Release 18.2 or 18.3 to Junos OS Release 17.3R3. [PR1355542](#)
- FPM board status is missing in SNMP MIB walk result. [PR1364246](#)
- The l2cpd process might crash when configuring MVRP with private VLAN and RSTP interface all. [PR1365937](#)
- MAC refresh packet might not be sent out from the new primary link after RTG failover. [PR1372999](#)
- FPC might crash when flapping the output interface of analyzer or sampling. [PR1374861](#)
- ARP request packets might be sent out with 802.1Q VLAN tag. [PR1379138](#)
- The dot1x does not work with Microsoft NPS server. [PR1381017](#)
- 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-limit with persistent MAC is not working after reboot. [PR1400507](#)
- The l2cpd might crash if the VSTP traceoptions and VSTP VLAN all commands are configured. [PR1407469](#)
- PEM alarm for backup FPC remains on master FPC though backup FPC is detached from VC. [PR1412429](#)

High Availability (HA) and Resiliency

- GRES might fail to start because of the missing state acknowledgment message from the Package Forwarding Engine. [PR1236882](#)

Infrastructure

- On an EX4300 egress VLAN-based firewall filter on a Q-in-Q interface, after a switch reboot, firewall counters might not increment as expected. [PR1165450](#)
- When **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](#)
- On EX4300 switches, the aggregated Ethernet interface goes down when interface member VLAN is P-VLAN and LACP is enabled. [PR1264268](#)
- Unable to provide management when em0 interface of FPC is connected to another FPC Layer 2 interface of the same Virtual Chassis. [PR1299385](#)
- The **monitor interface traffic** does not display incoming ICMP packets. [PR1303947](#)
- The file system might be corrupted multiple times during an image upgrade or while committing an operation. [PR1317250](#)
- The upgrade might fail if bad blocks occur in the flash memory device or file system. [PR1317628](#)
- The PFC feature might not work on EX4600. [PR1322439](#)
- The ifinfo process generates a core file on an EX4600 Virtual Chassis. [PR1324326](#)
- Need support for archiving dmesg file `/var/run/dmesg.boot*`. [PR1327021](#)
- EX4600 might be sending packet with incorrect destination mac-address in MPLS php scenario. [PR1334929](#)

Interfaces and Chassis

- The MAC address between aggregated Ethernet interface and the member port might be inconsistent in rare conditions. [PR1272973](#)
- The interface might not work properly after the FPC restarts. [PR1329896](#)
- On EX4300-VC platforms, the MAC address assigned to an aggregated Ethernet member interface is not the same as that of its parent aggregated Ethernet interface upon master Routing Engine halt. [PR1333734](#)
- Packets might drop on ICL of MC-LAG peer where MC-LAG is up. [PR1345316](#)
- MC-LAG peer does not send ARP request to the host. [PR1360216](#)

Junos Fusion Enterprise

- Mirrored packets are dropped if analyzer output extended port is reachable through the ICL link. [PR1211123](#)
- 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](#)
- DHCP snooping entry is deleted after l2ald restarts. [PR1281824](#)
- 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](#)
- Dot1x might crash in a Junos Fusion setup with dual AD. [PR1303909](#)
- In Junos Fusion environment SD displays U-Boot on the LCD screen. [PR1304784](#)
- Two to three seconds of packet loss is seen every 5 minutes on Junos Fusion. [PR1320254](#)
- A satellite device does not recover PoE after the device is offline for more than 10 minutes and rejoins the AD. [PR1356478](#)
- The l2ald might crash when issuing **clear ethernet-switching table persistent-learning** command. [PR1409403](#)

Layer 2 Ethernet Services

- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. [PR1373807](#)

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 flaps. [PR1389695](#)

Layer 3 Features

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

MPLS

- On a EX4600 switch, unified ISSU is not supported with MPLS configuration. [PR1264786](#)

Network Management and Monitoring

- Some parts of SNMP MIB jnxBoxAnatomy hierarchy related to chassis components might be missing. [PR1278197](#)

Platform and Infrastructure

- Ping does not go through device after WTR timer expires in ERPS scenario. [PR1132770](#)
- Layer 3 protocol packets are not being sent out from the switch. [PR1226976](#)

- On an EX4300 switch, PXE unicast ACK packets are dropped. [PR1230096](#)
- Traffic is not forwarded through the GRE tunnel on an EX4300 in some cases. [PR1254638](#)
- The PoE interfaces flap on an EX4300 when one PSU is removed in power redundancy N+N mode. [PR1258107](#)
- Unexpected pfex restart is seen when the Routing Engine switches over. [PR1258863](#)
- The mismatch of VLAN IDs between a logical interface and the VLAN configuration might result in traffic being silently dropped or discarded. [PR1259310](#)
- On an EX4300 Virtual Chassis, a 10-Gigabit Ethernet VCP might not get a neighbor after a system reboot. [PR1261363](#)
- The EX Series switch does not respond to SSH or Telnet. [PR1266045](#)
- The IRB interface does not go down when the master chassis is rebooted or halted. [PR1273176](#)
- The DHCP discover/offer packets might cause memory leaks and jdncpd core files might be generated. [PR1273452](#)
- CPU utilization for pfex_junos usage might go high if DHCP relay packets are received continually. [PR1276995](#)
- Traffic loss might be observed for about 10 seconds if the master member FPC reboots. [PR1283702](#)
- IGMP report packets might be dropped on EX4300-VC with persistent learning enabled. [PR1285807](#)
- The FBF might not work properly after the feature is activated or deactivated. [PR1293581](#)
- Some packets might be dropped after GRE encapsulation on EX4300. [PR1293787](#)
- On EX4300 switches, some functions of IPv6 Router Advertisement Guard do not work. [PR1294260](#)
- On EX4300 switches, when unknown unicast ICMP packets are received by an interface, packets are routed, so TTL is decremented. [PR1302070](#)
- The FRU PSU removal and insertion traps might not get generated [PR1302729](#)
- Unknown IPv6 multicast traffic is dropped if **mld-snooping** is enabled. [PR1304345](#)
- Inconsistent IEEE P-bit marking in 802.1Q header for OSPF packets. [PR1306750](#)
- Multicast receiver connected to the EX4300 switch might not be able to get the multicast streaming. [PR1308269](#)
- The **Traceroute** command is not working for routing instances on EX9200 devices running on Junos OS Release 17.1R3. [PR1310615](#)
- Autonegotiation is not working as expected between EX4300 and SRX5800. [PR1311458](#)
- IGMP snooping might not learn multicast router interface dynamically. [PR1312128](#)
- The interface with 1-gigabit SFP transceiver might go down if **no-auto-negotiation** is configured. [PR1315668](#)

- The l2cpd core files are generated if the interface is disabled under VSTP and enabled under RSTP. [PR1317908](#)
- High latency might be observed between the master Routing Engine and other FPCs. [PR1319795](#)
- The VLAN might not be processed, which leads to improper convergence of the STP. [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 EX4300 platform. [PR1325816](#)
- The L2cpd might generate core files if **set protocols layer2-control mac-rewrite** or **set protocols layer2-control bpdud-block** is configured on any of the child members of a LAG. [PR1325917](#)
- Extra EAP request packets might be sent unnecessarily. [PR1328390](#)
- The SNMP trap message are always sent out with log about "Fan/Blower OK" on EX4300-VC switch. [PR1329507](#)
- On EX4300 Series switch, when the TCAM is being exhausted, the TCAM table filter continues to be programmed. [PR1330148](#)
- On EX4300 platforms, storm control logs stopped after adding RTG configuration. [PR1335256](#)
- The IGMP packets are forwarded out of the RTG backup interface. [PR1335733](#)
- 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 statement **accept-source-mac** does not work if MAC move limit is configured. [PR1341520](#)
- MSTP might not work normally after a commit is performed. [PR1342900](#)
- A firewall filter might not be programmed in Packet Forwarding Engine even though TCAM entries are available. [PR1345296](#)
- Packet drop might be seen on the logical tunnel interfaces lt-x/2/x or lt-x/3/x. [PR1345727](#)
- On EX4300/EX4600s 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](#)
- Firewall syslog is not sent 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 SFP-T transceiver might be still up after system halt. [PR1354857](#)
- The FPC would crash because of the memory leak caused by the VTEP traffic. [PR1356279](#)
- Interface flapping is seen on EX4300 switch. [PR1361483](#)

- On EX4300/EX4600 platforms, the l2ald process might crash in dot1x scenario. [PR1363964](#)
- Packet Forwarding Engine might crash if encountering frequent MAC move. [PR1367141](#)
- Traffic drops on Packet Forwarding Engine as "invalid L2 token" when protocol changes from VPLS to EVPN. [PR1368802](#)
- The LLDP TLV with incorrect switch port capabilities might be sent. [PR1372966](#)
- Traffic might be dropped and discarded with indirect next hop and load balancing. [PR1376057](#)
- Packet drops on interface if the statement **gigether-options loopback** is configured. [PR1380746](#)
- IRB interface does not turn down when master of Virtual Chassis is rebooted or halted. [PR1381272](#)
- On the EX4300 switch, if a loss priority value of high is set for multicast packets by a classifier at the ingress interface, the configuration is overridden by the storm-control filter. [PR1382893](#)
- EX4300 device chooses incorrect bridge-id as RSTP bridge-id. [PR1383356](#)
- The dhcp-security binding table might not be updated because of the renew request with '0.0.0.0' value in 'ciaddr'. [PR1394341](#)

Routing Protocols

- Observed mcsnoopd core file at
`__raise,abort,__task_quit,__task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler_internal`
and (enable_slip_detector=true, no_exit=true) at
`../../../../src/junos/lib/libtask/base/task_scheduler.c:275`. [PR1305239](#)
- The OSPF routes cannot be installed to the routing table until the **lsa-refresh** timer expires. [PR1316348](#)
- BGP peer is not established after Routing Engine switchover when graceful-restart and BFD enabled. [PR1324475](#)
- The IGMP snooping feature might be enabled unexpectedly. [PR1327048](#)
- In Junos OS EX Series platform, the stateless firewall filter ignores IPv6 extension headers (CVE-2019-0005). See <https://kb.juniper.net/JSA10905> for more details. [PR1346052](#)
- The parity errors in Layer 3 IPv4 table in the Packet Forwarding Engine memory might silently drop and discard the traffic. [PR1364657](#)
- Host destined packets with filter log action might reach the Routing Engine. [PR1379718](#)
- EX4300 might drop incoming IS-IS hello packets when IGMP or MLD snooping is configured. [PR1400838](#)

User Interface and Configuration

- On EX4300, the J-Web allows the configuration of **source-address-filter**, which is not the expected behavior. [PR1281290](#)

Virtual Chassis

- On an EX4300 switch , FRU removal or insertion of trap is not generated for non-master members (the switch in backup or linecard role).. [PR1293820](#)

Resolved Issues: 17.1R2

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

- On an EX4300 switch or Virtual Chassis with 802.1X (dot1x) enabled, in a scenario with more than 254 clients (supplicants), plenty of clients might be going to the server-reject VLAN and have limited access to the server-reject VLAN although the clients have correct credentials. For a few authenticated clients, the authentication method might be displayed as "Server-Reject" although the client was authenticated in the correct VLAN---that is, the data VLAN. [PR1251530](#)
- After configuration change with "commit", "dot1x" radius authentication request may not be sent out when having the "wait-for-acct-on-ack" configuration option within "access profile" [PR1252456](#)

EVPN

- If an EX9200 switch is configured as a PE router connected to a multihomed site in an EVPN/MPLS network, RPD core files might be created on the EX9200 when more than 255 logical interfaces from the same physical interface/ESI are added to the virtual switch instance configuration. Then some logical interfaces are removed from the ESI (that is, rollback of the configuration). [PR1251473](#)

Infrastructure

- On EX/QFX Series switches, if the switch was power cycled then some process (like jdhcp/lacp/lldpd...could be any other process) might stop working after rebooting. [PR1222504](#)

Interfaces and Chassis

- MPC might crash during ISSU from Junos OS Release 15.1R1 to a later release when QSFP/CXP/CFP2 optics are present. [PR1216924](#)

Junos Fusion Enterprise

- On a Junos Fusion Enterprise, the EX4300 running Junos OS Release 17.1R2 cannot be added as a satellite. [PR1267767](#)
- On a Junos Fusion Enterprise, restarting satellite-related daemons and L2 learning result in some MAC entries getting stuck in DLR state. [PR1268619](#)

Network Management and Monitoring

- On EX9208 switches, after ISSU, storm control is taking effect only after deletion and re-creation. [PR1151346](#)
- The following system error is logged: **JAM: Plugin installed for %s PIC.** [PR1189100](#)
- After the reboot of the EX4600 Virtual Chassis, authentication of SNMPv3 users fails due to the change of the local engine ID. [PR1256166](#)

Platform and Infrastructure

- On EX4300 switches, Layer 2 traffic is dropped in some cases. [PR1157058](#)
- When a policer with the action of loss of priority is applied to the lo0 interface, all ICMP packets might be dropped. [PR1243666](#)
- SFP+ might not be recognized after EX4300 reboot. [PR1247172](#)
- On EX9200 switches, if ISSU is used to upgrade Junos, it is possible that an unnecessary thread will run on an FPC after the upgrade procedure. This thread can potentially enter into a loop and trigger a stop of forwarding traffic on that particular FPC. [PR1249375](#)
- The egress PE device (EX4300) sends out LLDP frames toward the CE device with the destination MAC address of 01:00:0c:cd:cd:d0 which is a duplicated frame and rewritten by ingress (PE) device. [PR1251391](#)
- On EX4300 switches, traffic is not forwarded through the GRE tunnel in some cases. [PR1254638](#)
- After you deactivate IPv6 RA and commit the configuration, the feature is not deactivated. [PR1257697](#)
- The filter applied to the lo0 interface with policer action might break the BGP session. [PR1258038](#)
- On the EX4300-VC, FPC crash and PFEX core file might occur. [PR1261852](#)

Port Security

- MACsec connections are deleted randomly in some scenarios. [PR1234447](#)
- High CPU usage caused by fxpc can lead to MACsec session drops. [PR1247479](#)
- After MACsec link flaps, traffic stops forwarding across the MACsec link. [PR1269229](#)

Routing Protocols

- The BGP session might flap during ISSU, resulting in 40-50 seconds of dropped traffic. [PR1247937](#)

Spanning Tree Protocols

- RSTP interface all edge with the BPDU block configures all interfaces to go into BPDU block even if an interface is explicitly disabled under RSTP. [PR1266035](#)

Subscriber Access Management

- The authd process generates core files continuously during RADIUS authentication. [PR1241326](#)

System Management

- On MX Series and EX9200 platforms, an enhancement for implementing sensor-specific temperature thresholds is needed. [PR1199447](#)

Virtual Chassis

- When you add the EX4300 to the VCF, the following error message is seen: **ch_opus_map_alarm_id alarm ignored: object 0x7e reason.** [PR1234780](#)

Resolved Issues: 17.1R1

Authentication and Access Control

- A dot1xd core file is observed during CoA with Juniper-Switching-Filter. [PR1219538](#)
- Security certificates are lost after reboot or upgrade, and the following error is seen: **Unable to derive certificate from input .** [PR1237732](#)

Infrastructure

- BGP sessions are dropped on the EX4300 when sending BGP host-inbound traffic. [PR1090033](#)
- GRE counters are incrementing very slowly after deactivating and activating the gr- interface. [PR1183521](#)
- DHCP return packets received across a GRE tunnel are not forwarded to clients. [PR1226868](#)
- A timeout error occurs when using the **request system snapshot slice alternate** command. [PR1229520](#)

Interfaces and Chassis

- MPC might crash during ISSU from Junos OS Release 15.1R1 to a later release when QSFP/CXP/CFP2 optics are present. [PR1216924](#)
- Restarting the interface process causes traffic loss in aggregate Ethernet (ae) bundle in MC-LAG scenario. [PR1229001](#)

MPLS

- Virtual Chassis/Virtual Chassis Fabric-I2ckt: FXPC core file is seen when deactivating core interface on MPLS I2ckt configuration using IRB interface. [PR1242203](#)

Platform and Infrastructure

- Firewall filter is getting deleted when a new bind point is added. [PR1214151](#)
- EBGp packets with ttl=1 and non-EBGP packets with ttl=1 go to the same queue on EX4300. [PR1215863](#)
- The dcd process might crash with configuration of **set vlans xxx interface all**. [PR1221803](#)
- Frame with CFI / DEI bit set to 1 dropped on ingress L3 interface on EX4300 in Junos OS Release 14.1X53-D40.8 [PR1237945](#)
- EX4300: Too many interfaces after >request system zeroize in default configuration. [PR1238848](#)
- Stale dot1x state leads to packet loss on trunk links if they are converted from access to trunk. [PR1239252](#)
- Certain multicast traffic might cause network impact on EX4300 switch. [PR1244351](#)
- EX4300 connectivity issue with 10/100M and full/half duplex interface. [PR1249170](#)
- On Junos Fusion Enterprise, Power over Ethernet (PoE) telemetries do not work. [PR1112953](#)
- Changes made in PoE configuration during SD Offline state are not getting reflected once the SD is back Online. [PR1154486](#)
- On a Junos Fusion Enterprise, issues with ARP traffic might occur if the Junos Fusion topology exceeds the documented limit of 6,000 extended port interfaces. [PR1186077](#)
- FF reject tcp-reset does not work on IRB interface. [PR1219953](#)
- Issue with the **show** command occurs in single supplicant mode captive portal. [PR1240259](#)
- ELS Style - There is no command to enable DHCP snooping without having to enable other FHS features. [PR1245559](#)

Routing Protocols

- Hops through GRE tunnel endpoints are seen in traceroute. [PR1236343](#)

Virtual Chassis

- Repeated log message kernel: **%KERN-5: tcp_timer_keep: Dropping socket connection due to keepalive timer expiration** is seen on EX4300. [PR1209847](#)

SEE ALSO

[New and Changed Features | 30](#)

[Changes in Behavior and Syntax | 36](#)

[Known Behavior | 40](#)

[Known Issues | 41](#)

[Documentation Updates | 59](#)

[Migration, Upgrade, and Downgrade Instructions | 60](#)

[Product Compatibility | 61](#)

Documentation Updates

IN THIS SECTION

- [Documentation Updates | 59](#)

This section lists the errata and changes in Junos OS Release 17.1R3 for the EX Series switches documentation.

Documentation Updates

- **Support for deletion of static routes when the BFD session goes down (EX Series)**—Starting with Junos OS 17.1R3, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.
- Starting with Junos OS Release 17.1R1, EX9200 Virtual Chassis is no longer supported, and the EX9200 Virtual Chassis documentation has been archived. See [EX Series Documentation Archives](#). For deployments with EX9200 switches, we recommend planning or moving to MC-LAG or Junos Fusion Enterprise architectures instead of using a Virtual Chassis.

SEE ALSO

[New and Changed Features | 30](#)

[Changes in Behavior and Syntax | 36](#)

[Known Behavior | 40](#)

[Known Issues | 41](#)

[Resolved Issues | 47](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 60

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 15.1, 16.1 and 16.2 are EEOL releases. You can upgrade from Junos OS Release 15.1 to Release 16.1 or even from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

SEE ALSO

Known Behavior 40
Known Issues 41
Resolved Issues 47
Documentation Updates 59
Product Compatibility 61

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 61

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 30
Changes in Behavior and Syntax 36
Known Behavior 40
Known Issues 41
Resolved Issues 47
Documentation Updates 59

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- New and Changed Features | 63
- Changes in Behavior and Syntax | 68
- Known Behavior | 68
- Known Issues | 71
- Resolved Issues | 73
- Documentation Updates | 75
- Migration, Upgrade, and Downgrade Instructions | 75
- Product Compatibility | 82

These release notes accompany Junos OS Release 17.1R3 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

New and Changed Features

IN THIS SECTION

- [Release 17.1R3 New and Changed Features](#) | 63
- [Release 17.1R2 New and Changed Features](#) | 63
- [Release 17.1R1 New and Changed Features](#) | 64

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

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

Release 17.1R3 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 17.1R3.

Release 17.1R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 17.1R2.

Release 17.1R1 New and Changed Features

Hardware

- **Satellite device support (EX2300 and EX3400)**—Starting with Junos OS Release 17.1R1, you can configure EX2300 and EX3400 switches as satellite devices in a Junos Fusion Enterprise topology. The satellite device in a Junos Fusion topology is managed and configured by the aggregation device. Junos Fusion Enterprise uses EX9200 switches in the aggregation device role.

[See [Junos Fusion Enterprise Overview](#).]

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

- **Authentication and access control features (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.1R1, Junos Fusion Enterprise supports controlling access to the network by using the following features:
 - 802.1X authentication
 - MAC RADIUS authentication
 - Server-fail fallback
 - TACACS+ authentication
 - Central Web authentication
 - RADIUS-initiated changes to an authorized user session (RFC 3576)
 - Flexible authentication order
 - RADIUS accounting interim updates
 - Dynamic filtering with multiple filter terms using VSAs
 - EAP-PAP protocol support for MAC RADIUS authentication
 - RADIUS accounting attributes Client-system-Name, Framed-MTU, Session-timeout, Acct-authentic, Nas-port-ID, and Filter-ID

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

Class of Service (CoS)

- **Class of Service support (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.1R1, Junos Fusion Enterprise supports the standard Junos CoS features and operational commands. Each extended port on a satellite device is a logical extension to the aggregation device. Therefore, the default CoS policy on the aggregation device applies to each extended port. An EX9200 aggregation device supports the following CoS features for each extended port:
 - BA classifier
 - Multifield classifier

- Input and output policer
- Egress rewrite

The satellite devices support the following CoS features for each extended port:

- BA classifier
- Queuing and scheduling

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

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

Layer 2 Features

- **Support for Layer 2 Features (Junos Fusion Enterprise)**—Starting in Junos OS Release 17.1R1, the following features are supported:
 - **Storm control**—Monitor traffic levels and take a specified action when a defined traffic level (called the *storm control level*) is exceeded, preventing packets from proliferating and degrading service. You can configure the switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs. [See [Understanding Storm Control for Managing Traffic Levels on Switching Devices](#).]
 - **Persistent MAC learning (Sticky MAC)**—Configure persistent MAC addresses (also called *sticky MAC addresses*) to help restrict access to an access port by identifying the MAC addresses of workstations that are allowed access to a given port. Secure access to these workstations is retained even if the switch is restarted. [See [Understanding Persistent MAC Learning \(Sticky MAC\)](#).]
 - **MAC limiting**—Configure MAC limiting on an interface or a VLAN, and specify the action to take on the next packet the interface or the VLAN receives after the limit is reached. Limiting the number of MAC addresses protects the switch from flooding the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). [See [MAC Limiting](#).]
 - **Loop detection on extended ports**—Enable downstream loop detection on the satellite device to prevent accidental loops caused by miswiring or misconfiguration on the extended ports.
- **Support for MAC/PHY features on Junos Fusion Enterprise**—Starting with Junos OS Release 17.1R1, the following MAC/PHY features are supported on Junos Fusion Enterprise:
 - **Digital optical monitoring (DOM)**—You can run the **show interfaces diagnostics optics *interface-name*** command to display the DOM information. The information includes diagnostics data and alarms for Gigabit Ethernet optical transceivers.
 - **Energy Efficient Ethernet (EEE)**—EEE reduces the power consumption of physical layer devices during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when a link is idle. You can run the **set interfaces *interface-name* ether-options**

ieee-802-3az-eee command at the **[edit]** hierarchy level to enable energy efficiency at the Ethernet ports. You can view the EEE status by using the **show interfaces interface-name detail** command. By default, EEE is disabled on EEE-capable ports.

- **Jumbo frames**—You can configure jumbo frames by using the **set interfaces interface-name mtu 9216** command at the **[edit]** hierarchy level.
- **Medium-dependent Interface (MDI)**—By default, the auto MDI/MDI-X feature is enabled on Junos Fusion Enterprise. This feature eliminates the need for a cross-over cable to connect the LAN port to a port on another device, as the crossover function is automatically enabled, when required.

Multicast

- **Support for multicast traffic forwarding (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.1R1, multicast traffic forwarding is supported in Junos Fusion Enterprise. Multicast replication is supported only on the aggregation device. The aggregation device performs ingress multicast replication to a set of extended ports. On the satellite device, multicast traffic is received for each of the extended ports. The following scenarios are supported for both IPv4 and IPv6 traffic: Layer 2 multicast with VLAN flooding and Layer 3 multicast.

[See [Understanding Multicast Forwarding on a Junos Fusion Enterprise](#).]

Network Management and Monitoring

- **Network monitoring and analysis (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.1R1, sFlow monitoring and port mirroring and analyzers are supported in Junos Fusion Enterprise:
 - sFlow technology, which is a monitoring technology for high-speed switched or routed networks, randomly samples network packets and sends the samples to a monitoring station. You can configure sFlow technology to continuously monitor traffic at wire speed on all interfaces simultaneously.
 - Port mirroring and analyzers facilitate analyzing traffic on switches at the packet level. You configure port mirroring on a switch to send copies of unicast traffic to an output destination such as an interface, a routing instance, or a VLAN. You can configure an analyzer to define both the input traffic and output traffic in the same analyzer configuration. The input traffic to be analyzed can be traffic that enters or exits an interface, or traffic that enters a VLAN.

[See [Understanding sFlow Technology on a Junos Fusion Enterprise](#) and [Understanding Port Mirroring Analyzers on a Junos Fusion Enterprise](#).]

Port Security

- **Media Access Control Security (MACsec) support on extended ports (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.1R1, MACsec is supported on extended ports in a Junos Fusion Enterprise topology. 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. Enabling MACsec on extended ports in a Junos Fusion Enterprise topology provides secure communication between the satellite device and connected hosts.

[See [Understanding Media Access Control Security \(MACsec\)](#).]

- **Access security support (Junos Fusion Enterprise)**—Starting with Junos OS Release 17.1R1, the following access security features are supported in Junos Fusion Enterprise:
 - **DHCP snooping**—DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information, which it uses to build and maintain a database of valid IP-address-to-MAC-address (IP-MAC) bindings called the DHCP snooping database. Clients on untrusted ports are only allowed to access the network only if they are validated against the database.
 - **DHCPv6 snooping**—DHCP snooping for DHCPv6.
 - **Dynamic ARP inspection (DAI)**—DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons.
 - **IP source guard**—IP source guard prevents IP address spoofing by examining each packet sent from a host attached to an untrusted access interface on the switch. The IP address, MAC address, VLAN, and interface associated with the host are checked against entries stored in the DHCP snooping database. If the packet header does not match a valid entry in the DHCP snooping database, the packet is discarded.
 - **IPv6 source guard**—IP source guard for IPv6.
 - **IPv6 neighbor discovery (ND) inspection**—IPv6 ND inspection mitigates attacks based on Neighbor Discovery Protocol; by inspecting neighbor discovery messages and verifying them against the DHCPv6 snooping table.

[See [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity](#).]

SEE ALSO

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

[Known Behavior | 68](#)

[Known Issues | 71](#)

[Resolved Issues | 73](#)

[Documentation Updates | 75](#)

[Migration, Upgrade, and Downgrade Instructions | 75](#)

[Product Compatibility | 82](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [System Management | 68](#)

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 17.1R3 for Junos Fusion Enterprise.

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the **peers** option at the `[edit system commit]` hierarchy level is not supported in batch configuration mode.

SEE ALSO

New and Changed Features 63
Known Behavior 68
Known Issues 71
Resolved Issues 73
Documentation Updates 75
Migration, Upgrade, and Downgrade Instructions 75
Product Compatibility 82

Known Behavior

IN THIS SECTION

- [Junos Fusion Enterprise | 69](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 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

- While applying a loopback filter on aggregation devices in a Junos Fusion Enterprise, Callback Control Protocol (CBCP) packets might be filtered, which might cause CBCP sessions to be dropped and one of the satellite devices in a redundant pair to be in the SplitBrainDn state. To work around this issue, you can add a filter similar to the following to the existing set of loopback filters:

```
set firewall family inet filter accept-icl term accept-icl from source-address
10.0.0.0/30
set firewall family inet filter accept-icl term accept-icl from
destination-address 10.0.0.0/30
```

[PR1183680](#)

- On a Junos Fusion, when using LLDP, the "Power via MDI" and "Extended Power via MDI" TLVs are not transmitted. [PR1105217](#)
- In a Junos Fusion Enterprise topology with dual aggregation devices, firewall statistics are not synchronized across the aggregation devices. [PR1105612](#)
- On a Junos Fusion Enterprise, 'show ethernet-switching table' takes a few minutes to show entries when an extended port receives with MAC count set to 150K. [PR1117567](#)
- In a Junos Fusion Enterprise, Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) fast start does not work. [PR1171899](#)
- On a Junos Fusion Enterprise, when the satellite devices of a cluster are rebooted, the output of the CLI command **show chassis satellite** shows the port state of the cascade ports as **Present**. [PR1175834](#)
- In a Junos Fusion Enterprise, conversion of EX2300 and EX3400 switches from satellite devices to Junos OS devices cannot be performed from the aggregation device using the command **request chassis satellite install junos-package-name fpc-slot slot-id**. As a workaround, use the following procedure:
 1. If automatic satellite conversion is enabled for the satellite device's FPC slot ID, remove the FPC slot ID from the automatic satellite conversion configuration.

```
[edit]
```

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

For example, to remove FPC slot ID 101 from the Junos Fusion.

```
[edit]
```

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

If automatic satellite conversion is enabled for the FPC slot ID, the Junos Fusion tries to convert the device back into a satellite device later in this procedure.

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

2. Log in to the satellite software (SNOS) on the switch to be converted back to Junos OS and use the following sequence of commands to install the Junos package:

```
#####  
dd bs=512 count=1 if=/dev/zero of=/dev/sda  
echo -e "o\nn\np\nl\n\n\nw" | fdisk /dev/sda  
mkfs.vfat /dev/sda1  
fw_setenv target_os  
reboot  
#####  
>>Get to the loader prompt  
#####  
loader> install --format tftp://<tftp server>/<Junos package name>
```

PR1213023

- In a Junos Fusion Enterprise, conversion of an EX2300 switch from Junos OS to satellite software (SNOS) takes 13-14 minutes. [PR1213853](#)
- In a Junos Fusion Enterprise, analyzer output is not supported for the aggregation device native interfaces. As a workaround, use RSPAN to capture analyzer output for the aggregation device. [PR1214596](#)
- In a Junos Fusion Enterprise, EX3400 and EX2300 operating as satellite devices might take longer time to re-converge from single-home to dual-home cluster due to a hardware limitation, compared to an EX4300 switch operating as a satellite device. [PR1226366](#)
- In a Junos Fusion Enterprise with dual aggregation devices, duplicate multicast packets are observed until L3 convergence happens between the aggregation devices, which might take a few seconds. [PR1231101](#)
- In a Junos Fusion Enterprise, a delay might result from moving a satellite device from cluster to non-cluster mode and vice versa. [PR1231678](#)
- Loss of connectivity of the link connecting the standalone switch might lead to conversion failure from Junos OS to satellite software (SNOS). As a workaround, reboot the standalone switch again to restart the conversion process. [PR1232798](#)

- In a Junos Fusion Enterprise, in order to use a non-default port as a clustering port in a clustering port policy, the policy must include at least one port that is a default uplink/clustering port for that platform. [PR1241808](#)
- In a Junos Fusion Enterprise, a satellite device might not come online when it is converted from cluster to non-cluster mode without accompanying topology changes. As a workaround, ensure the configuration of satellite devices matches the wiring topology: non-cluster devices should not be connected to other clustered devices by means of default or configured clustering/uplink ports. [PR1251790](#)
- In Junos Fusion Enterprise, when 802.1X authentication is configured in single-secure mode, a firewall counter is created for the default discard term in addition to the configured term. [PR1254503](#)
- In a Junos Fusion Enterprise it can take 6 to 30 seconds for the traffic to converge when the aggregation device is powered off or powered on. [PR1257057](#)
- During RE switchover on a Junos Fusion Enterprise, the BUM traffic is duplicated to indirectly connected satellite devices. This is because there is no current support to notify the GRES event to indirectly connected satellite devices. [PR1298434](#)

SEE ALSO

[New and Changed Features | 63](#)

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

[Known Issues | 71](#)

[Resolved Issues | 73](#)

[Documentation Updates | 75](#)

[Migration, Upgrade, and Downgrade Instructions | 75](#)

[Product Compatibility | 82](#)

Known Issues

IN THIS SECTION

- [Junos Fusion Enterprise | 72](#)

This section lists the known issues in hardware and software in Junos OS Release 17.1R3 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 a Junos Fusion Enterprise that has enabled PoE for all extended ports, the **show poe interface** command output displays the PoE administrative status as Enabled for non-PoE-capable interfaces. [PR1150955](#)
- On a Junos Fusion Enterprise, control packets from the aggregation device - including ping and DHCP packets - might not be forwarded to hosts connected to extended ports when the cascade ports on the aggregation device are down. [PR1173212](#)
- In a Junos Fusion Enterprise, restarting satellite processes from the aggregated device might not work. As a workaround, use the following commands to get the process ID and restart the process:

```
user@aggregated-device> request chassis satellite shell-command fpc-slot <slot-id> "ps -aef |
grep <process> | grep -v grep"
```

```
user@aggregated-device> request chassis satellite shell-command fpc-slot <slot-id> "kill -9
<process-id>"
```

Processes details:

amd—api-management-daemon

lcmd—chassis-management-daemon

dpd—discovery-and-provisioning-daemon

spfe—packet-forwarding-engine

ppman—ppman

ppman-lite—ppman-lite

[PR1244166](#)

- In a Junos Fusion Enterprise, backup link information might not be displayed in the output of the **show chassis satellite** command if cluster configuration is deleted and then added again on a single aggregated device. As a workaround, delete and then add configuration on both aggregated devices. [PR1247633](#)
- The ppm-lite process might generate a core file on the Fusion satellite devices. It is unexpectedly treating the IEEE PORT VLAN ID TLV on LLDP packets as a DCBXv1.01 TLV. [PR1364265](#)

SEE ALSO

Changes in Behavior and Syntax	68
Known Behavior	68
Resolved Issues	73
Documentation Updates	75
Migration, Upgrade, and Downgrade Instructions	75
Product Compatibility	82

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.1R3 | 73
- Resolved Issues: 17.1R2 | 74
- Resolved Issues: 17.1R1 | 74

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: 17.1R3

Junos Fusion Enterprise

- Mirrored packets are dropped if analyzer output extended port is reachable via the ICL link. [PR1211123](#)
- 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](#)
- DHCP Snooping entry is deleted after l2ald restart. [PR1281824](#)
- VRRP has a split-brain in dual autodiscovery Junos Fusion. [PR1293030](#)
- An aggregation device without a cascade port cannot reach hosts over the ICL link if they are authenticated by 802.1X in a different VLAN than the default (manually assigned) VLAN. [PR1298880](#)
- 802.1X authentication might fail in a Junos Fusion setup. [PR1299532](#)

- 802.1X authentication might fail in a Junos Fusion setup with dual aggregation devices. [PR1303909](#)
- In Junos Fusion environment SD displays U-Boot on the LCD screen. [PR1304784](#)
- Packet loss of 2-3 secs is seen every 5 minutes on Junos Fusion. [PR1320254](#)
- A satellite device does not recover PoE after the device is offline for more than 10 minutes and rejoins the aggregation device. [PR1356478](#)

Resolved Issues: 17.1R2

Junos Fusion Enterprise

- EX4300 with Junos OS Release 17.1R1 cannot be converted to satellite mode. [PR1267767](#)
- In Junos Fusion Enterprise, for **show ethernet-switching table**, a few entries are stuck in DLR state after **l2-learning** restart. [PR1268619](#)

Resolved Issues: 17.1R1

Junos Fusion Enterprise

- For Junos Fusion Enterprise, PoE telemetry is not working. [PR1112953](#)
- Changes made in PoE configuration during SD Offline state are not getting reflected once the SD is back Online. [PR1154486](#)
- Some ARPs are not resolving on Spirent when you exceed 6000 extended ports. [PR1186077](#)
- Traffic loss is seen after rebooting a satellite device in a satellite device cluster. [PR1168820](#)
- SNMP trap for satellite device reboot is not sent. [PR1182895](#)
- LLDP might stop working if manually deactivated and reactivated. [PR1188254](#)
- SDPD core files might be generated during conversion of EX2300/EX3400 cluster from JUNOS OS to SNOS. [PR1239915](#)

SEE ALSO

[New and Changed Features | 63](#)

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

[Known Behavior | 68](#)

[Known Issues | 71](#)

[Documentation Updates | 75](#)

[Migration, Upgrade, and Downgrade Instructions | 75](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R3 for Junos Fusion Enterprise documentation.

SEE ALSO

[New and Changed Features | 63](#)[Changes in Behavior and Syntax | 68](#)[Known Behavior | 68](#)[Known Issues | 71](#)[Resolved Issues | 73](#)[Migration, Upgrade, and Downgrade Instructions | 75](#)[Product Compatibility | 82](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 76](#)
- [Upgrading from Junos OS Release 16.1 to 17.1 in a JUNOS OS Fusion Enterprise System | 78](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 79](#)
- [Preparing the Switch for Satellite Device Conversion | 79](#)
- [Converting a Satellite Device to a Standalone Switch | 80](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 81](#)
- [Downgrading from Release 17.1 | 81](#)

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 [Installation and Upgrade Guide](#).

To download and install Junos OS:

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

9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **junos-install** package on the aggregation device.

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

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

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

All other customers, use the following commands.

```
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.

NOTE: After you install a Junos OS **junos-install** 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 **junos-install** package that corresponds to the previously installed software.

Upgrading from Junos OS Release 16.1 to 17.1 in a JUNOS OS Fusion Enterprise System

When the Junos Fusion Enterprise System includes clustered devices, use the following procedure to first upgrade the clustered devices to SNOS 3.0R1 and then upgrade the aggregation device from 16.1 to 17.1.

1. Enable hop-by-hop forwarding for control-traffic the on aggregation device using VTY commands.

- a. Start a shell on the aggregated device:

```
user@aggregation-device> start shell
```

- b. For each FPC which has cascade ports, start a VTY session. For example:

```
root@aggregation-device% vtty fpc1
```

- c. At the VTY prompt, enter the following command:

```
FPC1(aggregation-device vty)# set jnh ep stack-hostpath 0
```

2. Enable hop-by-hop forwarding for control-traffic on all satellite devices in a cluster.

```
user@aggregation-device> request chassis satellite shell-command vty -c 'test sd-cluster  
hop-to-hop enable' range fpc-start fpc-end
```

3. Update the satellite device cluster to the new image, which must be SNOS 3.0R1 or higher.

```
user@aggregation-device> request system software add upgrade-group cluster-upgrade-group  
image-location
```

4. Confirm all satellite devices are upgraded to the new image.

```
user@aggregation-device> show chassis satellite upgrade-group upgrade-group-name
```

5. Upgrade the aggregation device to the 17.1 image.

```
user@aggregation-device> request system software add aggregation-device-package-name
```

6. To complete the upgrade, reboot the system, including all satellite devices and aggregation device.

- To reboot the satellite devices:

```
user@aggregation-device> request chassis satellite reboot range fpc-start fpc-end
```

- To reboot the aggregation device:

```
user@aggregation-device> request system reboot
```

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

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

NOTE: The following conditions must be met before a Junos switch that is running Junos OS Release 17.1 can be converted to a satellite device when the action is initiated from the aggregation device:

- The Junos switch can only be converted to SNOS 3.0 and later.
- The Junos switch must be either set to factory default configuration using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

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

1. Log in to the device using the console port.

2. Clear the device:

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Switch

In the event that you need to convert a satellite device to a standalone device, you will need to install a new Junos OS software package on the satellite device and remove 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 15.1, 16.1 and 16.2 are EEOL releases. You can upgrade from Junos OS Release 15.1 to Release 16.1 or even from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Downgrading from Release 17.1

Junos Fusion Enterprise is first supported in Junos OS Release 16.1R1, although you can downgrade a standalone EX9200 switch to earlier Junos OS releases.

NOTE: It is not recommended to downgrade the aggregation device from 17.1 to 16.1 if there are cluster satellite devices in the setup.

To downgrade a Junos Fusion Enterprise from Junos OS Release 17.1 to 16.1, you must first downgrade the satellite software version on the satellite devices from 3.0R1 to 2.0R1.

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

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

After the satellite devices are downgraded to satellite software 2.0R1, they will not show as being online until the aggregation device is downgraded to 16.1.

2. Downgrade the aggregation device. Follow the procedure for upgrading, but replace the 17.1 **junos-install** package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

New and Changed Features 63
Changes in Behavior and Syntax 68
Known Behavior 68
Known Issues 71
Resolved Issues 73
Documentation Updates 75
Product Compatibility 82

Product Compatibility

IN THIS SECTION

- [Hardware and Software Compatibility | 82](#)
- [Hardware Compatibility Tool | 82](#)

Hardware and Software Compatibility

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 63
Changes in Behavior and Syntax 68
Known Behavior 68
Known Issues 71
Resolved Issues 73
Documentation Updates 75
Migration, Upgrade, and Downgrade Instructions 75

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

●	New and Changed Features 84
●	Changes in Behavior and Syntax 85
●	Known Behavior 86
●	Known Issues 86
●	Resolved Issues 87
●	Documentation Updates 88
●	Migration, Upgrade, and Downgrade Instructions 89
●	Product Compatibility 97

These release notes accompany Junos OS Release 17.1R3 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 17.1R3 New and Changed Features | 84](#)
- [Release 17.1R2 New and Changed Features | 84](#)
- [Release 17.1R1 New and Changed Features | 84](#)

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

Release 17.1R3 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 17.1R3.

Release 17.1R2 New and Changed Features

- There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 17.1R2.

Release 17.1R1 New and Changed Features

Junos Fusion

- **Support for satellite device clustering**—Starting in Junos OS Release 17.1R1, Junos Fusion Provider Edge supports satellite device clustering. Satellite device clustering enables you to connect up to 10 satellite devices into a single cluster, and to connect the satellite device cluster to the aggregation device as a single group instead of as individual satellite devices.

[See [Understanding Satellite Device Clustering in a Junos Fusion.](#)]

- **Support for LLDP-MED with VoIP integration**—Starting in Junos OS Release 17.1R1, Junos Fusion Provider Edge supports Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) with VoIP integration on the extended ports of satellite devices in a VoIP network. LLDP-MED with VoIP integration is an extension of LLDP that is used to support device discovery of VoIP telephones and to create location databases for these telephone locations.

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

SEE ALSO

Changes in Behavior and Syntax 85
Known Behavior 86
Known Issues 86
Resolved Issues 87
Documentation Updates 88
Migration, Upgrade, and Downgrade Instructions 89
Product Compatibility 97

Changes in Behavior and Syntax

IN THIS SECTION

- [System Management | 85](#)

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 17.1R3 for Junos Fusion Provider Edge.

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the **peers** option at the `[edit system commit]` hierarchy level is not supported in batch configuration mode.

SEE ALSO

New and Changed Features 84
Known Behavior 86
Known Issues 86
Resolved Issues 87
Documentation Updates 88
Migration, Upgrade, and Downgrade Instructions 89
Product Compatibility 97

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 for Junos Fusion Provider Edge.

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

SEE ALSO

New and Changed Features	 84
Changes in Behavior and Syntax	 85
Known Issues	 86
Resolved Issues	 87
Documentation Updates	 88
Migration, Upgrade, and Downgrade Instructions	 89
Product Compatibility	 97

Known Issues

There are no known issues in the Junos OS Release 17.1R3 for Junos Fusion Provider Edge.

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

SEE ALSO

New and Changed Features	 84
Changes in Behavior and Syntax	 85
Known Behavior	 86
Resolved Issues	 87
Documentation Updates	 88
Migration, Upgrade, and Downgrade Instructions	 89
Product Compatibility	 97

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.1R3 | 87](#)
- [Resolved Issues: 17.1R2 | 88](#)
- [Resolved Issues: 17.1R1 | 88](#)

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: 17.1R3

Junos Fusion Provider Edge

- Transit unicast traffic gets sometimes discarded when it passes through different logical interfaces of the same extended port on a satellite device. [PR1264900](#)
- In a Junos Fusion, the **show interfaces diagnostics optics satellite** command does not display any output. [PR1327876](#)
- SSH key-based authentication fails after reboot if **chassis satellite-management** is configured. [PR1344392](#)
- Laser receives power of the extended ports that is higher than the output power of the peer link. [PR1358007](#)

Resolved Issues: 17.1R2

Junos Fusion Provider Edge

- LACP on extended ports does not come up after GRES Routing Engine switchover event on MX104.[PR1262674](#)

Resolved Issues: 17.1R1

Junos Fusion

- Junos OS to satellite conversion initiated from aggregation device must use SNOS 3.0, SNOS 1.0R5, or SNOS 2.0R2.[PR1249877](#)

SEE ALSO

New and Changed Features 84
Changes in Behavior and Syntax 85
Known Behavior 86
Known Issues 86
Documentation Updates 88
Migration, Upgrade, and Downgrade Instructions 89
Product Compatibility 97

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R3 for Junos Fusion Provider Edge documentation.

SEE ALSO

New and Changed Features 84
Changes in Behavior and Syntax 85
Known Behavior 86
Known Issues 86
Resolved Issues 87
Migration, Upgrade, and Downgrade Instructions 89

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 89
- Upgrading an Aggregation Device with Redundant Routing Engines | 92
- Preparing the Switch for Satellite Device Conversion | 92
- Converting a Satellite Device to a Standalone Device | 93
- Upgrading an Aggregation Device | 96
- Upgrade and Downgrade Support Policy for Junos OS Releases | 96
- Downgrading from Release 17.1 | 96

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 [Installation and Upgrade Guide](#).

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

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

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

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

For upgrades from Junos Release 14.2 and earlier:

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

All other upgrades:

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.1R3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

Customers with EX4300 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-ex-4300-14.1X53-D30.3-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-qfx-5-14.1X53-D30.3-domestic-signed.tgz
```

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

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

```
[edit]
```

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

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

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

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

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

[edit]

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

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

[edit]

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

```
[edit]  
user@aggregation-device> request chassis satellite install  
/var/tmp/install-media-pxe-qfx-5-14.1X53-D30.3.tgz fpc-slot 101
```

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

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

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

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

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

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 17.1R3, you must also upgrade your satellite device to Satellite Device Software version 3.0R1.

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 15.1, 16.1 and 16.2 are EEOL releases. You can upgrade from Junos OS Release 15.1 to Release 16.1 or even from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Downgrading from Release 17.1

To downgrade from Release 17.1 to another supported release, follow the procedure for upgrading, but replace the 17.1 **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	 84
Changes in Behavior and Syntax	 85
Known Behavior	 86
Known Issues	 86
Resolved Issues	 87
Documentation Updates	 88
Product Compatibility	 97

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility](#) | 97

Hardware Compatibility

Hardware Compatibility

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

To determine the features supported on MX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. See the [Feature Explorer](#).

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features	 84
Changes in Behavior and Syntax	 85
Known Behavior	 86
Known Issues	 86
Resolved Issues	 87
Documentation Updates	 88
Migration, Upgrade, and Downgrade Instructions	 89

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

IN THIS SECTION

●	New and Changed Features	 99
●	Changes in Behavior and Syntax	 125
●	Known Behavior	 142
●	Known Issues	 150
●	Resolved Issues	 169
●	Documentation Updates	 232
●	Migration, Upgrade, and Downgrade Instructions	 234
●	Product Compatibility	 241

These release notes accompany Junos OS Release 17.1R3 for the MX Series routers. 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 17.1R3 New and Changed Features | 100](#)
- [Release 17.1R2 New and Changed Features | 101](#)
- [Release 17.1R1 New and Changed Features | 103](#)

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

Release 17.1R3 New and Changed Features

Interfaces and Chassis

- **Enhancement to increase the threshold of corrected single-bit errors (MPC7E, MPC8E, MPC9E on MX Series)**—In Junos OS Release 17.1R3, the threshold of corrected single-bit errors 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 on 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 errordetected in HMC 0 - Total count 25

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

[See [Alarm Overview](#).]

Restoration Procedures and Failure Handling

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (MX Series)**—In Junos OS Release 17.1R3, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, an automatic device recovery mode exists to help recover the system should it go into amnesiac mode. The system retries to boot with the saved rescue configuration. In this circumstance, the system that boots displays a banner **Device is in recovery mode** in the CLI (in both operational and configuration modes). In earlier releases, there is no automatic process to recover from amnesiac mode: Instead, a user with load and commit permission has to log in using the console and fix the issue in the configuration before the system can reboot.

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

Subscriber Management and Services

- **RADIUS attributes added to LNS messages (MX Series)**—Starting in Junos OS Release 17.1R3, 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)
- **Controlling search behavior for address allocation from linked pools (MX Series)**—Starting in Junos OS Release 17.1R3, 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 17.1R2 New and Changed Features

Interfaces and Chassis

- **Enhancement to ambient-temperature statement (MX Series)**—In Junos OS Release 17.1R2 and later, the default ambient temperature is set at 40° C on MX480, MX960, MX2010, and MX2020 Universal Routing Platforms. You can override ambient temperature by setting the temperature at 55° C or 25° C.

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

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

Routing Protocols

- **IGP cost calculation for next-hop-based dynamic tunnels(MX Series)**—Starting in Junos OS Release 17.1R2, IGP cost calculation is supported for next-hop-based dynamic tunnels. In multihoming networks with next-hop-based GRE or UDP tunnel, rpd chooses the best path by calculating IGP metrics. However, in single-homed networks, rpd installs the tunnel composite next hop in the Packet Forwarding Engine without any IGP cost calculation.

In earlier Junos OS releases, BGP preferred a path with the lowest router ID, which was not cost effective. When multiple PE devices advertise the same route, BGP did not take into account the IGP cost to those devices. This new feature allows BGP to choose an IGP path with the lowest metric and set up a tunnel to a PE device with the lowest cost. Note that in the absence of IGP connectivity, Junos OS does not install the advertised routes in the Packet Forwarding Engine or create a dynamic tunnel.

Subscriber Management and Services

- **Configurable grace period for unresponsive RADIUS servers (MX Series)**—Starting in Junos OS Release 17.1R2, you can use the **timeout-grace** statement at the **[edit access radius-options]** hierarchy level to configure a grace period that determines when an unresponsive RADIUS authentication server is marked as down or unreachable. When the server fails to respond to any of the attempts made for an authentication request, it times out, the time is noted, and the grace period begins. If the server is unresponsive for subsequent authentication requests, the grace period is checked each time the server times out. When the check determines that the grace period has expired, the server is marked as down or unreachable.

You can configure the grace period in the range 0 through 30 seconds; the default is 10 seconds. Use a short grace period to declare servers unavailable sooner and direct requests to available servers. Use a long grace period to give unresponsive servers more opportunities to respond.

In earlier releases, the grace period is 10 seconds and is not configurable.

- **Support for excluding tunnel attributes from RADIUS Access-Request messages (MX Series)**—Starting in Junos OS Release 17.1R2, you can use the **exclude** statement at the **[edit access profile profile-name radius attribute]** hierarchy level to exclude the following tunnel attributes from RADIUS Access-Request messages in addition to the previously supported Accounting-Start, and Accounting-Stop messages:
 - **acct-tunnel-connection**—RADIUS attribute 68, Acct-Tunnel-Connection
 - **tunnel-assignment-id**—RADIUS attribute 82, Tunnel-Assignment-Id
 - **tunnel-client-auth-id**—RADIUS attribute 90, Tunnel-Client-Auth-Id
 - **tunnel-client-endpoint**—RADIUS attribute 66, Tunnel-Client-Endpoint
 - **tunnel-medium-type**—RADIUS attribute 65, Tunnel-Medium-Type
 - **tunnel-server-auth-id**—RADIUS attribute 91, Tunnel-Server-Auth-Id
 - **tunnel-server-endpoint**—RADIUS attribute 67, Tunnel-Server-Endpoint
 - **tunnel-type**—RADIUS attribute 64, Tunnel-Type

Release 17.1R1 New and Changed Features

Hardware

- **Support for ODU path delay measurement for 100-Gigabit DWDM OTN MIC and 100-Gigabit DWDM OTN PIC (MX Series)**—Starting in Junos OS Release 17.1R1, Junos OS supports ODU path delay measurement for the 100-Gigabit DWDM OTN MIC (MIC3-100G-DWDM) on MPC3E (MX-MPC3E-3D) and MPC3E-NG (MPC3E-3D-NG) on MX Series routers and for the 100-Gigabit Ethernet DWDM OTN PIC (PTX-5-100G-WDM) on PTX3000 and PTX5000 routers. Delay is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (**DM**) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end (local interface).

To enable delay measurement, first enable looping of the delay measurement pattern at the remote interface by including the **remote-loop-enable** statement at the `[edit interfaces interfacename otn-options odu-delay-management]` hierarchy level. Then, measure the delay by including the **start-measurement** statement at the `[edit interfaces interfacename otn-options odu-delay-management]` hierarchy level. Use the **stop-measurement** statement to stop measuring the delay. To disable looping of the delay measurement pattern at the remote interface, use the **no-remote-loop-enable** statement.

- **1-port 100-Gigabit DWDM OTN MIC with CFP2 (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS release 17.1R1, support is provided for the 1-port 100-Gigabit Ethernet dense wavelength division multiplexing (DWDM) optical transport network (OTN) MIC (MIC3-100G-DWDM) with CFP2 analog coherent optical (CFP2-ACO) pluggable optics on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG). The 100-Gigabit Ethernet DWDM OTN MIC supports the following features:
 - Transparent transport of 100-Gigabit Ethernet signals with optical channel transport unit, OTU4 (V) framing
 - Dual-polarization quadrature phase shift keying (DP-QPSK) modulation with coherent receiver and soft-decision forward error correction (SD-FEC) for long-haul and metro applications
 - International Telecommunication Union (ITU)-standard OTN performance monitoring and alarm management
 - Extensive optical, digital signal processing (DSP), and bit error ratio (BER) performance monitoring statistics for the optical link

[See [100-Gigabit DWDM OTN MIC with CFP2-ACO](#) and [Configuring OTN Interfaces on MIC3-100G-DWDM MIC](#).]

Class of Service (CoS)

- **Copy ToS bits from incoming IP header to outer GRE IP header (MX Series with MPCs)**—Starting in Junos OS Release 17.1R1, you can set GRE tunnel interfaces to copy the ToS bits (DSCP value) from the incoming IPv4 header to the outer GRE IP header for transit traffic. You can set this at the individual GRE interface level by including the **copy-tos-to-outer-ip-header-transit** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level, or globally by including the **copy-tos-to-outer service-type ([gre] | [mt])** statement at the **[edit chassis]** hierarchy level.

You can also now rewrite the DSCP/IP precedence value in both the inner and outer headers with the **rewrite rules ([dscp] | [inet-precedence]) default protocol ([inet-both] | [inet-outer])** statement at the **[edit class-of-service interfaces interface-name]** hierarchy level.

[See [Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header.](#)]

EVPNs

- **Support for multihoming in an MSAN scenario with EVPN (MX Series routers with MPCs)**—Starting in Junos OS Release 17.1R1, the EVPN multihoming feature enables you to connect a customer site to two or more provider edge (PE) devices to provide redundant connectivity. A customer edge (CE) device can be multihomed to different PE devices or the same PE device. A redundant PE device can provide network service to the customer site as soon as a failure is detected. Thus, EVPN multihoming helps maintain EVPN service and traffic forwarding to and from the multihomed site in case of network failures such as:
 - Failure of the link between PE device to CE device
 - PE device failure
 - MPLS-reachability failure between the local PE device and a remote PE device

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

- **Support for VPWS with EVPN signaling mechanisms (MX Series)**—The Ethernet VPN (EVPN)-virtual private wire service (VPWS) network provides a framework for delivering the VPWS with EVPN signaling mechanisms. The VPWS with EVPN signaling mechanisms supports single-active or all-active multihoming capabilities and inter-autonomous system (AS) options associated with BGP-signaled VPNs. Starting with Junos OS Release 17.1R1, the **vpws-service-id** statement identifies the endpoints of the EVPN-VPWS network based on the **local** and **remote** identifiers configured on the provider edge (PE) routers in the network. These endpoints are autodiscovered by BGP and are used to exchange the service labels (learned from the respective PE routers) that are used by autodiscovered routes per EVPN instance (EVI).

Use the **show evpn vpws-instance** command to verify the routes and interfaces of the VPWS instance of the EVPN.

[See [Overview of VPWS Service with EVPN Signaling Mechanisms.](#)]

- **Support for inter-data center connectivity over pure Layer 3 network with EVPN (MX Series routers with MPCs)**—Starting in Junos OS Release 17.1R1, the control plane EVPN Type-5 supports IP prefix

for inter-subnet connectivity across data centers. The data packet is sent as the L2 Ethernet frame encapsulated in the VXLAN header over the IP network across the data centers to reach the tenant through the connectivity provided by the EVPN Type-5 IP prefix route.

[See [EVPN Type-5 Route with VXLAN encapsulation for EVPN/VXLAN.](#)]

- **Support for LACP in EVPN active-active multihoming (MX Series routers with MPCs)**—Starting with Junos OS Release 17.1R1, 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 an 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.
 - Include the **service-id number** statement at the **[edit switch-options]** hierarchy.

[See [Example: Configuring LACP for EVPN Active-Active Multihoming.](#)]

- **Support for IPv6 over IRB interfaces with EVPN (MX Series routers with MPCs)**—Starting in Junos OS Release 17.1R1, IPv6 addresses are supported on IRB interfaces with EVPN using the Neighbor Discovery Protocol (NDP). The following capabilities are introduced for IPv6 support with EVPN:
 - IPv6 addresses on IRB interfaces in master routing instances
 - Learning IPv6 neighborhood from solicited NA message
 - NS and NA packets on the IRB interfaces are disabled from network core
 - Virtual gateway addresses are used as Layer 3 addresses
 - Host MAC-IP synchronization for IPv6

You can configure the IPv6 addresses in the IRB interface at the **[edit interfaces irb]** hierarchy level.

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

- **Support for VLAN bundle service for EVPN (MX Series)**—Starting in Junos OS Release 17.1R1, Junos OS supports the VLAN bundle service for EVPN. The VLAN bundle service maps multiple VLAN IDs to one EVPN instance. Because a separate instance for each VLAN ID is not needed, this feature lowers the control plane overhead on the router by reducing the number of EVPN instances.

[See [VLAN Bundle Service for EVPN](#).]

General Routing

- **PHY timestamping support for MIC-3D-20GE-SFP-EH, MIC-3D-20GE-SFP-E, and built-in 10-Gigabit Ethernet ports (MX104)**—Starting with Junos OS Release 17.1R1, timestamping at the physical layer, also known as PHY timestamping, is supported on MIC-3D-20GE-SFP-EH, MIC-3D-20GE-SFP-E, and the built-in 10-Gigabit Ethernet ports on MX104 routers. PHY timestamping is the timestamping of the IEEE 1588 event packets at the physical layer. Timestamping the packet at the physical layer eliminates the noise or the packet delay variation (PDV) that is introduced by the Packet Forwarding Engine.

To enable PHY timestamping on MX104 routers, include the **phy-timestamping** statement at the **edit [protocols ptp]** hierarchy level.

[See [PHY Timestamping](#).]

- **Support for PTP over Ethernet, hybrid mode, and G.8275.1 profile (MPC5E and MX104)**—Starting in Junos OS Release 17.1R1, MPC5E and MX104 support the following features:
 - **PTP over Ethernet**—PTP over Ethernet enables effective implementation of packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks. PTP over Ethernet uses multicast addresses for communication of PTP messages between the slave clock and the master clock.
 - **Hybrid mode**—In hybrid mode, the synchronous Ethernet equipment clock (EEC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP.
 - **G.8275.1 profile**—G.8275.1 is a PTP profile for applications that require accurate phase and time synchronization. It supports the architecture defined in ITU-T G.8275 to enable the distribution of phase and time with full timing support and is based on the second version of PTP defined in IEEE 1588. You can configure the G.8275.1 profile by including the **profile-type g.8275.1** statement at the **[edit protocols ptp]** hierarchy level.

[See [Profile Type](#).]

High Availability (HA) and Resiliency

- **ISSU Feature Explorer**—The unified ISSU Feature Explorer is an interactive tool that you can use to verify your device's unified ISSU compatibility with different Junos OS releases.

[See [ISSU Feature Explorer](#).]

- **Support for unified ISSU on MX Series routers and MX Series Virtual Chassis with MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, MPC2E-3D-NG-Q, and MPC5E (MX240, MX480, MX960, MX2010,**

and MX2020)—Starting with Junos OS Release 17.1R1, unified in-service software upgrade (ISSU) is supported on MX Series routers and MX Series Virtual Chassis with MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, MPC2E-3D-NG-Q, and MPC5E.

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

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

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

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

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

[See [Protocols and Applications Supported by MX240, MX480, MX960, MX2010, and MX2020 MPC2E](#), [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC3E](#), and [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5Es](#).]

- **Unified in-service software upgrade support for 100-Gigabit DWDM OTN MIC (MX960)**—Starting with Junos OS Release 17.1R1, unified in-service software upgrade (unified ISSU) is supported for the 1-port 100-Gigabit Ethernet dense wavelength division multiplexing (DWDM) OTN MIC (MIC3-100G-DWDM) on MX960 routers with MPC3E (MX-MPC3E-3D) and MPC3E-NG (MX-MPC3E-NG).

Unified ISSU is a process to upgrade the system software with minimal disruption of transit traffic and no disruption of the control plane. You can use unified ISSU only to upgrade to a later version of the system software. When unified ISSU completes, the new system software state is identical to that of the system software when the system upgrade is performed through a cold boot.

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

- **New options for the show vrrp track command (MX Series)**—Starting with Junos OS Release 17.1R1, the **show vrrp track routes** command gives you the option to view all tracked routes. Another new option for the **show vrrp track** command, **all**, is equivalent to the already existing command **show vrrp track**.

[See [show vrrp track](#).]

Interfaces and Chassis

- **Getting load-balancing hash result information (MX Series)**—Starting in Junos OS Release 17.1R1, you can get the details for load-balancing hash results. You can get information for up to three levels of load balancing.

To get load-balancing results for routed IPv4, IPv6, and other L3 traffic, use the **show forwarding-options load-balance ingress-interface <interface-name> family <family-type> source-address <src-IP> destination-address <dest-IP> transport-protocol <transport-protocol> source-port <src-port> destination-port <dest-port> tos <TOS>** command. To get load-balancing results for raw packet dumps, use the **show forwarding-options load-balance ingress-interface <interface-name> family <family-type> packet-dump <pkt-dump>** command.

[See [show forwarding-options load-balance](#).]

- **Support for PPP-TCC encapsulation on MIC-3D-16CHE1-T1-CE (MX Series)**—Starting in Junos OS Release 17.1R1, Junos OS supports PPP-TCC encapsulation on channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE). PPP-TCC encapsulation is used for circuits with different media on either sides of the connection.
- **Removing the native VLAN ID from untagged traffic (MX Series)**—Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the **no-native-vlan-insert** statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.

[See [Sending Untagged Traffic Without VLAN ID to Remote End](#).]

- **Inline MultilinkPPP, Multilink FrameRelay, and Multilink FrameRelay End-to-End for time-division multiplexing WAN interfaces (MX Series)**—The ability to provide bundling services through the Packet Forwarding Engine without requiring a PIC or DPC by using inline Multilink PPP (MLPPP), Multilink Frame Relay (MLFR) FRF.16, and MLFR end-to-end FRF.15 for time-division multiplexing (TDM) WAN interfaces was first rolled out in Junos OS Release 14.1. Starting in Junos OS Release 17.1R1, this feature is also supported on the following MPCs: MPC5E (MX240, MX480, MX960, MX2010, and MX2020 routers) and MPC6E (MX2010 and MX2020 routers). Support includes multiple links on the same bundle as well as multiclass extensions for MLPPP. You can enable bundling services without additional DPC slots, freeing the slots for other MICs.

[See [Inline MLPPP for WAN Interfaces Overview](#), [Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End \(FRF.15\) for WAN Interfaces](#),] and [[Example: Configuring Inline Multilink Frame Relay \(FRF.16\) for WAN Interfaces](#).]

- **Enhancement to policer configuration (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure the MPC to take a value in the range 0 through 5 for the policer tick byte by using the **policer-limit** statement at the **[edit chassis]** hierarchy level. If this statement is not configured, the policer tick byte can take values up to 7, which is the default behavior. You can use the **set chassis policer-limit** command to enable this feature.

You must restart the MPC or the router for the changes to take effect.

- **Support for inline Two-Way Active Measurement Protocol (TWAMP) server and client on MPC7E (MX240, MX480, MX960)**—Starting in Junos OS Release 17.1R1, MX Series routers with MPC7E cards support the inline Two-Way Active Measurement Protocol (TWAMP) control-client and server for transmission of TWAMP IPv4 UDP probes between the session-sender (control-client) and the session-reflector (server). The TWAMP control-client and server can also work with a third-party server and control-client implementation.

TWAMP is an open protocol for measuring network performance between any two devices that support TWAMP. To configure the TWAMP server, specify the logical interface on the service PIC that provides the TWAMP service by including the `twamp-server` statement at the:[**edit interfaces si-fpc/pic/ port unit logical-unit-number rpm**] hierarchy level. To configure the TWAMP client, include the `twamp-client` statement at the:[**edit interfaces si-fpc/pic/ port unit logical-unit-number rpm**] hierarchy level.

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

- **Support for frame relay inverse ARP on MIC-3D-16CHE1-T1-CE (MX Series)**—Starting in Junos OS Release 17.1R1, Junos OS supports frame relay inverse ARP requests on channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE). You can configure MIC-3D-16CHE1-T1-CE to operate in either T1 or E1 mode. By default, all the ports operate in T1 mode.

[See [Configuring Inverse Frame Relay ARP](#).]

Layer 2 Features

- **Enhancement to MAC limit function (MX Series with MPCs)**—Starting in Junos OS Release 17.1R1, the handling of a burst of packets with new source MAC addresses is improved to reduce resource use and processing time. In earlier releases, new source MAC addresses are learned and placed in the MAC table even after the limit is exceeded. The Routing Engine later deletes the MAC address entries that are over the limit.

Now, the learning limit configured with the **interface-mac-limit** statement for new source MAC addresses is enforced at all levels: global, bridge domain, and VPLS. The MAC table is not updated with any new addresses after the limit has been reached. When any static MAC addresses are configured, the learning limit is the configured limit minus the number of static addresses.

[See [Limiting MAC Addresses Learned from an Interface in a Bridge Domain](#) and [Limiting the Number of MAC Addresses Learned from Each Logical Interface](#).]

Layer 2 VPN

- **Support for ETH-SLM and ETH-DM on aggregated Ethernet interfaces and LAG members on MPCs (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure ITU-T Y.1731 standard-compliant Ethernet synthetic loss measurement (ETH-SLM) and Ethernet delay measurement (ETH-DM) capabilities on aggregated Ethernet interfaces and LAG members on all MX Series MPCs. These ITU-T Y.1731 OAM services or performance-monitoring techniques can be measured in on-demand mode (triggered through the CLI) or proactive mode (triggered by the iterator application).

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. ETH-DM provides fine control to operators for triggering delay measurement on a given service and can be used to monitor service-level agreements (SLAs).

Management

- **Support for Junos Telemetry Interface sensor for queue depth statistics (MX Series)**—Starting with Junos OS Release 17.1R1, you can configure a Junos Telemetry Interface sensor that exports queue depth statistics for ingress and egress queue traffic. Telemetry data is exported directly from the line card. You can also apply one or more regular expressions to filter data. Include the **resource /junos/system/linecard/qmon/** statement at the **[edit system services analytics sensor sensor-name]** hierarchy level. Only UDP streaming of data is supported. gRPC streaming of queue depth statistics is not currently supported. Only MPC7E, MPC8E, and MPC9E are supported.

[See [sensor \(Junos Telemetry Interface\)](#).]

- **gRPC support for the Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.1R1, the Junos Telemetry Interface supports using a set of gRPC remote procedure call interfaces to provision sensors, subscribe to, and receive telemetry data. gRPC is based on an open source framework and provides secure and reliable transport of data. Use the **telemetrySubscribe** RPC to specify telemetry parameters and stream data for a specified list of OpenConfig commands paths. Telemetry data is generated as Google protocol buffers (gpb) messages in a universal key/value format. If your Juniper Networks device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download

the Network Agent package, which provides the interfaces to manage gRPC subscriptions. The package is available on the **All Junos Platforms** software download URL on the Juniper Networks webpage.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Support for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.1R1, the Junos Telemetry Interface enables you to export telemetry data from supported interface hardware. Sensor data, such as interface events, are sent directly to configured collection points without involving polling. On MX Series routers, only MPC1 through MPC9E are supported. For sensors that stream data through the User Datagram Protocol, all parameters are configured at the **[edit services analytics]** hierarchy level. For sensors that stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Not all hardware and sensors are supported in those previous releases.

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

MPLS

- **Support for subscriber management over MPLS pseudowire logical interface on virtual chassis (MX Series)**—Starting with Junos OS Release 17.1R1, MPLS pseudowire logical interface for subscriber management is supported on virtual chassis. The functionality of Ethernet interface types such as ae/ge/xe, works on virtual chassis.
- **Support for Layer 2 services provisioning on the services side of the pseudowire service logical interface (MX Series)**—Starting with Junos OS Release 17.1R1, Layer 2 services provisioning such as bridge domain or VPLS instance is possible on the services side of the pseudowire service logical interface anchored to logical tunnel interface.

Prior to Junos OS Release 17.1R1, Layer 2 encapsulations and features such as Spanning Tree Protocol (STP), VLAN and many more could not be configured on pseudowire service on the service logical interface.

[See [Layer 2 Services Provisioning on Services Side of Pseudowire Service Interface Overview](#).]

- **Support for port mirroring on pseudowire subscriber logical interface (MX Series)**—Starting with Junos OS Release 17.1R1, port mirroring is supported on the services side of an MPLS pseudowire subscriber logical interface.

You can configure pseudowire service interface in the same way as the logical interface or physical interface. The main purpose of port mirroring on pseudowire service interface is to allow configurations of pseudowire service interface as a mirrored interface at Layer 2 and Layer 3 levels as supported by firewall filters.

- **Support for LDP pseudowire auto-sensing (MX Series)**—Starting with Junos OS Release 17.1R1, Label Distribution Protocol (LDP) pseudowire auto-sensing addresses zero-touch provisioning. LDP pseudowire auto-sensing enables pseudowire headend termination to be dynamically provisioned rather than statically configured. Hence, it is referred to as zero-touch provisioning.

In Junos OS, pseudowire headend termination on service nodes is supported through the use of pseudowire service logical interfaces and physical interfaces. This approach is considered as superior in scalability to the old logical tunnel interface based approach, due to its capability of multiplexing and

demultiplexing subscribers or customers over a single pseudowire. Currently, the creation and deletion of the pseudowire service logical interfaces, pseudowire service physical interfaces, Layer 2 circuits, and Layer 2 VPNs for pseudowire headend termination rely on static configuration. This is not considered as ideal from the perspective of scalability, efficiency, and flexibility, especially in a network where each service node might potentially host a large number of pseudowires.

[See [LDP Pseudowire Auto-Sensing Overview](#).]

- **Order-aware abstract hops for MPLS LSPs (MX Series)**—Junos OS Release 17.1R1 introduces abstract hops, which are user-defined router clusters or groups that can be sequenced and used for setting up a label-switched path (LSP), similar to real-hop constraints.

The router groups are created using constituent lists that include constituent attributes, which is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs). Ordering among the router groups that satisfy the specified constituent attributes is achieved by using operational qualifiers in the abstract-hop definition.

A path can use a combination of real and abstract hops as constraints. To configure abstract hops, you need to create constituent lists with traffic engineering attributes, include the lists in the abstract-hop definition, and define path constraints that use the abstract hops.

[See [Abstract Hops For MPLS LSPs Overview](#) and [Example: Configuring Abstract Hops for MPLS LSPs](#).]

- **Support for extension of pseudowire redundancy condition to logical Interfaces (MX Series)**—Starting with Junos OS Release 17.1R1, pseudowire redundancy condition is supported on MPLS pseudowire subscriber logical interface. This is similar to the pseudowire redundancy feature for mobile backhaul by using the logical tunnel paired (lt-) interfaces.

The primary or backup pseudowire is terminated at the provider edge routers (ps0.0) and the corresponding pseudowire (ps0.1 to ps0.n) service logical interfaces connected to Layer 3 domain by configuring those service logical interfaces in the Layer 3VPN routing instances. There is a Layer 2 circuit across MPLS access node and provider edge with the pseudowire service on transport logical interface (ps0.0) as the local interface of Layer 2 circuit terminating at the provider edge device.

[See [Extension of Pseudowire Redundancy Condition Logic to Pseudowire Subscriber Logical Interface Overview](#).]

- **Increased scaling values for MPLS-over-UDP tunnels (MX Series routers with MPCs/MICs)**—The next-hop-based dynamic UDP tunnels are referred to as MPLS-over-UDP tunnels, and support the creation of a tunnel composite next hop for every dynamic tunnel created. Starting in Junos OS Release 17.1, the limit for the maximum number of next-hop-based dynamic MPLS-over-UDP tunnels that can be created on an MX series router with MPCs or MICs is increased. This provides additional scaling advantage for the total number of IP tunnels that can be created on the router.

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

Multicast

- **Rate sensitive upstream multicast hop (UMH) selection for multicast VPN source-active routes (MX Series)**—Starting in Junos OS Release 17.1R1, you can use the traffic rate on the ingress PE to trigger the egress PE to use an alternative UHM. Two new commands are introduced to support this feature, **min-rate** and **dampen**.

Use this feature, for example, to ensure that egress PEs only receive Source-Active A-D route advertisements from ingress PEs that are receiving traffic at or above a specified rate. Rather than advertising the Source-Active A-D route immediately upon learning of the S,G, the ingress PE waits the time specified in the **dampen** command for the traffic rate to remain above the **min-rate** before it sends Source-Active A-D route advertisements. If the rate drops below the threshold, the Source-Active A-D route is withdrawn. These new commands can be found at the **[edit routing-instancesinstance-name protocols mvpn mvpn-mode spt-only source-active-advertisement]** hierarchy level.

[See [min-rate](#) and [dampen](#).]

Network Management and Monitoring

- **Support for hrProcessorTable object (MX Series)**—Starting in Junos OS Release 17.1R1, support is provided for the **hrProcessorTable** object (object id: 1.3.6.1.2.1.25.3.3) described in the RFC2790, *Host Resources MIB*. The **hrProcessorTable** object provides the load statistics information per CPU for multi-core devices.

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

- **Get and walk support for SNMP Timing MIB objects (MX104)**—Starting in Junos OS Release 17.1R1, the get and walk functionality is supported for the following SNMP timing MIB objects:
 - jnxPtpClass
 - jnxPtpGmId
 - jnxPtpAdvClockClass
 - jnxPtpUtcOffset
 - jnxPtpUtcValid
 - jnxPtpOperationalSlaves
 - jnxPtpOperationalMaster
 - jnxPtpServoState
 - jnxPtpSlaveOffset
 - jnxTimingFrequencyTraceability
 - jnxTimingTimeTraceability
 - jnxClksyncQualityCode
 - jnxClksyncQualityCodeStr

- `jnxClksyncIIndex`
- `jnxClksyncIntfName`
- `jnxClksyncSynceQualityTable`
- `jnxClksyncSynceQualityIntfIndex`
- `jnxClksyncSynceQualityValue`
- `jnxClksyncSynceQualityIntfName`

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

- **Support for `mplsL3VpnIfConfTable` object (MX Series)**— Starting in Junos OS Release 17.1R1, support is provided for the `mplsL3VpnIfConfTable` object (object id: 1.3.6.1.2.1.10.166.11.1.2.1) described in RFC 4382, *MPLS/BGP Layer 3 Virtual Private Network (VPN) MIB*. The `mplsL3VpnIfConfTable` object represents the Layer 3 VPN enabled interfaces that are associated with a specific Virtual Routing and Forwarding (VRF) instance and shows the bitmask values of the supported protocols. The `mplsL3VpnIfConfTable` object creates entries for the interfaces that are associated with the VRF instances. If an interface is later removed from a VRF instance, the corresponding entry in the `mplsL3VpnIfConfTable` object gets deleted. To view details of the `mplsL3VpnIfConfTable` object, use the `show snmp mib walk mplsL3VpnIfConfTable` command.

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

- **Port mirroring enhancements (MX Series)**—Starting in Junos OS Release 17.1R1, the port mirroring feature supports several new enhancements:
 - Packet mirroring for both ingress and egress directions on subscriber IFLs
 - Support for the encapsulation of mirrored packets onto per-subscriber L2TP tunnels
 - Support for the removal of S-VLAN tags from mirrored packets

[See [Configuring Protocol-Independent Firewall Filter for Port Mirroring](#).]

OpenFlow

- **Destination MAC address rewrites for OpenFlow (MX80, MX240, MX480, and MX960)**—Some types of network equipment that function as routers accept and handle packets only if the destination MAC address in the packet is the same as the MAC address of the Layer 3 interface on which the packet is received. To interoperate with these routers, connected devices must also be able to rewrite the destination MAC address of an incoming packet. Starting with Junos OS Release 17.1R1, an OpenFlow controller can configure an MX Series router that supports OpenFlow to rewrite the destination MAC address of an incoming packet.

[See [Understanding How the OpenFlow Destination MAC Address Rewrite Action Works](#).]

Operation, Administration, and Maintenance (OAM)

- **Enhanced scale support for MIPs per chassis (MX Series with MPCs)**—Starting in Junos OS Release 17.1R1, Junos OS supports 8000 maintenance association intermediate points (MIPs) per chassis for bridge

domain and VPLS domain interfaces. Increasing the number of MIPs per chassis for specific domains enables effective Ethernet OAM deployment in scaling networks. To support the increased number of MIPs, configure the network services mode on the router as **enhanced-ip**. If you do not configure the network services mode, then Junos OS supports only 4000 MIPs.

[See [Configuring Maintenance Intermediate Points \(MIPs\)](#).]

- **Support for sender ID TLV**—Starting with Junos OS Release 17.1R1, you can configure Junos OS to send the sender ID TLV along with the packets. The sender ID TLV is an optional TLV that is sent in continuity check messages (CCMs), loopback messages, and Link Trace Messages (LTMs), as specified in the IEEE 802.1ag standard. The sender ID TLV contains the chassis ID, which is the unique, CFM-based MAC address of the device, and the management IP address, which is an IPv4 or an IPv6 address.

You can enable Junos OS to send the sender ID TLV at the global level by using the **set protocols oam ethernet connectivity-fault-management sendid-tlv** and the **set protocols oam ethernet connectivity-fault-management sendid-tlv send-chassis-tlv** commands. If the sender ID TLV is configured at the global level, then the default maintenance domain, maintenance association, and the maintenance association intermediate point (MIP) half function inherit this configuration.

The sender ID TLV, if configured at the hierarchy levels mentioned above, takes precedence over the global-level configuration.

NOTE: The sender ID TLV is supported only for 802.1ag PDUs and is not supported for performance monitoring protocol data units (PDUs).

[See [Junos OS Support for Chassis ID TLV](#).]

- **CFM enhancement for interoperability during unified ISSU (MX Series on MPC1, MPC2, MPC2-NG, MPC3-NG, MPC5, and MPC6 cards)**—Starting in Junos OS Release 17.1R1, Junos OS CFM works during a unified ISSU when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit CCM PDU (continuity check messages) during the unified ISSU upgrade.

To provide this interoperability, enable inline (Packet Forwarding Engine) keepalives with the **hardware-assisted-keepalives** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level. You must also configure the continuity-check interval to 1 second with the **interval** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]** hierarchy level. Interoperability during unified ISSU is not supported for any other interval value.

[See [Configuring Connectivity Fault Management for interoperability during Unified In-Service Software Upgrades](#).]

Platform and Infrastructure

- **Virtual broadband network gateway support on virtual MX Series router (vMX)**—Starting in Junos OS Release 17.1R1, vMX supports most of the subscriber management features available with Junos OS Release 17.1 on MX Series routers to provide a virtual broadband network gateway on x86 servers.

vBNG runs on vMX, so it has similar exceptions; the following subscriber management features available on MX Series routers 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 these licenses:

- vMX PREMIUM application package license with 1 Gbps, 5 Gbps, 10 Gbps, or 40 Gbps bandwidth
 - vBNG subscriber scale license with 1000, 10 thousand, 100 thousand, or 1 million subscriber sessions for one of these tiers: Introductory, Preferred, or Elite
- **Virtual MX Series router (vMX)**—Starting in Junos OS Release 17.1R1, you can deploy vMX routers on x86 servers. FreeBSD 10 is the underlying OS for Junos OS for vMX. vMX uses DPDK 2.2 to support improved performance.

vMX supports most of the features available on MX Series routers and allows you to leverage Junos OS to provide a quick and flexible deployment. vMX provides the following benefits:

- Optimizes carrier-grade routing for the x86 environment
- Simplifies operations by consistency with MX Series routers
- Introduces new services without reconfiguration of current infrastructure

Routing Protocols

- **IS-IS import policy and route prioritization (MX Series)**—Beginning with Junos OS Release 17.1R1, you can prioritize IS-IS routes that are installed in the routing table for better convergence. In a network with a large number of interior gateway protocol prefixes with BGP Layer 3 VPN or label-based pseudowire service established on top of some interior gateway protocol prefixes, it is important to control the order in which routes get updated in the forwarding table.

In previous releases, Junos OS installed IS-IS routes lexicographically in the routing table. Starting with Junos OS Release 17.1R1, you can configure an import policy to prioritize IS-IS routes as per your network requirements. Use a route tag, or filter the routes based on their prefix before setting a priority of **high**, **medium**, or **low**. Use the **reject** policy option to reject routes from a specific prefix or routes marked with a particular tag. The IS-IS protocol downloads routes to the rpd routing table based on the configured priority. If you do not configure an import policy, all routes are set to a medium priority by default.

[See [Example: Configuring a Routing Policy to Prioritize IS-IS Routes](#).]

- **Adjustable TCP MSS values (MX Series)**—Starting in Junos OS Release 17.1R1, you can use the **tcp-mss** statement to configure the maximum segment size (MSS) for transient TCP packets that traverse a router.

Adjusting the TCP MSS value helps reduce the likelihood of fragmentation and packet loss. The **tcp-mss** statement can be enabled on dynamic interfaces and supports protocols families **inet** and **inet6**.

[See [tcp-mss](#).]

- **BGP advertises multiple add-paths based on community value (MX Series)**—Beginning with Junos OS 17.1R1, you can define a policy to identify eligible multiple path prefixes based on community values. BGP advertises these community-tagged routes in addition to the active path to a given destination. If the community value of a route does not match the community value defined in the policy, then BGP does not advertise that route. This feature allows BGP to advertise not more than 20 paths to a given destination. You can limit and configure the number of prefixes that BGP considers for multiple paths without actually knowing the prefixes in advance. Instead, a known BGP community value determines whether or not a prefix is advertised.

[See [Example: Configuring a Routing Policy to Select and Advertise Multipaths Based on BGP Community Value](#).]

- **Selective advertising of BGP multiple paths (MX Series)**—Beginning with Junos OS Release 17.1R1, you can restrict BGP **add-path** to advertise contributor multiple paths only. Advertising all available multiple paths might result in a large overhead of processing on device memory and is a scaling consideration, too. You can limit and configure up to six prefixes that the BGP **multipath** algorithm selects. Selective advertising of multiple paths facilitates internet service providers and data centers that use route reflector to build in-path diversity in IBGP.

[See [Example: Configuring Selective Advertising of BGP Multiple Paths for Load Balancing](#).]

- **System performance enhancements for rpd, Packet Forwarding Engine, and kernel (MX Series)**—Beginning with Junos OS Release 17.1R1, performance of the routing protocol process (rpd), the Packet Forwarding Engine, and the kernel is enhanced to speed up the process with which the rpd learns the route states and changes, and reflects these changes in the ASIC-based Packet Forwarding Engine residing in the line cards. The key enhancements are faster route download rates when a router comes up after a reboot, or when you add a new line card, and faster update of the data plane in convergence scenarios. We recommend disabling daemons, such as Layer 2 address learning process (l2ald) and connectivity-fault management process (cfmd) —if they are not required— to improve system performance. Though these enhancements are mainly for the MX Series, other platforms might see some performance improvements as well.

Services Applications

- **Support for inline 6rd and 6to4 (MX Series routers with MPC5Es and MPC6Es)**—Starting in Junos OS Release 17.1R1, you can configure inline 6rd or 6to4 on MPC5Es and MPC6Es. You can use the inline capability to avoid the cost of using MS-DPCs for required tunneling, encapsulation, and decapsulation processes. Anycast is supported for 6 to 4 using next-hop service interfaces. Hairpinning is also supported for traffic between 6rd domains.

[See [Tunneling Services for IPv4-to-IPv6 Transition Overview](#), [show services inline software statistics](#), and [clear services inline software statistics](#).]

- **Support for IP reassembly on GRE tunnel interfaces (MX Series routers with MPCs)**—Starting in Junos OS Release 17.1R1, you can configure fragmentation and reassembly of generic routing encapsulation (GRE) packets on GRE tunnel interfaces on MX Series routers with the following Modular Port Concentrators: MPC2E-NGs, MPC3E-NGs, MPC5Es, and MPC6Es.

[See [Configuring Unicast Tunnels](#).]

- **Support for 464XLAT PLAT on MS-MPCs and MS-MICs (MX Series)**—Starting in Junos OS Release 17.1R1, the XLAT464 provider-side translator (PLAT) is supported on MS-MICs and MS-MPCs. The 464XLAT architecture provides a simple and scalable technique to provide IPv4 client-server connectivity across an IPv6-only network without having to maintain an IPv4 network and assign additional public IPv4 addresses on the customer side.

[See [464XLAT Overview](#).]

- **Logging and reporting framework (MX Series with MS-MPC and MS-MIC)**—Starting in Junos OS Release 17.1R1, the logging and reporting framework (LRF) enables you to log data for subscriber application-aware data sessions and send that data in an IP flow information export (IPFIX) format to an external log collector, using UDP-based transport. These data session logs can include subscriber information, application information, HTTP metadata, data volume, time-of-day information, and source and destination details. An external collector, which is not a Juniper Networks product, can then use this data to perform analytics that provide you with insights about subscriber and application usage.

[See [Logging and Reporting Function for Subscribers](#).]

- **Network attack protection for MS-MPCs and MS-MICs (MX Series)**—Starting in Junos OS Release 17.1R1, the MS-MPC and MS-MIC can detect and prevent network probing attacks, network flooding attacks, header anomaly attacks, and suspicious packet pattern attacks.

[See [Configuring Protection Against Network Attacks \(MS-MPCs and MS-MICs\)](#).]

- **Support for inline video monitoring on MPC7E, MPC8E, and MCP9E (MX Series)**—Starting in Junos OS Release 17.1R1, support for video monitoring using media delivery indexing (MDI) criteria is expanded to include the following Modular Port Concentrators: MPC7E, MPC8E, and MCP9E.

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

- **CLI command parity for carrier-grade NAT and stateful firewall (MX Series with MS-MPC)**—Starting in Junos OS Release 17.1R1, new operational commands and configuration options provide information previously available only when using the MS-DPC as the services PIC.

- To display information equivalent to that provided by **show services stateful-firewall flow-analysis** for the MS-DPC, use **show services sessions analysis** for the MS-MPC.
- To display information equivalent to that provided by **show services stateful-firewall subscriber-analysis** for the MS-DPC, use **show services subscriber analysis** for the MS-MPC.
- To drop sessions after a certain session setup rate is reached, include the new CLI option **max-session-creation-rate** at the **[edit services service-set service-set-name]** hierarchy level.

[See [max-session-creation-rate \(Service Set\)](#), [show services subscriber analysis](#), and [show services sessions analysis](#).]

- **Enhancements to stateful synchronization (MS-MIC, MS-MPC)**—Starting in Junos OS Release 17.1R1, stateful synchronization for long-running flows is enhanced for MS-MPC services PICs. These enhancements include:
 - Automatic replication of NAT flows for all service sets: NAT44 flows are automatically synchronized for all eligible service sets. You can selectively disable replication for individual service sets.
 - Checkpointing of IPv4 and IPv6 stateful firewall flows and NAPT-44 with address pooling paired (APP), with configurable timeout for checkpointing.

[See [Configuring Inter-Chassis Stateful Synchronization for Long Lived Flows \(MS-MPC, MS-MIC\)](#).]

- **Subscriber-aware and application-aware traffic treatment (MX Series with MS-MPC)**—Starting in Junos OS Release 17.1R1, Junos OS can perform subscriber-aware and application-aware policy enforcement for mobile or fixed-line subscribers. Junos OS determines the subscriber identity of traffic flow and applies the subscriber's policy rules to the flow. Application identification is performed through deep packet inspection (DPI) at Layer 7 and Layer 4. Subscriber policy actions can include:
 - Redirecting HTTP traffic to another URL or IP address
 - Forwarding packets to a routing instance to direct packets to external service chains
 - Setting the forwarding class
 - Setting the maximum bit rate
 - Performing HTTP header enrichment
 - Setting the gating status to blocked or allowed

[See [Subscriber-Aware and Application-Aware Traffic Treatment User Guide](#).]

- **Usage monitoring for subscribers (MX Series with MS-MPC)**—Starting in Junos OS Release 17.1R1, Junos OS can monitor the volume of traffic and the amount of time that a subscriber uses during a session if that subscriber's policy control rules are controlled by a policy and charging rules function (PCRF) server. The PCRF initiates this monitoring, and the MX Series sends the reports to the PCRF. Monitoring can take place for the entire subscriber session or for only specific data flows and applications. The PCRF provides threshold values to indicate when the Service Control Gateway sends a report to the PCRF, or the PCRF can request a report at any time.

[See [Understanding Usage Monitoring for TDF Subscribers](#).]

- **Traffic Load Balancer (MX Series with MS-MPCs)**—Starting in Junos OS Release 17.1R1, traffic load balancing is supported on MS-MPCs. The Traffic Load Balancer (TLB) application distributes traffic among multiple servers in a server group, and performs health checks to determine whether any servers should not receive traffic. TLB supports multiple VRFs.

[See [Traffic Load Balancer Overview](#).]

- **Support for H.323 gatekeeper mode for NAT on MS-MPC and MS-MIC (MX Series routers)**—Starting in Junos OS Release 17.1R1, H.323 gatekeeper mode is supported in NAPT44 and NAT64 rules and IPv4 stateful-firewall rules on the MX Series. H.323 is a legacy VoIP protocol.

[See [ALG Descriptions](#).]

- **Support for IKE and IPsec pass-through on NAPT44 and NAT64 (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 17.1R1, you can enable the passing of IKE and IPsec packets through NAPT44 and NAT64 rules between IPsec peers that are not NAT-T compliant by using the IKE-ESP-TUNNEL-MODE-NAT-ALG Application Layer Gateway (ALG) on MS-MPCs and MS-MICs. This ALG supports only ESP tunnel mode.

[See [ALG Descriptions](#).]

- **Class-of-service (Cos) marking and reclassification for the MS-MICs and MS-MPCs**—Starting with Junos Release 17.1R1, the MS-MIC and MS-MPC support CoS configuration, which enables you to configure differentiated services code point (DSCP) marking and forwarding-class assignment for packets transiting the MS-MIC or MS-MPC. You can configure the CoS service alongside the stateful firewall and NAT services, using a similar rule structure.

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

- **Services support for MPC7E (MX Series)**—Starting in Junos OS Release 17.1R1, the MPC7E (Multi-Rate) MPC supports the redirection of packets to the MS-MPC for the following services: carrier-grade NAT and stateful firewalls.
- **Support for distributing dynamic endpoint IPsec tunnels among AMS interfaces (MX Series routers with MS-MPCs)**—Starting in Junos OS Release 17.1R1, you can distribute IPsec tunnels with dynamic endpoints among aggregated multiservices (AMS) interfaces.

[See [Configuring Dynamic Endpoints for IPsec Tunnels](#).]

- **Enhancements to the RFC2544-based benchmarking tests (MX Series)**—Junos OS Release 17.1R1 extends support for the RFC2544 on MX Series routers with MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G) and the MPC6E (MX2K-MPC6E).

The RFC2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before activation of the service. The tests measure throughput, latency, frame loss rate, and back-to-back frames. Starting from Junos OS Release 17.1R1, RFC2544-based benchmarking tests on MX Series routers supports the following reflection function:

- Layer 2 reflection (ingress direction) for family **bridge**, **vpls**

To run the benchmarking tests on the MX Series routers, you must enable reflection feature on the corresponding MPC slot. To configure the reflector function on the MPC, use the **chassis fpc fpc-slot-no slamon-services rfc2544** statement at the **[edit]** hierarchy level.

[See [RFC2544-Based Benchmarking Tests Overview](#).]

- **Service redundancy daemon support for redundancy across multiple gateways (MX Series routers with MS-MPCs)**—Starting in Junos OS Release 17.1R1, you can configure redundancy across multiple service gateways. The redundancy actions are based on the results of monitoring system events, including:
 - Interface and link down events
 - FPC and PIC reboots
 - Routing protocol daemon (rpd) aborts and restarts
 - Peer gateway events, including requests to acquire or release mastership, or to broadcast warnings

[See [Service Redundancy Daemon Overview](#).]

Subscriber Management and Services

- **Support for access-line-identifier interface sets based on the Agent Circuit ID (ACI), the Agent Remote ID (ARI), or both (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure interface sets for dynamic subscriber VLANs based on the access-line identifiers (ALI) that are received in a DHCPv4, DHCPv6, or PPPoE discovery packet. The set can be created when the identifier received is the ACI, the ARI, both the ACI and the ARI, or when neither the ACI nor the ARI is received. These interface sets model subscriber identities in a 1:N S-VLAN access model, where a single VLAN exists per service, but more than one subscriber might be using the service. In earlier releases, only the ACI could create the interface sets (ACI sets); when it was not present, the discovery packet was dropped.

You can configure the creation of either ALI sets using this method or ACI interface sets using the legacy method, but not both. A CLI check prevents you from configuring both of these methods. The legacy ACI method might be deprecated in a future release.

[See [Access-Line-Identifier-Based Dynamic VLANs Overview](#).]

- **Static provisioning of unique subscriber ID including interface description (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure DHCP local server and DHCP relay agent to concatenate the interface description with the username during the subscriber or client authentication process. Use the **interface-description** statement to include either the logical interface description or the device interface description. The interface description is separated from the other username fields by the specified delimiter, or by the default delimiter “.” when you do not specify a delimiter. The specified delimiter must not be part of the interface description.

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

- **Flat file output for service filter-based accounting (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure service accounting statistics to be collected and reported in a local flat file as an

alternative to being collected and automatically reported to a RADIUS server. Statistics collection is initiated when the service profile is attached to the subscriber interface.

To configure local flat-file reporting:

1. Create a flat-file profile and specify the **service-accounting** option at the **[edit accounting-options flat-file-profile flat-file-profile-name fields]** hierarchy level.
2. Specify this profile with the **local** statement in the subscriber access profile.
3. Configure the access profile for local reporting by setting the accounting-order either to **local** or—if you plan to activate the service with a CLI configuration or command—to **activation-protocol** at the **[edit access profile profile-name service accounting-order]** hierarchy level.

[See [Configuring Service Accounting in Local Flat Files](#).]

- **Support for asymmetric DHCP leasing (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure an override to the DHCP configuration—typically on the relay agent—to send a shorter (asymmetric) lease to a DHCP client than the lease granted by the DHCP local server. When the local server sends a client an acknowledgment packet in response to the client's offer, the relay agent generates a new acknowledgment packet with the shorter time that you configured. When the client requests a lease renewal, the relay agent re-creates the short lease based on the original lease, rather than passing the request back to the local server. The relay agent continues to renew the shorter lease until the long lease renew time expires, at which time the asymmetric lease is no longer valid. Subsequent renewal requests from the client are forwarded to the server for consideration. If the client does not renew the lease before the short lease renew time expires, then the lease is considered to be abandoned by the client. The address is freed earlier than it would be if the granted lease was used. This feature is available for both DHCPv4 and DHCPv6 configurations.

[See [Configuring DHCP Asymmetric Leasing](#).]

- **shmlog support for CoS and firewall filter plug-ins (MX Series)**—Starting in Junos OS Release 17.1R1, you can use the **svc-sdb-id** filter option with the **show shmlog** command to display only the shmlog filter table entries associated with a service session identifier. For example, the following command displays only shmlog entries that include service session 3:

```
user@host> show shmlog entries logname all svc-sdb-id 3
```

Any client session can have multiple associated service sessions. When you specify only the client session ID, the output includes the entries for the client session in addition to entries for all the service sessions related to that client session:

```
user@host> show shmlog entries logname all sdb-id 2
```

Although you can specify multiple shmlog filters at the same time, inaccurate results are returned when you combine **svc-sdb-id** with any filter other than **sdb-id**. For example, if you combine **svc-sdb-id** with

vlan, the output does not display entries for the VLAN and service session. Instead, it displays no entries or only service session entries.

NOTE: The **svc-sdb-id** filter applies only to subscriber-based entries, because non-subscriber-based entries cannot be filtered. You can display those entries with the existing global commands. For example, for non-subscriber-based CoS and firewall entries, you can use the following commands:

```
user@host> show shmlog entries logname all
user@host> show shmlog entries logname *cos*
user@host> show shmlog entries logname *dfw*
```

- **LAC support for IPv6 address family and firewalls (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure the LAC to create the IPv6 address family (inet6) when tunneling the subscriber to the LNS. By default, the LAC requires only family inet to enable forwarding into an IP tunnel. It can apply IPv4 firewall filters to the session. Even when family inet6 is included in the dynamic profile, by default it is not created and IPv6 firewall filters cannot be applied.

Include the **enable-ipv6-services-for-lac** statement at the **[edit services l2tp]** hierarchy level to allow the IPv6 family to be created and IPv6 filters to be applied.

Use the **show services l2tp summary** command to display the current state, Disabled or Enabled, in the IPv6 Ssrvcis for LAC sessions field.

[See [enable-ipv6-services-for-lac](#).]

- **Dynamic subscriber and service management on statically configured interfaces (MX Series)**—Starting in Junos OS Release 17.1R1, enhanced subscriber management supports dynamic service activation and deactivation for static subscribers. These static subscribers work with the native Juniper Networks Session and Resource Control (SRC), or you can configure RADIUS to activate and deactivate the services with change of authorization (CoA) messages.

NOTE: However, with RADIUS, authentication failure does not prevent the underlying interface from coming up and forwarding traffic. Instead, it prevents the subscriber from coming up, and thus service activation or deactivation. Authorization parameters such as IP addresses, net masks, policy lists, and QoS are also not imposed when using RADIUS.

Use the following commands to provide administrative control of static subscribers:

- **request services static-subscribers login interface *interface-name***
- **request services static-subscribers logout interface *interface-name***

- **request services static-subscribers login group *group-name***
- **request services static-subscribers logout group *group-name***

Use the following commands to monitor static subscribers:

- **show static-subscribers**
- **show static-subscribers interface *interface-name***
- **show static-subscribers group *group-name***
- **Subscriber management and services feature parity (MX240, MX480, MX960)**—Starting in Junos OS Release 17.1R1, the MX240, MX480, and MX960 routers with the Routing Engine RE-S-X6-64G support all subscriber management and services features. These services include DHCP, PPP, L2TP, VLAN, and pseudowire.
- **Packet injection enhancements (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure packet injection by using the **packet-inject-enable** option and a reserved policy map named **packed-inject-flow**. When a packet marked with the **packet-inject-flow** policy map egresses out of a logical interface that has the **packet-inject-enable** option enabled, it is sent for packet injection.

The **show interfaces statistics** command output includes additional information about packet injection.

[See [packet-inject-enable](#).]

VPNs

- **Anti-spoofing protection for next-hop-based dynamic tunnels (MX Series Routers with MPCs)**—Starting in Junos OS Release 17.1R1, anti-spoofing capabilities are added to next-hop-based dynamic IP tunnels, where checks are implemented for the traffic coming through the tunnel to the routing instance using reverse path forwarding in the Packet Forwarding Engine.

Currently, when traffic is received from a tunnel, the gateway router does a destination address lookup before forwarding. With anti-spoofing protection, the gateway router does a source address lookup of the encapsulation packet IP header in the VPN to ensure that only legitimate sources are injecting traffic through their designated IP tunnels (strict mode). When a packet comes from a nondesignated tunnel, the reverse path forwarding check passes only in the loose mode. Traffic coming from nonexistent sources fails the reverse path forwarding check.

This feature is supported on virtual routing and forwarding (VRF) routing instances with strict mode as the default.

To enable anti-spoofing for dynamic tunnels, include the **ip-tunnel-rpf-check** statement at the **[edit routing-instances *routing-instance-name* routing-options forwarding-table]** hierarchy level.

[See [Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels](#) and [Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels](#).]

- **Increased scaling values for next-hop-based dynamic GRE tunnels (MX Series routers with MPCs/MICs)**—Starting in Junos OS Release 17.1R1, the limit for the maximum number of next-hop-based dynamic generic routing encapsulation (GRE) tunnels that can be created on an MX Series router with

MPCs or MICs is increased. This provides additional scaling advantage for the total number of IP tunnels that can be created on the router.

The increased scaling values of next-hop-based dynamic GRE tunnels benefits data center networks, where a gateway router is required to communicate with a number of servers over an IP infrastructure; for example, in Contrail networking.

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

SEE ALSO

[Changes in Behavior and Syntax | 125](#)

[Known Behavior | 142](#)

[Known Issues | 150](#)

[Resolved Issues | 169](#)

[Documentation Updates | 232](#)

[Migration, Upgrade, and Downgrade Instructions | 234](#)

[Product Compatibility | 241](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service \(CoS\) | 127](#)
- [General Routing | 127](#)
- [Interfaces and Chassis | 127](#)
- [Intrusion Detection and Prevention | 129](#)
- [Junos OS XML API and Scripting | 129](#)
- [Layer 2 VPN | 130](#)
- [Management | 130](#)
- [MPLS | 131](#)
- [Network Management and Monitoring | 132](#)
- [Operation, Administration, and Maintenance \(OAM\) | 134](#)
- [Routing Protocols | 134](#)
- [Security | 136](#)

- Services Applications | **136**
- Software Defined Networking | **137**
- Subscriber Management and Services | **137**
- System Management | **141**
- User Interface and Configuration | **141**
- VPNs | **141**

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 17.1R3 for MX Series routers.

Class of Service (CoS)

- **Support for 48 classifiers per family (MX Series)**—Starting with Junos OS Release 17.1R3, you can configure up to 48 classifiers per family at the `[edit class-of-service classifiers]` hierarchy level. In earlier releases, you could only configure up to 32 classifiers per family.

[See [CoS Features and Limitations on MX Series Routers](#).]

General Routing

- **Support for deletion of static routes when the BFD session goes down (MX Series)**—Starting with Junos OS 17.1R3, the default behavior of the static route at the `[edit routing-options static static-route bfd-admin-down]` hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

Interfaces and Chassis

- **Support for maximum queues configuration on MPC7E, MPC8E, and MPC9E (MX Series)**—Starting in Junos OS 17.1R1, you can configure the maximum number of queues per MPC on MPC7E, MPC8E, and MPC9E. By default, these MPCs operate in per port queuing mode.

You can use the **set chassis fpc slot-number max-queues queues-per-line-card** command to configure number of queues per MPC. The possible values for *queues-per-line-card* are **8k, 16k, 32k, 64k, 128k, 256k, 512k, or 1M**.

Per-unit scheduling and hierarchical queuing on MPC7E, MPC8E, and MPC9E are licensed features.

You cannot configure the **max-queues** and the **flexible-queuing-mode** statements at the same time.

You use the **flexi-queuing-mode** statement to configure a maximum of 32,000 queues per MPC.

If the **max-queues** statement is not configured, which is the default mode, the MPC starts with a message similar to the following:

FPC 0 supports only port based queuing. A license is required for per-VLAN and hierarchical features.

If the **max-queues** statement is configured and the value is less than or equal to 32,000, the MPC starts with a message similar to the following:

FPC 0 supports port based queuing and is configured in 16384 queue mode. A limited per-VLAN queuing license is required for per VLAN and hierarchical queuing features.

If the **max-queues** statement is configured and the value is greater than 32,000, the MPC starts with a message similar to the following:

FPC 0 supports port based queuing and is configured in 524288 queue mode. A full scale per-VLAN queuing license is required for per VLAN and hierarchical queuing features.

[See [Understanding Hierarchical Scheduling for MIC and MPC Interfaces](#) and [Flexible Queuing Mode Overview](#).]

- **Changes to show interfaces *interface-name* extensive output (MX Series)**—Starting in Junos OS Release 15.1R7, 16.1R5, 16.2R2, and 17.1R2, the **MAC Control Frames** field of the **show interface *interface-name* extensive** command for a specified 10-Gigabit Ethernet interface displays a value of zero. In previous releases, the value for this field was calculated. Because of continuous traffic and as a result of the calculations, the value displayed for this field changed continuously.
- **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.
- **Deprecated maximum transmission unit configuration option for virtual tunnel interfaces**—Starting In Junos OS Release 17.1R3, 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*.

Intrusion Detection and Prevention

- **Importing IS-IS tag value into LDP (MX Series)**—Starting in Junos OS Release 17.1R1, when a tag value is assigned to an IS-IS route, the IS-IS tag value is imported and used by LDP while installing the route in the inet.3 and mpls.0 routing tables if the **track-igp-metric** command is configured. This enables policy configuration to be applied on the inet.3 and mpls.0 routing tables based on the imported tag value.

Junos OS XML API and Scripting

- **XML output change for show subscribers summary port command (MX Series)**—Starting in Junos OS Release 17.1R1, the display format changed for the **show subscribers summary port** command to make parsing the output easier. The output is displayed as in the following example:

```
user@host> show subscribers summary port | display xml
```

```
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/17.1R1/junos">
  <subscribers-summary-information
    xmlns="http://xml.juniper.net/junos/16.1R2/junos-subscribers">
    <counters junos:style="port-summary">
      <port-name>ge-1/2/0</port-name>
      <port-count>1</port-count>
    </counters>
    <counters junos:style="port-summary">
      <port-name>ge-1/2/1</port-name>
      <port-count>1</port-count>
    </counters>
  </rpc-reply>
```

In earlier releases, that output is displayed as in the following example:

```
user@host> show subscribers summary port | display xml
```

```
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/16.1R2/junos">
  <subscribers-summary-information
    xmlns="http://xml.juniper.net/junos/16.1R2/junos-subscribers">
    <counters junos:style="port-summary">
      <port-name>ge-1/2/0</port-name>
      <port-count>1</port-count>
      <port-name>ge-1/2/1</port-name>
      <port-count>1</port-count>
```

```

    </counters>
  </rpc-reply>

```

Layer 2 VPN

- **Support for LSP on EVPN-MPLS**—Starting in Junos OS Release 17.1R3, 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.

Management

- **Enhancement to Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.1R1, Junos Telemetry Interface data streamed through gRPC no longer includes the phrase **oc-path** in the prefix field. For example, a physical interface sensor streaming data for interface et-0/0/0:0 now displays the following output: **str_value:/interfaces/interface[name='et-0/0/0:0']/**.
- **Enhancement to NPU memory sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.1R2, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is **/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/**

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

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 17.1R3, the format of telemetry data exported through gRPC for NPU memory and memory utilization implements prefix compression. This change reduces the payload size of data exported. The following example shows the new format:

key: __prefix__

str_value: /components/component[name='FPC0:NPU0']/properties/property

key: [name='mem-util-edmem-size']/value

uint_value: 12345

Telemetry data is exported in key-value pairs. Previously, the data exported included the component and property names in a single key string.

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

MPLS

- **Representation for OSPF DR node**—Up until version -10 of the BGP-LS draft, the OSPF DR node representation was ambiguous. One could represent DR node as 'AdvertisingRouterId-InterfaceIpAddress' or 'InterfaceIpAddress-1'. Junos OS used to follow 'InterfaceIpAddress-1' format. Starting with version '-11' of the BGP-LS draft, the representation for OSPF DR node must be 'AdvertisingRouterId-Interfaceaddress'. Junos OS now follows the latest format.
- **PPPoE subscribers do not bind over ps interfaces (MX Series)**—Starting with Junos OS Release 17.1R1, the termination of single, multiple, and dual-tagged service delimited VLANs are transported over a single Ethernet CCC pseudowire using ps virtual port devices. This feature provides scaled Layer 3 service application at the pseudowire head-end termination appliance. This behavior is as an extension and evolution for Ethernet pseudowire that is described in RFC 4448.
- **New field for LSP ping egress interface failure (MX Series)**—Starting in Junos OS 17.1R1, if an LSP ping is started and the chosen egress interface fails, pings are still sent to the failed interface and then dropped. The ping must be manually stopped and restarted to select a working interface to the destination (if one exists). To help detect this ping situation, a new field, **Packets dropped due to ifl down**, has been added to the output of the **show system statistics mpls** command.

[See [show system statistics mpls](#).]

- **Support for inet.0 and inet.3 labeled unicast BGP route for protocol LDP (MX Series)**--- Starting in Junos OS Release **17.1R3**, 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.
- Starting in Junos OS Release 16.1R7, 16.2R3, and 17.1R3, the previously hidden configuration statement, **session**, can be configured at the **[edit protocols ldp]** hierarchy level. This statement enables you to configure the LDP session parameters by specifying the session destination address.

[See [session](#).]

- **New option in show mpls lsp autobandwidth command (MX Series)**—Starting in Junos OS Release 17.1R3, a new option **—name lsp-name**— is introduced in the **show mpls lsp autobandwidth** command to specify the name of the LSP for which the autobandwidth information is displayed. With the **name** option, the autobandwidth information specific to the LSP name that has been provided can be obtained in the command output.

[See [show mpls lsp autobandwidth](#).]

- **Disable M-LDP from using RSVP-TE LSPs for tunneling (MX Series)**—Starting in Junos OS Release 12.3R1, Junos OS provides support for multipoint LDP for targeted LDP sessions with unicast replication,

in addition to link sessions. As a result, the current default behavior of multipoint LDP over RSVP tunneling is similar to unicast LDP.

However, because targeted LDP is chosen over LDP and link sessions to signal point-to-multipoint LSPs, you can enable LDP natively throughout the network, so the point-to-multipoint LSPs take the LDP paths.

[See [p2mp \(Protocols LDP\)](#).]

- **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.

Network Management and Monitoring

- **SNMP syslog messages changed (MX Series)**—Starting in Junos OS Release 17.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD —AgentX master agent failed to respond to ping. Attempting to re-register
NEW —AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD —NET-SNMP version %s AgentX subagent connected
NEW —NET-SNMP version %s AgentX subagent Open-Sent!

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

- **MIB buffer overruns only be counted under ifOutDiscard (MX Series)**—The change done via PR 1140400 introduced a CVBC where qdrops (buffer overruns) were counted under ifOutErrors along with ifOutDiscards. This is against RFC 2863 where buffer overruns should only be counted under ifOutDiscards and not under ifOutErrors. In Junos OS Release 17.1R1, this is now fixed.
- **Hard-coded RFC 3635 MIB OIDs updated (MX Series)**—In Junos OS Release 17.1R2, the following RFC 3635 MIB OIDs have been updated as default values:
 - dot3StatsFCSErrors and dot3HCStatsFCSErrors, framing errors
 - dot3StatsInternalMacReceiveErrors and dot3HCStatsInternalMacReceiveErrors, MAC statistics: Total errors (Receive)
 - dot3StatsSymbolErrors and dot3HCStatsSymbolErrors, code violations
 - dot3ControlFunctionsSupported, flow control
 - dot3PauseAdminMode, flow control
 - dot3PauseOperMode, auto-negotiation

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

- **Enhancement to SNMPv3 traps for contextName field (MX Series)**—Starting in Junos OS Release 17.1R2, the contextName field in SNMPv3 traps generated from a non-default routing instance, is populated with the same routing-instance information as is given in SNMPv2 traps. SNMPv2 traps provide the routing-instance information as context in the form of context@community. This information gives the network monitoring system (NMS) the origin of the trap, which is information it might need. But in SNMPv3, until now, the contextName field was empty. For traps originating from a default routing instance, this field is still empty, which now indicates that the origin of the trap is the default routing instance.

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

- **Update to SNMP support of apply-path statement (MX Series)**—In Junos OS Release 17.1R2, SNMP implementation for the apply-path configuration statement supports only two lists:

- **apply-path "policy-options prefix-list <list-name> <*>"**

This configuration has been supported from day 1.

- **apply-path "access radius-server <*>"**

This configuration is supported as of this release.

- **Juniper MIBs Loading Errors Fixed (MX Series)**—In Junos OS Release 17.1R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:

- jnx-gen-set.mib
- jnx-ifotn.mib
- jnx-optics.mib

[See [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 17.1, 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 default log level setting (MX Series)**—In Junos OS Release, 17.1R3, 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](#).]

- A decrease in the MPLS label-switched path (LSP) statistics pauses the SNMP MIB **mplsLsplInfoAggrOctets** count for one MPLS statistics gathering interval. In such cases, the **mplsLsplInfoAggrOctets** value is updated only after completing one more interval of the MPLS statistics gathering.

Operation, Administration, and Maintenance (OAM)

- **Change in behavior of the Ethernet OAM CFM process (MX Series)**—When you deactivate the connectivity fault management (CFM) protocol, the CFM process (cfmd) stops. When you activate CFM protocol, cfmd starts.

In releases before Junos OS Release 16.1R1, when you deactivate the CFM protocol, the CFM process continues to run.

Routing Protocols

- **Optimization of link-state packets (LSPs) flooding in IS-IS (MX Series)**—Starting in Junos OS Release 17.1R1, flooding of LSPs in IS-IS no longer occurs as a result of the commitment of configuration changes unrelated to IS-IS. Now, when the router is not in the restart state, every time a new LSP is generated after a CLI commit, the contents of the new LSP are compared to the contents of the existing LSP already installed in the link-state database (LSDB) between Intermediate Systems. When the contents of the two LSPs do not match, the system does not process the new LSP or install it in the LSDB, and consequently does not flood it through the IS-IS network. The new behavior does not affect the rebuilding of LSPs after they refresh in the LSDB. No configuration is required to invoke the new behavior.

In earlier releases, IS-IS generates new LSPs even when the configuration changes are not related to IS-IS. Because the new LSPs are flooded across the network and synchronized in the LSDB, this flooding process is time-consuming and CPU intensive in a scaled network environment.

- **Range of flow route rate-limit modified (MX Series)**—Starting with Junos OS Release 17.1R1, the range of flow route **rate-limit** has changed from [9600..1000000000000] to [0..1000000000000]. Earlier Junos OS releases had range restrictions for flow route **rate-limit** at the **[edit routing-options flow route flow then]** hierarchy level. Junos OS can now accept any configured **rate-limit** value. If the rate limit is set in the range of 0 through 999, the Packet Forwarding Engine discards the packets. For configured rate limit value between 1000 and 1000000000000, Junos OS sets the corresponding value in **kbps** as the rate limit.
- **Change in default behavior of router capability (MX Series)**—In Junos OS Release 17.1R1 and later releases, the router capability TLV distribution flag (S-bit), which controls IS-IS advertisements, will be reset, so that the segment-routing-capable sub-TLV is propagated throughout the IS-IS level and not advertised across IS-IS level boundaries.

- **Support for configuring higher PDU size for IS-IS hello packets (MX Series)**—Starting with Junos OS Release 17.1R1, you can configure the maximum protocol data unit (PDU) size of an IS-IS hello packet to up to 16000 bytes. You can achieve the maximum PDU size by configuring the **max-hello-size** configuration statement at **[edit protocol isis interface *interface-name*]** hierarchy and **[edit protocol isis]** hierarchy and by configuring the **hello-padding strict** configuration at the **[edit protocol isis]** hierarchy. The **max-hello-size** statement configured at the interface level has a higher precedence than the configuration at the **[protocol isis]** instance level.

NOTE: The maximum hello-size configuration at the **[protocol isis]** instance level must be less than or equal to the max-hello-size at the interface International Organization for Standardization (ISO) maximum transmission unit (MTU) level and not the interface MTU.

Previously, you could configure the **max-hello-size** configuration statement only at **[edit protocol isis]** hierarchy and the maximum size of IS-IS hello packets that were supported was 1492 bytes.

- **Weighted ECMP supports IS-IS SPRING next hops (MX Series)**—Starting in Junos OS Release 17.1R1, one hop weighted ECMP feature supports IS-IS SPRING based next hops. Currently weighted ECMP for SPRING routes does not support multiple next hop addresses.
- **MPLS configuration mandatory for indirect next-hop interfaces (MX Series)**—Starting in Junos OS Release 17.1R3, 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.
- **Modified output of show route forwarding-table (MX Series)**—Starting in Junos OS Release 17.1R3, the output of the **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](#).]

- For link-state distribution using an interior gateway protocol (IGP), ensure that OSPF is enabled on the donor interface for an unnumbered interface configuration, so the donor IP address is reachable to establish OSPF sessions.

[See [Configuring an Unnumbered Interface](#).]

Security

- **Packet types added for DDoS protection L2TP policers (MX Series routers with MPCs, T4000 routers with FPC5)**—Starting in Junos OS Release 17.1R1, the following eight packet types have been added to the DDoS protection L2TP protocol group to provide flexibility in controlling L2TP packets:

cdn	scccn
hello	sccrq
iccn	stopccn
icrq	unclassified

Previously, no individual packet types were available for this protocol group and all L2TP packets were policed the same based on the aggregate policer value. The default values for the bandwidth and burst policers for all packet types is 20,000 pps. The default **recover-time** is 300 seconds for each of the L2TP packet types.

[See [protocols \(DDoS\)](#).]

- **Global configuration for DDoS protection flow detection mode and flow level control (MX Series)**—Starting in Junos OS Release 17.1R1, you can configure the mode of operation (on, off, or automatic) for flow detection and tracking globally. You can also configure globally how traffic in culprit flows is handled (drop, keep, or police). Both configurations apply to all protocol groups and packet types in the traffic flow unless overridden by the configuration for a protocol group or packet type for all or some flow aggregation levels.

In earlier releases, you cannot configure the behavior globally; you can configure the behavior only for individual protocol groups or packet types, or at the individual flow aggregation levels: physical interface, logical interface, or subscriber.

See [Configuring How Flow Detection Operates Globally](#) and [Configuring How Traffic in a Culprit Flow Is Controlled Globally](#).

Services Applications

- **Deprecated security IDP statements (MX Series)**—In Junos Release 17.1R1 and later releases, **[edit security idp]** configuration statements are deprecated for the MX Series routers.
- **Device discovery with device-initiated connection (MX Series)**—In Junos OS Release 17.1R1 and later releases, when you configure statements and options under the **[system services ssh]** hierarchy and commit the configuration, make sure that the system reaches a stable state before you commit any **outbound-ssh** configurations.

You use the device discovery feature in the Devices workspace to add devices to Junos Space Network Management Platform. By default, Junos Space manages devices by initiating and maintaining a connection to the device.

[See [Device Discovery Overview](#).]

- **Change in enforcement of maintenance mode for changes to PCC action profiles (MX Series)**—Starting with Junos OS Release 17.1R1, a commit error occurs when you change the **logging-rule** or **steering** statements at the **[edit unified-edge pcef pcc-action-profiles *profile-name*]** hierarchy level if the TDF gateway is not in maintenance mode. Prior to Junos OS Release 17.1R1, a commit error was not displayed.
- **Change in error message displayed while fragmenting or de-fragmenting IPv6 GRE tunnel interface (MX Series routers)**—In Junos OS Release 17.1R3, 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 not supported

Software Defined Networking

- 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 17.1R3, the **locally-provisioned** option in the **show mpls lsp ingress** command is behaving as expected.

Subscriber Management and Services

- **Changes to the test aaa authd-lite user, test aaa dhcp user, and test aaa ppp user commands (MX Series)**—Starting in Junos OS Release 17.1R1, the following changes have been made to the **test aaa user** commands:
 - The Virtual Router Name and Routing Instance fields became the Virtual Router Name (LS:RI) field.
 - The Redirect VR Name field was renamed to Redirect VR Name (LS:RI).
 - The Attributes area in the CLI output header section was renamed to User Attributes.
 - The IGMP field was renamed to IGMP Enable.
 - The IGMP Immediate Leave and the MLD Immediate Leave default values changed from **disabled** to **<not set>**.
 - The Chargeable user identity value changed from an integer to a string.

- The Virtual Router Name field was added to the display for the DHCP client.
- The commands display only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise, the field displays **default:default**. The displayed value for all other attributes that are not received is **<not set>**.

[See [test aaa authd-lite user](#), [test aaa dhcp user](#), and [test aaa ppp user](#).]

- **interfaces statement restored for ESSM subscriber secure policy (MX Series)**—Starting in Junos OS Release 17.1R1, the **interfaces** statement was undeprecated at the **[edit services radius-flow-tap]** hierarchy level. When you use subscriber secure policies to mirror ESSM interfaces, you must configure the virtual tunnel (vt) interfaces that are used to send the mirrored packets to a mediation device. In some earlier releases, this statement was erroneously deprecated and hidden.

[See [interfaces \(Subscriber Secure Policy\)](#).]

- **New option to display all pending accounting stops (MX Series)**—Starting in Junos OS Release 17.1R1, the **brief** option is added to the **show accounting pending-accounting-stops** command. This option displays the current count of pending RADIUS accounting stop messages for subscribers, services, and total combined stops. The output is displayed as follows:

```
user@host> show accounting pending-accounting-stops brief
```

```
Total pending accounting stops: 4
  Subscriber pending accounting stops: 2
  Service pending accounting stops: 2
```

[See [show accounting pending-accounting-stops brief](#).]

- **Change to DHCP option 82 suboptions support to differentiate duplicate clients (MX Series)**—Starting in Junos OS Release 17.1R2, only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are considered when this information is used to identify unique clients in a subnet. Other suboptions, such as Vendor-Specific (suboption 9), are ignored.
- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 17.1R2, the **show subscribers extensive** command displays the **IPv6 Interface Address** field only when the dynamic profile includes the \$junos-ipv6-address predefined variable.

In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

- **Traffic shaping and L2TP tunnel switches (MX Series)**—Starting in Junos OS Release 17.1R1, when a dynamic profile attaches a statically configured firewall filter to an L2TP tunnel switch (LTS) session, the filter polices traffic from the LTS (acting as a LAC) to the ultimate endpoint LNS, in addition to the previously supported traffic from the LAC to the LTS (acting as an LNS). In previous releases, the firewall filter applied to only the traffic from the LAC to the LTS.

- **Default L2TP resynchronization method changed and statement deprecated (MX Series)**—Starting in Junos OS Release 17.1R2, the default resynchronization method for L2TP peers in the event of a control connection failure is changed to silent failover. In earlier releases, the default method is failover-protocol-fall-back-to-silent-failover. The silent failover method is preferred because it does not keep tunnels open without traffic flow, waiting for the failed peer to recover and resynchronize. You can use the new **failover-resync** statement at the **edit services l2tp tunnel** hierarchy level to specify either failover protocol or silent failover as the resynchronization method.

Because silent failover is now the default, the **disable-failover-protocol** statement is no longer needed and has been deprecated. If you upgrade to this release with a configuration that includes this statement, it is supported, but the CLI notifies you it is deprecated.

- **IPv6 Link Local Addresses Assigned to Underlying Static Demux Interfaces (MX Series)**—Starting in Junos OS Release 17.1R2, when you are using Router Advertisement for IPv6 subscribers on dynamic demux interfaces that run over underlying static demux interfaces, configure the software to use the same link-local address for both interfaces. In this case, the link-local address for the underlying interface should be based the MAC address of the underlying interface. The following statement causes the system to assign an address using the 64-bit Extended Unique Identifier (EUI-64) as described in RFC 2373:

```
system {
  demux-options {
    use-underlying-interface-mac
  }
}
```

- **Wildcard supported for show subscribers agent-circuit-identifier command (MX Series)**—Starting in Junos OS Release 17.1R3, 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.

- **Support for IPv6 all-routers address in nondefault routing instance (MX Series)**—Starting in Junos OS Release 17.1R3, 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 17.1R3, 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 \(L2TP\)](#).

- **Memory mapping statement removed for Enhanced Subscriber Management (MX Series)**— In Junos OS Release 17.1R3, 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.]

- **Source-specific multicast (SSM) CLI changes for dynamic IGMP and dynamic MLD (MX Series)**—Starting in Junos OS Release 17.1R3, the **ssm-map ssm-map-name** statement at the **[edit dynamic-profiles profile-name protocols (igmp | mld) interface interface-name]** hierarchy level is deprecated and is no longer supported. Instead, you define an SSM map policy with the **policy-statement** statement at the **[edit policy-options]** hierarchy level. Apply the policy for dynamic IGMP or dynamic MLD with the **ssm-map-policy ssm-map-policy-name** statement at the **[edit dynamic-profiles profile-name protocols (igmp | mld) interface interface-name]** hierarchy level.

If you upgrade from a release that does not support enhanced subscriber management (any release earlier than Junos OS Release 15.1R4) with a configuration that includes **ssm-map**, the configuration is allowed. However, the configuration has no effect and subscribers cannot log in.

- **Disabling a pseudowire underlying interface (MX Series)**—Starting in Junos OS Release 17.1R3, you cannot disable the underlying logical tunnel (lt) 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](#).]

- **Bandwidth options match for inline services and tunnel services (MX Series)**—Starting in Junos OS Release 17.1R3, 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 \(Tunnel Services\)](#).]

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the **peers** option at the `[edit system commit]` hierarchy level is not supported in batch configuration mode.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (MX Series)**—Starting in Junos OS Release 17.1R1, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 17.1R1, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the show system schema module juniper-command output directory (MX Series)**—Starting in Junos OS Release 17.1R1, when you issue the **show system schema module juniper-command** operational command in the Junos OS CLI, the device places the generated output files in the current working directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1R1, the generated output files are placed in the `/var/tmp` directory.
- **SLAX scripts included as part of the Junos OS image (MX Series)**—In Junos OS Release 17.1R1 and later releases, the Stylesheet Language Alternative Syntax (SLAX) scripts **services-oids-ev-policy.slax**, **services-oids.slax**, and **utils.slax** are included as part of the Junos OS image and automatically copied to the required location on the router when you install Junos OS.
- **Junos OS prohibits configuring ephemeral configuration database instances that use the name default (MX Series)**—Starting in Junos OS Release 17.1R3, 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

- **EVPN E-tree extended community (MX Series)**—In Junos OS Releases 17.1R2, and later releases, the E-tree leaf indication bit and leaf label in EVPN E-tree extended community follows the E-tree Extended Community as defined in the [E-TREE Support in EVPN & PBB-EVPN IET](#) IETF draft. A mixed network environment with routers running versions of Junos OS without this fix and routers with this fix would encounter unexpected forwarding behavior. Junos OS Release 16.1R4 has the incorrect label indication bit and leaf label encoding.
- **EVPN extended community and ISID using standard IANA value (MX Series)**—Starting in Junos OS Release 17.1R2, the router MAC extended community and service identifier (ISID) sub-type values have been corrected to use the Internet Assigned Numbers Authority (IANA) standardized value. In Junos OS Release 17.1R1, when you configure EVPN extended community using a pure type 5 routing mode with VXLAN encapsulation, you might encounter routing issues with the router from another vendor.
- **Support for ping on a virtual gateway address (MX Series)**—In Junos OS Release 17.1R2, Junos supports pinging an IPv4 or IPv6 address on the preferred virtual gateway interface. To set up support for ping,

you must include both the **virtual-gateway-accept-data** and the **preferred** statements at the **[edit interfaces irb unit]** hierarchy of the preferred virtual gateway. This enables the interface on the preferred virtual gateway to accept all packets for the virtual IP address, including ping packets.

SEE ALSO

[New and Changed Features | 99](#)

[Known Behavior | 142](#)

[Known Issues | 150](#)

[Resolved Issues | 169](#)

[Documentation Updates | 232](#)

[Migration, Upgrade, and Downgrade Instructions | 234](#)

[Product Compatibility | 241](#)

Known Behavior

IN THIS SECTION

- [Class of Service \(CoS\) | 143](#)
- [EVPN | 143](#)
- [Forwarding and Sampling | 144](#)
- [General Routing | 145](#)
- [High Availability \(HA\) and Resiliency | 146](#)
- [Interfaces and Chassis | 146](#)
- [MPLS | 147](#)
- [Platform and Infrastructure | 147](#)
- [Routing Protocols | 147](#)
- [Services Applications | 147](#)
- [Software Installation and Upgrade | 148](#)
- [Subscriber Management and Services | 149](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 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)

- **Filtering for Routing Engine sourced packets (MX Series)**—Starting in Junos OS Release 17.1R1, support is added for filtering on Differentiated Services Code Point (DSCP) and forwarding class for Routing Engine sourced packets. This includes IS-IS packets encapsulated in generic routing encapsulation (GRE). With this change comes a new order of precedence. When upgrading from a previous version of Junos OS where you have both a class of service (CoS) and firewall filter, and both include DSCP or forwarding class filter actions, the criteria in the firewall filter automatically takes precedence over the CoS settings. The same is true when creating new configurations; that is, where the same settings exist, the firewall filter takes precedence over the CoS, regardless of which was created first.

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 OS routing instance should be configured with **vlan-id none** so that the Ethernet tag in advertised MAC routes is set to zero. [PR945247](#)
- A PE device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE device. The IGP instance running in the VRF on the PE device might be able to discover the IGP instance running on the remote CE device through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- EVPN VPWS convergence and association with traffic loss is tied to the type of redundancy and the route exchange through BGP. In A/A this traffic loss is low because of the distribution of the traffic as well as protocols that can be used on the CE-PE link to steer the traffic away from the failed link as soon as the failure occurs. Here is the data for AA and AS. The number for AS are higher and are due to inherent limitations of this redundancy scheme. AA: a) ESI Goes DOWN : <10 msec. b) ESI comes UP: <50msec (for Traffic Items corresponding to 80RIs ? 1VPWS CKT per RI) = 350 msec approx. (For Traffic item corresponding to 2000CKTs in one RI) AS: a) ESI goes Down: 4950msec (Approx.) b) ESI Comes UP: 2100 msec (Approx.) [PR1181523](#)
- In scaled up EVPN VPWS configuration (approximately 8000 EVPN VPWS), during a Routing Engine switchover, rpd scheduler slip messages might be seen. [PR1225153](#)

- An incorrect PE router is attached to an ESI when the router receives two copies of the same AD/ESI route (for example, one through eBGP and another one received from an iBGP neighbor). This causes a partial traffic that is discarded without notification and with stale MAC entries. You can confirm the issue by checking the members of the ESI: `user@router> show evpn instance extensive ...` **Number of ethernet segments: 5 ESI: 00:13:78:00:00:00:00:00:01 Status: Resolved Number of remote PEs connected: 3 Remote PE MAC label Aliasing label Mode 87.233.39.102 0 0 all-active 87.233.39.1 200 0 all-active <<<< this PE is not part of the ESI 87.233.39.101 200 0 all-active** [PR1231402](#)
- When you activate and deactivate the Route Target per bridge domain in EVPN, the rpd process might crash resulting in traffic loss. We recommend that you do not toggle to this configuration in Junos OS Release 17.1R1. [PR1244956](#)

Forwarding and Sampling

- We have an interface for a direct route starting in ifdown condition. The remote side is then brought up, so I/F goes to ifup. Since it is a direct route, rpd does not install the route or nexthop. It receives that info from the kernel, and just updates a next hop in rpd local storage. The route and next hop for the interface are taken care of in the kernel. There is no route change in rpd. The route_record depends on route flash to find out about updates. Because there is no route change, there is no route flash, so route_record is unaware. In order to change this, we would need to decide that we want a route flash for this case. Currently, for direct and local routes/next hops, these are "don't care" in rpd, as far as route updates go. We just update our next hop information, without marking for any other notifications. A complication for the solution is a change that was done for PR 1002287, where if the NOTINSTALL flag is set, do not send the update to srrd. That flag is set for direct and local routes. Incidentally, this is day 1 operation. If the interface is up at startup, it should all work correctly. FIB table can provide OIF/GW only. SRC_MASK, DST_MASK, SRC_AS and DST_AS are not available in PFE FIB Table. So SRRD connection is required. Listening to both SRRD and FIB table, and consolidating information will complicate implementation. Scanning the entire FIB Table just for the few such routes will have performance impact and will complicate present implementation. This is day 1 implementation for SRRD/Sampled. Workarounds: There are two possible workarounds. 1) Have the far end interface up when the DUT interface is brought up. In cases where that is not happening, a recovery would be to disable the DUT interface, then enable it again. At that point, everything should be initially brought up in the state we are looking for. 2) Enable the **nexthop-learning** command. [PR1224105](#)
- FreeBSD 10.x based Junos OS is not supported on 32-bit Routing Engines in Junos OS Release 17.1R1. [PR1252662](#)

General Routing

- The **rpd** process might crash if ECMP routes have more than 38 IS-IS IPv6 next hops—If the **maximum-ecmp 64** statement is enabled and ECMP routes have more than 38 IS-IS IPv6 next hops, then the **rpd** process might crash because the next hop gateway addresses get overwritten and stored in a circular buffer.

NOTE: If all the next-hop IP addresses are IPv6 addresses, you can configure only 38 ECMP next-hop addresses for IS-IS.

- **Support for simultaneous PTP over Ethernet and PTP over IPv4 master streams is not available for G.8275.1 profile (MPC5E and MX104)**—In Junos OS Release 17.1R2, support for simultaneous PTP over Ethernet encapsulation and PTP over IPv4 master clock interface is not available on MPC5E and MX104, for a G.8275.1 PTP profile.
- 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 150,000 cps). This issue is not seen with up to 50,000 control connections at 10,000 through 30,000 cps. [PR1087561](#)
- NAT64: Source-prefix filtering and protocol filtering of the CGNAT sessions are incorrect. For example, the **show services sessions extensive protocol udp source-prefix <0:7000::2>** command displays incorrect filtering of the sessions. [PR1179922](#)
- Chef for Junos OS supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev resources. The netdev resource developed for interface configuration has a limitation to 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](#)
- When LACP is used and considers the member link to be ready to forward traffic, the member link must not be used by the load balancer until all the micro-BFD sessions of the particular member link are in the up state. [PR1192161](#)
- In certain interface scaling scenarios, during configuration commit or 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](#)
- PIC gets rebooted without generating a core file, in spite of having dump on flow control configured. [PR1217167](#)

- The ptp master streams on IP and Ethernet are not supported simultaneously. [PR1217427](#)
- There is no unified ISSU from a Junos OS Release with NPU image size less than 60 MB to a Junos OS Release with NPU image size greater than 60 MB. [PR1222540](#)
- In this feature(next-hop-based tunnel), GRE/UDP tunnel does not have any tunnel liveness protocol, which can state tunnel up/down event. With the current implementation, tunnel will be up if the next hop is installed in the kernel or the Packet Forwarding Engine. This next hop will be withdrawn only if the BGP router is removed from the bgp.l3vpn.o table. [PR1223727](#)
- OSPF is used as routing protocol between the clients and the DEP router with TD configured. The OSPF protocol traffic brings the IPsec up on spokes and DEP router. The IPsec SAs are distributed on the DEP router. The neighbor state between the OSPF peers move to full but it does not stay in that state. States change to init, 2-way, ex-start, and to full again. As a result, the data traffic between the routers is getting dropped. Thus tunnel distribution with protocol traffic is not supported. [PR1232277](#)
- vMX does not detect interface link state correctly in SR-IOV mode with i40e driver. [PR1271902](#)
- CFM is not supported for L2-over-GRE tunnel. CCM can pass through as transit traffic through GRE interfaces transparently using datapath. Link trace functionality uses MAC-learning and re-injecting LTM on GRE interface in case the bridge is configured with CFM. This is not supported feature. [PR1275833](#)
- With Junos OS Release 16.2R1 or later, the error message about jlaunchd, **jlaunchd: %AUTH-1: commit-batch is thrashing, not restarted**, might be seen after a system reboot or a Routing Engine switchover. [PR1284271](#)

High Availability (HA) and Resiliency

- **Unified ISSU incompatibility with VPLS dynamic profiles (MX Series)**—Using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 17.1R1 does not work if VPLS dynamic profiles are configured and enhanced subscriber management is not configured.

[See [ISSU System Requirements](#)]

Interfaces and Chassis

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

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

- 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](#)

MPLS

- When Flow-Label (FL) is enabled for pseudowire, the OAM packets are not sent with Flow-Label because rpd is not aware of the Flow-Label values assigned by the Packet Forwarding Engine software. Hence, the packets were getting dropped by the Packet Forwarding Engine at the egress. The remote PE was expecting the packet with Flow-Label and pseudowire label. [PR1217566](#)

Platform and Infrastructure

- The FPC reports the following errors and it is not able to connect any subscriber: **Pkt Xfer:** WEDGE DETECTED IN PFE 0 TOE host packet transfer: %PFE-0: reason code 0x1"** Also, the MQ FI may be wedged and the following log can be seen: **Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) FI Reorder cell timeout Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) FI Enqueuing error, type 1 seq 404 stream 0 Apr 11 12:09:11.945 2013 NSK-BBAR3 fpc7 MQCHIP(0) MALLOC Pre-Q Reference Count underflow - decrement below zero** [PR873217](#)

Routing Protocols

- When a Junos OS aggregation gateway uses an IPv6 address as next hop for IPv4 aggregates announced to downstream, it might attract traffic prematurely before Packet Forwarding Engines are programmed with more specific IPv4 routes. This happens when the IPv6 address is advertised in BGP inet6-labeled-unicast family. [PR1220235](#)
- With GRES+ NSR enabled, the master Routing Engine replicates kernel states and protocol states on the backup Routing Engine. Both kernel state (ifstates) and protocol state replication are independent processes. The ksyncd takes care of ifstates replication. The rpdinfra takes care of replication (mirror) connection between two Routing Engines. And NSR-supported protocols have their own mechanism to replicate their database using mirror connection. According to PIM/MVPN NSR design, on the backup Routing Engine, it walks through the replication database (RDB). Once a PIM/MVPN state is processed on the backup Routing Engine, the associated RDB is deleted. If kernel replication is restarted, it can lead to interface deletions and additions only on the backup Routing Engine. PIM states on the backup Routing Engine go out of synchronization. [PR1224155](#)

Services Applications

- Account Session ID, Interface Identifier, and Subscriber User Name trigger attributes are optimized for a scaled subscriber management environment. If you include any of the other, non-optimized, trigger attributes in a scaled subscriber management environment, a significant delay might be observed between

the time when the DTCP ADD message is sent and the time when forwarding starts for the mirrored traffic. For example, if there are 10,000 subscriber sessions on the router, forwarding of the mirrored traffic might be delayed for 20 minutes. This delay occurs when you specify any non-optimized attribute, with or without any optimized attribute. The delay occurs regardless of the order of attributes in the DTCP packet. [PR1269770](#)

- Broadband-edge platforms do not support service-set integration with dynamic profiles when the service set is representing a carrier-grade NAT configuration. As a workaround, you can use next-hop service set configurations and routing options to steer traffic to a multiservices (ms) interface where NAT functionality can be exercised. The following configuration snippet shows the basics of statically configuring the multiservices interface next hop and a next-hop service set. Traffic on which the service is applied is forced to the interface inside the network by configuring that interface as the next hop. This configuration does not show other routing-options or NAT configurations relevant to your network.

```
routing-options {
  static {
    route 0.0.0.0/0 {
      next-hop ms-3/0/0.1;
      preference 0;
    }
  }
  ...
}
services {
  service-set CGN {
    nat-rules CGN_SAMPLE;
    next-hop-service {
      inside-service-interface ms-3/0/0.1;
      outside-service-interface ms-3/0/0.2;
    }
  }
  nat {
    ...
  }
}
```

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

Software Installation and Upgrade

- On a broadband network gateway (BNG) that is running enhanced subscriber management, you must take the service cards offline before you can perform an in-service software upgrade (ISSU) to Junos OS Release 17.1 from a Junos OS release that includes the application-aware policy control feature (16.1R4 and later).

- **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 17.1 release. The warning **ISSU is not supported for RPM configuration** appears.

Subscriber Management and Services

- If a graceful Routing Engine switchover (GRES) is triggered by an operational mode command, the state of aggregated services interfaces (ASIs) are not preserved. For example:

```
request interface <switchover | revert> asi-interface
```

However, if GRES is triggered by a CLI commit or FPC restart or crash, the backup Routing Engine updates the ASI state. For example:

```
set interface si-x/y/z disable
commit
```

Or:

```
request chassis fpc restart
```

- 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.

SEE ALSO

[New and Changed Features | 99](#)

[Changes in Behavior and Syntax | 125](#)

[Known Issues | 150](#)

[Resolved Issues | 169](#)

[Documentation Updates | 232](#)

[Migration, Upgrade, and Downgrade Instructions | 234](#)

Known Issues

IN THIS SECTION

- [Authentication and Access Control](#) | 151
- [Class of Service \(CoS\)](#) | 151
- [EVPN](#) | 151
- [Forwarding and Sampling](#) | 152
- [General Routing](#) | 153
- [High Availability \(HA\) and Resiliency](#) | 159
- [Infrastructure](#) | 159
- [Interfaces and Chassis](#) | 159
- [Layer 2 Ethernet Services](#) | 160
- [Layer 2 Features](#) | 160
- [MPLS](#) | 161
- [Network Management and Monitoring](#) | 163
- [Platform and Infrastructure](#) | 163
- [Routing Protocols](#) | 165
- [Services Applications](#) | 167
- [Subscriber Access Management](#) | 167
- [User Interface and Configuration](#) | 168
- [VPNs](#) | 168

This section lists the known issues in hardware and software in Junos OS Release 17.1R3 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.

Authentication and Access Control

- Platform-specific callbacks are not getting initialized. [PR1354855](#)

Class of Service (CoS)

- If the **interface-set** statement is configured for CoS, then the FPC might crash when the CoS configuration tries to access an interface-set-related pointer that was freed after the FPC or device was reloaded. [PR1224046](#)
- The cosd might crash during commit through NETCONF if **xcess-priority** is configured. It is a timing issue. [PR1403147](#)

EVPN

- When unknown unicast traffic is received on the egress EVPN PE device, the input packet and rate will be counted as twice as large as the actual input packet and rate. [PR830535](#)
- **Show evpn vpws-instance** SID NNN is not supported. [PR1122695](#)
- The l2ald process might generate a core file a scaled Layer 2 setup with bridge domain, VPLS, and EVPN. The core file generation usually follows a kernel page fault. In most cases the issues resolves on its own after the generation of the l2ald core file. In some cases, you need to manually restart the process. **Logs:**
/kernel: %KERN-3-BAD_PAGE_FAULT: pid 69719 (l2ald), uid 0: pc 0x88beb5ce got a read fault at 0x6ca, x86 fault flags = 0x4 /kernel: %KERN-6: pid 69719 (l2ald), uid 0: exited on signal 11 (core dumped) init: %AUTH-3: l2-learning (PID 69719) terminated by signal number 11. Core dumped!
[PR1142719](#)
- EVPN VPWS convergence and association with traffic loss is tied to the type of redundancy and the route exchange through BGP. In A/A this traffic loss is low because of the distribution of the traffic as well as protocols that can be used on the CE-PE link to steer the traffic away from the failed link as soon as the failure occurs. Here is the data for 1 and 2. The numbers for 2 are higher and are due to inherent limitations of this redundancy scheme.
1: a) ESI goes down: <10 ms.
b) ESI comes up: <50 ms (for traffic items corresponding to 80RIs ? 1VPWS CKT per RI) = 350 ms approx.
(For traffic item corresponding to 2000CKTs in one RI)
2: a) ESI goes down: 4950 ms (approx.)
b) ESI comes up: 2100 ms (approx.)
[PR1181523](#)
- In an EVPN scenario with static MAC configured on the EVPN instance, the remote EVPN instance can see the MAC route information. However, after the static MAC in the EVPN instance is deactivated and

then activated, and the MAC route information in the remote EVPN instance is checked, no such MAC route is found in the EVPN route table. [PR1193754](#)

- This issue is applicable to MAC-in-MAC PNN-EVPN and does not affect any other scenario. When a PBB-EVPN configuration is reloaded on MX Series routers, error logs are seen while deleting interfaces related to backbone bridge component. These errors do not result in any functional issues. [PR1323275](#)
- PBB-EVPN cannot flood traffic toward a core layer. Recover traffic by running **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 do not affect any other scenario. [PR1323503](#)
- When EVPN PE (RR) is configured as single home without ESI, evpn bgp routes from table " bgp.evpn.0 " might leak into default evpn table (__default_evpn__.evpn.0) causing label leak. Leak might lead to all label exhaustion and result in to rpd core. [PR1333944](#)
- In a device running Junos OS platform, the l2ald process might crash during MAC address processing. As a result, the MAC learning process is impacted; however, the l2ald process recovers on its own. [PR1347606](#)
- When EVPN is configured with class-of-service-based forwarding (CBF), traffic might be lost for the CBF services. [PR1374211](#)
- In an EVPN (Ethernet VPN) scenario, if the router receives a Type 2 MAC+IP route advertisement having 2 MPLS labels, and then withdrawal of the same route with only 1 label occurs, the withdrawal will not be processed and that route will be stuck. [PR1399726](#)

Forwarding and Sampling

- When a policing filter is applied to an active LSP carrying traffic, the LSP resignals and drops traffic for approximately 2 seconds. It can take up to 30 seconds for the LSP to come up under the following conditions: (1) Creation of the policing filter and its application to the LSP through configuration occurs in the same commit sequence and (2) Load override of a configuration file that has a policing filter and policing filter application to the LSP are followed by a commit. [PR1160669](#)
- The firewall filter family **any** configured with the **shared-bandwidth-policer** statement on an MC-aggregated Ethernet interface does not reconfigure bandwidth or carve up the policer when the standby device becomes the active device after A/S switchover; it drops all packets. [PR1232607](#)
- After the **show firewall** command is executed, the **dfwinfo: tvptest:dfwlib_owner_create tvp driven policer_byte_count support 0** message is seen in messages logs. This message is a cosmetic issue and it can be ignored safely. This message can be seen with the following sample configuration: << **sample config** >> **set interfaces ge-0/0/0 unit 0 family inet filter input test_filter set interfaces ge-0/0/0 unit 0 family inet address 100.100.100.1/24 set firewall family inet filter test_filter term policer then policer policer_test set firewall policer policer_test if-exceeding bandwidth-limit 100m set firewall policer**

policer_test if-exceeding burst-size-limit 125k set firewall policer policer_test then loss-priority low
[PR1248134](#)

- In some stress test conditions, the sampled process crashes and generates a core file when connecting to Layer 2 Bitstream Access (L2BSA) and EVPN subscribers aggressively. [PR1293237](#)
- Remote CE1 MAC address might take a long time to clear after running **clear mac**. [PR1304866](#)
- In an EVPN A/A scenario with an MX Series router or EX Series switch acting as a PE device, flood next hops 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 are reported to RADIUS. [PR1383354](#)

General Routing

- This issue occurs when the configured global-MAC limit is less than the interface MAC limit and the same interface is configured with packet action. When the traffic is sent with a higher packet rate, all the MAC entries are learned by the Packet Forwarding Engine. The Routing Engine later trims this to the configured global-MAC limit. When the traffic is sent with a lower packet rate, the Routing Engine learns somewhat more than the configured global-MAC limit and subjects the remaining packets (with newer MACs) to the configured drop action. [PR1002774](#)
- ICMP echo_reply traffic with applications such as IPsec does not work with the MS-MIC and MS-MPC cards in an asymmetric traffic environment, because these cards employ a stateful firewall by default. The packet is dropped at the stateful firewall because it acknowledges an ICMP reply that has no matching session. [PR1072180](#)
- An intermittent issue occurs when the **bypass-queuing-chip** statement is configured on an aggregated Ethernet interface. The follow-up configuration changes are such that, removing a child link from an aggregated Ethernet bundle and configuring the **per-unit-scheduler** statement 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](#)
- For the translation types napt-44 and deterministic-napt44, a few NAT sessions are seen to be stuck when service sets or corresponding applications are deactivated or activated with traffic running. [PR1183193](#)
- Reporting fabric self ping blackhole got decoupled from Fabric Hardening events and will execute disable-pfe action per default and will raise a Major Alarm. [PR1184761](#)
- AMS redundant interfaces are not listed under possible completions for operational mode commands. [PR1185710](#)
- IR-mode configuration statement commit failure is seen with MPC7, MPC8, and MPC9. [PR1192228](#)
- GUMEM errors for the same address might continually be logged if a parity error occurs in a locked location in GUMEM. Because GUMEM utilizes ECC memory, any error is self-correcting and has no

impact on the operation of the router. In a rare case, such a parity error might appear repeatedly at a specific location. As a workaround, the error can be cleared by rebooting the FPC. [PR1200503](#)

- When ppm deviation exceeds 10 ppm, do not display off-frequency if the clock source is still being locked. Display **in-use#** instead. This displayed value indicates that it is still locked to the source, although the clock has a considerably large ppm deviation. [PR1202327](#)
- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, consider the following displayed message: **fpc7 KERNEL/PFE APP=NH OUT OF SYNC: error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none**. This issue does not seem to have any service impact on MPC2 and MPC3 line cards. [PR1205593](#)
- In an OAM environment with a GRE interface, if the physical interface is brought down, then an OAM keepalive brings down the GRE interface. But within a few second, GRE interface comes up and the OAM keepalive is down. The traffic forwarding might be affected. [PR1207017](#)
- When APS is configured on already present interfaces, without a maintenance window, the rpd might lose the protect interface state due to some sequence of events. This will cause the local route of the protect interface to go into an inconsistent state and any protocol using this local route might also see issues. [PR1210951](#)
- Several files are copied between Routing Engines during the FFP synchronize phase of the commit (such as `/var/etc/mobile_aaa_ne.id` and `/var/etc/mobile_aaa_radius.id`). These files are copied even if there was no corresponding change in the configuration, thus unnecessarily increasing commit time. [PR1210986](#)
- When you issue an operational mode clear command, the queuing monitor sensor counters on the Junos Telemetry Interface (JTI) server are not reset. As a result, after such **clear** CLI commands are issued, the QMON sensor statistics on the JTI server do not match those in the output of CLI and vty commands. [PR1226948](#)
- The normal discard count in the output of the **show pfe statistics traffic** command continuously increases, even without any user traffic. This issue occurs because internal control traffic that is expected to be dropped silently is unexpectedly being counted as normal discard traffic. There is no impact on user traffic. [PR1227162](#)
- When a configuration that moves the Packet Forwarding Engine offline and another configuration that brings the Packet Forwarding Engine back online are committed in quick succession, out-of-synchronization syslog errors might occur. Most of the time these are benign errors, but sometimes these errors might result in Packet Forwarding Engine crash. [PR1232178](#)
- The following error messages occur during GRES and unified ISSU: **syslog errors @ agentd_rts_async_rtbm_msg : FLM : Failed to create private.**[PR1232636](#)
- When the 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](#)

- To distinguish between flow and kernel logical interfaces for VLAN-OOB, subscribers use the option **idl-arch-type**: **router> show interfaces ge-1/0/3.3221225476 ifl-arch-type ? Possible completions:**
flow Display flow ifls rtsock Display rtsock ifls [PR1236713](#)
- When the IPv4 or IPv6 address configured as the local-gateway for the IPSec VPN service is not actually assigned to any interface in the up state (not present a local/direct route in the routing-table), the system still sends ISAKMP packets for IKE exchange. An address of the outgoing interface is selected as the source address for these packets. [PR1238112](#)
- On MX Series routers with the rpd process in ASYNC mode and distributed IGMP configured, rpd might generate a core file and then crash. [PR1238333](#)
- In a BGP or an MPLS scenario, if the next-hop type of label route is indirect, then the following changing events about the next-hop interface family **mpls** might cause the route to be in the dead state, and the route remains dead even when the family **mpls** is again activated: Deactivating and activating the interface family **mpls**. Deleting and adding back the interface family **mpls**. Changing maximum labels for the interface. [PR1242589](#)
- For ANCP subscribers in Idle state the previously reported speed in the ANCP Port UP message is not applied. [PR1242992](#)
- ANCP neighbors are reinitialized (and could go down) after an ephemeral commit of any ANCP-related configuration. [PR1243164](#)
- Sensors are not reused when the subscriptions have no common paths. When subscribed from multiple servers for the same subscription, sensors are not reused. [PR1245902](#)
- After you connect 1000 L2BSA subscribers and run the CLI command **show ancp subscriber detail | match Aggregate Circuit Identifier Binary**, the output stops at a certain point and gets stuck for minutes. Even pressing Ctrl-C does not help terminate the CLI output. In some cases, pressing Ctrl-C causes the ancpd process to crash. [PR1250996](#)
- The MX104 Routing Engine might be stuck in boot loop after disabling interface fxp0 in the configuration. [PR1253155](#)
- In a Junos Telemetry Interface (JTI) scenario using Junos OS Release 16.1R3 and later releases with non-upgraded freeBSD, if openconfig is used, then the na-grpcd process experiences memory leak or memory increase continuously (that is, continuous subscribing and unsubscribing or aggressive timers for interfaces for about 2 seconds or other conditions), eventually causing the na-grpcd process to crash due to memory exhaustion. As a result, the collector does not get the streaming data during the na-grpcd crash. [PR1254794](#)
- VPLS MAC table is not being populated properly as verified by using the **show vpls mac-table** command, although all subscribers have traffic. This is considered a cosmetic issue. [PR1257605](#)
- In earlier releases, the code does not contain strict enforcement of checking the targeting configuration syntax. Thus having targeting-distribution only in the dynamic VLAN profile but not in the client profile is allowed. This leads to confusion and potentially unexpected behavior. With this fix, strict checking is introduced. Targeting-distribution is required at all levels to bring up the client. [PR1258955](#)

- Errors like **mshpmand[190]: msvcs_session_send**: Plugin id 3 not present in the svc chain for session are usually cosmetic. [PR1258970](#)
- On MX Series routers, in a rare case the backup Routing Engine is slow to process replication. Replication on the master Routing Engine continues too long under a purge condition and results in logic problems and smgd process crash on the backup Routing Engine. [PR1261268](#)
- 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](#)
- If the dynamic VLAN profile does not have IFF configuration (for example, family PPPoE or family inet), but has a firewall filter configuration, firewall filter indexes are not released after the dynamic VLAN is removed. This eventually leads to depletion of available firewall filter indexes. [PR1265973](#)
- If the rpd process crashes on a device that has the **switchover-on-routing-crash** statement enabled, on the device, live VM core files might be seen on both Routing Engines without an impact on the system. [PR1267796](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- On all MX Series devices with the Point-to-Point Protocol over Ethernet (PPPoE) subscriber scenario, the unexpected log message **VTAG not found in uflow** might be seen when a PPPoE subscriber logs on a static virtual local area network (VLAN) logical interface (IFL, i.e., unit ge-0/0/0.40) [PR1284966](#)
- After a Layer 2 routing instance is renamed, A10-NSP interfaces attached to the old routing interface do not get moved to the new routing instance. [PR1287070](#)
- PPPoE cannot dial in because of all padi dropped as "unknown iif" when deactivated or activated aggregated Ethernet configuration which a aggregated Ethernet child leaves and joins bundle in quick succession. And lead to out of order for substructs msg. And the Fix is to process all substructs hanging off the parent logical interface ifstate in the order in which they were enqueued. Restore-only is to reboot FPC. [PR1291515](#)
- iLatency (calculated by differing producer timestamp and gRPC server timestamp) can sometimes be negative for Packet Forwarding Engine related telemetry packets because drift in Routing Engine and Packet Forwarding Engine NTP servers. [PR1303376](#)
- Customer reported that after running a RADIUS disaster backup procedure (disable access interface, change the RADIUS server and enable interfaces), VLAN-OOB sessions came up, but no new ANCP session was established. **user@router> show ancp neighbor Version IP Address PartID State Time Subscriber Capabilities Count * 0x32 172.16.4.1 0 Not Estblshd 21:17 0 Topo, OAM * 0x32 172.16.5.1 0 Not Estblshd 21:18 0 Topo, OAM** **user@router> show subscribers summary Subscribers by State Active: 2988 Total: 2988 Subscribers by Client Type VLAN-OOB: 2988 Total: 2988** The customer was running the following steps to reproduce this issue: 1. Bring up 21,000 subscribers with a session mix over all products (DCIP, EVPL, PPP, L2TP, and L2BSA). 2. Run RADIUS backup procedure using the script. Then the procedure waits until all subscribers are logged out of the BNG (duration of about 10 minutes). 3. After the access interface are enabled within the procedure, many VLAN-OOB sessions come up again, even though no new ANCP session gets established. [PR1306872](#)

- The XML request and reply do not have the same prefix. This is an oversight and might cause confusion as the prefixes normally match. However, fixing the reply now would break current installations. This PR serves as a reminder about the mismatch. [PR1312364](#)
- The following harmless logs will be noticed on router's with next generation Routing Engine's **Oct 1 22:17:25 abc vhclicent.9947.daemon: vhclicent instantiated by /bin/sh Oct 1 22:17:25 abc vhclicent.9947.daemon: rsh -JU __juniper_private5__ 192.168.1.2 export PATH=\$PATH:/usr/sbin:/sbin/ ; date -s '2017-10-01 22:17:25' Oct 2 01:30:08 abc vhclicent.23832.daemon: vhclicent instantiated by /bin/sh Oct 2 01:30:08 abc vhclicent.23832.daemon: rsh -JU __juniper_private5__ 192.168.1.2 export PATH=\$PATH:/usr/sbin:/sbin/ ; date -s '2017-10-02 01:30:08' Oct 2 01:30:08 abc vhclicent.23845.daemon: vhclicent instantiated by /bin/sh Oct 2 01:30:08 abc vhclicent.23845.daemon: rsh -JU __juniper_private5__ 192.168.1.2 export PATH=\$PATH:/usr/sbin:/sbin/ ; hwclock -w.** These logs are harmless however might fill up messages file. [PR1315128](#)
- Making changes in services traffic-load-balance instance for one instance, can lead to refresh of existing instances. [PR1318184](#)
- When certain MPC (Modular Port Concentrator) model like MPC4E has very specific hardware failure and it fails to boot up because of FPC (Flexible PIC Concentrator) internal I2C error, other FPCs might go offline. [PR1319560](#)
- With regards to FPC restarts/Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections and the reactions to failure situation might not get handled gracefully. The TCP connection timeout occurs because of jlock hog crossing boundary value (5 seconds) causing bad consequences in MX Series Virtual Chassis. Currently, the only solution would be to enable marker infrastructure during MX Series Virtual Chassis setup. [SPR1332765](#)
- Some vmhost commands are missing on Zero Touch Provisioning (ZTP) for MX Series platforms with VM host support (that is, next-generation Routing Engines, such as RE-S-X6-64G). This might cause the ZTP for vmhost images to fail on this kind of platform. [PR1343338](#)
- Changed the implementation of AES-GCM-256 to improve compatibility with other vendors' MACSec 256-bit implementation. [PR1336834](#)
- On MX Series platform with 100M SFP used on MIC-3D-20GE-SFP-E/MIC-3D-20GE-SFP-EH, SFP might not work if it's not from Fiberxon or Avago. [PR1344208](#)
- ancpd might generate a core file when ANCP subscribers are cleared in a scaled scenario with **enhanced-ip** configured. [PR1344805](#)
- During stress conditions, error log messages regarding route add, change, or delete might be incorrect. [PR1350713](#)
- When an ephemeral database instance is configured, if committing changes that are unrelated to IGMP/MLD (such as **set interfaces ge-0/0/1.0 description**), and the number of ephemeral commits reaches the maximum number, the ephemeral database might purge all commits and roll over. 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 them

back again. This might cause PIM to prune the groups on those interfaces and send join messages again. Finally, multicast traffic flapping and drop might be seen. [PR1352499](#)

- In a L2BSA subscriber scenario, if there is a misconfiguration on the RADIUS profile for the L2BSA subscriber (for example, the routing instance returned from RADIUS is not configured as VPLS) or the authentication part is missing from the physical interface configuration, the bbe-smgd process might crash during L2BSA subscriber log in. [PR1367472](#)
- An FPC restart or FPC core file under heavy traffic load might lead to generation of a bbe-smgd core file. The core file is created due to cleanup issues with the VLAN creations in flight. [PR1371926](#)
- On MX Series routers enabled with enhanced subscriber management, if the subscriber profile initiates a filter service for each subscriber, and there are a very large number of broadband edge subscribers (for example, 10,000) logging in and out repeatedly, the filter service might fail to get installed for the subscriber. In some rare conditions, it might also lead to an FPC crash. [PR1374248](#)
- On MX Series routers in a subscriber scenario, if a large number of subscribers (for example, more than 1000) set up connections simultaneously, the setup rate might be 30 percent lower than expected. [PR1384722](#)
- If **commit fast-synchronize** is enabled, the device with more than five IP addresses configured in the DHCP server group might go into amnesiac mode after reboot. But in practice it should not allow more than five IP addresses based on the implementation, and this validation for "commit check" is skipped when **fast-synchronize** is configured. [PR1385902](#)
- In low-end 32-bit systems, rpd has a lower level of available memory. It is desired to have a log message to alert customers when the average memory usage or transient memory usage exceeds thresholds. [PR1387465](#)
- After upgrading to Junos OS Release 17.2 or later releases, the **chained-composite-next-hop ingress l3vpn extended-space** configuration statement cannot be configured any longer on a logical system. [PR1402390](#)
- On MX Series and ACX Series platforms, when you offline and then online the MIC-3D-16CHE1-T1-CE-H card, the related FPC might crash. [PR1402563](#)
- When **auto-bandwidth** is configured for RSVP LSP, when timeout occurs during LSP statistics query, large bandwidth might be incorrectly reserved for the LSP. If there are no sufficient resources (for example, bandwidth or alternative path) in the network, other LSPs might be torn down, or might not go up. [PR1406822](#)
- The process rpd might crash after a non forwarding route (that is, a route to an indirect next hop association is a non forwarding indirect next hop) that is received from multiple protocols is resolved again by using the non forwarding path. [PR1407408](#)

High Availability (HA) and Resiliency

- Ksyncd might return wrong address if vccpd is slow in setting protocol mode which lead to VC-BM can't sync with VC-MM when VC-B split and reforming Virtual Chassis. Restart ksyncd daemon in VC-BM to restore. [PR1361617](#)

Infrastructure

- Starting from Junos OS Release 14.2R3 the **show class-of-service fabric statistics** command might fail, displaying the **Error = Operation timed out** message in some cases (especially if there are many FPCs in the chassis). This is because data structures used to query fabric statistics became significantly larger in later releases. Thus when multiple FPCs start transmitting data to the Routing Engine at the same time, some packets might get dropped in the internal Ethernet switch on the Control Board. If re-transmission does not happen within the timeout, the **Operation timed out** error is seen. [PR1228293](#)
- The **set system ports console log-out-on-disconnect** command logs the user out from the console and closes the console connection. If the **set system syslog console any warning** command 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 using a Telnet connection to the console. As a workaround, remove the later configuration by using the **set system syslog console any warning** command, which solves the issue. [PR1230657](#)
- System log (syslog) messages are observed when one of the following CLI commands is executed: **system syslog file messages kernel any** or **system syslogfile messages any any**. These syslog messages do not indicate any functionality, breakage, or impact. If you need to enable **any-any**, then you must skip these logs with an appropriate match condition. [PR1239651](#)

Interfaces and Chassis

- During the configuration change and reuse of the VIP address on an interface, you need to stop the configuration, do a commit and then add the interface address configuration in the next commit. [PR1191371](#)
- In a VPLS multihoming scenario, the CFM packets are forwarded over the standby PE device link, resulting in duplicate packets or a loop between the active and standby links. [PR1253542](#)
- Out-of-sequence packets are seen with the LSQ interface. [PR1258258](#)
- In Junos OS BNG solutions, after commit event, when configuration contains duplicate **vlan-id** configured on aggregate and demux interfaces, the MX Series router might go into database prompt mode and a kernel core file is generated. [PR1274038](#)
- Junos OS upgrade involving Junos OS Release 14.2R5 and later 14.2 maintenance releases, and Junos OS Release 16.1 maintenance releases with CFM configuration can cause cfmd process crash after upgrade. This issue is due the old version of **/var/db/cfm.db**. [PR1281073](#)

- Y.1731 delay measurement is not supported on MPC6. [PR1303672](#)
- This PR is to suppress the unnecessary cfmd logs such as the following: `Mar 9 11:30:51.614 2018 MX cfmd[28796]: %DAEMON-3: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0 Mar 9 11:30:51.614 2018 MX cfmd[28796]: %DAEMON-3: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x8d69160 Mar 9 11:30:51.614 2018 MX cfmd[28796]: %DAEMON-3: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0 Mar 9 11:30:51.614 2018 MX cfmd[28796]: %DAEMON-3: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x8d69160 Mar 9 11:30:51.614 2018 MX cfmd[28796]: %DAEMON-3: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0 Mar 9 11:30:51.614 2018 MX cfmd[28796]: %DAEMON-3: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0`. [PR1347650](#)
- With the following configuration present, the interface flaps after a commit in which an aggregated Ethernet interface is being added: `set interfaces <interface-name> otn-options trigger oc-tsf hold-time up <> down <> set interfaces <interface-name> otn-options trigger odu-bei hold-time up <> down <>`
- In a subscriber management environment, the subscriber (for example, subscriber A) might not access the device (A can get IP address x.x.x.x but then the connection will be terminated), because the address x.x.x.x was previously assigned to another subscriber B and then reassigned to A before confirming whether the respective access route for address x.x.x.x was removed. [PR1405055](#)

Layer 2 Ethernet Services

- When MSTP is configured under a routing instance, both the primary and standby VPLS pseudowires get stuck in ST state because of a bug in the software. This issue has been fixed and now the pseudowire status is set correctly. [PR1206106](#)
- After the underlying physical interface for a static VLAN demux interface, the **NAS-Port-ID** formed is based on the previous physical interface. [PR1255377](#)
- DHCP core file might be generated after deleting and adding the VPLS/BD related configuration in one commit. [PR1267810](#)
- When a configuration change adds an existing interface to a new routing instance or logical system and the same configuration change is used to enable BBE DHCP subscriber functionality on that routing instance, the client creation might fail. [PR1294274](#)
- MX Series routers might display the false positive CB alarm **PMBus Device Fail**. [PR1298612](#)
- Port-extender (RJ-45 ports only) LAG interfaces are not up after SNOS 3.1R1.4 upgrade. [PR1354718](#)

Layer 2 Features

- Because of MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes on a router functioning as a VPLS PE device and equipped with one of the following line card:

T4000-FPC5-3D, MX-MPC3E-3D, MPC5E-40G10G, MPC5EQ-40G10G, MPC6E, MX2K-MPC6E. [PR1148971](#)

- On routers running Junos OS with Routing Engine GRES enabled, if VPLS is configured with a dynamic profile association, some traffic loss is observed when the Routing Engine switches from master to standby. This issue is due to a change in the underlying database that handles the dynamic profile sessions. As a result, the VPLS connection is destroyed and re-created after a Routing Engine switchover. [PR1220171](#)
- If an LDP-VPLS routing instance is configured with active and backup neighbors, and flow label capability is enabled on the active neighbor but not on the backup neighbor, upon switching to the pseudowire to the backup neighbor, Junos OS on the VPLS PE device will continue to send traffic with flow label based on the capability learned from the previously active neighbor. [PR1393447](#)

MPLS

- When using **mpls traffic-engineering bgp-igp-both-ribs** with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit ABR when there are two or more such 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 ABRs 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 performed 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 doing so in 64-bit mode. The issue is seen when you use Junos OS Release 13.3 or later with the configuration statement **auto-64-bit** configured, or 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, use the CLI command **set system processes routing force-32-bit**. [PR1141728](#)
- If the **minimum-bandwidth** and **bandwidth** options are both present in the configuration, the bandwidth selection of the LSP is inconsistent. [PR1142443](#)
- When flow-label (FL) is enabled for pseudowire, the OAM packets are not sent with flow-label because rpd is not aware of the flow-label values assigned by the Packet Forwarding Engine. Hence, the packets are getting dropped by the Packet Forwarding Engine at the egress PE device. The remote PE device was expecting the packet with FL and pseudowire label. [PR1217566](#)
- In a CE-CE setup, traffic loss might be observed over the secondary LSP when the primary LSP fails over. [PR1240892](#)
- On MX Series and PTX Series platforms, the rpd might crash when the RSVP bypass undergoes reoptimization and the reoptimized instance encounters failure before it becomes the main instance. [PR1250253](#)
- A new configuration, **protocols mpls traffic-engineering bgp-igp-both-ribs**, in the routing instance is required to make a channelized optical carrier (cOC) work. [PR1252043](#)

- Because of the current way of calculating bandwidth, you see a minimal discrepancy between MPLS statistics and the adjusted bandwidth reported. The algorithm is enhanced so that both values match 100 percent. [PR1259500](#)
- The throughput measurement might be inaccurate when the performance of an MPLS LSP is measured. [PR1274822](#)
- 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](#)
- After the RSVP MPLS LSP link flaps (link goes down and comes back up), RSVP tries to create a second MPLS LSP instance, if Resv/PathErr message drops for the second MPLS LSP instance, then the second MPLS LSP instance is stuck, and no further optimizations are possible. [PR1338559](#)
- If an inet address is not configured for the gr- interface, the gr- interface will borrow an address from the loopback interface. From Junos OS Release 16.1R1 onwards, the RSVP creates a node-neighbor by default. There are duplicate neighbors with the same IP address because the gr- interface is an borrowing address from the 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 an address from the loopback interface. [PR1340950](#)
- Executing a **restart chassisd** in an MX Series Virtual Chassis router with the following elements configured might result in a core file being generated. (1) IGP: OSPF/OSPF3 (area 0, LFA), IS-IS (Level 2, LFA) LDP synchronization, IPv4 and IPv6 (2) IBGP: dual, redundant route reflection, IPv4 and IPv6 (3) MPLS: LDP (IGP synchronization, track IGP metric), RSVP (node link protection, adaptive, auto bandwidth, refresh reduction) (4) L3VPN: OSPF, OSPF3, BGPv4, BGPv6, RIPv2, static, MBGP, next-generation MVPN, L3VPN CNH with ext space, any-to-any, hub and spoke, MPLS access, Ethernet access, multicast extranet, per-VPN and per-prefix labels, SRX Series based network address translation, SRX based firewall (5) Direct Internet Access: EBGP (6) CoS: BA/MF classification, policing/shaping, queuing/scheduling, hierarchical queuing/shaping/scheduling, eight traffic classes (7) BFD/OAM/CFM: liveness detection (8) Load balancing: L2 aggregated Ethernet, IP ECMP , and MPLS ECMP (9) High Availability GRES/NSR, ISSU, fabric redundancy, tail end protection, and BGP prefix independent convergence edge (10) Security: loopback filter, ARP policers, control plane traffic policers, unicast RPF check with all feasible paths, TTL filtering, J-Flow/IPFIX export only, and SRX Series based DDoS. [PR1352227](#)
- Traceroute MPLS from Juniper to Huawei routers does not work as expected due to unsupported TLV. [PR1363641](#)
- If RSVP is disabled and reenabled globally, and in a rare situation, the new RSVP task tries to access memory allocated by the old RSVP task during a particular RSVP Path State Control Block changed path, then the rpd might crash. [PR1366243](#)
- When **protocols ldp dual-transport inet-lsr-id** is not the same as router-id, LDP fails to advertise Layer 2 circuit label mapping to its neighbor. Thus, the Layer 2 circuit will not come up properly. [PR1405359](#)

Network Management and Monitoring

- While polling the Ethernet connectivity fault management protocols statistics, the SNMP process might crash. [PR1364001](#)

Platform and Infrastructure

- When there is huge logical interface scaling on aggregated Ethernet interfaces (500 or more) with more than 32 member links and when all FPCs are restarted one by one, followed by member link addition to the link aggregation group (LAG), the state dependency evaluation in the kernel takes a long time given the scale involved. As a result, the FPCs do not get all the states from the Routing Engine. This is an uncommon sequence of events or conditions. [PR938592](#)
- When TCP authentication is enabled on a TCP session, the TCP session might not use the selective acknowledgment (SACK) TCP extensions. [PR1024798](#)
- When using **show | compare** method to commit, part of the configuration might be treated as noise and return syntax error. [PR1042512](#)
- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- On MX Series routers, parity memory errors might occur in pre-classifier engines within an MPC. Packets are silently discarded because such errors are not reported and hence harder to diagnose. CM errors such as syslog messages and alarms should be raised when parity memory errors occur. [PR1059137](#)
- CoS error messages might appear when a nonexistent path for a database file is configured for CoS. These messages do not affect any service or traffic. [PR1158127](#)
- Access to a stale or invalid pointer causes a particular check based on the pointer structure field to unpredictably fail, resulting in the assert later in the code. The issue occurs when a sequence of events related to firewall filters results in the filter structure getting deleted and re-created. [PR1205325](#)
- Several files such as **/var/etc/mobile_aaa_ne.id**, **/var/etc/mobile_aaa_radius.id** are copied between Routing Engines during the **ffp synchronize** phase of the commit. These files are copied even if there is no corresponding change in the configuration, thus unnecessarily increasing commit time. [PR1210986](#)
- In Junos OS Release 17.1R1 and earlier, MGD with extend-db feature supports a database of 2.5 GB size (maximum) on 64-bit platforms, which is a problem solved through this PR. After this PR, the maximum configurable database size supported with the extend-db feature is 1.5 GB on i386 platforms (both 32 bit and 64 bit). [PR1228629](#)
- Starting In Junos OS Release 13.3, SRX Series clusters need to run auditd on both nodes. However, on MX-VC Bm and TXP all LCC also add auditd. Because LCC and VC-BM do not have a route for the accounting server, the following message is generated: **813 unreachable infor. user@router> show system processes extensive | match "-re|audit" sfc0-re0:**

```
----- 2565 root 1 96 0 3304K 2620K RUN
```

0:01 0.00% auditd lcc0-re0: ----- 2398 root
1 96 0 3240K 2536K select 0:01 0.00% auditd lcc1-re0:

----- 2791 root 1 96 0 3244K 2544K select
0:01 0.00% auditd %DAEMON-3: auditd[2398]: sendmsg to 10.233.225.78(10.233.225.78).1813 failed:
Network is down %DAEMON-3: auditd[2398]: AUDITD_RADIUS_REQ_SEND_ERROR: auditd_rad_send:
sendto/sendmsg: Network is down [PR1238002](#)

- When certain hardware transient failures occur on an MQ-chip-based MPC, traffic might be dropped on the MPC, and syslog errors **Link sanity checks** and **Cell underflow** are reported. There is no major alarm or self-healing mechanism for this condition. [PR1265548](#)
- This issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have XE/GE links toward the downstream switch and LAG bundles as uplinks towards upstream routers. The XE/GE link is part of the physical loop in the topology. Spanning tree protocols such as VSTP/RSTP/MSTP are used for loop avoidance. Some MAC addresses are not learned on DUT when LAG bundles that are part of such bridge domains are flapped and other events such as spanning tree root bridge change occur. [PR1275544](#)
- Even though multicast appears to be active with **show multicast route extensive** command, it is not forwarded to the subscriber interface. [PR1277744](#)
- When **chassis control restart** is done with the CoS rewrite rule configured on the aggregated Ethernet interface, the **Platform failed to bind rewrite** message could be seen in syslog. This issue is specific to aggregated Ethernet interfaces. It is a timing issue that can occur when the logical interface deletion is delayed due to high scale and when the logical interfaces come up again after restart they have different indexes. [PR1315437](#)
- Output policing action for EVPN-VXLAN might not be applied to an interface despite configuration on the IRB interface. [PR1348089](#)
- On MX Series routers enabled with next-generation subscriber management, if subscribers are enabled with distributed IGMP, and there are some stressful operations (for example, subscribers log in or log out as well as join or leave IGMP groups repeatedly) some line cards might crash due to the timing issue. [PR1355334](#)
- In a layer 3 VPN topology, traceroute to a remote PE for a CE-facing network see the ICMP TTL expired reply with a source address of only one of the many CE-facing networks. In 15.1R5, 16.1R3, and 16.2R1+ there is a kernel sysctl value, icmp.traceroute_l3vpn. Setting this to 1 will change the behavior to selected an address based on destination specified in the traceroute command. This PR adds the option to the configuration. [PR1358376](#)
- Sometimes the OSPF flaps while performing unified ISSU from Junos OS Release 16.2R2 to Release 17.2R3. [PR1371879](#)

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 nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- When only the default routing instance is present, the **show bgp summary** command 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 the master routing instance, inet.0), the BGP establish state is showed properly. This issue occurs for IPv4 BGP sessions only. On IPv6, we always see all the BGP states as default. [PR600308](#)
- Soft core files might be continuously generated because of the **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, which is automatically created after the **soft-assert-core** file is collected, without any impact to the traffic or service. [PR815146](#)
- In rare cases, rpd might generate a core file with the **rt_notbest_sanity: Path selection failure on ...** error. The core is soft, which means there should be no impact to traffic or routing protocols. [PR946415](#)
- During interoperation with other vendors in a draft-rosen multicast VPN, by default the Junos OS attaches a route target to multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities will be excluded from propagating if the BGP route target filtering is enabled on the device running Junos OS. [PR993870](#)
- On MX Series routers, when an instance type is changed from VPLS to EVPN, and in the same commit operation an interface is added to the EVPN instance, the newly added EVPN interface might not be able to come up. [PR1016797](#)
- When LDP is deactivated, in a rare case, the result of remote loop-free alternate (remote LFA) might be computed to go through the deactivated LDP node. The situation is self-recovered in the next SPF calculation. [PR1202392](#)
- JTASK_SCHED_SLIP for rpd might be seen on doing restart routing or ospf protocol disable with scaled bgp routes in MX104 router [PR1203979](#)
- In the context of a large number of configured VPNs, routes changing during a BGP path-selection configuration change can sometimes lead to the generation of an rpd core file. The core file has been seen to be generated after the **always-compare-med** option is removed. [PR1213131](#)
- The rpd process leaks memory as a result of topology and configuration. However, adding or deleting static flowspec routes in isolation does not cause any memory leak. The exact configuration that causes the leak is currently unknown. [PR1213959](#)

- 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](#)
- A few Bidirectional Forwarding Detection (BFD) protocol sessions flap while coming up after an FPC reboot. The flapping does not impact the system, because it occurs during the bring-up phase. The issue occurs because of a race condition in PPMAN code. [PR1274941](#)
- When **route-distinguisher-id** is configured and VRF with a route distinguisher is automatically assigned with the **auto-rd** feature configured, the MX Series BNG allows the configuration to be committed but this is followed by rpd process crash. [PR1278582](#)
- In an MVPN (multicast VPN) scenario, if **routing-instances <instance name> protocols pim static** is configured, the rpd might crash when the **deactivate routing-instances instance name protocols pim static** to deactivate the routing instance of PIM static. [PR1284760](#)
- Backup Routing Engine scheduler slips when Cisco Rosen7 PE with MDT-SAFI is enabled. However, the MDT-SAFI update does not include the route-target extended community attribute, NSR is enabled, policies are set to import or export the inet-mdt table, but Rosen is not configured. [PR1295712](#)
- An MX104 is connected to an SRX1500. IS-IS is running between these device and BFD has been configured between the IS-IS peers. Unfortunately, BFD does not come up between these devices successfully. [PR1312298](#)
- The rpd process might crash and generate core files in a distributed IGMP environment. [PR1314679](#)
- In RPKI (Resource Public Key Infrastructure) scenario, the validation replication database might have much more entries than the validation database after restarting RPKI cache server and the validation session is reestablished. [PR1325037](#)
- When route target filtering (RTF) is configured for VPN routes and multiple BGP sessions 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 BGP, LDP, and IS-IS configurations, deleted IS-IS routes might still be present in the routing table. The PR does not affect or have any impact on route selection or other functionality of rpd. Just that deleted IS-IS routes do not get removed with specific configurations. [PR1329013](#)
- In a large-scale OSPF network (for example, there are more than 500 devices in an area), OSPF remote loop-free alternate default PQ node selection algorithm does not provide the proper protection paths. [PR1335570](#)
- On all platforms running Junos OS and enabled with GRES and NSR, if Routing Engine switchover is executed, the BGP peers in the new master Routing Engine might flap due to hold-timer expiry after GRES. [PR1390113](#)
- With GRES and NSR enabled, if executing switchover, all the BGP sessions might flap. [PR1391084](#)

Services Applications

- On MX Series routers with L2TP configured, the L2TP packet in the ICRQ re-transmission message is set to an incorrect value, and this causes frequent L2TP session flapping. [PR1206542](#)
- It is not recommended to configure an ms- interface when the ams bundle in one-to-one mode has the same member interface. [PR1209660](#)
- The NAT auto-injected routes might fail to install when back-to-back commits with changes made to service sets or NAT rules are performed. This issue occurs with a unique configuration where thousands of routes are added by the service PIC process (spd), which manages installation of NAT return routes and destination routes. [PR1223729](#)
- On a Layer 2 Tunneling Protocol (L2TP) access concentrator (LAC) router where Access Node Control Protocol (ANCP) is used for bandwidth adjustment, the L2TP Connect Speed Update Notification (CSUN) message might be sent to the L2TP network server (LNS) after a short delay after the ANCP Port Up message with updated access line parameters was received. This delay is caused by the current interaction scheme between ANCP and L2TP daemons and can last up to 5 seconds. In a production network scenario, this delay should not be visible as the L2TP daemon checks for state updates each time there is an L2TP packet to be sent or received. [PR1234674](#)
- If an L2TP subscriber has a static pp0 interface on the LAC side, LCP renegotiation is configured on the LNS side, and CPE has been changed, an issue with successful negotiation of the PPP session between LNS and the CPE device can occur. [PR1235554](#)

Subscriber Access Management

- On MX Series routers with the subscriber management feature enabled, after GRES switchover the **show network-access aaa statistics radius** command displays only zeros and the **clear network-access aaa statistics radius** command does clear statistics as it should. This is a cosmetic issue and communication with the RADIUS server is working fine. However, the affected CLI commands do not work as expected. [PR1208735](#)
- In Point-to-Point Protocol over Ethernet (PPPoE) subscribers scenario with a large number of subscribers (for example, 3000), during login and logout operations, some subscribers might be stuck in an error state of "Terminated". This issue impacts the traffic for these subscribers. [PR1262219](#)
- Authd generates a core file when **gx-plus** is enabled in getDynamicRequestSource. [PR1277137](#)
- Usage-monitoring-information AVP as part of PCRF gx-plus provisioning is causing service accounting activation. [PR1391411](#)
- In a subscriber scenario, the authd might crash multiple times due to a memory corruption issue. [PR1402012](#)

User Interface and Configuration

- When frequent load replace operations are being performed on the router, commit might take longer. [PR1029477](#)
- When **persist-groups-inheritance** is configured, doing configuration changes and issuing rollback might cause persist-groups tree corruption and eventually cause improper configuration propagation after commit. This situation might lead to mgd process (daemon) crash as well. [PR1214743](#)
- The log messages **ffp[52861]: %DAEMON-3: "dynamic-profiles": Profiles are being modified** can appear after a configuration is change but the dynamic profiles are not modified. The issue was reproduced by changing the configuration by using NETCONF but without commit so the configuration changes were cleared after ending the NETCONF session. The next configuration change with commit generates the log messages **Profiles are being modified**. [PR1234446](#)

VPNs

- In a next-generation MVPN scenario, when **forwarding-cache timeout never non-discard-entry-only** is configured for an MVPN instance, even though the cache lifetime is shown as forever in the output of the CLI command **show multicast route instance X extensive**, the route disappears after 7–8 minutes. [PR1212061](#)
- In an MVPN setup with the **SPT-only** option, if the source or receiver is connected directly to the candidate RP PE router and the MVPN data packet arrives at the candidate RP PE router before its transition to SPT, the MVPN data packet will be dropped. [PR1223434](#)
- Starting in Junos OS Release 15.1F5, under the next-generation MVPN environment, when multicast production data is stopped, VRF S,G entry and MVPN/BGP routes might persist, whereas they should be deleted. [PR1236733](#)
- In a multicast VPN with Border Gateway Protocol (next-generation MVPN) scenario with only SPT mode configuration, under certain conditions the PIM register-stop packet might be sent before the Source Tree Join (Type-7) packet, which might cause some multicast packets to drop. [PR1238916](#)
- In an MVPN scenario with I-PMSI tunnels and multihomed source, if the link between Source and PIM-DR PE1 goes down, then the second PE2 takes the PIM-DR role and starts to advertise Type-5 prefixes. Then as the link between the Source and PE1 comes back up and PE1 takes the PIM-DR role back, PE1 might not generate Type-5 BGP prefixes for active sources in some multicast groups. Without Type-5 prefixes from the ingress PE device, the receivers' PE device do not generate Type-6 or Type-7 and the ingress PE device does not send multicast traffic. Workaround: Clear PIM joins in the affected instance. `PE1> clear pim join instance _MVPN_instance_name_.` [PR1242493](#)
- When a C-multicast route (Type 7 or Type 6) for inter-as non-segmented option C topology is sent with the originator's IP address, Junos source PE does not accept this route and hence the PIM join fails. [PR1327439](#)

SEE ALSO

New and Changed Features	99
Changes in Behavior and Syntax	125
Known Behavior	142
Resolved Issues	169
Documentation Updates	232
Migration, Upgrade, and Downgrade Instructions	234
Product Compatibility	241

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.1R3 | [169](#)
- Resolved Issues: 17.1R2 | [215](#)
- Resolved Issues: 17.1R1 | [224](#)

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: 17.1R3

Application Layer Gateways (ALGs)

- IPsec IKEv2 negotiation fails when IKE ALG is enabled. [PR1300448](#)
- IKEv2 negotiation might fail when IKE ESP ALG is enabled in an IKEv2 redirection scenario. [PR1329611](#)

Authentication and Access Control

- MAC move might occur in a DHCP security scenario. [PR1369785](#)

Class of Service (CoS)

- The Routing Engine level **scheduler-hierarchy** command misses a forwarding class when the **per-unit-scheduler** mode is configured. [PR1281523](#)
- CoS wildcard configuration is applied incorrectly after router restart. [PR1325708](#)
- The Routing Engine might get into amnesiac mode after restarting if **excess-bandwidth-share** is configured. [PR1348698](#)
- The 802.1P rewrite might not work on inner VLAN. [PR1375189](#)
- The FPC might reboot when changing CoS mode from **hierarchical-scheduler** to **per-unit-scheduler**. [PR1387987](#)

EVPN

- In an EVPN scenario, the rpd might crash during MAC moving in Active-Active and Active-Standby multihomed PE devices. [PR1216144](#)
- MAC entry is incorrectly programmed in the Packet Forwarding Engine, leading to some traffic to be discarded with notification. [PR1231402](#)
- In VXLAN-EVPN, the VLAN tag of egress ARP reply is removed if egress interface is lt-. [PR1252522](#)
- Rpd might crash with signature similar at **evpn_mirror_mac_process_update_instance**. [PR1258835](#)
- An FPC or MPC might crash in an EVPN/MPLS or EVPN/VXLAN environment. [PR1274976](#)
- Ethernet A-D route per Ethernet segment (Type-1 Per es) is not generated with a new route target after the vrf-target is changed. [PR1279529](#)
- EVPN traffic mapping to specific LSPs is not working. [PR1281415](#)
- Local preference for EVPN type-5 route might cause unexpected results if BGP multipaths are configured. [PR1292234](#)
- Traffic might be dropped after updated ARP route update packets are received from the peer Layer3 gateway in an EVPN/VXLAN scenario. [PR1306024](#)
- In an EVPN environment the rpd might crash on QFX10002 after the rpd process is restarted. [PR1320408](#)
- Discard EVPN route is installed on the local PE device after connection flaps on a remote PE device in a multihomed EVPN topology. [PR1321125](#)
- The rpd might crash during EVPN-VXLAN configuration changes. [PR1321839](#)
- In an EVPN/NSR scenario, the rpd process crashes and generates a core file on the backup Routing Engine when any configuration is changed on the master Routing Engine. [PR1336881](#)

- The rpd might crash if the IRB interface and routing instance are deleted together in the same commit operation. [PR1345519](#)
- The rpd might crash if the EVPN instance refers to a vrf-export policy that does not have “then community”. [PR1360437](#)

Forwarding and Sampling

- Firewall daemon might leak some memory when a filter-based forwarding (FBF) configuration is committed. [PR1157714](#)
- Accounting Interim Interval is reset after GRES. [PR1261472](#)
- Unexpected messages might be seen in logs. [PR1270686](#)
- Sampled stops collecting data on Routing Engine based sampling supported platforms. [PR1270723](#)
- Firewall filter might not be matched when a wildcard (*.*) is specified as a matching condition. [PR1274507](#)
- Unicast traffic is forwarded out of the logical interface even after the interface is disabled. [PR1277697](#)
- The sampled route reflector process (srrd) might crash in the large routes churn situation. [PR1284918](#)
- The mib2d process might crash if an SNMP walk is active when a configuration is committed or rolled back. [PR1286448](#)
- sampled core at, strlen.c:100 when sampling is enabled or modified on an interface. [PR1289530](#)
- Observing pfed core files 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](#)
- Some accounting files might be missed in case the remote archive sites are unreachable. [PR1300764](#)
- There is memory leak on mib2d when firewall MIBs are polled. [PR1302553](#)
- The dfwd process might crash during the execution of the **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 commit might fail when the **nexthop-learning** configuration statement is enabled for J-Flow v9. [PR1316349](#)
- The FPC CPU usage might continue to be at 100% if **shared-bandwidth-policer** is configured. [PR1320349](#)
- DHCP service crashes after switch/router 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](#)

- The L2ALD daemon might crash if a duplicate MAC is learned by two different interfaces. [PR1338688](#)
- In EVPN-VXLAN, the **clear ethernet-switching table** command might not work correctly. [PR1341328](#)
- Junos OS allows firewall filters with the same name under **edit firewall** and **edit firewall family inet** hierarchy levels. [PR1344506](#)
- Commit fails when you attempt to delete any demux0 unit numbers that are greater than or equal to 1000,000,000. [PR1348587](#)
- Backup Routing Engine writing dummy interface accounting records. [PR1361403](#)
- The filter counter is not written to the accounting file when accounting is enabled on the bridge firewall filter. [PR1392550](#)

General Routing

- Enhanced IP/enhanced-Ethernet and MS-DPC compatibility. [PR1035484](#)
- ICMP reply traffic might get dropped on MS-MPC line cards. [PR1059940](#)
- In timing hybrid mode MX MPC2 cards are not working with ACX with VLAN (native-vlan-id.). [PR1076666](#)
- Memory leak on when an Layer 3 VPN configuration is committed for a Layer 3 VPN scaling test. [PR1115686](#)
- The **show storm_cntl help database** command on the FPC might cause a crash. [PR1127870](#)
- No warning is raised when the bridge family is configured with **interface-mode** trunk but without **VLAN-tagging** or **flexible-VLAN-tagging** . [PR1154024](#)
- The ksyncd process might crash because of transient replication errors between Routing Engines. [PR1161487](#)
- Unexpected MobileNext Gateway Activation license alarm is raised when TDF gateway is configured. [PR1162518](#)
- Kernel displays I2C bus timeout errors when there are multiple commit processes. [PR1174001](#)
- The KRT queue might be stuck if the rpd sends two deleted requests to the kernel for the same next hop. [PR1186334](#)
- SNMP trap sent for **PEM Input failure** alarm is not generated when a single input feed fails on MX960. [PR1189641](#)
- The replacement PIC might bounce when the 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 the configuration is committed. [PR1190569](#)
- The CLI commands **request vmhost zeroize** and **request vmhost zeroize both** might work only on the local Routing Engine. [PR1197152](#)
- The **chassisd[1825]: pvidb_get_root_node: Error(2) retrieving rootnode value** error messages might be seen. [PR1198817](#)

- The rpd might crash after the configuration is committed. [PR1200174](#)
- The RSVP auto-mesh is flapping every 15 minutes for BGP peers that have only EVPN address family enabled and data is going across the PE device. [PR1202926](#)
- Stale VBF states occur without sdb sessions. [PR1204369](#)
- SMID daemon stopped responding to the management requests. [PR1205546](#)
- IPsec phase2 soft lifetime calculation is different between Junos OS Release 11.4R12 and Junos OS Release 14.2R6. [PR1209883](#)
- TACACS access does not work after upgrade. [PR1220671](#)
- CMIC:CMIC(0/1): Unable to deregister sub error (131072) for error(0x1b0001) for module MIC. Error messages are seen on MPC5E. [PR1221337](#)
- mqtt Routing Engine scope : continuous pdb_open error messages for Routing Engine scope MQTT broker. [PR1224705](#)
- CoS service with Reflexive CoS-rule should modify CoS values for reverse flow. [PR1227021](#)
- 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](#)
- MPC2E-NG and MPC3E-NG generate core files with a specific MIC because of a tight loop of PCI express critical exceptions. [PR1231167](#)
- False AC PEM failure(status bits: 0xff) alarm/SNMP trap is seen with MX5, MX10, MX40, and MX80 routers. [PR1231893](#)
- FPC might crash and PTP might remain in **INITIALIZING** state after configuration commit. [PR1232740](#)
- Power OK SNMP trap (jnxPowerSupplyOk) is not raised when PEM and SCB are inserted in the chassis. [PR1232885](#)
- Major errors related to XQ-chip L4NP parity errors might be reported on the MPC. [PR1232952](#)
- The MS-MPC might crash when OSPFv3 IPv6 traffic goes through it. [PR1233459](#)
- FPCs go offline with **FPC Incompatible with SCB** during system restart. [PR1235132](#)
- No DNS information in the output of the **show subscribers extensive** command for DHCP subscribers. [PR1237525](#)
- The **multicast-replication** setting cannot be reflected in the redundancy environment after rebooting both Routing Engines. [PR1240524](#)
- Tracking PR for enabling mobiled for an MX-VC environment. [PR1241857](#)
- JDI-RCT:vRCT: chassisd[9132]: LIBJSNMP_NS_LOG_NOTICE: NOTICE: netsnmp_ipc_client_connection: unix connection error: socket(-1) main_session(0x9812f80) error messages are seen after chassis-control restart. [PR1243364](#)
- XM chip-based line card might drop traffic under high temperature. [PR1244375](#)

- MX104 router does not report the HSL2 CMERROR alarm upon HSL2 CRC errors. [PR1247707](#)
- The error messages about "jlock hog" might be seen after restarting routing in large scale of routes. [PR1248246](#)
- The RPT-PHY-RIAD:shm-rtssdbd.core@0x080529e3 in rtssdb_async_msg (state=0x88715844, rtssb_addr=...) core is seen at ../../../../src/junos/lib/libshm-rtssdb/rtssdb_server.c:1729. [PR1249116](#)
- MACsec session fails with the generation of dot1x core files. [PR1251508](#)
- The EOAM LFM adjacency on an MX Series MPC or EX9200 might flap when an unrelated MIC, which is in the same MPC slot, is brought online. [PR1253102](#)
- KRT queue remain stuck with "EINVAL -- Bad parameter in request" in the new master Routing Engine after unified ISSU. [PR1254519](#)
- The **validation-state:unverified** routing entry might not be shown with proper location in show route. [PR1254675](#)
- Prolonged flow-control core observed for the TFTP ALG traffic (10,000 simulated users). [PR1255973](#)
- Dynamically injected routes for the concentrator and the IPv6 prefix can be programmed with the wrong precedence. [PR1256672](#)
- MIB walk ascii jnxFabricMib provides wrong DropBytes statistics. [PR1257569](#)
- The rpd might crash during a next-hop change, if unicast reverse path forwarding (uRPF) is used. [PR1258472](#)
- The rpd might crash during the next-hop change, if unicast RPF is used. [PR1258472](#)
- The device control process (dcd) crashes during an ATM-related configuration commit. [PR1258744](#)
- MX/QFX/VMX Licensing: license keys entered through the configuration **system license keys** can be lost (not effective anymore) after certain events or changes. [PR1259460](#)
- **HEAP: Free at interrupt level /Free interrupt violation!** syslog messages might be seen when an interface is down. [PR1259757](#)
- Many LCP Term Req and PADT messages are not processed during PPPoE subscriber logout. [PR1260626](#)
- The wrong XML RPC command output is shown for **show route bgp advertise-protocol/receive-protocol**. [PR1261421](#)
- Traffic drops when an MPC has a high rate of cell underflow drops after link sanity check. [PR1262868](#)
- vMX FPC core - panic (format_string=format_string@entry=0x9e509c4 "Thread %s attempted to %s with irq priority at %d\n"). [PR1263117](#)
- Duplicate sensor resources are created when the difference is trailing "/" . [PR1263446](#)
- Extra link transitions might be seen after an MPC is restarted. [PR1264039](#)
- BGP hold time might expire after a GRES or NSR switchover. [PR1264436](#)
- The rpd might crash after some VRF instances are deleted, if **vrf-table-label** is configured. [PR1264464](#)

- PTP is lost when the master line card is restarted. [PR1264530](#)
- All traffic received from specific fabric streams is dropped with only XMCHIP FI: cell underflow error syslog event. [PR1264656](#)
- PCC-controlled LSP metric does not get updated on the controller-PCE-delegated LSPs do not come up. [PR1265864](#)
- The first and last addresses are not translated. [PR1266774](#)
- MPC might report a parity error with the when the **fast-lookup-filter** statement is configured. [PR1266879](#)
- ISSU-related limitation under highly scaled scenarios. [PR1267680](#)
- Junos OS: bbe-smgd process denial of service while processing VLAN authentication requests/rejects (CVE-2018-0006). [PR1268129](#)
- RSI output is augmented by detailed "nhinfo" dumps. [PR1268460](#)
- The openflowd process might get stuck at 100% CPU usage when an OpenFlow filter is deleted and queried at the same time. [PR1268527](#)
- On MX Series, 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 might cause the SIP ALG to generate a core file. [PR1268891](#)
- Wrong "Voltage Threshold Crossed" alarms are seen if keep suspends/re-starts Junos. [PR1269157](#)
- MIC error interrupts excessively loads the CPU when MIC/fpc restart is initiated. [PR1270420](#)
- Multicast traffic silently dropped without notification when uplink is flapping with MoFRR enabled. [PR1270939](#)
- The management daemon (MGD) might crash after you invoke a specific RPC. [PR1271024](#)
- Messages related to **Logical Addr xxxxxxxx is invalid** seems when FPC restart also passing traffic. [PR1271810](#)
- Virtual forwarding plane failed to load files from the virtual control plane if the interconnection has an MTU less than 1500. [PR1273365](#)
- The mspm and log messages about memory zone level are generated incorrectly. [PR1273901](#)
- Some received packets might be incorrectly dropped by DA rejects after a 40-Gigabit Ethernet or 100-Gigabit Ethernet port is configured under a LAG. [PR1274073](#)
- L2-over-GRE tunnel uses the underlying physical interface MTU directly without deducting the IP/GRE header length. [PR1274203](#)
- The IPv6 ping might fail after route leaking policy deployment is done between two Layer 3 VPN routing instances. [PR1274339](#)

- The **show subscribers detail**, **show subscribers extensive**, **show subscribers count client-type <>**, and other commands fail to execute because the subscriber management database is unavailable. [PR1274464](#)
- The FPC might crash when a route is received through BGP and sampled through J-Flow. [PR1275021](#)
- Link stays down after a flap on next-generation MPCs with QSFP+40G direct attach copper (DAC) cable. [PR1275446](#)
- Fixing the default behavior of the configuration statement added for static route's dependency on BFD_ADMIN_DOWN, through PR 1070477. [PR1275973](#)
- When the static link protection mode is configured and the backup state is down, the primary port goes to down state instead of the secondary port, and the secondary port remains in up state. [PR1276156](#)
- Fabric input stream might drop all packets upon sustained oversubscription or when CRC errors are injected on single plane. [PR1276301](#)
- Junos OS does not use the complete TCP window size and slows the connection when JET application over grpc is installed on Junos. [PR1276443](#)
- FPC connections might drop with syslog messages: **CHASSISD_MAIN_THREAD_STALLED: main chassis-control thread stalled for XXX sec -- exiting**. [PR1276605](#)
- Spd memory leak might be observed after the **service-set** statement is added or removed. [PR1276809](#)
- The KRT asynchronous queue might be stuck, which might impact the synchronization for RIB and FIB. [PR1277079](#)
- The l2ald memory might leak for every IPv6 ND message it receives from peer MC-LAG and it is not freeing the memory allocated. [PR1277203](#)
- L2C BUS stuck causes SFP+ thread hogging and MPC restart. [PR1277467](#)
- IS-IS adjacencies over MLPPP links do not connect to the LSQ bundle interface. [PR1278377](#)
- bbe-smgd might generate core files in certain cases when using logical interface sets in Universal Call Admission Control policy mode. [PR1278543](#)
- **jnh_vbf_flow_get_oif_index: Rollback cmd not found for flow** syslog messages generated by MPC during subscriber login. [PR1278580](#)
- The routing protocol process (rpd) might be stuck at 100% when the same BGP prefix routes are learned in different routing instances with multipath and auto-export configured. [PR1279260](#)
- On MX104 with GRES enabled, the chassis network-services might not get set as "Enhanced-IP". [PR1279339](#)
- BBE-smgd core files are generated when the packet is received with unexpected TPID. [PR1279402](#)
- VLAN out-of-band subscriber session fails in auto configured mode. The physical interface goes down even if it is physically up. [PR1279612](#)
- CoS attachment might be attached to the wrong link if issuing some changes to the aggregated Ethernet bundle. [PR1279788](#)

- The temperature value is being displayed as "Testing" in the output of **show chassis fpc detail** after GRES. [PR1280030](#)
- After a MS-MPC or MS-PIC is offline or brought 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](#)
- MTU for a Layer 2 over GRE gr- interface should be unlimited. [PR1281173](#)
- The ingress **service-accounting-deferred** statement is not providing the correct IP traffic statistics for for L2BSA subscribers. [PR1281201](#)
- Subscribers might get stuck in Init state if there is an SDB access error during their login. [PR1281896](#)
- The subscribers might fail to bind after FPC restart followed by bbe-smgd restart. [PR1281930](#)
- Optics levels are not sent in Junos Telemetry Interface for down interfaces. [PR1281943](#)
- Buffer overflow in sockets library (CVE-2017-2344). [PR1282562](#)
- The kernel might crash in a rare corner case. [PR1282573](#)
- 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](#)
- The rpd process might crash if dynamic interfaces are used by multiple applications. [PR1282854](#)
- VBF flows are not programmed correctly on aggregated Ethernet interfaces. [PR1282999](#)
- GRE OAM fails to come up when the GRE tunnel source and family inet address are the same. [PR1283646](#)
- Junos: bbe-smgd process denial of service while processing VLAN authentication requests/rejects (CVE-2018-0006). [PR1284213](#)
- Bad MPLS encoding in Junos Telemetry Interface. [PR1284317](#)
- 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 enhancement of reporting total SBE errors when the corrected single-bit errors threshold of 32 is exceeded for MPC7E, MPC8E, MPC9E. [PR1285315](#)
- LC, PFH, and Packet Forwarding Engine interface are not coming up on RE1. [PR1285606](#)
- 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 paths. [PR1285979](#)
- Internal latency increases overtime for Packet Forwarding Engine sensors with streaming telemetry. [PR1286286](#)

- The missing statement **Shared bandwidth policer not supported for interface ge-x/x/x** is seen, during a failed commit in Junos OS Release 16.1R3. [PR1286330](#)
- Unified ISSU is not supported from Junos OS Release 15.1 onward, because the source release includes one or more BBE features such as logical interface (IFL) options, CoS fragmentation map, MLPPP, advisory options, advanced services, and multicast distribution. [PR1286507](#)
- DDoS culprit flows are not reported by CLI or during login to an MX Series router with a single Packet Forwarding Engine. [PR1286521](#)
- Framed routes might get stuck in the KRT queue. [PR1286849](#)
- The one-set/leaf-list configuration might not get deleted with the delete operation through JSON. [PR1287342](#)
- The LTS to LNS connection is not working 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](#)
- The bbe-smgd process might crash and generate a core file on the standby Routing Engine during a reboot upgrade with active locally terminated PPPoE subscribers. [PR1288121](#)
- During unified ISSU (FRU upgrade) micro BFD flap is observed. [PR1288433](#)
- Performance issues can be seen when nontranslated traffic is introduced to a service set using a large number of NAT terms. [PR1288510](#)
- After GRES smid was declared thrashing and was not restarted after fatal SDB error. [PR1288871](#)
- Kernel 'rtdat' 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 be in down state after GRES. [PR1289493](#)
- NAT-T and DPD functionality do not work for aggressive mode. [PR1290689](#)
- Incorrect temperature is displayed for MPCP5 and MPC7 in the **show chassis fpc** command output. [PR1290771](#)
- Memory leak in bbe-smgd process on subscriber logout for subscribers who have joined any multicast group. [PR1290918](#)
- LSP traffic might silently drop and get discarded after a link goes down in the bypass path. [PR1291036](#)
- The routing protocol process (rpd) might generate a core file when the process is restarted. [PR1291110](#)
- The switch might wrongly learn its own IRB MAC address. [PR1291184](#)

- JDI-RCT-RPD: Device going to the DB prompt `db@jsr_jsm_send_ka_after_merge,send_proto_keepalive` was observed on the master Routing Engine. [PR1291247](#)
- L2TP ICCN fast retransmission occurs after tunnels go down. [PR1291557](#)
- Kernel is not installing the route and throwing an error. [PR1291917](#)
- The `bbe-smgd` process might crash and subscribers might get stuck when a large group of different types of subscribers log in or log out. [PR1291969](#)
- Recursive lookup in Packet Forwarding Engine might happen over a dynamic tunnel. [PR1292425](#)
- An error in `vbf_filter_add_orphan_check` might be seen when the subscribers using filters log out or log in. [PR1292582](#)
- An error message might be seen while bringing up the subscriber in a subscriber management environment. [PR1293057](#)
- **DDR3 TEMP ALARM** messages are logged in `chassisd` log. [PR1293543](#)
- CPCDD core files are generated when Routing Engine based HTTP-redirect is used. [PR1293553](#)
- Performing load replace terminal and attempting to replace the interface stanza might terminate the current CLI session and leave the user session hanging. [PR1293587](#)
- The **show extensible-subscriber-services sessions** command displays an incorrect timestamp after a unified ISSU. [PR1293800](#)
- MX Series router might not honor the do-not-fragment bit in subscriber environment. [PR1294282](#)
- The flow export rate remains lower than the configured export rate in an inline sampling scenario. [PR1294296](#)
- Loss of DHCP/PPPoE subscribers is observed during unified ISSU from 16.1-20170718_161_r4_s5.0 to 16.1-20170718_161_r4_s5.0. [PR1294709](#)
- During PPPoE subscriber login, errors such as [`vbf_flow_src_lookup_enabled`] and [`failed to find iff structure, ifl`] were seen on the FPC. [PR1294710](#)
- The `rpd` might crash if the interface or BGP flaps. [PR1294957](#)
- The KRT queue might be stuck with the error **RPD_KRT_Q_RETRIES: chain nexthop add: Unknown error: 0**. [PR1295756](#)
- `xmlproxyd` generate core files during telemetry streaming. [PR1295831](#)
- The service profile's CoS might be overridden by the client profile's CoS when second family DHCP sessions are added in a dual-stack subscriber scenario. [PR1296002](#)
- The `mspmmand` process might crash if you use SCG services on MS-MPC or MS-MIC. [PR1296422](#)
- The `jdhcpd` might crash when using **dhcp-security** related command in enhanced subscriber management mode. [PR1296461](#)
- The kernel might crash continuously when a lot of terms are configured for firewall filters. [PR1296884](#)

- In ECMP fast reroute scenario, traffic might get silently dropped or discarded because of a next hop in "hold" state. [PR1297251](#)
- The mgd process might consume high Routing Engine CPU when certain **show** commands are executed. [PR1297728](#)
- Some random number of ports on MPC7E-10G card might not come up after the remote system or line card restarts or interface flaps. [PR1298115](#)
- The log message about shutdown time is wrong when the system exceeds chassis over temperature limit. [PR1298414](#)
- The bbe-smgd process might crash when traceoption is enabled due to an invalid username character. [PR1298667](#)
- MX Series BNG does not respond to PADI after GRES on some ports/VLANs. [PR1298890](#)
- The error messages about PEM might be seen in an MX Series router with AC PEM. [PR1299284](#)
- 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](#)
- bbe-smgd can generate core files after a Routing Engine mastership switch. [PR1299812](#)
- Subscriber database is stuck in "not-ready" state after GRES. [PR1299940](#)
- Chassisd core is seen after insertion of REMX2K-X8-64 in MX2000 platform along with older RE-S-1800x4. [PR1300083](#)
- 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](#)
- ICMP/ICMPv6 error messages might be discarded while forwarding through an AMS interface. [PR1301188](#)
- The rpd might crash when executing the **show route extensive** command during deletion of the IS-IS configuration. [PR1301849](#)
- The rpd might crash when NSR is enabled and routing-instance specific configurations are committed. [PR1301986](#)
- Continuous interface flapping might lead to unwanted MIC reset. [PR1302246](#)
- Service cookie data that is sent from Packet Forwarding Engine to service PIC can be corrupted and might lead to unexpected behavior. [PR1302493](#)
- The rpd might crash when toggling the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statements. [PR1302504](#)
- The chassisd crashes if ISSU is aborted in FRU upgrade phase. [PR1303086](#)

- The **multicast resolve-rate** value might go back to the default after system upgrade or reboot. [PR1303134](#)
- 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 in the output of the **show subscribers routing-instance** command. [PR1303199](#)
- The inline-ka PPP echo requests are not generated for aggregated Ethernet interfaces. [PR1303249](#)
- Fan speed changes frequently on MX Series router after an upgrade to Junos OS with the change introduced by PR:1244375. [PR1303459](#)
- The kernel log GENCFG messages with Severity 1 (Alert) might be seen. [PR1303637](#)
- If MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic with an error message **Fatal trap 12: page fault** while in kernel mode. [PR1303798](#)
- MX Series MIB polling returns a value that has "sdg". Polling result should include the "svc" generic value. [PR1303848](#)
- Truncated output appears for the **show pppoe lockout** CLI command. [PR1304016](#)
- The fabric planes might go into "check" state after restarting the line cards with SFB2 used on MX2010 or MX2020 platform. . [PR1304095](#)
- Effective rate of E3 in framed mode is limited to 30 Mbps on certain channelized MICs. [PR1304344](#)
- DSCP value changes are not reflected in LLDP PDU. [PR1304627](#)
- RPF-check strict mode causes 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 report SNMP index value of LSI interface instead of 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 experience the generation of rpd core file when traceoptions are enabled. [PR1305440](#)
- Repeated log messages are seen on the backup Routing Engine when **set system internet-options no-tcp-reset drop-all-tcp** and NSR option are enabled. [PR1305729](#)
- **start shell pfe network fpc** command is not working on MX960. [PR1306236](#)
- Bbe-smgd might fail to properly add access-internal routes when the router is extremely busy. [PR1306650](#)
- Smihelperd generates core files when SNMP is polling for JUNIPER-SUBSCRIBER-MIB::jnxSubscriberGeneral.7.0. [PR1306966](#)

- The kmd process error **UI_DBASE_OPEN_FAILED** is seen because of too many open files. [PR1308380](#)
- License is lost during Routing Engine switchover in scale-subscriber scenario. [PR1308620](#)
- CoS applied to a subscriber demux logical interface is not working. [PR1308671](#)
- FPC syslog errors with **pfeman_inline_ka_steering_gencfg_handler: nh not found** could mean that steering rules are not installed correctly. [PR1308884](#)
- All the MICs on one FPC, with PWHT subscribers configured, might go offline during the restart of an FPC in another slot. [PR1308995](#)
- Error messages might be often seen after an MPC restarts. [PR1309013](#)
- Incorrect values are found in the Event-timestamp of RADIUS Accounting-Stop packets for L2BSA subscribers. [PR1309212](#)
- MX2020/MX2010: After smooth upgrade from SFB to SFB2, if one plane/SFB is restarted, link training fails between those planes and MPC6 cards. [PR1309309](#)
- bbe-mibd might generate core files after a Routing Engine mastership switch. [PR1309341](#)
- First access-request fails for L2BSA subscribers when changing the MTU of LACP aggregated Ethernet A10-NSP interface. [PR1309599](#)
- Subscribers might not be able to access the device if dynamic VLAN is used. [PR1309770](#)
- 90% percent subscribers might go down after ISSU from 16.1 to 17.3. [PR1309983](#)
- In next-generation subscriber management release, bbe-smgd process memory leak is seen after deleting or adding the address pool. [PR1310038](#)
- The MS-MIC or MS-MPC memory utilization might stay at a 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](#)
- Some harmless syslog messages might be seen. [PR1310678](#)
- Local IPv6 interface from NDRA prefix not removed from service interface when subscriber dual-stack session is removed. [PR1310752](#)
- Performing a commit check just after setting the master password can trigger improper decoding of configuration secrets. [PR1310764](#)
- After BSYS reboot, rpd is sometimes unresponsive on one GNFs. [PR1310765](#)
- An incorrect error number might be reported for syslog messages with the prefix of **%DAEMON-3-RPD_KRT_Q_RETRIES**. [PR1310812](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- Rpd core files observed after multiple session flaps on a scaled setup. [PR1312169](#)
- PEM alarms and I2C failures are observed on MX240/MX480/MX960/EX92/SRX5K series. [PR1312336](#)

- MIC-MRATE might restart after port speed change. [PR1312504](#)
- Counter at PPPoE session logical interface is incremented wrongly because the accounting packet contains the wrong Acct-input-packets value and the wrong Acct-input-octets value. [PR1312998](#)
- False overtemperature SNMP trap could be seen when using MPC5, MPC6, MPC7, MPC8, MPC9 on MX2020. [PR1391](#)
- MX-VC: BNG: IPv6 RS (router-solicit) packets are dropped in nondefault RI, for default RI it is working. [PR1313722](#)
- **show version detail** gives severity error log **traffic-dird[20126]: main: swversion pkg: 'traffic-dird' name: 'traffic-dird' ret: 0**. [PR1313866](#)
- mspmand core file due to flow-control seen while clearing CGNAT+SFW sessions. [PR1314070](#)
- 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 rpd might crash in an MoFRR scenario. [PR1314711](#)
- MPC7E- IR-mode knob commit failure. [PR1314755](#)
- RPC error while committing **system services subscriber-management enable** through NETCONF. [PR1314968](#)
- The L2TP LAC might drop packets that have an incorrect payload length while sending packets to the LNS. [PR1315009](#)
- Too many logs are generated after executing many Vhclient related commands. [PR1315128](#)
- The RIB and FIB might get out of synchronization if the KRT asynchronous queue is stuck. [PR1315212](#)
- FPC crash is observed when a route has unilist next hops in an RSVP scenario. [PR1315228](#)
- The **show version detail** generates the severity error log **mobiled: main Neither BNG LIC nor JMOBILE package is present,exit mobiled**. [PR1315430](#)
- The command of **show version detail** might generate the severity error log **main: name: SRD ret: 0**. [PR1315436](#)
- The fan speed might frequently keep changing between normal and full for MX Series Router. [PR1316192](#)
- Demux interface sends neighbor solicitation with source link MAC address with all zeros. [PR1316767](#)
- The output of the **show configuration <> | display json** command might not be properly enclosed in double quotation marks. [PR1317223](#)
- Linux-based microkernel might panic because of concurrent update on mutable objects. [PR1317961](#)
- The rpd might crash when the link flaps on an adjacent router. [PR1318476](#)
- The daemon bbe-smgd might crash after GRES is performed. [PR1318528](#)
- The FPC crashes on configuration change for Packet Forwarding Engine sensors. [PR1318677](#)
- The bbe-smgd process might crash multiple times and does not recover in a rare scenario. [PR1318887](#)

- MS-MPC and MS-MIC might crash after a new IPsec tunnel is added. [PR1318932](#)
- The task replication might not be complete to certain network protocols after multiple GRES. [PR1319784](#)
- The **MIB2D_COUNTER DECREASING: pfes_stats_delta: counter** error message might be seen on VMX. [PR1319996](#)
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX-VC ISSU. [PR1320370](#)
- Various types of boards might crash while performing unified ISSU. [PR1320683](#)
- The **show subscriber summary** command output displays incorrect terminated subscriber count. [PR1320717](#)
- PPP inline keepalive does not work fine as expected when the CPE device aborts the subscriber session. [PR1320880](#)
- MX Series routers send the IPv6 router advertisements and DHCPv6 advertisements before sending IPCPv6 ACK from the CPE device. [PR1321064](#)
- The logical interface bind changes are taking more time, and many log messages like **IFL TCP (38) Bind change notify ran for** are generated by the FPC. [PR1321086](#)
- MX-VC CoS is not applied to Packet Forwarding Engine when VCP link is added. [PR1321184](#)
- The bbe-smgd process generates core files after a large number of clients log out and log in a PPPoE dual stack subscriber scenario. [PR1321468](#)
- There is CoA-NAK with **Error-Cause = Invalid-Request** sent back to Radius server if a drop policy is applied under the **radius-flow-tap** configuration in an L2TP subscriber scenario. [PR1321492](#)
- The rpd might crash when two next hops are installed with the same next hop index. [PR1322535](#)
- MS-MIC logical interfaces remain down after many iterations of taking them offline and bringing them back online. [PR1322854](#)
- Line card might crash upon receipt of specific MPLS packet. [PR1323069](#)
- Memory leaks in MGD-API daemon 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](#)
- Memory leak is seen in mosquito-nossd daemon. [PR1324531](#)
- SNMP interface filter does not work when **interface-mib** is part of the dynamic-profile configuration. [PR1324573](#)
- The VLAN rewrite function might use the wrong VLAN ID when Ethernet OAM is configured on DPCE cards. [PR1325070](#)
- SNMP values might not be increased monolithically. [PR1325128](#)
- MPC cards might drop traffic under high temperature. [PR1325271](#)
- IS-IS adjacency fails to establish because of packets drop on Packet Forwarding Engine. [PR1325311](#)

- Denial-of-Service vulnerability in MS-PIC, MS-MIC, MS-MPC, MS-DPC and SRX Series flow daemon (flowd) is related to the SIP ALG. [PR1326394](#)
- The VLAN demux interface does not respond to the ARP request in a subscriber scenario with subscriber-management enabled on MX Series routers running Junos OS releases after Junos OS Release 15.1. [PR1326450](#)
- In an MX Series BNG, a CoS service object is not deleted properly for TCP and scheduler. [PR1326853](#)
- GRE interface might not come up after deactivating/activating the routing instances. [PR1327099](#)
- Some of the show commands were issued twice when request support information is executed. [PR1327165](#)
- With auto-installation usb configured, interface related commits might not take effect due to dcd error. [PR1327384](#)
- Add error message for AMS load-balancing support. [PR1329049](#)
- MS-MIC or MS-MPC might restart when sampling the MPLS traffic. [PR1329189](#)
- When an AMS bundle has a single aggregated multiservices member interface (mams-) added to it, the subinterfaces do not recover after they are disabled. [PR1329498](#)
- On MX platform in dynamic subscriber over PS interface scenario, if CoS host-outbound-traffic is configured for ieee-802.1p rewrite, it might not work correctly for the packet bit. [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](#)
- '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](#)
- Rpd core files are generated on the new backup Routing Engine at task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler after NSR and GRES are disabled. [PR1330750](#)
- FPC wedge with fragmented packets on LSQ interface - PT1: Head and tail out of sync. [PR1330998](#)
- Non-NEBS compliant optics might be disabled when chassis temperature exceeds non-nebs-optics-overheat-trigger. [PR1331186](#)
- The FPC might crash due to logical interface index corruption when IPv6 traffic goes through the IRB interface. [PR1331911](#)
- On all Junos OS products, the local DHCPv6 server might incorrectly respond to "Confirm" messages from clients. [PR1331995](#)
- The rpd generates core files in a Layer 2 circuit or a Layer 2 VPN environment. [PR1332260](#)
- Inaccurate J-Flow records might be seen for output interface and next hop. [PR1332666](#)
- The dot1xd might crash if ports in multi-supplciant mode flaps. [PR1332957](#)

- The subinfo process might crash and cause the PPPOE subscribers to get disconnected. [PR1333265](#)
- The MX Series router might not be able to learn the global IPv6 neighbor address of its DHCPv6 subscriber client. [PR1333392](#)
- In AA Multihoming EVPN VXLAN, some race conditions can trigger constant high CPU on the backup Routing Engine. [PR1334235](#)
- The UID limit is reached in a large-scale subscriber scenario. [PR1334886](#)
- When using **show subscribers** and if 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](#)
- The RIP route updates might be partially dropped when NSR is enabled. [PR1335646](#)
- The MAC_STUCK message might be seen on MS-MPC or MS-MIC. [PR1335956](#)
- Subscriber might experience SDB DOWN event and drop the clients' connections when issuing **show subscribers** commands. [PR1336388](#)
- On MX2000 with SFB card installed, high traffic volume on or heavy traffic on MPC7E, MPC8E or MPC9E might cause traffic drops with cell underflow messages. [PR1336446](#)
- With certificate hierarchy, where intermediate CA profiles are not present on the device, in some corner cases, the PKI daemon can become busy and stop responding. [PR1336733](#)
- The hash value generated for 256-bit key length of AES-GCM-256 algorithm is incorrect. [PR1336834](#)
- BBE-SMGD might core when you configure CoS on logical interface sets. [PR1336852](#)
- The link flaps or stays down due to an interoperability issue between MX Series routers or EX9200 switches and a and transport device. [PR1337327](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- Error log message **sdb_db_interface_remove: del ifl:si- <index> with licnese cnt non zero on** can be seen on LTS during subscriber logout. [PR1337000](#)
- Very few of subscribers show wrong accounting values in large-scale subscribers scenario. [PR1340512](#)
- There might be traffic loss on some subscriber sessions when more than 32000 L2TP subscriber sessions are anchored in an ASI interface. [PR1341659](#)
- With discard interfaces (configured with IGMPv3), the KRT queue gets stuck while deleting multicast next hop with error **EPERM -- Jtree walk in progress**. [PR1342032](#)
- SNMP walk might failed for LLDP related OIDs. [PR1342741](#)
- MX Series routers send the IPv6 router advertisements and DHCPv6 advertisements before sending IPCPv6 ACK from the CPE device. [PR1344472](#)

- The Framed-route "0.0.0.0/0" won't be installed in MX Series platform with Junos enhanced subscriber management releases. [PR1344988](#)
- Dot1x reauthentication issue. [PR1345365](#)
- An rpd crash might be seen if **no-propagate-ttl** is set in a routing instance that has a specific route. [PR1345477](#)
- New PPPoE users might fail to log in. [PR1346226](#)
- **AC system error** counter in **show pppoe statistics** is not working. [PR1346231](#)
- VCCP-ADJDOWN detection is delayed on VC-Bm when deleting one VCP link on VC-Mm. [PR1346328](#)
- NAT might not work and the spd might crash. [PR1346546](#)
- Statistics daemon PFED might generate a core file on an upgrade between certain releases. [PR1346925](#)
- **Twice-napt-44** sessions does not sync to the backup SDG when stateful sync is configured. [PR1347086](#)
- IPv6 MAC resolve might fail if the DHCPv6 client uses a non-EUI64 link-local address. [PR1347173](#)
- The rpd might crash when the dynamic-tunnels next hop that is resolving migrates to a more specific IGP route. [PR1348027](#)
- Issue with handling the community_action ("add") in RPC call. [PR1348082](#)
- MIC-3D-20GE-SFP-E might generate core files due to ISR 2 MIC error interrupt hogging. [PR1348107](#)
- The authd and smgd might crash and create a core file. [PR1348727](#)
- The per-service accounting statistic value is not accurate. [PR1348796](#)
- The chassisd might crash after MPC6E or MPC7E is replaced with MPC9E. [PR1348834](#)
- DHCPv6 Solicit dropped on L2TP LNS in MX-VC when incoming interface is on VC-master and both anchor si- interface and VCP port on VC-backup on MPC2 NG or MPC2 NG. [PR1348846](#)
- Routing Engine mastership keepalive timer is not updated after the GRES configuration is removed. [PR1349049](#)
- Major alarm:"Major PEM 0 Input Failure" might be observed for DC PEM. [PR1349179](#)
- The MPC might crash when the MIC is removed. [PR1350098](#)
- Pseudowire subscriber over redundant logical tunnels function does not work on MPC7 and MPC9. [PR1350115](#)
- The pccd might crash after a delegated LSP is removed in a PCEP scenario. [PR1350240](#)
- Multicast traffic gets dropped as Invalid policy ID exception. [PR1350380](#)
- The VCP port might not come back up after removing and adding it again. [PR1350845](#)
- PPE Errors async xtxn error when FPC is restart or removal. [PR1350909](#)
- The pfed process is consuming 80-90% CPU running subscriber management on PPC-based routers. [PR1351203](#)

- After GRES, the BGP neighbors at Master RE might reset and the BGP neighbors at Backup RE take long time to establish [PR1351705](#)
- Offlining MIC6-100G-CFP2 MIC through CLI command might trigger FPC card to crash. [PR1352921](#)
- Rpd might permanently hog the CPU due to Logical System configuration commit. [PR1353548](#)
- Syslog error: dfw_bbe_filter_bind:1125 BBE Filter bind type 0x84 index 167806251 returned 1. [PR1354435](#)
- The rpd generates core files when adding an inter-region template in a routing-instances. [PR1354629](#)
- Newly provisioned IPsec tunnel could not forward traffic. [PR1354757](#)
- The static-subscribers do not properly update firewall information on the Packet Forwarding Engine when dynamic configuration changes are made to active subscribers. [PR1354774](#)
- Memory leak is found in agentd when running valgrind. [PR1354922](#)
- Changing the **ipv4-flow-table-size** does not change the amount of available IPv4 flow table memory in an inline J-Flow scenario. [PR1355095](#)
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100 Gbps). [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 GRES scenario. [PR1355463](#)
- Rpd crash might be seen when issuing CLI **show dynamic-tunnels database terse** and when the system have RSVP tunnels configured. [PR1356254](#)
- I2c messages from PEM/PSM are reported if SNMP is enabled. [PR1356259](#)
- The **show pppoe underlying-interfaces** command in a scaled environment might cause bbe-smgd memory leak. [PR1356428](#)
- The bbe-smgd generates core files in recursive loop between functions bbe_autoconf_if_I2_input and bbe_if_I3_input. [PR1356474](#)
- DHCP subscribers fail after reconfiguration of port from tagged to untagged mode. [PR1356980](#)
- Routing Engine switchover that occurs before the backup Routing Engine is not GRES ready might cause a linecard restart. As a result, the Routing Engine kernel crashes and multiple chassisd crashes occur. [PR1357427](#)
- MPC/FPC might be unable to reply request messages to the Routing Engine in a high subscriber scale scenario. [PR1358405](#)
- Multiple bbe-smgd crashes might be seen when multiple subscribers log in simultaneously. [PR1358868](#)
- The **show chassis fpc** command might show **Bad Voltage** message for FPC powered off by configuration or CLI command after the **show chassis environment fpc** command is executed. [PR1358874](#)
- The IPv6 subscriber might fail to access the network. [PR1359520](#)

- The bbe-smgd might fail to add members to some of the aggregated Ethernet interfaces randomly when there are many aggregated Ethernet interfaces in the access configuration. [PR1359986](#)
- The rpd generates core files 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 performing a Routing Engine switchover with SRTE routes. [PR1360354](#)
- Mirrored traffic is not going out through the LT interface. [PR1360489](#)
- FPC core might be observed after GRES switchover. [PR1361015](#)
- An rpd scheduler slip might be seen when you frequently delete, modify, or add groups that are applied on the top level. [PR1361304](#)
- IP over VPLS traffic is affected by EXP rewrite rule on the core-facing MPLS interface. [PR1361429](#)
- MX Series BNG does not generate ESMC/SSM quality level failed SNMP trap. [PR1361430](#)
- The rpd is struck at 100 percent after the **clear bgp neighbor** operation. [PR1361550](#)
- Spontaneous bbe-smgd core files might be seen on the backup Routing Engine. [PR1362188](#)
- The MS-MPC might reset continuously on MX Series routers. [PR1362271](#)
- Route installation failure might be seen after the BGP neighbor and route flaps. [PR1362560](#)
- Executing the **show route prefix proto ip detail** command during route churn in a route scale scenario might lead the FPC to crash. [PR1362578](#)
- The non-default routing instance is not supported correctly for NTP packets in a subscriber scenario. [PR1363034](#)
- Select CLI functions are not triggering properly (**set security ssh-known-hosts load-key-file**, **set system master-password**). [PR1363475](#)
- MX Series Virtual Chassis: Request to record VCCP heartbeat state change in the syslog by default. [PR1363565](#)
- Some error logs might be seen on MX2010 and MX2020 routers equipped with SFB2. [PR1363587](#)
- The multicast route update might be stuck in KRT queue and the rpd might crash if rpd and kernel go out of synchronization. [PR1363803](#)
- FPM board status is missing in SNMP MIB walk result. [PR1364246](#)
- A traffic loop might occur even though that port is blocked by RSTP in a ring topology. [PR1364406](#)
- Configuration commit might be delayed by 30 seconds. [PR1364621](#)
- Default adapter type is changed from E1000 to VMXNET3. [PR1365337](#)
- MPC7E: A ukern crash and FPC reboot are seen with the vty **show agent sensors verbose** command. [PR1366249](#)
- MS-MPC/MS-PIC might crash in a NAT scenario. [PR1366259](#)
- The next hop of MPLS path might be stuck in hold state, which could cause traffic loss. [PR1366562](#)

- The **show system resource-monitor fpc** command might display a nonexistent Packet Forwarding Engine. [PR1367534](#)
- RTG interface status will be shown as incorrect status with **show interface**. [PR1368006](#)
- In BBE configurations, receipt of a crafted IPv6 exception packet causes denial of service (CVE-2018-0058). [PR1368599](#)
- SNMP MIB walk causes KMD errors. [PR1369938](#)
- Kernel crash might be seen after committing a demux related configuration. [PR1370015](#)
- The rpd might crash after a Routing Engine switchover is performed or the rpd is restarted if interface-based dynamic GRE tunnel is configured. [PR1370174](#)
- Packets exceeding 8000 bytes might be dropped by MS-MPC in an ALG scenario. [PR1370582](#)
- FPC causes high CPU utilization or crash during a hot-banking condition. [PR1372193](#)
- Image installation on SD fails with error **Unable to read reply from software add command to re1; error 1**. [PR1372877](#)
- The Routing Engine might crash after a non-GRES switchover. [PR1373079](#)
- LDP convergence delay might be seen after IGP metric change with **bgp-igp-both-ribs** configured. [PR1373855](#)
- Cosmetic log **warning: [---] is protected, 'protocols ---' 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 a dynamic VLAN subscriber scenario. [PR1374478](#)
- The bbe-smgd generates core files continuously while deleting a multicast group node from the tree. [PR1374530](#)
- A bbe-smgd core file might be seen after performing GRES. [PR1376045](#)
- MS-MPC might have performance degradation under scaled fragmented packets. [PR1376060](#)
- Interface optic output power is not zero when the port has been disabled. [PR1376574](#)
- Packets might be dropped on the data plane in the inline J-Flow scenario. [PR1377500](#)
- The ICMPv6 packets larger than 1024 might be dropped if **icmp-large-packet-check** is configured on the IDS. [PR1378852](#)
- PCS statistics (bit errors and errored blocks) could not increment in the Routing Engine CLI output. [PR1379147](#)
- The Routing Engine might crash with various core files due to the deadlock issue on the SDB STS. [PR1380231](#)
- Memory leak observed in MS-MPC card. [PR1381469](#)
- Subscribers not able to log in after double GRES, after reboot, or after configuration. [PR1382050](#)

- The MPC6E might crash while fetching PMC device states. [PR1382182](#)
- Flows are getting exported before the expiration of the configured active timeout value. [PR1382531](#)
- The kmd crashes with a core file after bringing up the IPsec connection. [PR1384205](#)
- Missing **interface-description** configuration statement for static subscribers. [PR1384421](#)
- IPsec VPN traffic might fail when passing through MS-MPC of MX Series router with CGNAT enabled. [PR1386011](#)
- Output of the **show class-of-service interface** command incorrectly shows adjusting application as PPPoE IA tags for DHCP subscribers. [PR1387712](#)
- 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](#)
- IGMP group threshold exceed log message appears and prints an incorrect demux logical interface. [PR1389457](#)
- The jnxFruState might show incorrect PIC state after replacing an MPC with another MPC having less PICs. [PR1390016](#)
- The CoS **adjustment-control-profile** configuration for application DHCP tags does not get applied. [PR1390101](#)
- The bbe-smgd process might crash after committing configuration changes. [PR1391562](#)
- The spd might crash when **any-ip** is configured in the from clause of the NAT rule with the static translation type. [PR1391928](#)
- If FPGA on the new master CB has a specific hardware failure, the chassis might keep crashing after a GRES switchover. [PR1393884](#)
- The MS-MPC might generate a core file when mspmand receives a non-synchronized TCP packet. [PR1396785](#)
- IPsec tunnel cannot be established because the tunnel SA and rule are not installed in the PIC. [PR1398849](#)
- The bbe-smgd process might crash when executing the **show pppoe lockout** command. [PR1398873](#)
- Smg-service can become unresponsive. [PR1403480](#)
- The FPC might crash in a CoS scenario. [PR1404325](#)
- Fabric performance drop occurs on MPC7, MPC8, MPC9E and SFB2 based MX2000 routers. [PR1406030](#)

High Availability (HA) and Resiliency

- GRES might fail to start because of missing state ack message from the Package Forwarding Engine. [PR1236882](#)
- The backup Routing Engine might go to the database prompt after removing or restoring the configuration. [PR1269383](#)

- The ksyncd might crash. [PR1275022](#)
- Line card reboots after GRES. [PR1286393](#)
- After interface flapping occurs, the GNFs on the server CB ports show the message **Switchover Status: Not Ready**. [PR1306395](#)
- The ksyncd process might crash continuously on the new backup Routing Engine after GRES is performed. [PR1329276](#)
- Insufficient available space on hard disk is caused by the crashinfo files that is generated by ksyncd when GRES is configured in large-scale configuration scenario. [PR1332791](#)
- VC-Bm cannot synchronize with VC-Mm when the Virtual Chassis splits and then reforms. [PR1361617](#)

Infrastructure

- The **show system users** command output displays users that are not using the router. [PR1247546](#)
- The **show interface** command does not return any values and sometimes the command is completely stuck. [PR1250328](#)
- When **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](#)
- SNMP MIB walk IfHighSpeed returns instable values for em logical interface when no bandwidth is configured for it. [PR1257566](#)
- Vmcore is created because of mbuf leak. [PR1261996](#)
- Some of the syslog records of CGNAT session might have incorrect time. [PR1295442](#)
- The device might fail to upgrade. [PR1298749](#)
- The **syscalltrace.sh** might create a huge output file, which could cause the router to run out of storage space. [PR1306986](#)
- Kernel crash (vmcore) occurs during broadcast storm after enabling **monitor traffic interface fxp0**. (CVE-2018-0029) [PR1322294](#)
- Cleanup at thread exit in FreeBSD kernel is causing memory leaks. [PR1328273](#)
- On all platforms running Junos OS, on a port configured with both **dot1x static mac by-pass** and **normal authentication**, the hosts configured for static MAC bypass might not be able to send traffic. [PR1335125](#)
- The kernel might crash and the system might reboot in SNMP query reply scenario. [PR1351568](#)
- The **show system virtual-memory | display xml validate** command displays errors. [PR1356423](#)

Interfaces and Chassis

- The output value is incorrect when querying the optical power of OTN interfaces in the router. [PR1216153](#)
- On an Enhanced DPC, the dcd process might increase the CPU usage to a very high level after **commit check** is executed. [PR1236088](#)

- VRRP mastership does not change after priority is changed. [PR1242243](#)
- Iterator adjacency is removed, leading to inability to display **sla-iterator-statistics** within CFM performance monitoring. [PR1244525](#)
- RL-dropped packets are not displayed in the output of **show interfaces <ifl or interface-set ifl> detail/extensive** commands. [PR1249164](#)
- At a high logical interface scale, an ifinfo process (daemon) generates a core file when the command **show interfaces extensive | no-more** is executed. [PR1254189](#)
- The SNMP-set on the supported configuration in the jnxOpticsConfigTable fails if the FPC slot is 10 and above or the port number is 10 and above. [PR1259155](#)
- The MRU of an aggregated Ethernet interface might be reset to the default value. [PR1261423](#)
- The error messages about **RLIMIT_STACK** and **RLIMIT_SBSIZE** might be seen in a PPP scenario after issuing **show version detail**. [PR1262629](#)
- Some error messages might be seen when setting or deleting VCP port for MPC7, MPC8, and MPC9E cards. [PR1271089](#)
- BERT test shows the elapsed time "in progress" but gets stuck after a few seconds and never gets completed. [PR1274896](#)
- MTU configuration for vt- interface causes the vt- interfaces to be removed because the MTU on this interface is already set to unlimited. [PR1277600](#)
- The **PPP Chap Challenge-Length** option is not initialized with the default value. [PR1280263](#)
- The line card hosting an Ethernet OAM LFM session might reboot during a unified ISSU. [PR1283280](#)
- The monitor interface on aggregated Ethernet logical interfaces displays an incorrect bps value compared to that shown in the **show interface** output. [PR1283831](#)
- Interface flapping is observed when Routing Engine switchover is performed if the member links of an aggregated Ethernet 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](#)
- Family **inet** is displayed as not-configured after the loopback address is added or deleted. [PR1294267](#)
- In VRRP scenario, when tracked interface or route goes down, the mastership switchover is delayed for a longtime. [PR1294417](#)
- L2TP subscribers might not be cleared if the access-internal routes fail to install. [PR1298160](#)
- An absolute value can be configured for the **delay-buffer-rate** option on an inline LSQ interface. [PR1300281](#)
- IRB interface shows incorrect bandwidth value. [PR1302202](#)
- VRRP could not support logical interfaces using the same group ID in VRRP delegated-process mode. [PR1305327](#)

- AFEB might not come up when LFM is deactivated. [PR1306707](#)
- After the **request system reboot both** CLI command is executed, the PPP daemon might become unresponsive. [PR1310909](#)
- The PPPoE subscriber might not log in correctly after authentication failure in a subscriber scenario. [PR1311113](#)
- MPC CPU might reach 100% when **otn ufec** statement is configured. [PR1311154](#)
- The ifinfo process might crash and generate a core file when the **show interfaces name** command is executed with the name greater than 128 characters. [PR1313827](#)
- Invalid interface-set configuration might get committed and result in continuous dcd and chassisd crash. [PR1316976](#)
- There is no route to an IP address from a directly connected route. [PR1318282](#)
- The **show interfaces interface-set** command is displaying an incorrect logical interface. [PR1319682](#)
- **IPv6 Framed Interface Id** field (from **show subscribers extensive** output) is not properly matching the negotiated one. [PR1321392](#)
- IPCP negotiation might fail for dual stack PPPoE subscribers. [PR1321513](#)
- Subscribers might fail to access the device after deleting the needless aggregated Ethernet configuration. [PR1322678](#)
- 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 that supports subscriber management. [PR1328251](#)
- Traffic loss might be seen after the aggregated Ethernet bundle unit 1 is deleted. [PR1329294](#)
- The interface might not work properly after FPC restarts. [PR1329896](#)
- The dcd process might crash due to memory leak and causing commit failure. [PR1331185](#)
- The last logical interface digit is sometimes truncated in jpppd trace logs. [PR1332483](#)
- The transportd might crash when SNMP query on jnxoptIfOChSinkCurrentExtTable with unsupported interface index. [PR1335438](#)
- A momentary dip in traffic occurs when a GRES is performed. [PR1336455](#)
- Restarting chassisd with GRES is disabled might cause FPC to restart and some demux sessions to be deleted. [PR1337069](#)
- VRRP virtual MAC addresses disappear, which causes VRRP virtual IP to be not reachable. [PR1338277](#)
- The 100-Gigabit DWDM interface might go down for 15 seconds after a loss-of-signal event. [PR1343535](#)
- The PPPoE subscribers might fail to log in for authd running on 100 percent utilized CPU with a high frequency of on-demand IP address allocation requests. [PR1348578](#)

- The link-degrade-monitor configuration might cause commit synchronization failure on the backup Routing Engine. [PR1350192](#)
- The jpppd core file is generated on the backup Routing Engine in longevity test at `../..../src/junos/usr/sbin/jpppd/pppMain.cc:400`. [PR1350563](#)
- Native VLAN ID support is needed on ps-interface. [PR1352933](#)
- 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 aggregated Ethernet interface might flap when the link speed of the aggregated Ethernet bundle is configured to 10G. [PR1355270](#)
- Clients might not get IPv4 address in PPPoE dual-stack scenario. [PR1360846](#)
- Many PPPoE subscribers might be lost after unified ISSU/GRES. [PR1360870](#)
- Error messages like `ifname [ds-5/0/2:4:1] is candidate` are seen during a commit operation. [PR1363536](#)
- In case of MPLS, DMR packets are sent with different MPLS expiration bits if the MX Series router receives CFM DMM packets with varying expiration values on the MPLS header. [PR1365709](#)
- In rare cases, there might be L2TP subscribers stuck in the terminated state. [PR1368650](#)
- ISSU could be aborted at **Timed out Waiting for protocol backup chassis master switch to complete** with MX Virtual Chassis configuration. [PR1371297](#)
- The dcd process might go down when `vlan-id none` is configured for the interface. [PR1374933](#)
- 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 routers. [PR1378912](#)
- The dcd is 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](#)
- A dcd core file might be seen after a 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](#)
- The backup Routing Engine might get stuck in amnesiac mode after reboot. [PR1398445](#)
- All dcd operations might be blocked if profile-db is corrupt. [PR1399184](#)

J-Web

- Denial of service occurs in J-Web (CVE-2018-0062). [PR1264695](#)

- Unauthenticated remote code execution through J-Web interface. [PR1269932](#)

Layer 2 Ethernet Services

- The IPv4/IPv6 packets originating from the Routing Engine might be corrupted when the bridge domain has VLAN ID set to **none**, but the outgoing Layer 2 interface for the packet is tagged and CoS is enabled. [PR1263590](#)
- DHCP is not using the configured IRB MAC address as the source MAC address in DHCP offer unicast replies. [PR1272618](#)
- DHCPv6 client bound to IA_PD prefix on reception of DHCPv6 Request for IA_NA, MX deletes the existing binding. [PR1286359](#)
- The jdhcpd process crashes, generating a core file, and restart. [PR1288475](#)
- ARP requests are not generated for IRB configured in VPLS over GRE tunnel. [PR1295519](#)
- PPPoE/DHCP clients cannot log in to PPPoE/DHCP dual-stack subscriber scenario. [PR1298976](#)
- A parameter-handling problem might cause the kernel to panic when a neighbor discovery message arrives on an IRB interface. [PR1303415](#)
- Multiple jdhcpd core files are observed in jdhcpd_update_groups at `../..../src/junos/usr/sbin/jdhcpd/jdhcpd_config.c:2290`. [PR1311569](#)
- DHCPv6 traffic might be dropped in a subscriber scenario. [PR1316274](#)
- jdhcpd generates core files after DHCP configuration is changed/modified. [PR1324800](#)
- The snmpget for OID: dot3adInterfaceName might not work. [PR1329725](#)
- The l2cpd process leaks memory if the Layer 2 learning process is disabled. [PR1336720](#)
- When DHCP subscribers are in BOUND (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state, if dhcp-service is restarted then the subscribers in this state are logged out. [PR1350710](#)
- DHCP relay agent discard the DHCP request message silently if the requested IP address has been allocated to the other client. [PR1353471](#)
- On MX Series routers, restart of the FPC that hosts micro-bfd link might cause a LACP core file. [PR1353597](#)
- A jdhcpd crash is observed while processing the DHCPv6 Information-Request. [PR1368377](#)
- 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](#)

Layer 2 Features

- The mbuf leaks because of processing of MPLS packets in the VPLS network. [PR1272898](#)
- In scaling VPLS scenario, convergence time is taking more than 10 minutes. [PR1279192](#)

- A misconfiguration that adds an aggregated Ethernet bundle and its member link to a VPLS instance might cause 100% routing protocol process (rpd) utilization. [PR1280979](#)
- On MX Series routers with MPCs or MICs, packets received on the IRB interface in VPLS gets double-tagged. [PR1295991](#)
- The rpd process memory leak is observed when VPLS configuration changes such as deleting or re-configuring VPLS interfaces occur. [PR1335914](#)
- VPLS instance stays in NP state after LDP session flaps. [PR1354784](#)
- The unicast traffic from an IRB interface toward LSI might be dropped due to Packet Forwarding Engine mismatch at egress processing. [PR1381580](#)

MPLS

- The rpd generates core files in `rio_hello_timeout_cb` under high CPU load. [PR1138190](#)
- RSVP P2MP sub-LSPs with more than one sub-LSP in down state might not get re-optimized after the transit path goes down. [PR1174679](#)
- Automatic bandwidth underflow is not being registered following the first bandwidth adjustment when there is no traffic flowing over an LSP. [PR1233293](#)
- The rpd might crash when moving a static LSP from one routing instance to another. [PR1238698](#)
- The rpd might crash after performing **restart routing** or Routing Engine switchover in MPLS environment. [PR1239102](#)
- Potential issues with policy-based selection of RSVP LSPs [PR1261739](#)
- A rpd crash might be seen if egress-policy is configured in LDP. [PR1266358](#)
- The **created time** value in **show mpls lsp extensive** drifts by a second when the **show** command is issued multiple times. [PR1274612](#)
- The ingress RSVP LSP fails to come up after the **clear rsvp lsp all** command is run on the egress router. [PR1275563](#)
- The rpd might crash in LDP L2circuit scenario. [PR1275766](#)
- The rpd might crash on the egress LER of a fast-reroute protected LSP. [PR1276748](#)
- A crafted MPLS packet might lead to a kernel crash. [PR1276786](#)
- The Routing Engine might crash during next-hop addition in a race condition. [PR1284850](#)
- MPLS I2ckt ping packet incorrectly parsed by the output loopback filter. [PR1288829](#)
- LDP egress policy is not advertising the label for the inet.3 BGP labeled-unicast route. [PR1289860](#)
- The routing protocol process (rpd) crashes because of LDP defect during NSR-enabled Routing Engine switchover. [PR1290789](#)
- Received MTU might not get updated in RSVP MTU signaling. [PR1291533](#)

- In an RSVP environment, a stale LSP might get created after a Routing Engine switchover with nonstop routing (NSR) enabled. [PR1292526](#)
- The rpd might crash when the MPLS LSP path changes. [PR1295817](#)
- The process rpd might crash when MPLS traceroute is performed. [PR1299026](#)
- The traffic in a 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 kysncd process might crash after removing and inserting the backup Routing Engine in an analytics and MPLS sensor scenario. [PR1303491](#)
- The **explicit-null** feature might block incoming host-bound traffic from LSPs. [PR1305523](#)
- 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 the interface is down when UHP-based LSPs are configured. [PR1309397](#)
- The rpd process might crash if LDP updates the label for the BGP route. [PR1312117](#)
- The **install-nexthop lsp/lsp-regex** statement in the policy does not work with dynamic LSPs (RSVP automesh). [PR1313185](#)
- Delayed **show mpls container-lsp** output. [PR1314960](#)
- RSVP node-neighbor found even when node-hello has been disabled. [PR1317241](#)
- The rpd might crash after the primary link failure of link protection. [PR1317536](#)
- With dynamic tunnels configured, the rpd might crash when it is restarted or Routing Engine switchover is executed. [PR1319386](#)
- The IPv4/IPv6 multicast traffic might get dropped in MX Series Virtual Chassis when the traffic comes in through the Layer 2 circuit and goes out through an aggregated Ethernet member interface across Virtual Chassis members. [PR1320742](#)
- With dynamic-tunnels configured, the rpd might crash when the rpd is restarted or Routing Engine switchover is executed. [PR1319386](#)
- The rpd might crash because of a memory leak in an RSVP scenario. [PR1321952](#)
- Receipt of specially crafted UDP packets over MPLS might bypass stateless IP firewall rules (CVE-2018-0031). [PR1326402](#)
- SNMP OID counters for mplsLspInfoAggrOctets show a constant value in **show mpls lsp statistics** for some LSPs even though traffic constantly increases. [PR1327350](#)
- Rpd daemon crashes upon receipt of specific MPLS packet (CVE-2018-0043). [PR1328058](#)
- Packets loss might be observed when auto-bandwidth is enabled for CCC connections. [PR1328129](#)
- The rpd might crash on the backup Routing Engine because of memory exhaustion. [PR1328974](#)
- The rpd might crash with MPLS traceoption configured. [PR1329459](#)

- Whenever there is a decrease in the statistics value across an LSP, the `mplsLsplInfoAggrOctets` value takes two intervals to get updated. [PR1342486](#)
- LDP label is generated for serial interface subnet route unexpectedly. [PR1346541](#)
- The MPLS LSP does not come up after changing admin-group mapping. [PR1348208](#)
- The rpd might crash in an RSVP setup-protection scenario. [PR1349036](#)
- In a very rare scenario, rpd might crash when LDP fails to allocate self-id for the P2MP FEC. [PR1349224](#)
- Nondeterministic load balancing of Routing Engine generated traffic is observed. [PR1354738](#)
- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. [PR1359956](#)
- The 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 LSP might remain UP even if no path is acceptable due to CSPF failure. [PR1365653](#)
- The route prefixes with an assigned label might be missed in the LDP database. [PR1366619](#)
- The rpd might crash in a BGP LU and LDP scenario. [PR1366920](#)
- RSVP authentication might fail between some Junos OS releases and cause traffic loss during local repair. [PR1370182](#)
- 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 Networks device to another device that does not support RFC6424. [PR1372924](#)
- The traffic might not be load-balanced equally across LSPs with **ldp-tunneling** configured. [PR1373575](#)
- The rpd process might crash continuously if the **nsr-synchronization** statement or **all** flag is used in RSVP traceoptions. [PR1376354](#)
- Receipt of a specifically crafted malicious MPLS packet leads to a Junos OS kernel crash (CVE-2018-0049). [PR1380862](#)
- The rpd might crash on the backup Routing Engine after switchover. [PR1382249](#)

Multicast

- DHCPv6 relay is not working unless DHCP is restarted [PR1316210](#)
- Multicast traffic is not forwarded on the newly added P2MP branch/receiver. [PR1317542](#)

- Some IGMP groups might have wrong upstream interface because an incorrect discard route is installed in the PIM. [PR1337591](#)
- With discard interfaces (configured with IGMPv3), the KRT queue gets stuck while deleting the multicast next hop with the error **EPERM -- Jtree walk in progress**. [PR1342032](#)

Network Management and Monitoring

- Denial-of-service vulnerability in SNMP MIB-II subagent daemon (mib2d) (CVE-2018-0019). [PR1241134](#)
- jnxDomCurrentLaneTxLaser* SNMP MIBs used for tracking the Tx and Rx power values are not working for P3-15-U-QSFP28 PIC. [PR1265412](#)
- Command Esc-q does not work when the system log is disabled, and syslog messages continue to be displayed. [PR1269274](#)
- SNMP MIB hierarchy is missing. [PR1278197](#)
- mib2d-related syslog messages **MIB2D_RTSLIB_READ_FAILURE: rtslib_iflm_snmp_pointchange** are seen during the removal and restoration of configurations. [PR1279488](#)
- snmpd denial of service upon receipt of crafted SNMP packet (CVE-2017-2345). [PR1282772](#)
- The mib2d process might crash in the SNMPv3 environment. [PR1286005](#)
- The mib2d process 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 an logical interface of lo0 is deleted. [PR1286351](#)
- The **show arp no-resolve interface X** command for the nonexistent interface X shows unrelated static ARP entries. [PR1299619](#)
- After SNMP configuration, activation of the snmpd process starts to consume a lot of CPU time. [PR1300016](#)
- The syslog might generate duplicate entries of hostname and timestamp. [PR1304160](#)
- The mib2d might crash during SNMP polling on interface MIBs and meanwhile the FPC restarts or the interface flaps. [PR1318302](#)
- The jnxDomLaneAlarmSet trap is sent with an empty interface description. [PR1318913](#)
- SNMP stops or becomes very slow after a very long period of time. [PR1328455](#)
- With **interafce-mib**, the MX Series router is responding with **type : NoSuchInstance** for OIDs when multiple OIDs are polled in one SNMPGET request. [PR1329749](#)
- jnxDcuStatsEntry and jnxScuStatsEntry OIDs are missing post interface config change. [PR1354060](#)

Platform and Infrastructure

- Traffic drop might occur under a large-scale firewall filter configuration. [PR1093275](#)
- Kernel might crash on issuing **show arp** or **clear arp** if there is an IPv4 255.255.255.255 address. [PR1120114](#)
- The **LIBJSNMP_NS_LOG_WARNING** messages are observed continuously in **/var/log**. [PR1159551](#)
- 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](#)
- Unable to roll back to a certain configuration version when using admin users with restricted permissions. [PR1206074](#)
- Packet Process Engine UCODE rebalancing is getting enabled by default. [PR1207532](#)
- Unexpected scheduler queue ID mapping might be seen if an aggregated Ethernet interface is configured with **scheduler-map** in enhanced-ip mode. [PR1236541](#)
- ISSU might fail, displaying the message **Backup Routing Engine not ready**. [PR1240788](#)
- With a commit script configured, the mgd process might crash when any feature (statement/option) is configured in private configuration mode. [PR1244015](#)
- The **commit complete** message is displayed three times on every commit. [PR1244031](#)
- XM chip-based line card might drop traffic under high temperature. [PR1244375](#)
- The MX104 router does not report HSL2 CMERROR alarm upon HSL2 CRC errors. [PR1247707](#)
- The error messages about "jlock hog" might be seen after restarting routing in large scale of routes. [PR1248246](#)
- In some scenarios, certain interface configuration change (otn-options, wan-phy, etc) might trigger shm-rtssdb to generate core files because of unexpected internal messages size exchanged between Kernel and shm-rtssdb daemon. [PR1249116](#)
- One of the processes (dcd, rpd, dfwd, pfed, cosd, sampled) might generate a core file in a large-scale 8000 ESSM login or logout with an ephemeral database. [PR1249979](#)
- Kernel crash might be observed with the panic string "rn_clone_unwire no ifclone parent". [PR1253362](#)
- Error message **rn timer_delete_nh: no pat-node** might be seen when the subscriber logs out. [PR1263983](#)
- Configuration changes under the logical-system statement for a logical system (LSYS) user does not take effect after a single commit with fast-synchronize enabled. [PR1265139](#)
- Transient hardware problem causes high fabric traffic drops. [PR1265385](#)
- An error message might be seen if a new line card or service card is brought comes online. [PR1266336](#)

- Dropping the TCP RST packet incorrectly on the Packet Forwarding Engine might cause traffic drop. [PR1269202](#)
- The queued statistics of interface are not correct for CoS on MX Series routers. [PR1271055](#)
- The RPM Loss percentage values for "over all tests" through SNMP might be incorrect. [PR1272566](#)
- Every few seconds syslog prints messages related to **luss_cassxr_hotbank_check CASS XR Heavy Bank Mask**: seen on MPCE FPCs. [PR1273439](#)
- The **show ddos-protection protocols arp culprit-flows** command displayed the wrong source MAC address. [PR1274457](#)
- EVPN-VXLAN traffic gets dropped as **Incorrect vxlan fw path executed** because of a sampling configuration on the core interface. [PR1280539](#)
- The MPC might crash after an IRB interface is deleted or any other change is made on an IRB interface. [PR1281107](#)
- Password might be required when you issue the **request routing-engine login other-routing-engine** command. [PR1283430](#)
- Error messages might be observed with MPC5E card. [PR1283850](#)
- 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 **set system commit delta-export** command is run. [PR1284788](#)
- Administratively disabling an interface might cause high FPC CPU usage. [PR1285673](#)
- 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](#)
- The output values of the **show system resource-monitor** command are not accurate. [PR1287592](#)
- There might be memory leak on MPC if the next-hop address that is defined in the next-hop-group is reachable through multiple interfaces. [PR1287870](#)
- Unauthenticated remote root access is possible when RSH service is enabled (CVE-2018-0052). [PR1288932](#)
- The source MAC address learned from Packet Forwarding Engines across aggregated Ethernet interfaces might bounce between aggregated Ethernet member and Packet Forwarding Engines for a long time and might cause an MLP-ADD storm. [PR1290516](#)
- The rmopd might get stuck at sbwait upon receiving a specific response from the HTTP agent. [PR1292151](#)
- Transient flow control is asserted by XLP MAC after MX Series router is upgraded 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 2 to 4 times a day, which raises an alarm when polled through jnxRpmResSumPercentLost MIB. [PR1300049](#)
- Packet corruption with EVPN MPLS double label push with 3 or more ieee 802.1Q VLAN tags. [PR1300211](#)
- Traffic might be dropped in the egress Packet Forwarding Engine because of a hashing mismatch. [PR1300789](#)
- Packet Forwarding Engine might crash after an MPC reset in a firewall filter scenario. [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 MX Series routers with DPC, on which CoS is configured. [PR1301723](#)
- MX Series FPC wedges when creating more than 4000 logical tunnel interfaces per Packet Forwarding Engine. [PR1302075](#)
- The **interface-MAC-limit** might fail for aggregated Ethernet interface. [PR1303293](#)
- MQSS parcel error might result in performance degradation or the forwarding through the Packet Forwarding Engine might stall [PR1303529](#)
- The Two-Way Active Measurement Protocol (TWAMP) Request-TW-Session message's Type-P Descriptor format is not RFC-compliant. [PR1305752](#)
- When an "auditd" child process is terminated, the **System reaching processes ceiling <low or high or critical> watermark** error message might be seen. [PR1305964](#)
- 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](#)
- Service cookie opaque data reset wrongly leading data sent to service pic getting corrupted. [PR1310904](#)
- The built-in MPC in MX5/10/40/80 might crash due to CPU hogging after the chip fails to initialize. [PR1312286](#)
- AMS ICMP error handling forwarding to the correct service PIC-Packet Forwarding Engine. [PR1313668](#)
- 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 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 **fabric-header-crc-enable** feature is enabled. [PR1320874](#)
- Traffic with more than two VLAN tags might be incorrectly rewritten and sent out. [PR1321122](#)

- In MX104 router, the **sdk-vmmd: %USER-3: is_platform_rainier: Platform could not be detected** syslog is logged with the wrong severity level. [PR1321622](#)
- The **no-propagate-ttl** might not take effect if **chained-composite-next-hop ingress l3vpn extended-space** is configured. [PR1323160](#)
- MAC addresses might not be learnt on MX Trio-based card due to the negative value of the bridge MAC table limit counter. [PR1327723](#)
- The packet might get dropped in LSR if MPLS pseudowire payload does not have control word and its destination MAC starts with '4' or '6'. [PR1327724](#)
- Traffic loss might be observed on lt- interface. [PR1328371](#)
- Directories and files under **/var/db/scripts** do not have execution permission or the jet directory is missing under **/var/db/scripts** causing the **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](#)
- Denial of service occurs in telnetd (CVE-2018-0061). [PR1331234](#)
- Router opens a database prompt at netisr_process_workstream_proto. [PR1332153](#)
- RPM mib pingResultsMinRtt, pingResultsMaxRtt, pingResultsAverageRtt response as "1" while target address is unreachable, 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](#)
- Backup Routing Engine kernel crash is observed on committing **set system management-instance**. [PR1335903](#)
- The MPC might crash after setting max-queues to a very large number. [PR1338845](#)
- Route corruption in Packet Forwarding Engine with connectivity fault management enabled for Layer 2 circuit. [PR1338854](#)
- While downgrading a Junos OS platform from a later release, the router goes into amnesiac state. [PR1341650](#)
- Configuring the same DHCP server in different routing instances is not supported in a DHCP relay scenario. [PR1342019](#)
- Transition of VRRP backup to master might result in dead next hops. [PR1342707](#)
- Route corruption in Packet Forwarding Engine with connectivity fault management enabled for Layer 2 circuit. [PR1342881](#)
- Junos OS: Multiple vulnerabilities NTP. [PR1343195](#)
- The rpd might crash when performing a 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](#)
- Junos OS: cURL: Multiple vulnerabilities in multiple cURL versions. [PR1347361](#)
- The IPv4 GRPS traffic over aggregated Ethernet interface might be dropped if `gtp-tunnel-endpoint-identifier` is configured. [PR1347435](#)
- FPC CPU utilization with `lt-` interfaces is pegged continuously at 100%. [PR1348840](#)
- ICMP error messages are not generated if 'don't fragment' packets exceed the MTU of the multiservice interface. [PR1349503](#)
- JNH memory leak is seen with VTEP traffic. [PR1356279](#)
- Traffic is dropped without notification in a large-scale scenario. [PR1357707](#)
- On Junos OS, the next-hop index allocation fails and private index space gets exhausted through incoming ARP requests to the management interface (CVE-2018-0063). [PR1360039](#)
- Junos OS: Multiple vulnerabilities in libxml2. [PR1364019](#)
- The **Disconnected after ISSU and before switchover** error message might be seen and FPC is restarted during unified ISSU. [PR1364514](#)
- Subscribers over aggregated Ethernet interface might have tail drops that will affect the fragmented packets due to the QXCHIP buffer getting filled up. [PR1368414](#)
- Forwarding is broken after adding protocol EVPN extended-vlan-id. [PR1368802](#)
- The host outbound traffic might get dropped when `class-of-service host-outbound-traffic ieee-802.1 rewrite-rules` is configured. [PR1371304](#)
- Traffic might drop on newly added interfaces on MX Series router after unified ISSU. [PR1371373](#)
- JNH memory leaks occur in multicast scenario with MoFRR enabled. [PR1373631](#)
- FPC crash might be seen after the FPC restarts. [PR1380527](#)
- Packet drops on interface if `gigether-options loopback` is configured. [PR1380746](#)
- MAC learning might get stuck on MX Series routers with DPC and MPC. [PR1383233](#)
- The RVT interface might flap. [PR1399102](#)

Routing Policy and Firewall Filters

- Condition-based policy fails to take action even though the condition is matched. [PR1300989](#)
- The `rpd` might crash if `vrf-target auto` is configured for a routing instance. [PR1301721](#)
- 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 `set metric multiplier offset` command might cause overflow or underflow. [PR1349462](#)

Routing Protocols

- The **show bgp summary** command displays an incorrect result while assisting GR. [PR1045151](#)
- Multipath does not recalculate after enabling the **AS-PATH-IGNORE** option, and clearing the session triggers rebuilding of the multipath. [PR1163945](#)
- BGP extended communities with sub-type 4 erroneously displayed at LINK_BANDWIDTH. [PR1216696](#)
- The routing protocol process (rpd) on the backup Routing Engine might restart unexpectedly upon the addition of a new L2VPN routing instance. [PR1233514](#)
- The rpd process generates core files in the ASBR when BGP is deactivated in the ASBR before all stale labels have been cleaned up. [PR1233893](#)
- The routes learned from a BGP peer might not be advertised to others if **add-path** is configured. [PR1246349](#)
- The stale BFD session might remain up on the previous anchor FPC. [PR1246363](#)
- 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](#)
- BGP-LU label might go into "dead" state in the forwarding table after the MPLS address family on the next-hop interface is removed and added again. [PR1262180](#)
- IPv6 BFD session(s) configured under IS-IS might not come up after interfaces comes up. [PR1266211](#)
- MPLS over UDP tunnel creation fails in the absence of a VRF table. [PR1270955](#)
- The rpd might crash after BGP is deactivated or activated. [PR1272202](#)
- PIM is stuck in the "InProgress" state when NSR is enabled. [PR1273538](#)
- BGP-ORR not working correctly in an IS-IS overload scenario. [PR1274802](#)
- The BFD down for BGP might cause customer traffic to be dropped without notification. [PR1276497](#)
- Error messages might be seen when receiving BGP update messages with UNREACH NLRI. [PR1276758](#)
- After bfdd restart, the issue is seen with next-generation MVPN and Layer 2 VPN route exchange causing MVPN and VPLS traffic drop. [PR1278153](#)
- With NSR enabled, rpd might generate core files in the master Routing Engine when there is change in kernel id. [PR1278741](#)
- The rpd core files are generated because of BGP update with malformed optional transitive attributes. [PR1279204](#)
- OSPF neighbors might not come up during router reload under high load if PIM is also configured. [PR1279682](#)
- IS-IS LSPs might be dropped during interoperability with a Cisco device in a segment routing scenario. [PR1280522](#)

- Routing loops might be seen after configuring BGP Prefix-Independent Convergence (BGP PIC). [PR1282520](#)
- 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 (RP) router. [PR1282848](#)
- Some BGP-related traceoptions flag settings are not effective immediately after the configuration commit, until the BGP sessions are flapped. [PR1285890](#)
- The rpd might crash if the dynamic rendezvous point goes down in ECMP topology and also if **PIM join-load-balance automatic** is configured. [PR1288316](#)
- With BGP traceoption enabled, executing the **rollback** and **load merge** commands for the configuration might cause rpd to crash. [PR1288558](#)
- BGP-RR sends full route updates to its RR-Clients when any family **mpls** interface gets bounced due to any fiber cut or manual events causing high CPU spike. [PR1291079](#)
- Multihop BFD sessions flap continuously. [PR1291340](#)
- BGP Monitoring Protocol (BMP) might send malformed route-monitoring messages. [PR1292848](#)
- Graceful Restart helper might lose capabilities during peering establishment. [PR1293174](#)
- Rpd crashes upon receipt of malformed PIM packet (CVE-2019-0013). [PR1293306](#)
- Multicast flow reset might occur on OIF for RPT joined branch when PIM prune comes on another interface. [PR1293900](#)
- Impd (link mangement protocol daemon) crashes repeatedly when **logical-system** is configured on the same router. [PR1294166](#)
- The rpd might crash if BGP flap happens. [PR1295062](#)
- ISSU might take more time to complete and the FPC might go offline during ISSU reboot. [PR1298259](#)
- The rpd process might crash because of the AS PATH check error that occurs when RIB groups are added first and later the routing instances are added. [PR1298262](#)
- Inline-BFD on IRB will be broken after GRES/NSR switchover, and the anchor FPC subsequent goes offline. [PR1298369](#)
- MSDP sessions might flap due to data replication stuck between backup and master Routing Engines with huge SA burst between peers. [PR1298609](#)
- The rpd process might crash on the backup Routing Engine. [PR1298711](#)
- Junos OS: The rpd might crash due to a malformed BGP update packet (CVE-2018-0020). [PR1299199](#)
- BGP might send incorrect AS path when alias is enabled and multiple peers are under the BGP group. [PR1300333](#)
- IBGP route damping does not take effect on IBGP **inet-vpn** address family. [PR1301519](#)

- The rpd process might crash with a core file while deleting a multipath route. [PR1302395](#)
- BGP sessions established without SYNC flag. [PR1302426](#)
- Multicast traffic might be pruned for random groups following a designated router failover. [PR1303050](#)
- 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](#)
- The BFD session might flap when querying interface statistics through SNMP or executing a show command through CLI in vMX. [PR1305308](#)
- BGP traceoption logs are still written when it is deactivated. [PR1307690](#)
- 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](#)
- With resource public key infrastructure (RPKI) enabled, rpd successive crashes are seen during route validation database processing. [PR1309944](#)
- BGP labeled-unicast protection might break multicast reverse path forwarding (RPF). [PR1310036](#)
- The BGP session might flap when the connection between the master Routing Engine and the backup Routing Engine keeps flapping, with NSR configured. [PR1311224](#)
- The rpd might crash when the neighbor IS-ISv6 router is restarted, causing route churn. [PR1312325](#)
- BGP route age gets refreshed when the secondary path goes down, with BGP PIC enabled. [PR1312538](#)
- The IS-IS SPF might be triggered by LSP updates containing changes only in **reservable bandwidth**. [PR1313147](#)
- The rpd might crash if RIP neighbor is configured with the local interface IP address. [PR1313712](#)
- BGP prefixes with three levels of recursion for resolution get stuck with a stale next hop at the first level after a link down event. [PR1314882](#)
- The rpd might constantly consume high CPU in a BGP setup. [PR1315066](#)
- On a chassis with BMP configured, the rpd might crash when the rpd process is gracefully terminated. [PR1315798](#)
- OSPF routes cannot be installed to the routing table until the lsa-refresh timer expires. [PR1316348](#)
- The primary path of MPLS LSP might switch to another address. [PR1316861](#)
- Isdb entry cleanup might cause rpd crash, if loop free alternative is configured. [PR1317023](#)
- The inactive route cannot be installed in multipath next-hop after disabling and enabling the next hop interface in a Layer 3 VPN scenario. [PR1317623](#)

- The MPLS labels next hop for IPv4 labeled unicast route are incorrect if some changes are made to the active LDP route. [PR1317800](#)
- BGP-LU update oscillates with BGP-PIC. [PR1318093](#)
- Remove syslog message that got added to code unintentionally. [PR1318458](#)
- IS-IS might choose a sub-optimal path after the metric change in ECMP links. [PR1319338](#)
- Traffic might get blackholed temporarily when BGP GR is triggered and the direct interface flaps. [PR1319631](#)
- The rpd process might crash when deactivating the static route if the next-hop interface is of P2P type. [PR1323601](#)
- When prefix limit is reached, increasing **maximum-prefixes** does not take effect. [PR1323765](#)
- BGP peer is not established after routing engine switchover when graceful-restart and BFD enabled. [PR1324475](#)
- Process mcsnoopd memory leak occurs. [PR1326410](#)
- IGMP snooping might be enabled unexpectedly. [PR1327048](#)
- Multiple next hops might not be installed for IBGP multipath route after an IGP route update. [PR1327904](#)
- The rpd might crash on the backup Routing Engine after the BGP peer is deleted. [PR1329932](#)
- Manual GRES with MX-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 backup selection policy. [PR1333198](#)
- The discard next hop is being installed when the primary LSP interface drops. When the primary interface returns, discard next hop remains until BGP LU neighbor is cleared. This impacts only the cloned route (S=0.) [PR1333570](#)
- IGMP joins are not processed with **passive allow-receive** configured on IGMP interfaces. [PR1334913](#)
- BGP sessions get stuck in active state after the device is restarted at the remote end (Cisco). [PR1335319](#)
- The rpd might crash if SRLG information is in the protocol IS-IS. [PR1337849](#)
- The rpd might crash after the remote BGP peer closes the TCP session. [PR1340379](#)
- The rpd crashes due to receipt of crafted BGP NOTIFICATION messages (CVE-2018-0037). [PR1340689](#)
- The rpd crash might occur when receiving BGP updates. [PR1341336](#)
- The VRF static route might not be exported when **route-distinguisher-id** is used on route reflector in a BGP Layer 3 VPN scenario. [PR1341720](#)
- Changes to the displayed value of AIGP in the **show route ... extensive** command. [PR1342139](#)
- Traffic is discarded without notification if local DUT receives BFD-down. [PR1342328](#)

- The rpd might crash when EBGp neighbor flaps. [PR1342481](#)
- The rpd might crash when deleting or deactivating the VRF routing instance in a BGP Layer 3 VPN environment. [PR1343578](#)
- The rpd process might crash after GRES when multipath is configured. [PR1346954](#)
- The rdp might crash if a route for RPF uses a qualified next hop. [PR1348550](#)
- A rpd crash might be seen after executing a Routing Engine switchover. [PR1349167](#)
- IGMP snooping over LDP VPLS might lose an Ipi-bound downstream snooping next hop after certain multicast topology changes. [PR1349388](#)
- Traffic loss might be seen after the upstream interface shifts from one to another during receipt of the PIM prune packet. [PR1350806](#)
- The source-as community is not appended to RP (display issue in "show route" detail output). [PR1353210](#)
- Ukern memory leak and core crash are seen in BGP environment. [PR1366823](#)
- Static route gets unexpectedly refreshed on commit when configured with **resolve**. [PR1366940](#)
- About 10 minutes of traffic loss is caused by BGP flap during MX Series ISSU. [PR1368805](#)
- The static route might persist even after its BFD session goes down. [PR1385380](#)
- The rpd might crash after issuing the **show route detail** command for RIP route. [PR1386873](#)
- Penultimate-hop router does not install BGP LU label, causing traffic to be discarded without notification. [PR1387746](#)
- IGMPv3/MLD membership requests could not work normally. [PR1389119](#)
- An rpd core file might be dropped due to a soft assert if a non-BGP protocol route with an AS_PATH is used. [PR1391767](#)
- The pcmd on the Routing Engine might run with high CPU utilization after Routing Engine switchover. [PR1392704](#)
- The rpd generates core files on the backup Routing Engine during neighborship flap when using an authentication key with more than 20 characters. [PR1394082](#)
- The best and the second-best routes might have the same weight value if BGP PIC is enabled. [PR1395098](#)
- A rpd soft core file might be seen when Layer 2 VPN is used. [PR1398685](#)

Services Applications

- DTCP non-optimized trigger attributes can delay mirrored traffic forwarding in scaled environments. [PR1269770](#)
- Lawful intercept: ingress control packets from the subscriber are mirrored to the mediation device twice. [PR1275592](#)
- Business service fails to get deactivated after Routing Engine switchover. [PR1280074](#)

- Backup Routing Engine goes to the database prompt with a **vmcore** file if the configuration for the ASI interface that has gone down is deleted. [PR1281882](#)
- PCP mappings cannot be manually cleared when a NAT pool is shared between PCP and standard NAT. [PR1284261](#)
- TLVs in ICRQ for actual-rate-downstream/actual-data-rate-upstream do not reflect PPPoE-IA value. [PR1286583](#)
- One of the internal HA queues get corrupted , which results in mspmand generating a core file on the backup SDG. [PR1291664](#)
- L2TP subscribers are down after a GRES while verifying framed IPv6 route support for an 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 GRES switchover. [PR1295248](#)
- L2TP subscribers might get stuck in terminating state during login. [PR1298175](#)
- [OC/ST] Continuous generation of *jl2tpd_era_Ins* log files occurs even though l2tp is not configured. [PR1302270](#)
- LTS clients experience packet drop for large packets due to fragmentation in LTS. [PR1312691](#)
- L2TP Tunnel Tx and Rx Bytes counts sometimes decrease when subscriber sessions are reduced within the tunnel. [PR1318133](#)
- SNMP MIBs do not yield data related to sp- interfaces. [PR1318339](#)
- The MRU might be changed to 1492 instead of the default 1500 in an L2TP scenario. [PR1319252](#)
- A long route remains in forwarding table after subscriber session goes down. [PR1322197](#)
- L2TP LTS might drop the first "CHAP Success" packet from an LNS due to delayed programming of /136 route on Packet Forwarding Engine. [PR1325528](#)
- The jl2tpd might crash if the RADIUS server returns 32 tunnel-server-endpoints. [PR1328792](#)
- Not all CSURQ replied. [PR1330150](#)
- Crash at ../src/junos/lib/libjuniper/mgmt-sock/mgmt_sock_select_info.c:35. [PR1337406](#)
- Command **show services stateful-firewall flows count** shows incorrect flow count after services configuration change. [PR1338704](#)
- Internal termination code and RADIUS Acct-Terminate-Cause in RADIUS Acct-Stop for a tunneled PPP session might be incorrect. [PR1339911](#)
- 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](#)

- UDP checksum inserted by MS-DPC after NAT64 is not valid when incoming IPv4 packet has UDP checksum set to 0. [PR1350375](#)
- The **show services stateful-firewall** flows counter shows ridiculously high numbers. [PR1351295](#)
- The JI2tpd process might crash shortly after one of the L2TP destinations becomes unavailable. [PR1352716](#)
- IPsec tunnels might flap when SNMP walk is executed if IPsec is configured with DPD enabled. [PR1353240](#)
- L2TP access concentrator (LAC) tunnel connection request packets might be discarded on the LNS device. [PR1362542](#)
- Some tunneled PPPoE subscriber stuck in terminating state in corner case. [PR1363194](#)
- Accounting stop message is not sent to RADIUS server after bringing down the L2TP subscriber. [PR1368840](#)
- IPsec-VPN IKE security-associations might get stuck in "Not Matured" state. [PR1369340](#)
- NAT64 does not translate ICMPv6 Type 2 packet (packet is too big) correctly when MS-DPC is used for NAT64. [PR1374255](#)
- Twice NAT is not supported on FTP ALG and causes MS-PIC crash. [PR1383964](#)
- L2TP subscribers might be stuck in initialization state in a corner case. [PR1391847](#)
- 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 the IPv4 ICMP packet in a NAT64 with MS-DPC deployment. [PR1402450](#)

Software Installation and Upgrade

- New versions of Junos OS do not have the tool for accessing the aux port - `/usr/libexec/interposer`. [PR1329843](#)
- Commit might fail in single-user mode. [PR1368986](#)

Subscriber Access Management

- The DNS might not be assigned when **authentication-order none** is used for subscribers. [PR1273034](#)
- The DHCP subscriber might not get an IP address if the address pool utilization is tight. [PR1274870](#)
- The bbe-smgd might crash after it is restarted in a scaled subscriber management scenario. [PR1277099](#)
- Some RADIUS attributes might not be filtered out of the accounting-on/accounting-off message on an MX Series router. [PR1279533](#)
- The authd might crash when deleting RADIUS configurations in a subscriber environment. [PR1283109](#)

- The IP addresses of subscribers assigned by RADIUS might be counted within the local pool incorrectly after Virtual Chassis switchover. [PR1286609](#)
- The authd process generates a core file at DynamicRequestEntry::addHistory authd_aaa_dyn_req. [PR1289215](#)
- A few IP addresses might be stuck on a policy and charging rules function (PCRF) router. [PR1302509](#)
- Service interim for DHCP subscribers is not working in a 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](#)
- Subscriber might be stuck in "Init" state when **test aaa xxx** command is executed. [PR1311263](#)
- Memory leak might happen after clearing subscriber with script or manually. [PR1312517](#)
- Service interim missing for random users in a JSRC scenario. [PR1315207](#)
- The delegated prefix from RADIUS is incorrectly parsed when the prefix length is fewer than 20 bytes long. [PR1315557](#)
- The unified ISSU is allowed to proceed when the account is suspended. [PR1320038](#)
- IP addresses are assigned discontinuously from the linked IP pools. [PR1323829](#)
- multiple-radius-servers having different dynamic-request-port is not supported. [PR1330802](#)
- In dual stack subscribers scenario with NDRA pool configured, the linked pools are not used when the first NDRA pool is exhausted. [PR1351765](#)
- The subscriber might fail to bind, and some processes might restart in a large-scale subscriber environment due to a rare timing issue. [PR1358339](#)
- The authd process might not be started after executing a Routing Engine switchover on the backup Routing Engine 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](#)
- CoA updates subscriber with the original dynamic-profile if RADIUS has returned a different dynamic-profile name. [PR1381230](#)
- Some subscribers fail to get SRL service as provided in the 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 might consist of additional AVP **Destination-Host** even it is not configured for Gx-Plus session. [PR1384011](#)

- The **authd: gx-plus: logout: wrong state for request session-id <xyz>** log message is seen when a subscriber manually logs out. [PR1384599](#)
- JSRC used RADIUS service accounting protocol instead of JSRC for the SRC installed service. [PR1403835](#)

User Interface and Configuration

- The commitd process might generate a core file when removal of certain configuration is followed by a commit operation. [PR1267433](#)
- CLI session might die while issuing command **show configuration | compare rollback 1**. [PR1331716](#)

VPNs

- An rpd memory leak in processing L2CKT/L2VPN configuration leads to its crash as it is out of memory. [PR1220363](#)
- The rpd crashes and generates core files on the backup Routing Engine. [PR1258595](#)
- Next-generation MVPN IPv6 RP bootstrap type 3 S-PMSI AD route prefix ff02::d persists after BSR data stop. [PR1269234](#)
- The routing protocol process (rpd) crashes after a Layer 2 VPN configuration change followed by **ping mpls l2vpn**. [PR1272612](#)
- Memory leak in rpd in Rosen7 MVPN scenario. [PR1276041](#)
- 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](#)
- Layer 2 circuits stitched through logical tunnel peer interfaces might be stuck in "LD" (local site signaled down) state. [PR1305873](#)
- The rpd might crash on the standby Routing Engine during a Routing Engine switchover if l2circuit is configured. [PR1310934](#)
- Non-optimal route to source might be selected for next-generation MVPN with **unicast-umh-election** enabled. [PR1315011](#)
- Un-hide **set protocols pim mvpn family inet6 disable** configuration to allow users to disable **inet6** on MVPN. [PR1317767](#)
- The rpd might crash after ISSU in a large scale scenario with PIM configuration. [PR1322530](#)
- Moving an MC-LAG from an LDP-based pseudowire to a BGP-based pseudowire might cause rpd crash. [PR1325867](#)
- MVPN sender-site configuration not allowed with S-PMSI. [PR1328052](#)
- A rpd core file is on the backup Raouting Engine with NG-MPVPN and NSR configuration. [PR1328246](#)
- The rpd might crash on the backup Routing Engine when changing the **l2circuit virtual-circuit-id** in an NSR scenario. [PR1345949](#)

- The rpd might crash on the backup Routing Engine when changing the **virtual-circuit-id** in an l2circuit scenario. [PR1345949](#)
- The process rpd might crash after configuration change in a Layer 2 VPN scenario [PR1351386](#)
- In dual-homed next-generation MVPN, the receipt of type 5 withdrawal removes downstream join states for some routes. [PR1368788](#)
- High rpd CPU utilization on the backup Routing Engine might be observed in an MVPN+NSR scenario. [PR1392792](#)
- Downstream interface is not removed from multicast route after receiving a PIM prune. [PR1398458](#)

Resolved Issues: 17.1R2

Class of Service (CoS)

- The cosd process might crash when you execute the command **show class-of-service queue-consumption**. [PR1066009](#)

Forwarding and Sampling

- In proto file AccessListObjBind message the structure needs to change. [PR1230587](#)
- J-Flow v9 is sending the flows with the source-address inverted in show firewall log. [PR1249553](#)
- In MX Series subscriber management environment, l2ald daemon might crash during EVPL subscriber login logout loop. [PR1258853](#)
- Service stats reported in the wrong direction. [PR1262876](#)
- Routing-instances information is not updated in the flat accounting file. [PR1275225](#)

General Routing

- Temp Sensor Fail alarm might be raised incorrectly while an AS-MCC PIC is coming up. [PR1036412](#)
- ICMP reply traffic might get dropped on MS-MPC line cards. [PR1059940](#)
- MPLS traffic might not route through MX Series platform for ingress LSP with channelized E1/T1 circuit emulation MIC interface as the outgoing interface [PR1064515](#)
- Log message jnh_if_get_input_feature_list(9723): Could not find ifl state. [PR1140527](#)
- Port block efficiency and Unique pool users stats shows negative and INFINITY value respectivity in the NAT pool which is being used by the sessions, upon adding address into the NAT pool which is not being used by the sessions, both NAT pools are used under the same SS. [PR1177244](#)
- The destination-prefix-list support list is added for NAT rule with twice-napt-44 translation. [PR1177732](#)
- Interfaces on the MIC-3D-4XGE-XFP installed in MPC2E-3D-NG or MPC3E-3D-NG might flap when they are connected to a DWDM device. [PR1180890](#)
- MS MIC crash might be seen in some instances when there is a service configuration. [PR1183828](#)

- Syslog "JAM: Plugin installed for %s PIC" logged as ERROR level. [PR1189100](#)
- NAT IP pools information split between AMS members is incorrect after rebooting the FPC/ PIC. [PR1190461](#)
- The CPU of processes might get near 100% and messages are repeatedly logged into syslog when restarting the agentd process several times. [PR1192366](#)
- On MX Series and EX9200 platforms, an enhancement is needed for implementing sensor specific temperature thresholds. [PR1199447](#)
- The command **show subscribers summary port extensive** output might have the wrong tunneled/terminated sessions count. [PR1206208](#)
- The ppsman based sessions might be flapping when executing offline/online MIC-3D-20GE-SFP MIC inserted into MPC2E-NG/MPC3E-NG. [PR1211702](#)
- Syslog message : **fpc_pic_process_pic_power_off_config:xxxx :No FPC in slot y** is displayed on empty FPC slots with no PIC power off configured. [PR1216126](#)
- The routers equipped with NG-REs might raise memory size mismatch alarm after upgrade. [PR1220061](#)
- CoS service with Reflexive cos-rule should modify CoS values for reverse flow. [PR1227021](#)
- **vbf_ifl_bind_change_var_walker:377: ifl.demux.22698 (1073764522): IFL TCP (38) Bind change notify ran for 1480 us** log messages are often seen. [PR1229967](#)
- Optional service with blanks in a service string causes session termination. [PR1232287](#)
- High MPC5 CPU on a scaled setup with 64 - 128 K subscribers. [PR1233452](#)
- Dynamic-profile service with service-volume (VSA 67) data collecting interval is not 5 minutes. [PR1234887](#)
- PIC-based MPLS J-Flow not working with MPLS packet sampling at egress side. [PR1236892](#)
- LI enabled subscribers might experience packet drops because of MAC validation failures. [PR1237519](#)
- Junos Telemetry Interface: Frequent disconnects seen in MQTT when IFL sensor is provisioned for longer duration. [PR1238803](#)
- MPC9E might generate FPC core file on Junos OS Release 16.1R2.11, when configured with "mixed-rate AE bundles" and "adaptive load balancing". [PR1238964](#)
- MIB ifJnxTable is not supported. [PR1240632](#)
- Session database synchronization might fail in certain scenarios. [PR1241162](#)
- Untagged bridged traffic might not be mirrored on the second port of the mirrored group. [PR1241403](#)
- **ms90 kernel: kern.maxfiles limit exceeded by uid 0, please see tuning(7)** message seen after injecting more than 2M routes. [PR1243581](#)
- MXVC-Some VBF flows are missing after FPC restart. [PR1244832](#)
- Route Target per bridge domain for EVPN is not supported. [PR1244956](#)

- MX2010/MX2020 (AC & DC) PSMs goes to Present State whenever there is a feed failure even though the PSM properly gives output power. [PR1245459](#)
- The jsd process might crash while subscribing for telemetry data with 2 seconds frequency. [PR1247254](#)
- The rpd process might crash and restart when a MAC address is learned from a given PE on a different ESI. [PR1247338](#)
- PADI dropped due to duplicate client. [PR1248282](#)
- The bbe-smgd process might crash if duplicate variable names are used for different purposes in the dynamic-profile configuration. [PR1248725](#)
- **telemetry_start_polling_fd: evSelectFD failed, errno: 9** messages are continuously seen in the log. [PR1248813](#)
- Only one IA-NA dhcpv6 (without PD request) can establish in case two or more subscribers are provided with the same PD from RADIUS. [PR1249837](#)
- Syslog "JAM:PL: Registered attributes for c23" should be logged as INFO. [PR1250091](#)
- MPC5E/MPC2E-NG/MPC3E-NG/MPC7/MPC8/MPC9 might crash due to a software defect. [PR1250335](#)
- Ukern process crash on Linux based FPC due to a scheduler issue. [PR1250691](#)
- smihelperd core file is generated during subscriber logout process. [PR1250760](#)
- RADIUS Accounting Stats of subscribers get doubled after unified ISSU. [PR1250919](#)
- The rpd might crash when some interfaces go down and some peers go down. [PR1250978](#)
- Cosmetic issue occurs on MS-MIC-16G when you enable it online. [PR1251400](#)
- KRT queue stuck on Routing Engine causes RIB and FIB to go out of sync. [PR1251556](#)
- When a non-0 slot MIC is re-inserted or replaced, the MIC might fail to come online and MIC0 info might disappear. [PR1252998](#)
- **show pfe statistics traffic** displays 2^64 counter for packets output. [PR1253299](#)
- The Routing Protocol process (rpd) might restart unexpectedly when waiting for an acknowledgment from kernel (with "indirect-next-hop-change-acknowledgements" configuration option). [PR1254735](#)
- Interface is not coming up on MPC3E-NG/MPC2E-NG line cards between third party switches. [PR1254795](#)
- After switchover, KRT queue might get stuck on the new master RE with the error "ENOENT -- Item not found". [PR1254980](#)
- Incorrect data in the output of 'show subscribers extensive '. [PR1255029](#)
- MX Series FPC crash due to out of memory condition when an IRB is part of a L3 multicast group. [PR1255290](#)
- Multiple Riot core files might be seen in VMX platform. [PR1255866](#)

- The messaged **krt_decode_comp** read a non specific nh from kernel nhid is constantly seen after upgrading to Junos OS Release to 16.2R1-S1. [PR1256197](#)
- Core files are constantly were observed when NAT term calls application-set with no active applications. [PR1258060](#)
- Unable to run "show subscribers extensive" and some other CLI commands after GRES because the subscriber-management database is unavailable. [PR1258238](#)
- na-grpc log handling needs to be fixed. [PR1258484](#)
- DCD daemon crashes during the ATM related configuration commit. [PR1258744](#)
- When using an AMS interface and running the show interfaces extensive command the sub-interfaces will only show 0 for the packet counters. [PR1258946](#)
- QSFPP-40GBASE-LR4 might remain down after fiber link flap. [PR1259930](#)
- Incorrect egress classification of L3 multicast traffic from ingress VLAN bridge interface after configuration change. [PR1260413](#)
- MPC going offline during unified ISSU. [PR1260714](#)
- A Packet Forwarding Engine saves only the first multicast IPv4 packet when waiting for a resolve request. [PR1260729](#)
- Deviation in dynamic profile service accounting. [PR1260898](#)
- During multicast activation of dynamic subscribers via a service profile, the bbe-smgd daemon in backup Routing Engine could sometimes crash. [PR1261285](#)
- GRPC physical interfaces *-pkts fields zero suppressed by its own counter. [PR1261589](#)
- Dynamic VLAN is removed after 30 seconds if there are no subscribers on it and remove-when-no-subscribers is set regardless of its idle-timeout. [PR1262157](#)
- ICMP network unreachable message is not sent back when the subscriber is terminated in vrf. [PR1263094](#)
- Dynamic VLAN interface is logged out upon reaching idle-timeout even though there is a client session (PPPoE or DHCP) above it. [PR1263131](#)
- CoS Service Profile without line rate adjust needs to use "adjust-always" for proper revert behavior. [PR1263337](#)
- Socket for JSD is not listening randomly after router reboot or JSD process crash. [PR1263748](#)
- smg-service subsystem is not responding to management requests. [PR1264038](#)
- In the Ethernet frames with more than 2000 bytes of payload, the mspmand process might crash. [PR1264712](#)
- MX LAC does not send packets in the l2tp tunnel for some static ppp subscribers. [PR1265414](#)
- PRPD/JET API: BgpRouteMonitorRegister() might not send end-of-rib operation. [PR1265427](#)

- After high subscriber churn BBE_DFW_FINDEX_EXHAUSTED: Filter index space exhausted error prevented subscribers from connecting. [PR1265973](#)
- BNG accepts IGMPv3/MLDv2 membership reports sent to non-standard multicast addresses. [PR1266309](#)
- Unified ISSU failure might be seen with Junos OS Release 16.1R4-S1. [PR1266317](#)
- ARP requests are hitting AE_RESERVED_IFL_UNIT (AEx.32767) when VSTP is enabled on double tagged AE IFL. [PR1267238](#)
- bbe-smgd core file is generated after following subscriber login/logout on backup Routing Engine under certain boundary conditions. [PR1267646](#)
- The CLI configuration command **set chassis effective-shaping-rate** is enabled for the MX104. [PR1267829](#)
- **HALP-lbnh_xlate_cntr_db_get_stats:250counter id 1573873: Unable to find lbnh xlate counter** is flooding the syslog. [PR1268452](#)
- Rpd crash and BGP session flapping might be seen during flapping interfaces or when changing configurations. [PR1269116](#)
- xnm:error in rpc-reply in show arp interface | display xml. [PR1269170](#)
- Router MAC extended community is not using standardized value. [PR1269236](#)
- Log message **sdk-vmmd: %USER-3: is_platform_rainier: Platform found as rainier** is logged with error severity. [PR1271134](#)
- The Routing Engine might stop all services after GRES or unified ISSU. [PR1271306](#)
- Some received packets might be incorrectly dropped after 40GE/100GE port is configured under a LAG. [PR1274073](#)

High Availability (HA) and Resiliency

- **Vmcore** file were generated on both VCMm and VCBm at the same time. [PR1274438](#)

Infrastructure

- Smartd **Offline uncorrectable sectors** critical logs keep reporting every 30 minutes. [PR1233992](#)
- A ksyncd crash might be seen on the backup Routing Engine due to stale next hops on the master Routing Engine. [PR1250880](#)
- Kernel core file is generated with `userland_sysctl / sysctl_root / sysctl_kern_proc_env / panic_on_watchdog_timeout`. [PR1254742](#)
- Device is rebooting due to watchdog timeout. [PR1259616](#)

Interfaces and Chassis

- Configuring ODU FRR related otn-options might crash the FPC without producing a core file. [PR1038551](#)
- MPC might crash during unified ISSU from Junos OS Release 15.1R1 to a later release when QSFP/CXP/CFP2 optics are present. [PR1216924](#)

- LCP packets might still be sent after PADT is sent. [PR1234027](#)
- t3 interface not coming up due to incorrect substrate. [PR1238395](#)
- AE target distribution will need "manual" keyword in configuration. [PR1239724](#)
- MX Series can calculate MTU value incorrectly on pp0 interface. [PR1240257](#)
- DT_LNS: NCP is not responding and gets stuck in ncpResponseBufferDelayed. [PR1241946](#)
- Static PPPoE session cannot be established after GRES. [PR1245465](#)
- The cfmd might crash when CFM filter refers to a firewall policy. [PR1246822](#)
- Need **send-chassis-tlv** configuration statement help text. [PR1248583](#)
- IPv6 ND does not work for DHCPv6 sessions when using static Demux VLAN with RA. [PR1250313](#)
- SNMP reporting ifHCInUcastPkts counter value is equivalent to $(2^{64})-1$. [PR1252716](#)
- Daemon cfmd memory leak upon commits if bridge-domain is configured. [PR1255584](#)
- For CFM over AE, incorrect Anchor fpc is selected. [PR1258490](#)
- I2C BUS timeout causes SFP thread hogging and MPC restart. [PR1260517](#)
- IPCP/IPv6CP re-negotiation is terminated by MX Series BNG. [PR1260829](#)
- Jpppd might crash when traceoptions is enabled over PPPoE. [PR1264000](#)
- Message appears: MXVC CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC 0. [PR1264647](#)
- Malformed PPP Echo Reply causing keepalive failure. [PR1273083](#)
- dot1agCfmFaultAlarm with dot1agCfmMepHighestPrDefect="-1". [PR1273278](#)

Layer 2 Ethernet Services

- STP status gets wrong after changing outer vlan-tags [PR1121564](#)
- The MAC address might not be learnt due to spanning-tree state "discarding" in kernel table after RE switchover [PR1205373](#)
- The IPv4/IPv6 packets originating from RE might be corrupted when the bridge domain has 'vlan-id' set to none, but the outgoing L2 interface for the packet is tagged and CoS is enabled [PR1263590](#)
- DHCP is not using the configured IRB MAC as the source MAC in DHCP offer unicast replies. [PR1272618](#)
- The IA_PD prefix might be deleted when MX receives a DHCPv6 IA_NA request [PR1286359](#)
- jdhcpd process core and restart [PR1288475](#)
- JDHCPD memory leak during dhcp/pppoe login / logout loop [PR1289780](#)
- 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](#)

- Kernel panic using irb and neighbor-discovery secure security-level default [PR1303415](#)
- On EX9200 , log messages related to DHCP snooping prints IP address in reverse order. [PR1310003](#)
- Multiple jdncpd core files are observed in jdncpd_update_groups at `../..../src/junos/usr/sbin/jdncpd/jdncpd_config.c:2290`. [PR1311569](#)
- DHCPv6 traffic might be dropped in subscriber scenario [PR1316274](#)
- jdncpd core dump after making DHCP config changes [PR1324800](#)
- The snmpget for OID: dot3adInterfaceName might not work [PR1329725](#)
- JSA10868 2018-07 Security Bulletin: Junos OS: A malicious crafted IPv6 DHCP packet might cause the JDHCPD daemon to core (CVE-2018-0034) [PR1334230](#)
- The jdncpd process might spike to 100% from less than 10% when DHCPv6 is used. [PR1334432](#)
- The 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](#)
- When DHCP subscribers are in BOUND (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state, if dhcp-service is restarted then the subscribers in this state are logged out [PR1350710](#)
- DHCP relay agent will discard DHCP request message silently if the requested IP address has been allocated to the other client [PR1353471](#)
- Restart FPC which homing micro-bfd link causes LACP core [PR1353597](#)
- JSA10889 2018-10 Security Bulletin: Junos OS: The jdncpd process crash during processing of specially crafted DHCPv6 message (CVE-2018-0055) [PR1368377](#)
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. [PR1373807](#)

MPLS

- The rpd might crash while making static LSPs go up. [PR1084736](#)
- RSVP LSP might not honor TE metric change. [PR1205996](#)
- Entropy label calculation might not provide good load sharing result. [PR1235258](#)
- The LDP routes are not installing with matched L-IS-IS routes in inet.3 route table. [PR1248336](#)
- RPD on backup Routing Engine might consume excessive CPU time if it cannot connect to the RPD on the master Routing Engine. [PR1250941](#)
- When the configured metric for one of the LSPs used in ECMP is removed, other LSPs with configured metric might not honor the configured metric value. [PR1261961](#)
- Traffic loss is seen during auto-BW MBB on ingress router as "invalid fabric token". [PR1264089](#)
- When "explicit-null" is configured for LDP, label 0 is assigned as IPv6 explicit null label. [PR1264753](#)
- Remote targeted LDP session might remain up, though it should not be up. [PR1266802](#)

- TE++ Container LSP statistics are showing the same 10 LSPs and looping. [PR1267774](#)
- FRR bypass tunnel does not appear to be working; the bypass label looks incorrect. [PR1270877](#)
- The CLI command **show route extensive** might cause RPD to crash. [PR1272993](#)

Network Management and Monitoring

- Empty responses for SNMPv3 bulk-get requests if SNMP max message size is lower than OID value. [PR1207683](#)
- Eventd process stops sending syslog message to a configured syslog server. [PR1246712](#)
- SNMPv3 trap does not contain routing-instance information in contextName field. [PR1265288](#)

Platform and Infrastructure

- NPC generated core file with reference to [0x41490f64 in trinity_policer_free (result_ptr=0x5d671f64, nh_ptr=0x5d671f78). [PR1071040](#)
- MPC cell packing wedge might occur with multicast or bridge flood traffic. [PR1180397](#)
- The "rdd" process is restarted in get_mview_root() during GRPC JVISION activation while chassis Packet Forwarding Engines are coming up. [PR1225086](#)
- MAC entry aging is not updated with Source MAC refresh on MPC3E/MPC4E line card at slow traffic rate. [PR1230516](#)
- The apply-path functionality might get broken after you change it. [PR1232299](#)
- The FPC crash or only traffic loss might be seen on MPC1E/2E/3E/4E or MPC-3D-16XGE-SFPP during ISSU. [PR1241729](#)
- Minimum buffer value programmable in the Packet Forwarding Engine changed from 4096 bytes to 1568 bytes. [PR1246197](#)
- MPC or FPC cards report LUCHIP EDMEM errors during ISSU. [PR1249395](#)
- The configuration database is locked when a user that was configure exclusive is logged out unexpectedly. [PR1250305](#)
- The auditd might crash when RADIUS accounting is configured and the RADIUS accounting server becomes unreachable. [PR1250525](#)
- Unexpected flooding for a known unicast VPLS or BRIDGE traffic ingress MPC5 or MPC6 might be observed intermittently toward remote Packet Forwarding Engines. [PR1255073](#)
- GRE tunnel traffic gets dropped after you disable and re-enable the gr- interface. [PR1255706](#)
- FPC might crash and generate a core file during unified ISSU because memory is not properly recycled. [PR1258795](#)
- mgd might crash after you execute the command **show ephemeral-configuration | display inheritance**. [PR1258823](#)

- Mismatching in/out pps value is shown with **show pfe statistics traffic detail**. [PR1259427](#)
- Routed traffic going out via irb/I2 interface with VXLAN EVPN is getting dropped after I2 interface switch. [PR1259551](#)
- DHCP/BOOTP reply packet for an unnumbered interface might trigger FUD process failure. [PR1260623](#)
- WRED drop occurs on one VLAN when the other VLAN is congested. [PR1260951](#)
- DDRIF checksum error might lead to a traffic black hole. [PR1260983](#)
- On a MX Series Virtual Chassis running as a MVPN bud node, traffic is not being forwarded to the local receiver. [PR1261172](#)
- FPC might crash with interface-specific firewall filters with policers configured. [PR1267908](#)

Port Security

- Traffic drops are seen on MPC7E cards after rekeying of MACsec. [PR1257041](#)

Routing Protocols

- The RPD might crash in large-scale BGP routes environment with multipath configured. [PR1209695](#)
- The bgpPeerState/bgpPeerTable returns an invalid value when there is an IPv6 peer. [PR1233790](#)
- BGP-LU add-path in combination with per-prefix-label can result in incorrect MPLS.0 routing/forwarding swap state. [PR1238119](#)
- Session uptime in **show bfd session detail** output omits seconds if uptime is longer than 24 hours. [PR1245105](#)
- The RPD process might crash if static rt-constrain feature is configured but family route-target is not present on any BGP. [PR1247625](#)
- OSPF nex thop might keep flapping, if multi-area rLFA along with policy is configured. [PR1248746](#)
- LLGR feature does not work between Juniper PE to other vendor's RR. [PR1248823](#)
- The configuration statement **learn-pim-router** not working properly. [PR1251439](#)
- BGP peers remain stuck in idle state after unified ISSU. [PR1261902](#)
- Routing protocol process (rpd) might restart unexpectedly with a reference to `ioth_session_delete_internal()` routine. [PR1261970](#)
- The rpd might crash if the IS-IS segment routing is configured but a certain interface is not configured with RSVP. [PR1262612](#)
- MPLS label entry for direct route as BGP-LU route is permanently stuck in KRT queue when vrf-table-label is configured in CoS VRF. [PR1263291](#)
- When applying import policy to a BGP neighbor, the rpd might crash continuously. [PR1265224](#)
- "Nexthop AFI=3" is observed in BGP open message after you configure **family inet unicast extended-nexthop**. [PR1272807](#)

Services Applications

- Backup SDG reported memory-usage zone in RED. [PR1202872](#)
- L2TP tunnels might get stuck in "Terminating" state on MX Series LNS. [PR1249768](#)
- Traffic is dropped when changing the source-address under a NAT rule term for BASIC-NAT translation. [PR1257801](#)
- L2TP Congestion Window set to 128 instead of 1 when tunnel is created. [PR1265001](#)
- KMD process might crash because of apply-group configuration. [PR1265404](#)
- Kernel crash might be seen after performing the CLI command commit. [PR1273357](#)

Subscriber Access Management

- The auth request does not cause the router to send the RADIUS REQUEST message, "Failed to queue the request, will be queued in authd internal queue". [PR1178813](#)
- Configuration statement **set access radius-options timeout-grace** should be unhidden. [PR1249847](#)
- Need option to exclude tunnel attributes in access-request on LNS. [PR1264024](#)
- Possible CPS degradation for scaled dhcpv4/v6 and pppoev4 subscribers. [PR1264052](#)
- Incorrect number of messages in the queue to RADIUS server in the output **show network-access aaa statistics radius detail**. [PR1267307](#)

VPNs

- IoT issue between Juniper and third party for SSM Rosen 07 based Inter-AS MVPN. [PR1238807](#)
- The L2circuit does not switch based on the APS status. [PR1239381](#)
- Rpd memory leak is observed in NG-MVPN environment. [PR1259579](#)

Resolved Issues: 17.1R1

Class of Service (CoS)

- Incorrect CoS rewrite for L3VPN traffic when chained-composite-next-hop is enabled. [PR1062648](#)
- QMON - Queue 3 in both ingress and egress do not have the correct maximum depth values, in the **show interfaces queue ...** CLI commands. [PR1226558](#)
- The cosd might crash after you activate/deactivate the CoS configuration. [PR1236866](#)
- The error message of **cos_check_temporal_buffer_status** might be observed when configuring Hierarchical CoS with strict-high scheduling. [PR1238719](#)

Forwarding and Sampling

- Local backup for accounting flat files might not perform after transfer to archive site fails. [PR1198095](#)
- The policer on Trio based card allows more traffic when packet size is less than 128 bytes. [PR1207810](#)

- Commit fails after applying bandwidth-percent policer on ps interface. [PR1225977](#)
- Configuration for ipv4-flow-table-size and ipv6-flow-table-size does not propagate to FPC after reboot if sampling instance is not associated. [PR1234905](#)
- J-Flow version 9 cannot get TCP flag information from IPv6 fragment packets. [PR1239817](#)

General Routing

- The MS-MPC/MS-MIC card might crash after the NAT session is removed. [PR1117662](#)
- Trace-route does not work on Services PIC. [PR1163472](#)
- MX240 DC power shows abnormal electrical current value even its external DC power sources circuit breaker is at off position. [PR1177536](#)
- DNS Query fails for fragmented DNS traffic. [PR1182910](#)
- Error messages are reported during unified ISSU on MX Series router. [PR1200045](#)
- Login/logout of PPPoE subscriber causing link up/down traps if **no-traps** command is configured. [PR1204949](#)
- With local source, Continuous iif-mismatch is reported on MoFRR backup interface. [PR1206121](#)
- FPC might crash with any inline feature enabled. [PR1210060](#)
- AMS interface works incorrectly in warm-standby mode. [PR1216030](#)
- Memory allocation might fail in Trio-based FPC due to memory fragmentation. [PR1216300](#)
- RPD consumes high CPU when VPLS instances are configured for the first time or a system with VPLS instances is rebooted. [PR1216332](#)
- Replacing an MQ FPC with an XM one might cause all other MQ-based cards to report "FI Cell underflow at the state stage" on MX Series platform. [PR1219444](#)
- Packet loss might occur when multicast traffic enters and exits the Packet Forwarding Engine in a different FPC. [PR1219962](#)
- On an MX Series Virtual Chassis environment traffic loss might be observed due to incorrectly programmed Aggregated Ethernet interfaces. [PR1220934](#)
- RPD might crash after offlining or onlining FPC/MPC or doing GRES. [PR1221183](#)
- Continuous login and logout PPPoE/DHCP subscribers might cause some subscribers to fail to bind. [PR1221690](#)
- "Show chassis hardware detail" shows ada0 and ada1 entries in reverse order. [PR1222330](#)
- The subscribers are unable to connect due to "uifl inactive issue" error. [PR1222829](#)
- "unnumbered-address" under dynamic profile shows the incorrect value. [PR1222975](#)
- The bbe-smgd process memory might leak in the backup Routing Engine. [PR1223625](#)
- A pfed core file is observed after deleting apply-groups. [PR1223847](#)

- **early/opDel: bad stored heap** messages seen on sending traffic using captive-portal-content-delivery service. [PR1226782](#)
- The chassisd might crash with **show chassis ucode-rebalance** command on MX Series platform. [PR1227445](#)
- Openflow: Flowstat reply has incorrect DL type. [PR1228383](#)
- Different behavior might be observed for TCP and non-TCP RE-generated traffic when the route pointing to indirect next-hop is not subjected to 'load-balance per-packet'. [PR1229409](#)
- Unequal load balance over LSP does not work if destination route is IPv6. [PR1230186](#)
- Interface statistics are not restored on MX Series VC after unified ISSU, which causes the RADIUS volume accounting stats value to remain unchanged. [PR1230524](#)
- The dynamic-profile service filter matches the traffic that is not defined in the prefix-list applied to the filter. [PR1230997](#)
- ICMP identifier is not translated back to expected value during ICMP traceroute for TTL exceeded packets on NAT using Multiservice MPC. [PR1231868](#)
- IPsec SAs are not cleared after disabling the ms interface inside a logical interface IFL. [PR1232276](#)
- Optional service with blanks in a service string causes session termination. [PR1232287](#)
- Some Packet Forwarding Engine statistics counters do not work in MPC7/8/9. [PR1232540](#)
- Packet Forwarding Engine statistics input packets pps counter has a large error. [PR1232547](#)
- Input Framing errors are incrementing on interfaces connected to MPC2E-NG with 4x10G MIC. [PR1232618](#)
- Some error messages might be seen during offlining/onlining FPC or link flap. [PR1232686](#)
- RPD core file is generated with mem_assert , rta_route_session_ref_free, rta_parse_session_delete, task_module_dyn_config_server. [PR1232742](#)
- LSP-ping might fail and IP packets with options will not get mirrored in port-mirror environment. [PR1234006](#)
- SNMP trap description does not match the trap signal. [PR1234083](#)
- offlining/onlining SFB2 can trigger another fabric plane to go to check state. [PR1234224](#)
- After the backup Routing Engine is replaced, the new Backup Routing Engine cannot synchronize with Master Routing Engine if 'dynamic-profile-options versioning' is configured. [PR1234453](#)
- With **show route forwarding table *** enabled protocols field additional flags. [PR1234501](#)
- False login attempts might be seen on MPC7E/8E/9E for receiving noise. [PR1234712](#)
- VLNS(VBNG) - Commit generated a "warning: requires 'l2tp-inline-lns' license" but a valid license is installed. [PR1235697](#)

- The Aggregated Ethernet interface with per-packet load sharing configured might drop packets unexpectedly. [PR1235866](#)
- The outer source MAC in ARP reply packet for IRB interface is different than the inner virtual MAC. [PR1236225](#)
- A stale route is present in inetflow.0 rib after deleting rib-group and deactivating static flow route. [PR1236636](#)
- PIC-based MPLS J-Flow not working with MPLS packet sampling at the egress side. [PR1236892](#)
- Offlining/onlining SFB2 can trigger another fabric plane to go to check state. [PR1237134](#)
- The MS-MPC might crash when receiving internally corrupted frames from another FPC. [PR1237667](#)
- High Routing Engine CPU usage might be seen with router-advertisement configured. [PR1237894](#)
- "Empty license directory copied from the master" logs are seen on backup Routing Engine when the number of licenses for scale-subscriber is exceeded. [PR1238615](#)
- MX Series is sending accounting interim without the update-interval configuration statement. [PR1239273](#)
- Total traffic loss for BGP-PIC learned prefixes occurs on link failure. [PR1239357](#)
- Traceroute will not resolve VRF loopback address where SI and pseudointerface exist. [PR1240221](#)
- Incorrect CoS adjustment and missing adjustment application occur for PPPoE session with dynamic-profile services. [PR1241201](#)
- Delay in PTP clock class changes. [PR1241211](#)
- With IPsec dynamic endpoints (DEP) over IPv6, the ARI IPv6 routes might be missing after GRES with NSR. [PR1242503](#)
- The FPC might crash when adding physical interface sensor. [PR1243411](#)
- The rpd process might crash and restart when a MAC address is learned from a given PE on a different ESI. [PR1247338](#)
- When an IPv6 node receives an ICMPv6 PTB (Packet Too Big) message with MTU < 1280, the node will emit atomic fragments. This behavior might result in denial of service attack. [PR1250832](#)

High Availability (HA) and Resiliency

- Connection might be broken between master and backup Routing Engine after ISSU. [PR1234196](#)

Infrastructure

- The gdb can be exploited to allow execution of unsigned binary. [PR968335](#)
- Continuous kernel logs and LDP stats timeout error occurs when you run **show ldp traffic-statistics**. [PR1215452](#)
- SMART ATA Error Log Structure error: invalid SMART checksum logs are seen after upgrade. [PR1222105](#)

Interfaces and Chassis

- ARP entry learned through Aggregated Ethernet interface does not expire when the ARP IP is no longer reachable. [PR1211757](#)
- MPC might crash during unified ISSU from Junos OS Release 15.1R to later release when QSFP/CXP/CFP2 optics are present. [PR1216924](#)
- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces but in the same routing-instance. Only one logical interface was assigned with the identical address after commit. There was no warning during the commit but just syslog messages indicating incorrect configuration. This issue is fixed and it is now not allowed to configure the same IP address (the length of the mask does not matter) on different logical interfaces. [PR1221993](#)
- RPT MMX Regression: During firewall script run, switchover is performed. The new master takes ownership and stays up but the old master goes to db>. [PR1222582](#)
- Stuck L2TP session remains after session/tunnel termination. [PR1228802](#)
- Interface is not coming up after de-activating and activating "protocols oam ethernet connectivity-fault-management maintenance-domain". [PR1231315](#)
- Commit failure, error: Bandwidth on IFL <static vlan demux interface> cannot be greater than that of its IFD. [PR1232598](#)
- The MX Series routers might fail to send the IPCP Configure-Ack packet to the subscriber. [PR1235261](#)
- NCP is not responding and gets stuck in ncpResponseBufferDelayed. [PR1241946](#)
- JPPPD core file is generated during scaled login/logout. [PR1245848](#)
- VRRP might be stuck in (state: unknown, VR State: bringup) when VRRP is configured on one IFL without VLAN and the lower-unit-number logical interface in same physical interface has VLAN configured. [PR1247050](#)

Layer 2 Ethernet Services

- The MPC might power back on from offline state after you commit the configuration if it is configured to be offline when detecting major errors. [PR1218304](#)
- MX Series is not including Delegated-IPv6-Prefix in accounting interim. [PR1231665](#)
- MX Series BNG waits 30 seconds before replying to any rapid commit option set DHCPv6 Solicit retransmissions messages. [PR1234009](#)
- After upgrading to Junos OS Release 15.1F2-S13 "/var partition is full" alarm might be seen. [PR1237218](#)
- LACP might time out during unified ISSU when LACP is configured in fast periodic along with the **fast-hello-issu** configuration statement. [PR1240679](#)

MPLS

- Both **load-balance-label-capability** and **no-load-balance-label-capability** could be configured under forwarding-options. [PR1126439](#)
- The command **no-install-to-address** not always honored for PCC-delegated LSPs. [PR1169889](#)
- The rpd process might crash when dynamic-tunnel is configured but RSVP signaling is disabled. [PR1213431](#)
- FPC sockets disconnects and various scheduling slips occur when executing the **show ldp traffic-statistics** command with many ECMP links and L3VPN routes. [PR1214961](#)
- Carrier-over-carrier VPN PE router "protocol mpls" under RI breaks existing "protocol connection". [PR1222570](#)
- RPT RIAD VMX Regressions : rsvp-lsp-enh-lp-upstream-status is taking more time for synchronization on the backup Routing Engine on egress. [PR1242324](#)

Multicast

- Kernel: %KERN-3: fmbb_uc_pfes_pre: rnh_get_pfe_id failed with ENOTSUP 45. This error is not fatal; it just means that FMBB cannot be done. [PR1230465](#)

Network Management and Monitoring

- The statistics of OID ifOutError incorrectly include ifOutDiscards. [PR1243071](#)

Platform and Infrastructure

- The junos:key attribute is not emitted when the configuration is emitted in JSON format. [PR1195928](#)
- Blank firewall log is generated for IPv6 packets with nexthead hop-by-hop. [PR1201864](#)
- The firewall filters are incorrect after GRES. [PR1230954](#)
- The scripts process might crash when some special combination of jcs:printf(...) and some special characters at the boundary of the buffer are used. [PR1232418](#)
- With non-Ethernet frame payload, traffic might not be correctly load-balanced. [PR1232943](#)
- The scale-subscriber license count might increase to an invalid license state with L2TP/LTS clients. [PR1233298](#)
- NTP.org and FreeBSD have published security advisories for vulnerabilities resolved in ntpd (NTP daemon). Server-side vulnerabilities are only exploitable on systems where NTP server is enabled within the [edit system ntp] hierarchy level. [PR1234119](#)
- Flow-tap-dtcp service login via SSH with key-based authentication fails. [PR1234464](#)
- ADC based line card might fail to boot up on the FPC slot that was previously used for MPC6E. [PR1235861](#)
- J-Flow cannot sample multicast traffic in multi-copy scenario in MX2010/MX2020. [PR1237164](#)
- FPC and Routing Engine might stuck in high CPU usage when DDoS SCFD is turned on. [PR1237486](#)

- FPC might crash during unified ISSU. [PR1239304](#)
- Low temporal buffer configuration is not honored. [PR1240756](#)

Provider Edge Satellite Software

- Traffic forwarding is not working from AD to SD. [PR1231227](#)

Routing Protocols

- The rpd process on the backup Routing Engine might crash because of a memory leak with the PIM configuration. [PR1155778](#)
- The rpd process might crash during MSDP instance deletion. [PR1216078](#)
- The rpd process might crash after performing BGP flapping. [PR1222554](#)
- The rpd might crash when BGP add-path is configured and the same prefix is received from multiple peers with different source AS. [PR1223651](#)
- Rpd core could be seen if MPLS goes down. [PR1228388](#)
- Junos OS 15.1 and later releases might be impacted by the receipt of a crafted BGP UPDATE which can lead to an rpd (routing process daemon) crash and restart. Repeated crashes of the rpd daemon can result in an extended denial of service condition. [PR1229868](#)
- Rpd crash might be seen if ISIS LSP is purged. [PR1235504](#)
- RSVP bandwidth load-balancing is not working after LSPs are advertised in the IS-IS or IS-IS TE shortcuts are configured. [PR1237531](#)
- Rpd generates a core file due to an assertion condition related to changing a policy for a BGP neighbor. [PR1239990](#)
- After doing some configuration modification related to sham-link, the sham-link might not be able to be brought up anymore. [PR1240391](#)
- Multicast route leaking does not work correctly. [PR1240656](#)
- The rpd process might crash if static rt-constrain feature is configured but family route-target is not present on any BGP. [PR1247625](#)

Services Applications

- LNS-Tunnel/session establishment get stalled when the LNS is flooded by high rate L2TP messages. [PR990081](#)
- FTP ALG on MX fails to translate the PORT command when the FTP client uses Active Mode and requests AUTH(SSL-TLS) but the FTP server does not use AUTH. [PR1194510](#)
- The kmd process might consume excessive CPU resources during continuous polling for IKE related data through SNMP. [PR1209406](#)
- Traffic black holes occur due to service-set programming on MS-MPC. [PR1223302](#)

- PPPoE - L2TP subscribers might get stuck in Terminating state in longevity login/logout test. [PR1235996](#)
- MS-DPC - Performance degradation in CGNAT scaling occurs during memory stress. [PR1242556](#)

Subscriber Access Management

- Syslog is not generated when RADIUS server is marked “dead”. [PR1207904](#)
- Gy support is seen for the 3GPP-SGSN-MCC-MNC AVP in CCR messages. [PR1233847](#)
- The DHCPv6 solicits are ignored instead of being responded to with an advertise packet with status code NoPrefixAvail(6) when no delegated prefix is available. [PR1234042](#)
- The authd daemon might generate a core file when traceoption filters are configured during GRES not-ready state. [PR1234395](#)

User Interface and Configuration

- The rpd memory leak might be triggered when configuring or reconfiguring IS-IS interface. [PR1243702](#)
- Uncommitted lines are displayed right after commit with "delta-export". [PR1245187](#)

VPNs

- After issue "clear pim join" on source PE the multicast flow stops in an NG-MVPN scenario with the **asm-override-ssm** configuration statement for the SSM group. [PR1232623](#)
- The rpd might crash on backup Routing Engine when changing the I2circuit neighbor in an NSR scenario. [PR1241801](#)

SEE ALSO

New and Changed Features 99
Changes in Behavior and Syntax 125
Known Behavior 142
Known Issues 150
Documentation Updates 232
Migration, Upgrade, and Downgrade Instructions 234
Product Compatibility 241

Documentation Updates

IN THIS SECTION

- [Subscriber Management Access Network Guide | 232](#)
- [Subscriber Management Provisioning Guide | 232](#)
- [Subscriber Management VLANs Interfaces Guide | 233](#)

This section lists the errata and changes in Junos OS Release 17.1R3 documentation for MX Series.

Subscriber Management Access Network Guide

- The “Configuring a Pseudowire Subscriber Logical Interface Device” and “anchor-point (Pseudowire Subscriber Interfaces)” topics have been updated to state that you cannot dynamically change an anchor point that has active pseudowire devices stacked above it. Both topics describe the steps to follow when you must change such an anchor point.
- 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

- Support for the packet-triggered subscribers and policy control rule base (PTSP) feature was discontinued starting in Junos OS Release 13.1R1, but this was not reflected in the documentation. Text exclusive to PTSP has been removed from the *Broadband Subscriber Sessions User Guide*. This includes all CLI topics and the following chapters:
 - “Configuring the PTSP Feature to Support Dynamic Subscribers”
 - “Configuring the PTSP Partition to Connect to the External Policy Manager”

- “Configuring PTSP Services and Rules”
- “Monitoring and Managing Packet-Triggered Subscribers”

Topics for other features that refer to PTSP are updated to report the end of support.

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

[See [Suspending AAA Accounting and Baseline Accounting Statistics Overview](#).]

- 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 99
Changes in Behavior and Syntax 125
Known Behavior 142
Known Issues 150
Resolved Issues 169
Migration, Upgrade, and Downgrade Instructions 234
Product Compatibility 241

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.1 | 235](#)
- [Procedure to Upgrade to FreeBSD 10.x based Junos OS | 235](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 237](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 239](#)
- [Upgrading a Router with Redundant Routing Engines | 240](#)
- [Downgrading from Release 17.1 | 240](#)

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

Starting with Junos OS Release 15.1, in some of the devices, FreeBSD 10.x is the underlying OS for Junos OS instead of FreeBSD 6.x. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.x remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see [Understanding Junos OS with Upgraded FreeBSD](#).

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.x-based Junos OS	FreeBSD 10.x-based Junos OS
MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 17.1

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.

[Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).

Procedure to Upgrade to FreeBSD 10.x based Junos OS

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install FreeBSD 10.x based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.

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

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

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.1R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.1R3.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.1R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.1R3.9-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 10.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 10.x, and Junos OS (FreeBSD 6.x) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 17.1 **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: MX80, and MX104.

To download and install FreeBSD 6.x based Junos OS:

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

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

- All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/jinstall-ppc-17.1R3.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot  
source/jinstall-ppc-17.1R3.x-limited-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.

- For software packages that are downloaded and installed from a remote location:

- `ftp://hostname/pathname`
- `http://hostname/pathname`
- `scp://hostname/pathname`

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 17.1 **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 15.1, 16.1 and 16.2 are EEOL releases. You can upgrade from Junos OS Release 15.1 to Release 16.1 or even from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

For more information about EEOL releases and to review a list of EEOL releases, see <https://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 17.1

To downgrade from Release 17.1 to another supported release, follow the procedure for upgrading, but replace the 17.1 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 125
Known Behavior 142
Known Issues 150
Resolved Issues 169
Documentation Updates 232
Product Compatibility 241

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 241](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features 99
Changes in Behavior and Syntax 125
Known Behavior 142
Known Issues 150
Resolved Issues 169
Documentation Updates 232
Migration, Upgrade, and Downgrade Instructions 234

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- New and Changed Features | 242
- Changes in Behavior and Syntax | 252
- Known Behavior | 257
- Known Issues | 259
- Resolved Issues | 262
- Documentation Updates | 270
- Migration, Upgrade, and Downgrade Instructions | 270
- Product Compatibility | 274

These release notes accompany Junos OS Release 17.1R3 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 17.1R3 New and Changed Features | 243
- Release 17.1R2 New and Changed Features | 243
- Release 17.1R1 New and Changed Features | 243

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

Release 17.1R3 New and Changed Features

- There are no new features or enhancements to existing features for PTX Series in Junos OS Release 17.1R3.

Release 17.1R2 New and Changed Features

- There are no new features or enhancements to existing features for PTX Series in Junos OS Release 17.1R2.

Release 17.1R1 New and Changed Features

Hardware

- **P3-10-U-QSFP28 PIC (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the P3-10-U-QSFP28 is supported on PTX3000 and PTX5000 routers that have third-generation FPCs installed. The P3-10-U-QSFP28 PIC has ten ports that are configurable as 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet ports. The interface speeds are configured by port group—ports 0 through 4 and ports 5 through 9. To configure the port speed, use the following command:

```
[edit chassis]
user@host# set fpc slot-number pic pic-number port port-number port-speed (10G | 40G | 100G)
```

[See the [PTX Series Interface Module Reference](#).]

- **Upgrade of FPCs in an operational PTX5000**—Starting in Junos OS Release 17.1R1, you can upgrade the first-generation FPCs or second-generation FPCs to third-generation FPCs in an operational PTX5000. You might need to upgrade the following components before you can upgrade the FPCs in a PTX5000:
 - SIBs
 - Fan tray
 - Power distribution unit
 - Power supply module

[See the [PTX5000 Packet Transport Router Hardware Guide](#).]

- **New PIC P3-24-U-QSFP28 (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the PIC P3-24-U-QSFP28 is supported on PTX3000 and PTX5000 routers. The P3-24-U-QSFP28 PIC has 24 ports configurable as either 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.

To install the P3-24-U-QSFP28 PIC, you must have a third-generation FPC installed on your system.

[See the [PTX Series Interface Module Reference](#).]

- **New SIB SIB3-PTX5K (PTX5000)**—Starting in Junos OS Release 17.1R1, the SIB3-PTX5K SIB is supported on PTX5000 routers.

[See the [PTX5000 Packet Transport Router Hardware Guide](#).]

- **New FPCs FPC3-PTX-U1-L, FPC3-PTX-U1-R, FPC3-PTX-U2-L, FPC3-PTX-U2-R, FPC3-PTX-U3-L, and FPC3-PTX-U3-R (PTX5000)**—Starting in Junos OS Release 17.1R1, the FPC3-PTX-U1-L, FPC3-PTX-U1-R, FPC3-PTX-U2-L, FPC3-PTX-U2-R, FPC3-PTX-U3-L, and FPC3-PTX-U3-R FPCs are supported on PTX5000 routers. The FPCs provide the following throughput:
 - FPC3-PTX-U1-L and FPC3-PTX-U1-R—1.0 Tbps
 - FPC3-PTX-U2-L and FPC3-PTX-U2-R—2.0 Tbps
 - FPC3-PTX-U3-L and FPC3-PTX-U3-R—3.0 Tbps

When installing these third-generation FPCs on the PTX5000 chassis, you might need to install the following components:

- SIB3-PTX5K SIBs
- FAN3-PTX-H fan tray
- PDU2-PTX-DC power distribution unit
- PSM2-PTX-DC power supply module

NOTE: Some new features provided by these third-generation FPCs can be accessed only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

[See the [PTX5000 Packet Transport Router Hardware Guide](#).]

- **New horizontal fan tray FAN3-PTX-H (PTX5000)**—Starting in Junos OS Release 17.1R1, the FAN3-PTX-H horizontal fan tray is supported on PTX5000 routers.

[See the [PTX5000 Packet Transport Router Hardware Guide](#).]

- **Third-generation FPCs (PTX3000)**—Starting in Junos OS Release 17.1R1, third-generation FPCs are supported on PTX3000 routers. FPC3-SFF-PTX-U1 FPCs (model numbers FPC3-SFF-PTX-U1-L and FPC3-SFF-PTX-U1-R) support 1.0 Tbps of throughput. FPC3-SFF-PTX-U0 FPCs (model numbers FPC3-SFF-PTX-U0-L and FPC3-SFF-PTX-U0-R) support 500 Gbps of throughput.

Third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) are supported only in a PTX3000 with SIB3-SFF-PTX SIBs. Third-generation FPCs and FPC-SFF-PTX-P1-A first-generation FPCs can interoperate with each other in the same system.

NOTE: Some features provided by these third-generation FPCs can be accessed only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

[See the [PTX3000 Packet Transport Router Hardware Guide](#).]

- **SIB3-SFF-PTX SIBs (PTX3000)**—Starting in Junos OS Release 17.1R1, SIB3-SFF-PTX SIBs are supported on PTX3000 routers. The SIB3-SFF-PTX SIBs are required with third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1). The SIB3-SFF-PTX SIBs also support FPC-SFF-PTX-P1-A first-generation FPCs—third-generation FPCs and FPC-SFF-PTX-P1-A first-generation FPCs can interoperate with each other in the same system.

[See the [PTX3000 Packet Transport Router Hardware Guide](#).]

- **Upgrading to third-generation FPCs and SIBs in an operational router (PTX3000)**—Starting in Junos OS Release 17.1R1, you can upgrade to third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) and SIB3-SFF-PTX SIBs in an operational PTX3000.

[See the [PTX3000 Packet Transport Router Hardware Guide](#).]

- **Support for P2-10G-40G-QSFPP and P2-100GE-OTN PICs on third-generation FPCs (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the P2-10G-40G-QSFPP PIC is supported on PTX Series routers that have third-generation FPCs installed.

[See the [PTX Series Interface Module Reference](#).]

- **The P1-PTX-24-10G-W-SFPP PIC is supported on third-generation FPCs (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the P1-PTX-24-10G-W-SFPP PIC is supported on PTX Series routers that have third-generation FPCs installed.

[See the [PTX Series Interface Module Reference](#).]

- **5-port 100-Gigabit DWDM OTN PIC with CFP2 (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F6 and 17.1R1, the 5-port 100-Gigabit dense wavelength division multiplexing (DWDM) optical transport network (OTN) PIC (PTX-5-100G-WDM) with CFP2 analog coherent optical (CFP2-ACO) pluggable optics on third-generation FPCs is supported on the PTX3000 and PTX5000 series routers. The 5-port 100-Gigabit DWDM OTN PIC supports the following features:
 - Transparent transport of five 100-Gigabit Ethernet signals with optical channel transport unit, OTU4 (V) framing.
 - Dual-polarization quadrature phase shift keying (DP-QPSK) modulation with coherent receiver and soft-decision forward error correction (SD-FEC) for long-haul and metro applications.

- International Telecommunication Union (ITU)-standard OTN performance monitoring and alarm management
- Extensive optical, digital signal processing (DSP) and bit error ratio (BER) performance monitoring statistics for the optical link.

[See the [PTX5000 Packet Transport Router Hardware Guide](#).]

- **New Routing and Control Board RCB-PTX-X6-32G (PTX3000)**—Starting in Junos OS Release 17.1R1, the Routing and Control Board (RCB) is supported on PTX3000 routers. The RCB combines the functionality of a Routing Engine, Control Board, and Centralized Clock Generator (CCG) in a single FRU. Although the functionality is combined in a single FRU, you must install an RCB companion card in the **RE0** and **RE1** slots adjacent to each RCB to enable the RCBs to communicate through the backplane.

Class of Service (CoS)

- **Support for shaping of traffic exiting third-generation FPCs on PTX3000 and PTX5000 routers (PTX Series)**—Beginning with Junos OS Release 17.1R1, you can shape the output traffic of an FPC3 physical interface on a PTX3000 or PTX5000 packet transport router so that the interface transmits less traffic than it is physically capable of carrying. Shaping on all PTX Series packet transport router interfaces has a minimum rate of 1 Gbps and an incremental granularity of 0.1 percent of the physical interface speed after that (for example, 10 Mbps increments on a 10 Gbps interface). You can shape the output traffic of a physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces interface-name]** or **[edit class-of-service traffic-control-profiles profile-name]** hierarchy level and applying the traffic control profile to an interface.

[See [shaping-rate \(Applying to an Interface\)](#).]

- **ISSU Feature Explorer**—Starting in Junos OS Release Feature Explorer, an interactive tool, to verify your device's unified ISSU compatibility with different Junos OS releases.

[See [ISSU Feature Explorer](#).]

Interfaces and Chassis

- **Aggregated Ethernet Statistics Enhancements (PTX Series Routers)**—Starting in Junos OS Release 17.1R1, multicast and broadcast counters from individual links are supported for aggregated Ethernet interfaces and are displayed in the **show statistics ae interfaces** command.
- **Support for different Ethernet rates in aggregated Ethernet interfaces (PTX5000)**—Starting in Junos OS Release 17.1R1, the **mixed** statement is supported for the **link-speed** configuration statement on aggregated Ethernet interfaces. The **mixed** configuration statement is configured at the **[edit interfaces interface-name aggregated-ether-options link-speed (speed | mixed)]** hierarchy level.

[See [link-speed \(Aggregated Ethernet\)](#).]

- **Support for configuring the port speed (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the **speed** configuration statement is used to configure the port speed on interface modules that support multiple port speeds. The **speed (10G | 40G | 100G)** configuration statement is configured at the **[edit chassis fpc slot-number pic pic-number port port-number]** hierarchy level.

[See [speed.](#)]

- **Support for configuring interface loopback (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the **loopback (local | remote)** configuration statement is used to specify whether local or remote loopback is enabled. This allows you to test the transceiver cable connection from the far end to the retimer interface without changing the cable. The **loopback (local | remote)** configuration statement is configured at the **[edit interfaces interface-name together-options]** hierarchy level.

See [loopback \(Local and Remote\).](#)

- **Support for configuring the LED on a port to flash (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, the **led-beacon** command causes the LED for the specified port to flash green. This enables you to physically locate a specific optic port on the PIC. The **led-beacon** configuration statement is configured at the **[edit interfaces interface-name (with port number)]** hierarchy level.

[See [led-beacon.](#)]

- **Synchronous Ethernet clock synchronization on third-generation FPCs (PTX3000)**—Starting in Junos OS Release 17.1R1, Synchronous Ethernet clock synchronization is supported on third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) on the PTX3000.

[See [Synchronous Ethernet Overview.](#)]

- **Integrated photonic line card (IPLC) (PTX3000)**—Starting in Junos OS Release 17.1R1, the PTX3000 can provide a fully integrated photonic line system for converged core and metro core packet optical networks running point-to-point and ring topologies. The following optical components are available for the PTX3000:
 - Integrated photonic line card (IPLC) base module—Provides the combined functionality of a 32-port reconfigurable optical add/drop multiplexer (ROADM), optical amplifier, optical equalizer, and optical channel monitor on a single card.
 - IPLC expansion module—Increases the channel capacity of the IPLC node to 64 channels.

The standalone optical inline amplifier (ILA) provides periodic amplification of the optical line signal to enable long-distance transmission.

To complete the optical solution, you can use Juniper Networks 100G Coherent transponders, along with the IPLC, optical ILA, and Connectivity Services Director (CSD), which runs on the Junos Space Network Management platform to provide an end-to-end, fully managed packet optical solution.

You can configure, manage, and monitor the IPLC through Junos Space Connectivity Services Director 2.0, the Junos CLI, or your SNMP management system.

[See [PTX3000 Integrated Photonic Line Card User Guide.](#)]

- **Support for configuring and managing Juniper Networks optical inline amplifier (ILA) through Junos OS CLI**—Starting with Junos OS release 17.1R1, you can configure and manage certain capabilities of the optical inline amplifiers (ILA)s over the optical supervisory channel (OSC) of the PTX3000 integrated photonic line system, including authentication, performing resets, software upgrades, and performance monitors thresholds.

[See [Understanding Optical Supervisory Channel Communication in the Amplifier Chain.](#)]

Management

- **gRPC support for the Junos Telemetry Interface (PTX Series)**—Starting with Junos OS release 17.1R1, the Junos Telemetry Interface supports using a set of gRPC remote procedure call interfaces to provision sensors, subscribe to, and receive telemetry data. gRPC is based on an open source framework and provides secure and reliable transport of data. Use the **telemetrySubscribe** RPC to specify telemetry parameters and stream data for a specified list of OpenConfig commands paths. Telemetry data is generated as Google protocol buffers (gpb) messages in a universal key/value format. If your Juniper Networks device is running a version of Junos OS with an upgraded FreeBSD kernel, you must download the Network Agent package, which provides the interfaces to manage gRPC subscriptions. The package is available on the **All Junos Platforms** software download URL on the Juniper Networks webpage.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface.](#)]

- **Support for Junos Telemetry Interface (PTX Series)**—Starting in Junos OS Releases 17.1R1, you can use the Junos Telemetry Interface to export telemetry data from supported interface hardware. Sensor data, such as interface events, are sent directly to configured collection points without involving polling. FPC1, FPC2, and FPC3 are supported. For sensors that stream data through the User Datagram Protocol, all parameters are configured at the **[edit services analytics]** hierarchy level. For sensors that stream data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Not all hardware and sensors are supported in previous releases.

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

- **Support for adding non-native YANG modules to the Junos OS schema (PTX Series)**—Starting in Junos OS Release 17.1R1, you can load custom YANG models on devices running Junos OS to add data models that are not natively supported by Junos OS but can be supported by translation. Doing this enables you to extend the configuration hierarchies and operational commands with data models that are customized for your operations. The ability to add data models to a device is also beneficial when you want to create device-agnostic and vendor-neutral data models that enable the same configuration or RPC to be used on different devices from one or more vendors. You can load custom YANG modules by using the **request system yang add** operational command.

[See [Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS.](#)]

MPLS

- **Egress peer engineering of service labels (BGP, MPLS) and egress peer protection for BGP-LU (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, you can enable traffic engineering of service traffic, such as MPLS LSP traffic between autonomous systems (ASs), using BGP labeled unicast for optimum utilization of the advertised egress routes. You can specify one or more backup devices for the primary egress AS boundary router. Junos OS installs the backup path in addition to the primary path in the MPLS forwarding table, which enables MPLS fast reroute (FRR) when the primary link fails. It provides support for the FRR protection backup scheme to do an IP lookup to determine a new egress interface.

[See [Egress Peer Traffic Engineering Using BGP Labeled Unicast Overview.](#)]

- **Order-aware abstract hops for MPLS LSPs (PTX Series)**—Starting in Junos OS Release 17.1, support is provided for abstract hops, which are user-defined router clusters or groups that can be sequenced and used for setting up a label-switched path (LSP). They resemble real-hop constraints.

The router groups are created using constituent lists that include constituent attributes, which is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs). Ordering among the router groups that satisfy the specified constituent attributes is achieved by using operational qualifiers in the abstract-hop definition.

A path can use a combination of real and abstract hops as constraints. To configure abstract hops, you need to create constituent lists with traffic engineering attributes, include the lists in the abstract-hop definition, and define path constraints that use the abstract hops.

[See [Abstract Hops For MPLS LSPs Overview](#) and [Example: Configuring Abstract Hops for MPLS LSPs](#).]

Multicast

- **Support for next generation MVPN and Internet multicast (PTX5000 and PTX3000)**—Starting in Junos OS Release 17.1R1, the **mpls-internet-multicast** routing instance type uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP (or Next Gen) MVPN. Next generation MVPN is available only for PTX Series routers that have third-generation FPCs installed.

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

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

Network Management and Monitoring

- **Support for hrProcessorTable object (PTX Series)**—Starting in Junos OS Release 17.1R1, support is provided for the **hrProcessorTable** object (object id: 1.3.6.1.2.1.25.3.3) described in the RFC2790, *Host Resources MIB*. The **hrProcessorTable** object provides the load statistics information per CPU for multi-core devices.

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

- **Support for mplsL3VpnIfConfTable object (PTX Series)**— Starting in Junos OS Release 17.1R1, support is provided for the **mplsL3VpnIfConfTable** object (object id: 1.3.6.1.2.1.10.166.11.1.2.1) described in RFC 4382, *MPLS/BGP Layer 3 Virtual Private Network (VPN) MIB*. The **mplsL3VpnIfConfTable** object represents the Layer 3 VPN enabled interfaces that are associated with a specific Virtual Routing and Forwarding (VRF) instance and shows the bitmask values of the supported protocols. The **mplsL3VpnIfConfTable** object creates entries for the interfaces that are associated with the VRF instances. If an interface is later removed from a VRF instance, the corresponding entry in the **mplsL3VpnIfConfTable**

object gets deleted. To view details of the `mplsL3VpnIfConfTable` object, use the `show snmp mib walk mplsL3VpnIfConfTable` command.

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

Routing Policy and Firewall Filters

- **Optimized performance for DSCP and traffic-class firewall filter match conditions (PTX Series with third-generation FPCs)**—Starting in Junos OS Release 17.1R1, the `promote dscp` and `promote traffic-class` indicators are supported in firewall filters for IPv4 and IPv6 traffic. When either are applied to a filter, the entire filter is compiled in a way that optimizes its performance for the `dscp` or `traffic-class` match condition. The indicators are configured at the `[edit firewall family (inet | inet6) filter filter-name]` hierarchy level.

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

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

- **Support for firewall feature matching on gre-key (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1 on PTX3000 and PTX5000, the `promote gre-key` statement is supported to configure gre-key as one of the matches in a filter. When `promote gre-key` is configured and gre-key is used in any of the terms in a filter, the entire filter is compiled in a way that optimizes its performance for gre-key matching. The `promote gre-key` configuration statement is configured at the `[edit firewall family family-name filter filter-name]` hierarchy level.

[See [promote gre-key](#).]

- **Support for filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling (PTX Series with third-generation FPCs)**—Starting in Junos OS Release 17.1R1, filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling uses firewall filters to provide decapsulation of GRE traffic. The filter-based GRE decapsulation also supports routing-instance as an action.

NOTE: Configuring filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling is supported only when the **enhanced-mode** statement is configured at the `[edit chassis network-services]` hierarchy level.

- **Support for configuring the GTP-TEID field for GTP traffic (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, 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 hash-key family inet layer-4]` or `[edit forwarding-options hash-key family inet6 layer-4]` hierarchy level.

[See [gtp-tunnel-endpoint-identifier](#).]

Routing Protocols

- **Support for BGP to carry flow-specification routes (PTX5000 and PTX3000)**—Starting in Junos OS Release 17.1R1, BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX Series routers that have third-generation FPCs installed. Propagating firewall filter information as part of BGP enables you to propagate firewall filters against denial-of-service (DOS) attacks dynamically across autonomous systems.

[See [Example: Enabling BGP to Carry Flow-Specification Routes](#).]

- **Support for Bidirectional Forwarding Detection protocol intervals (PTX3000 and PTX5000)**—Starting in Junos OS Release 17.1R1, longer configuration ranges for Bidirectional Forwarding Detection (BFD) protocol intervals are supported on PTX Series routers that have third-generation FPCs installed.

NOTE: The longer configuration ranges are supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

Security

- **Support for Secure Boot (PTX3000)**—Starting in Junos OS Release 17.1R1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement Secure Boot.

Services Applications

- **Support for inline-jflow (PTX Series routers with third-generation FPCs)**—Starting in Junos OS Release 17.1R1, you can use inline-jflow's export capabilities with IP Flow Information Export (IPFIX) to define a flow record template suitable for IPv4 or IPv6 traffic on PTX Series routers that have third-generation FPCs installed.

[See [Monitoring Network Traffic Flow Using Inline Flow Monitoring on PTX Series Routers](#).]

User Interface and Configuration

- **Monitoring, detecting, and taking action on degraded physical 100-Gigabit Ethernet links to minimize packet loss (PTX3000 and PTX5000)**—Starting with Junos OS Release 17.1R1, you can monitor physical link degradation (indicated by bit error rate (BER) threshold levels) on Ethernet interfaces, and take corrective actions if the BER threshold value drops to a value in the range of 10^{-13} to 10^{-5} .

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

- To monitor physical link degrade on Ethernet interfaces, configure the **link-degrade-monitor** statement.
- To configure the BER threshold value at which the corrective action must be triggered on or cleared from an interface, use the **link-degrade-monitor thresholds (set value | clear value)** statement.
- To configure the link degrade interval value, use the **link-degrade-monitor thresholds interval value** statement. The configured interval value determines the number of consecutive link degrade events that are considered before any corrective action is taken.
- To configure link degrade warning thresholds, use the **link-degrade-monitor thresholds (warning-set value | warning-clear value)** statement.
- To configure the link degrade action that is taken when the configured BER threshold level is reached, use the **link-degrade action media-based** statement.
- To configure the link degrade recovery options, use the **link-degrade recovery (auto interval value | manual)** statement.

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

SEE ALSO

[Changes in Behavior and Syntax | 252](#)

[Known Behavior | 257](#)

[Known Issues | 259](#)

[Resolved Issues | 262](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

[Product Compatibility | 274](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [General Routing | 253](#)
- [Interfaces and Chassis | 253](#)
- [Management | 254](#)
- [MPLS | 254](#)

- Network Management and Monitoring | 255
- Routing Protocols | 256
- Services Applications | 256
- System Management | 256
- User Interface and Configuration | 256

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 17.1R3 for the PTX Series.

General Routing

- **ECMP next hops supported for IS-IS IPv6**—When maximum-ecmp 64 option is enabled and if an IS-IS route has multiple next hops or if it is above the maximum limit, then the rpd crashes because the next hop gateway addresses are overwritten and stored in a circular buffer. Note: In the worst case (if all the next hops are IPv6), only 38 ECMP next hops are fully supported for IS-IS IPv6 instead of 64.
- **Support for deletion of static routes when the BFD session goes down (PTX Series)**—Starting with Junos OS Release 17.1R3, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session-down message.

[See [Enabling BFD on Qualified Next Hops in Static Routes for Route Selection.](#)]

Interfaces and Chassis

- **Message now displayed when SIB autohealing is complete (PTX3000 and PTX5000)**—In Junos OS Release 17.1R1 and later, the output of **show chassis fabric errors autoheal** displays a message when SIB autohealing is complete, as shown in the following example:

```
user@host> show chassis fabric errors autoheal
Mar 30 01:43:00
Time                               Error log of first 100 errors
2016-03-29 23:46:23 PDT             Req: sib 0
2016-03-29 23:46:23 PDT             Action: SIB 0 (autohealing)
2016-03-29 23:54:52 PDT             Completed: SIB 0 (autoheal)
```

Management

- **Enhancement to Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.1R1, Junos Telemetry Interface data streamed through gRPC no longer includes the phrase **oc-path** in the prefix field. For example, a physical interface sensor streaming data for interface et-0/0/0:0 now displays the following output: `str_value:/interfaces/interface[name='et-0/0/0:0']/`.
- **Enhancement to NPU memory sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.1R2, the path used to subscribe to telemetry data for network processing unit (NPU) memory and NPU memory utilization through gRPC has changed. The new path is `/components/component[name="FPC<fpc-id>:NPU<npu-id>"]/`

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

- **Enhancement to NPU memory sensors for Junos Telemetry Interface (PTX Series)**—Starting with Junos OS Release 17.1R3, the format of telemetry data exported through gRPC for NPU memory and memory utilization implements prefix compression. This change reduces the payload size of data exported. The following example shows the new format:

key: `__prefix__`

str_value: `/components/component[name='FPC0:NPU0']/properties/property`

key: `[name='mem-util-edmem-size']/value`

uint_value: `12345`

Telemetry data is exported in key-value pairs. In releases before Junos OS Release 17.1R3, the data exported includes the component and property names in a single key string.

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

MPLS

- **Representation for OSPF DR node**—Up until version -10 of the BGP-LS draft, the OSPF DR node representation was ambiguous. One could represent DR node as 'AdvertisingRouterId-InterfacelpAddress' or 'InterfacelpAddress-1'. JUNOS used to follow 'InterfacelpAddress-1' format. Starting with version '-11' of the BGP-LS draft, the representation for OSPF DR node must be 'AdvertisingRouterId-Interfacelpaddress'. Junos OS now follows the latest format.
- **PPPoE subscribers do not bind over ps interfaces (MX Series)**—Starting with Junos OS Release 17.1R1, the termination of single, multiple, and dual-tagged service delimited VLANs are transported over a single Ethernet CCC pseudowire using ps virtual port devices. This feature provides scaled Layer 3 service application at the pseudowire head-end termination appliance. This behavior is as an extension and evolution for Ethernet pseudowire that is described in RFC 4448.

Network Management and Monitoring

- **SNMP syslog messages changed (PTX Series)**—Starting in Junos OS Release 17.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - Old message—**AgentX master agent failed to respond to ping. Attempting to re-register**
New message—**AgentX master agent failed to respond to ping, triggering cleanup!**
 - Old message— **NET-SNMP version %s AgentX subagent connected**
New message— **NET-SNMP version %s AgentX subagent Open-Sent!**

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

- **Update to SNMP support of apply-path statement (PTX Series)**—Starting in Junos OS Release 17.1R2, SNMP implementation for the **apply-path** configuration statement supports only two lists:
 - **apply-path "policy-options prefix-list <list-name> <*>"**
This configuration has been supported from day 1.
 - **apply-path "access radius-server <*>"**
This configuration is supported as of Junos OS Release 17.1R2.
- **MIB loading errors fixed (PTX Series)**—Starting in Junos OS Release 17.1R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:
 - jnx-gen-set.mib
 - jnx-ifotn.mib
 - jnx-optics.mib

[See [MIB Explorer](#).]

- **Change in default log level setting (PTX Series)**—Starting in Junos OS Release, 17.1R3, 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 that come from a nondefault routing instance and a nondefault logical system (PTX Series)**—In Junos OS Release

17.1, a new option, **context-oid**, for the **trap-options** statement enables you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional variable binding, or varbind.

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

Routing Protocols

- **Change in default behavior of router capability (PTX Series)**—In Junos OS Release 15.1F7, 16.1R4, 16.1X65, and 17.1R1 and later, the router capability TLV distribution flag (S-bit), which controls IS-IS advertisements, will be reset, so that the segment-routing-capable sub-TLV is propagated throughout the IS-IS level and not advertised across IS-IS level boundaries.

Services Applications

- **Device discovery with device-initiated connection (PTX Series)**—In Junos OS Release 17.1R1 and later, when you configure statements and options under the [**system services ssh**] hierarchy and commit the configuration, make sure that the system reaches a stable state before you commit any **outbound-ssh** configurations.

You use the device discovery feature in the Devices workspace to add devices to Junos Space Network Management Platform. By default, Junos Space manages devices by initiating and maintaining a connection to the device.

[See [Device Discovery Overview](#).]

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the **peers** option at the [**edit system commit**] hierarchy level is not supported in batch configuration mode.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (PTX Series)**—Starting in Junos OS Release 17.1R1, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 17.1R1, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the show system schema module juniper-command output directory (PTX Series)**—Starting in Junos OS Release 17.1, when you issue the **show system schema module juniper-command** operational command in the Junos OS CLI, the device places the generated output files in the current working directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1, the generated output files are placed in the **/var/tmp** directory.

SEE ALSO

[New and Changed Features | 242](#)[Known Behavior | 257](#)[Known Issues | 259](#)[Resolved Issues | 262](#)[Documentation Updates | 270](#)[Migration, Upgrade, and Downgrade Instructions | 270](#)[Product Compatibility | 274](#)

Known Behavior

IN THIS SECTION

- [General Routing | 258](#)
- [Interfaces and Chassis | 258](#)
- [High Availability \(HA\) and Resiliency | 258](#)

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 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

- The na-grpcd process generates core files after subscribing to gRPC data for the interfaces sensor for queue statistics in the gRPC stack. This issue occurs rarely with the gRPC beta stack. The process should work fine starting in Junos OS Release 17.2 because the gRPC stack has been upgraded to 1.0.0.[PR1255206](#)

Interfaces and Chassis

- 1. Delay Measurement support for 5-port 100G DWDM PIC and 5-port 100G DWDM MIC is **ONE TIME Delay Measurement**. This means that customers intending to measure Delay 2 points should ensure that the link is up on both sides and then conduct this test one time. The result value is valid one time once the test is finished. The test result on CLI is not valid after one time measurement because the old result might show up on the Routing Engine CLI. 2. **Remote-loop-enable** should be configured first on the remote end. 3. Each time a customer wants to verify this, the test has to be **repeated**. 4. Processing delays in each mode are different: HGFEC [For 5-port 100G DWDM MIC] being the highest, SDFEC in the interim and GFEC being the lowest for the same cable length. 5. In summary, any breakage in Transmit/Receive path during the Delay Measurement test will hinder delay measurement. This is true for all FEC modes - GFEC, SDFEC, HGFEC. 6. Currently SNMP walk is not available for Delay Measurement. [PR1233917](#)

High Availability (HA) and Resiliency

- **Unified ISSU incompatibility with VPLS dynamic profiles (PTX Series)**—Using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 17.1R1 does not work if VPLS dynamic profiles are configured and enhanced subscriber management is not configured.
[See [ISSU System Requirements](#).]

SEE ALSO

[New and Changed Features | 242](#)

[Changes in Behavior and Syntax | 252](#)

[Known Issues | 259](#)

[Resolved Issues | 262](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

[Product Compatibility | 274](#)

Known Issues

IN THIS SECTION

- General Routing | 259
- Interfaces and Chassis | 261
- MPLS | 261
- Platform and Infrastructure | 261
- Routing Protocols | 261
- User Interface and Configuration | 261

This section lists the known issues in hardware and software in Junos OS Release 17.1R3 for the PTX Series.

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

General Routing

- On PTX3000 and PTX5000, the MIB jnxDomMib does not return the proper result from the SNMP query. [PR1045804](#)
- The PTX 100GbE-LR4 interfaces might flap when the reference clock switches over from being the line clock to operating as the holdover. This switchover and the ensuing interface flapping is initiated when the PIC on which the line clock sources reside is taken offline. When the router uses line clock sources and when it does not have any external clocks from BITS-a or BITS-b, taking the PIC offline brings the line clock down and the reference clock is switched from line clock to holdover. This reference clock transition might cause a large clock phase-shift in the 100GbE-LR4 CFP modules, and this phase-shift might cause distortion in the output optical pulse waveform on the associated interfaces. This distortion results in interface flap. This issue cannot be fixed by software because of a hardware limitation. [PR1130403](#)
- 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, the chassisd might crash and the upgrade or downgrade might fail. This issue is not seen if both base and target images are from Junos OS Release 15.1F or from Junos OS Release 16.x and later. [PR1171652](#)
- On PTX Series with third-generation FPC line cards, in rare scenarios, while restarting the FPC, a PIC index mismatch issue might lead the FPC to crash if it is configured with **inline-JFlow**. [PR1183215](#)

- If more than 10 FO CRC errors are seen in an interval of 30 seconds, then CMERROR infra raises an alarm and appropriate action needs to be taken. [PR1197865](#)
- Power budget values for a PTX5000 chassis, FPC, and PICs have been revised. For routers operating on limited power, this can change the point where alarms for power-over-budget or insufficient power are raised or cleared. [PR1216404](#)
- On PTX1000 and QFX10002, some random ports that use 100G Lumentum optics might not come up after a reboot. This is a timing issue because of failures during optics read on some ports. As a workaround, remove and insert the optics, which might bring up the ports. [PR1227029](#)
- Sometimes l2cpd core files are generated when LLDP neighbors are cleared. [PR1270180](#)
- A memory leak in the Packet Forwarding Engine's clone route storage could lead to memory exhaustion for MPLS routes. [PR1289311](#)
- On PTX5000/PTX3000/PTX10016/PTX10008/PTX1000 line cards, if there is an error due to HMC (Hybrid Memory Cube) internal error conditions, Junos OS software resiliency performs the default **disable-pfe** action and shuts down all WAN interfaces. However, the system continues to install active routes and their next hops into the faulty HMC memory. The Junos OS software resiliency has been enhanced to skip route out next-hop programming and queue statistics collections when the **disable-pfe** action is enabled to prevent those symptoms mentioned above. This does not prevent the HMC fatal error, but only reduces the operational impact upon such failure. [PR1300180](#)
- Interfaces might go down when the Packet Forwarding Engine encounters **TOE::FATAL ERROR** TOE is a module in Packet Forwarding Engine. The fatal error can be caused by either software issues or hardware issues such as memory parity errors). You can reboot the line card to recover the service when the issue arises. [PR1300716](#)
- The iLatency value (calculated by differing producer timestamp and gRPC server timestamp) can sometimes be negative for telemetry packets related to Packet Forwarding Engines because of a drift in the Routing Engine and Packet Forwarding Engine NTP servers. [PR1303376](#)
- Disabling of et-0/0/5:2 also disables et-0/0/5:0. This issue occurs because of incorrect QSFP28 optics channel and TQ-chip retimer lane mapping. [PR1337975](#)

Interfaces and Chassis

- Junos OS upgrade involving Junos OS Release 14.2R5 (and subsequent Junos OS Release 14.2 maintenance releases) and Junos OS Release 16.1R2 (and subsequent maintenance releases) with CFM configuration can cause the cfmd to crash after the upgrade. This is because of the presence of an old version of `/var/db/cfm.db`. [PR1281073](#)

MPLS

- On PTX Series platforms, the rpd might crash when the RSVP bypass undergoes reoptimization and the reoptimized instance encounters failure before it becomes the main instance. [PR1250253](#)

Platform and Infrastructure

- On PTX Series platforms with **chassis network-services enhanced-mode** configured, the default policy **junos-ptx-series-default** is not loaded correctly during some configuration operations, which causes the BGP routes not to be installed in the forwarding table as expected. As a workaround, reboot the router after any configuration operation on the network services. [PR1204827](#)

Routing Protocols

- A few Bidirectional Forwarding Detection (BFD) protocol sessions flap while coming up after the FPC restarts or reboots. This does not impact the system, because the flap is seen during the bring-up phase. This occurs because of a race condition in the PPMAN code. [PR1274941](#)

User Interface and Configuration

- When **persist-groups-inheritance** is configured, doing configuration changes and issuing rollback might cause persist-groups tree corruption and eventually cause improper configuration propagation after commit. As a result, the mgd process might crash as well. [PR1214743](#)

SEE ALSO

[New and Changed Features | 242](#)

[Changes in Behavior and Syntax | 252](#)

[Known Behavior | 257](#)

[Resolved Issues | 262](#)

[Documentation Updates | 270](#)

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.1R3 | [262](#)
- Resolved Issues: 17.1R2 | [267](#)
- Resolved Issues: 17.1R1 | [269](#)

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

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

Resolved Issues: 17.1R3

General Routing

- The **request vmhost zeroize** and **request vmhost zeroize both** commands might work only on the local Routing Engine. [PR1197152](#)
- The rpd process might crash after a configuration is committed. [PR1200174](#)
- On PTX Series routers, the chassisd thread might not get CPU resources for 200 seconds generating multiple chassisd core files. [PR1226992](#)
- The SNMP MIB walk is not working for optical MIB jnxDomCurrentRxLaserPower. [PR1249262](#)
- An FPC major alarm might be seen with the following error message: **DLU: ilp memory cache error** and **DLU: ilp prot1 detected_imem_even error**. [PR1251154](#)
- The **validation-state:unverified** routing entry might not be displayed with the proper location in the output of the **show route** command. [PR1254675](#)
- The PTX5000 third-generation FPCs (FPC3-PTX-U3 and FPC3-PTX-U2) contain new hardware components that require a Junos OS upgrade. [PR1258693](#)
- SPMB ukern panic occurs during ASIC error recovery. [PR1268253](#)

- Error messages might be seen on a PTX5000 router with third-generation FPCs with P2-10G-40G-QSFPP(10-Gigabit Ethernet/40-Gigabit Ethernet LAN/WAN OTN PIC with QSFP+) PIC inserted. [PR1273575](#)
- 100GBase-ER4 (740-045420) is shown as UNKNOWN when the **show chassis hardware** command is executed in Junos OS Release 15.1R5. [PR1280089](#)
- FPCs might go offline because of fabric healing on a PTX3000 with the SIB-SFF-PTX-240-S SIB. [PR1282983](#)
- The PTX SPMB might crash after the FPC replacement followed by a SIB restart. [PR1283553](#)
- The MPLS TTL might be reset to 255 on third-generation PTX Series FPCs if the **protocols mpls no-propagate-ttl** statement is configured. [PR1287473](#)
- A memory leak in the Packet Forwarding Engine's clone route storage might lead to memory exhaustion for MPLS routes. [PR1289311](#)
- LSP traffic might silently drop and get discarded after a link goes down in the bypass path. [PR1291036](#)
- The rpd process might generate a core file while restarting. [PR1291110](#)
- Incorrect SNMP OID values are sent in SNMP traps for removal or insertion of front panel display on PTX Series routers. [PR1294741](#)
- The **LINK** LED is red when the port is disabled on PTX Series routers. [PR1294871](#)
- When interface flaps or BGP session flaps, the system will receive a response from the kernel for a generic next-hop object. After the response is received for the current child object, all parent objects are walked. In some situations, the parent object might also be of type generic next hop and might also be waiting for a kernel response. If both the parent and the child object are being programmed to the kernel simultaneously, the rpd might crash. [PR1294957](#)
- Alarms and syslog errors are seen with priority strict-high on an AF4 queue, in oversubscription cases (1X100G egress to 1X10G egress setup). [PR1297343](#)
- "Link errors" alarm messages might be seen after migrating to FPC3 on PTX3000 routers. [PR1298841](#)
- PTX Series FPC3 drops MPLS packets when the maximum transmission unit is less than the MPLS packet size on the outgoing interface with IPv4 traffic. [PR1302256](#)
- Heap memory leak might be observed on PTX Series FPCs during a multicast route installation into the Packet Forwarding Engine. [PR1302303](#)
- On PTX3000 platforms, the powering on a FPC (PTX-IPLC-B-32) line card might cause the other FPC line cards to reboot. [PR1302304](#)
- The third-generation FPC (FPC3-SFF-PTX) might not boot on a PTX3000 with a Control Board or Routing Engine. [PR1303295](#)
- On PTX3000 and PTX5000 platforms, the 100-Gigabit Ethernet interfaces might not come up. [PR1303324](#)
- If MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic with an error message: **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)

- Repeated log messages `%PFE-3 fpcX expr_nh_index_tree_ifl_get` and `expr_nh_index_tree_ipaddr_get` are observed when the sampling packet is discarded with `log(or syslog)` configuration statement under the firewall filter. [PR1304022](#)
- PTX3000 with RCB-PTX Routing Engine might not come online or recognize the integrated photonic line cards (IPLCs). [PR1304124](#)
- On PTX Series devices with EVPN-VXLAN setup, some error messages might be observed. They have no service impact. [PR1307014](#)
- The FPC1/2/3 line card on PTX Series router might crash due to a timing issue. [PR1307067](#)
- The "interface hold-time down" timer does not take effect on PTX5000 routers with optical interface. [PR1307302](#)
- The rpd might crash and generate core files after multiple session flaps on a scale setup. [PR1312169](#)
- When a next-generation Routing Engine is used on the router, a log is generated after executing a VHclient (Vmhost Client) related command (for example, `show vmhost version/uptime`). These commands are generated at "info" level, so they are generated by default. Too many logs are generated after executing many VHclient related commands. These logs are harmless but might fill up the log file. [PR1315128](#)
- The RIB and FIB might get out of synchronization because the KRT asynchronous queue might get stuck. [PR1315212](#)
- When the interface on PTX3000/PTX5000 third-generation Flexible PIC Concentrator (FPC) or PTX10000 line card goes down (for example, FPC on remote end shutdown or restart), the packets already in the application-specific integrated circuit (ASIC) pipeline need to be drained. If the packets are not drained fast enough, the Packet Forwarding Engine connecting to the interface might be disabled with the error of **stream drain problem**. [PR1315823](#)
- The line card might crash upon receipt of a specific MPLS packet. The affected line cards include MPC7E, MPC8E, and MPC9E on MX Series platform, the third-generation FPC on PTX3000 (FPC3-SFF-PTX), the third-generation FPC on PTX5000 (FPC3-PTX-U2 that is, FPC-P1; FPC3-PTX-U3 that is, FPC-P2) and the built-in FPC on PTX1000. Refer to <https://kb.juniper.net/JSA10864> for more details. [PR1323069](#)
- The MPLS traceroute fails across the PTX Series platform. [PR1327609](#)
- Link instability is observed after a link-down event on PTX Series routers. [PR1330708](#)
- In an Equal-Cost Multipath (ECMP) scenario on PTX Series platforms, the unilist next-hop might be incorrectly programmed on the Packet Forwarding Engine when a member link flaps, resulting in traffic getting silently dropped and discarded. [PR1333274](#)
- On a PTX Series routers, the FPC might reboot in certain rare scenarios when a flapping interface configuration is committed. [PR1335161](#)
- A member of IPv4 unilist next hops might get stuck in "Replaced" state after the interface flaps. [PR1336201](#)

- On a PTX Series platform system with P3-10-U-QSFP28 PIC, because of the incorrect channel and lane mapping, disabling one breakout 10G port on et-0/0/5 also disables another breakout 10G port on et-0/0/5 (for example, disabling et-0/0/5:2 also disables et-0/0/5:0). [PR1337975](#)
- On PTX Series platform with FPC/FPC2/FPC E, if multirate PIC is used, the FPCs might not forward the traffic. [PR1339524](#)
- When PTX5000 and PTX3000 routers with P3-24-U-QSFP28 PIC installed are rebooted, the interface might flap continuously. [PR1342681](#)
- On PTX Series platforms with FPC3 cards, if point-to-multipoint (P2MP) MPLS LSP is configured with link-protection, then traceroute of MPLS LSP might cause the FPC to crash. [PR1348314](#)
- This issue applies to PTX3000 FPC-SFF-PTX-P1-A/FPC-SFF-PTX-T. When BFD is configured for BGP, any Layer 3 packet injects from the line card might result in BFD sessions not coming up on the PTX3000. [PR1352112](#)
- This issue occurs in a BGP multipath scenario in which a route installation failure is not handled properly. Traffic loss might occur. [PR1362560](#)
- On PTX Series routers with third-generation FPCs, if optics that are not certified by Juniper Networks are used and there is a specific traffic pattern with congestion, traffic might be dropped. [PR1378392](#)
- On PTX Series platforms, the Layer 3 VPN traffic might be dropped if one core-facing interface goes down in the Layer 3 VPN multipath scenario. [PR1380783](#)
- A single Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in an error or wedge condition. [PR1400716](#)

Infrastructure

- The **show system users** command displays users who are not using the router. [PR1247546](#)
- The PTX Series router might get to abnormal state because the protection mechanism for F-Label malfunctions. [PR1336207](#)

Interfaces and Chassis

- Interface flap occurs while executing a Routing Engine switchover, if the member links of an aggregated Ethernet interface are configured with framing settings. [PR1287547](#)
- The 100-Gigabit Ethernet interfaces might not come up when **otn-options laser-enable** is configured on PTX Series platforms. [PR1297164](#)
- The transportd process might crash when an SNMP get query on jnxoptIfOChSinkCurrentExtTable with an unsupported interface index is executed. [PR1335438](#)
- PTX Series routers with FPC3 that encounter oversize_drop errors will trigger an FPC major alarm with each error seen. The alarm should clear on its own. The severity of this alarm is being reduced because there is no lasting impact to the FPC. [PR1375030](#)

Layer 2 Ethernet Services

- On PTX Series platforms, the following message is filling up the syslog: **l2cpd[2486]: task_connect: task MVRP l2ald ipc./var/run/l2ald_control addr /var/run/l2ald_control: No such file or directory.** [PR1278189](#)

MPLS

- When RSVP is configured, RSVP neighbors are present, and the system is under very high CPU load conditions, the rpd process might crash in rare cases. [PR1138190](#)
- In an RSVP environment, a stale LSP might get created after a Routing Engine switchover with nonstop routing (NSR) enabled. [PR1292526](#)
- The rpd process might crash when the MPLS LSP path change occurs. [PR1295817](#)
- The rpd process might crash on the backup Routing Engine due to memory exhaustion. [PR1328974](#)
- The rpd process might crash while tracing LSP events when MPLS traceoption is configured. [PR1329459](#)
- When using an LSP to forward traffic, the statistics are not displayed in the command **show mpls lsp ingress statistics** output, whereas the interface displays the traffic sent out properly. This behavior can be seen when you have the logical system on the same router used as a provider where the kernel will be in synchronization with the Self ID allocation between master and logical system to display the statistics properly. The CLI command **show mpls lsp ingress statistics lose** MPLS LSP statistics in the output. [PR1344039](#)
- If the LDP route with indirect next-hop exists (for example, the LDP egress-policy is used to advertise BGP route into LDP), the rpd might crash when the LDP route is deleted. [PR1398876](#)

Platform and Infrastructure

- Continuous log messages occur. For example, **tftpd[23724]: Timeout #35593 on DATA block 85.** [PR1315682](#)
- In a large-scale setup, even when rpm/twamp is not configured, in case of network churn and lots of interface flaps, a traffic black hole might be observed. [PR1357707](#)
- A vulnerability in the IP next-hop index database in Junos OS Release 17.3R3 might allow a flood of ARP requests, sent to the management interface, to exhaust the private internal routing interfaces (IRIs) next-hop limit. Once the IRI next-hop database is full, no further next hops can be learned and existing entries cannot be cleared, leading to a sustained denial of service (DoS) condition. [PR1360039](#)

Routing Protocols

- BGP-LU does not react to family MPLS up/down and invalidate BGP-LU routes received with label-operations. In this situation, BGP-LU label might go into "dead" state in the forwarding table after the MPLS address family on the next-hop interface is removed and re-added. [PR1262180](#)
- The rpd process crashes and generates core files multiple times when you receive an open message from an existing BGP peer. [PR1299054](#)
- On all platforms with BGP configured, if BGP routes are queued to send to any peer (this could be caused by simply having network churn), the rpd might constantly consume high CPU (98 percent). [PR1315066](#)

- The primary path of an MPLS LSP might switch to another address. [PR1316861](#)
- The rpd process might crash after the passive interface under IS-IS is deactivated. [PR1318180](#)
- The rpd process might crash if SRLG information is in the protocol IS-IS. [PR1337849](#)
- The rpd process might generate a core file on the backup Routing Engine during neighborship flap when using an authentication key with more than 20 characters. When using "security authentication-key-chains key-chain <*> key <*> secret <*> " with hmac-sha-1 algorithm and the secret key length is more than 20 characters, causes memory corruption in the rpd, and later the rpd process crashes on backup Routing Engine. [PR1394082](#)

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](#)

Resolved Issues: 17.1R2

General Routing

- The routers equipped with NG-REs might raise memory size mismatch alarm after upgrade. [PR1220061](#)
- User-configured TPID is not applying on a single-tagged VLAN interface. [PR1237687](#)
- Junos Telemetry Interface: Frequent disconnects seen in MQTT when IFL sensor is provisioned for longer duration. [PR1238803](#)
- Tx rate not guaranteed for an extreme case scheduler with traffic from multiple ingress Packet Forwarding Engines to a single egress Packet Forwarding Engine. [PR1241291](#)
- Add **set** parameter to CLI **request system software add** command. [PR1246675](#)
- "telemetry_start_polling_fd: evSelectFD failed, errno: 9" are continuously seen in log. [PR1248813](#)
- cs605x_otu_defect_active: cs605x_get_otu_alarms failed. Messages are logged when only the first two ports are not configured on 4x100G OTN PIC. [PR1250707](#)
- While processing lookup results, IRP block raises an interrupt upon detecting an error condition. The interrupt is active until the trap code error is read. Under certain conditions, software is not reading this trap code error upon IRP interrupt. This causes the following syslog messages:

```
fpc5 INTR: throttle 60sec PECHIP[2]:pe.irp.intr.status:ap0_trap(0): (Count:1)
fpc5 INTR: throttle 3630sec PECHIP[2]:pe.irp.intr.status:ap0_trap(0): (Count:3434)
fpc5 INTR: throttle 3600sec PECHIP[2]:pe.irp.intr.status:ap0_trap(0): (Count:6841)
```

[PR1256736](#)
- The following log message may be printed frequently during normal operation on backup Routing Engine of PTX5000 router when the Routing Engine type is RE-DUO-C2600. The message is cosmetic and does not indicate any service impact or Routing Engine mastership loss:

/kernel: mastership: sent other Routing Engine mastership loss signal

[PR1260884](#)

- On a PTX Series platform, if running interfaces on QSFP28 PIC in 10G mode, some of the interfaces on the QSFP28 PIC may not come up after a system reboot or PIC restart. [PR1263413](#)
- This only affects only PTX5000 and PTX3000 platforms with Third-generation FPCs. Software periodically monitors voltages on the FPCs to check if they are within the proper range. This change adjusts the expected values for voltages on certain power rails of the FPCs. In rare cases it is possible that a marginal FPC was operating inside the older limits but outside the new limits, in which case a new chassis alarm will be raised for that FPC. [PR1263675](#)
- In PTX routers equipped with Next Generation Routing Engine (RE-S-X6-64G, REMX2K-X8-64G, RE-PTX-X8-64G/CB2-PTX), the following log messages might be displayed as an error messages after a **commit** command is executed:

```
sdk-vmmd: %USER-3: is_platform_rainier: Platform found as rainier
```

[PR1271134](#)

Interfaces and Chassis

- Configuring ODU FRR related otn-options might crash the FPC without producing a core file. [PR1038551](#)
- 5-port 100G DWDM PIC: Unsupported CFP is initialized even when part number is not valid. [PR1174080](#)

MPLS

- The RPD might crash while making static LSPs up. [PR1084736](#)
- RPD crash in MPLS OAM environment. [PR1233042](#)
- The LDP routes are not installing with matched L-ISIS routes in inet.3 route table. [PR1248336](#)
- RPD core when rpd is terminating while there are a large number of RSVP LSPs. [PR1257367](#)

Platform and Infrastructure

- The "rdd" process restarted in get_mview_root() during GRPC JVISION activation while chassis PFEs are coming up. [PR1225086](#)
- mgd might crash after executing command **show ephemeral-configuration | display inheritance**. [PR1258823](#)

Resolved Issues: 17.1R1

Class of Service (CoS)

- The error message `cos_check_temporal_buffer_status` might be observed when configuring hierarchical CoS with strict-high scheduling. [PR1238719](#)

General Routing

- False login attempts might be seen on MPC7E/8E/9E for receiving noise. [PR1234712](#)
- NGRE: Routing Engine switchover resulting when trying to connect to FPC through `cty`. [PR1235761](#)
- PTX Series router might send wrong packets if MPLS LSPs have protection configured. [PR1239634](#)
- 'oc-path' to be removed from prefix for IFD sensor (both FreeBSD 10.x-based Junos OS and FreeBSD 6.1-based Junos OS). [PR1244658](#)

Infrastructure

- Continuous kernel logs and LDP stats timeout error when executing `show ldp traffic-statistics`. [PR1215452](#)

Interfaces and Chassis

- 5-port 100G DWDM PIC: `chassisd` logs are flooded with power related messages. [PR1184415](#)
- ARP entry learned through Aggregated Ethernet interface does not expire when the ARP IP is no longer reachable. [PR1211757](#)

Routing Protocols

- The `rpd` process might crash after performing BGP flapping. [PR1222554](#)
- An `rpd` core file could be seen if MPLS goes down. [PR1228388](#)
- Kernel crashes in the chassis after FPC reset. [PR1242362](#)

SEE ALSO

[New and Changed Features | 242](#)

[Changes in Behavior and Syntax | 252](#)

[Known Behavior | 257](#)

[Known Issues | 259](#)

[Documentation Updates | 270](#)

[Migration, Upgrade, and Downgrade Instructions | 270](#)

[Product Compatibility | 274](#)

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R3 documentation for PTX Series.

SEE ALSO

New and Changed Features 242
Changes in Behavior and Syntax 252
Known Behavior 257
Known Issues 259
Resolved Issues 262
Migration, Upgrade, and Downgrade Instructions 270
Product Compatibility 274

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 17.1 | 270](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 273](#)
- [Upgrading a Router with Redundant Routing Engines | 274](#)

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

Basic Procedure for Upgrading to Release 17.1

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 [Installation and Upgrade Guide](#).

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 17.1R3:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.

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

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

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-17.1R3.9.tgz
```

Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-17.1R3.9-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

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

NOTE: After you install a Junos OS Release 17.1 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 15.1, 16.1 and 16.2 are EEOL releases. You can upgrade from Junos OS Release 15.1 to Release 16.1 or even from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

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

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

SEE ALSO

New and Changed Features 242
Changes in Behavior and Syntax 252
Known Behavior 257
Known Issues 259
Resolved Issues 262
Documentation Updates 270
Product Compatibility 274

Product Compatibility

IN THIS SECTION

- [Hardware Compatibility | 275](#)

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features	 242
Changes in Behavior and Syntax	 252
Known Behavior	 257
Known Issues	 259
Resolved Issues	 262
Documentation Updates	 270
Migration, Upgrade, and Downgrade Instructions	 270

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- [New and Changed Features](#) | [276](#)
- [Changes in Behavior and Syntax](#) | [296](#)
- [Known Behavior](#) | [302](#)
- [Known Issues](#) | [304](#)
- [Resolved Issues](#) | [309](#)
- [Documentation Updates](#) | [326](#)

- Migration, Upgrade, and Downgrade Instructions | 326
- Product Compatibility | 339

These release notes accompany Junos OS Release 17.1R3 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 17.1R3 New and Changed Features | 277
- Release 17.1R2 New and Changed Features | 277
- Release 17.1R1 New and Changed Features | 277

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

NOTE: The following QFX Series platforms are supported in Release 17.1R3: QFX5100, QFX10002, QFX10008, and QFX10016.

Release 17.1R3 New and Changed Features

MPLS

- **Order-aware abstract hops for MPLS LSPs (QFX Series)**—Junos OS Release 17.1 introduces abstract hops, which are user-defined router clusters or groups that can be sequenced and used for setting up a label-switched path (LSP), similar to real-hop constraints.

The router groups are created using constituent lists that include constituent attributes, which are a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs). Ordering among the router groups that satisfy the specified constituent attributes is achieved by using operational qualifiers in the abstract-hop definition.

A path can use a combination of real and abstract hops as constraints. To configure abstract hops, you need to create constituent lists with traffic engineering attributes, include the lists in the abstract-hop definition, and define path constraints that use the abstract hops.

[See [Abstract Hops For MPLS LSPs Overview](#) and [Example: Configuring Abstract Hops for MPLS LSPs](#).]

Restoration Procedures Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (QFX Series)**—Starting in Junos OS Release 17.1R3, 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 the 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](#).]

Release 17.1R2 New and Changed Features

- There are no new features or enhancements to existing features for QFX Series in Junos OS Release 17.1R2.

Release 17.1R1 New and Changed Features

Hardware

- **QFX10008 switch**—Starting with Junos OS Release 17.1R1, the Juniper Networks QFX10000 line of Ethernet switches provides cloud builders and data center operators scalable solutions for both core and spine data center deployments. The QFX10008 switch is an 8-slot, 13 U chassis that supports up to eight line cards. This switch was previously supported in an “X” release of Junos OS.

[See [QFX10008 Switch Hardware Guide](#).]

- **QFX10016 switch**—Starting with Junos OS Release 17.1R1, the Juniper Networks QFX10016 modular data center spine and core Ethernet switch provides cloud and data center operators with high-level scale and throughput. The largest of the QFX10000 line of switches, the QFX10016 can provide 96 Tbps of throughput and 32 Bpps of forwarding capacity in a 21 rack unit (21 U) chassis. The QFX10016 has 16 slots for line cards that allow for a smooth transition from 10-Gigabit Ethernet and 40-Gigabit Ethernet networks to 100-Gigabit Ethernet high-performance networks. This switch was previously supported in an “X” release of Junos OS.

[See [QFX10016 Switch Hardware Guide](#).]

- **QFX10000-60S-6Q line card (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.1R1, the QFX10000-60S-6Q line card provides 60 SFP+ ports and six flexible configuration ports for 100Gbps and 40Gbps. Note that as of Release 17.1R1, the SFP+ ports do not support 1-Gbps.

Of the six flexible configuration ports, two ports have QSFP28 sockets that support either 100-Gbps or 40-Gbps speeds. The remaining four ports have QSFP+ sockets that can be configured as either a native 40-Gbps port or four 10-Gbps ports using a breakout cable. With breakout cables, the line card supports a maximum of 84 logical 10-GbE ports.

[See [QFX10000-60S-6Q Line Card](#).]

Class of Service (CoS)

- **Support for class-of-service-based forwarding (QFX 10000 Series)**—CoS-based forwarding (CBF) enables the control of next-hop selection based on a packet's class-of-service field. Starting with Junos OS Release 17.1R1, QFX 10000 Series switches support CBF. You can implement CBF by creating a **next-hop-map** at the **[edit class-of-service forwarding-policy]** hierarchy level and then applying the **next-hop-map** to a **policy-statement** at the **[edit policy-options]** hierarchy level. CBF can only be configured on a device with eight or fewer forwarding classes plus a default forwarding class.

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

- **Support for data center bridging quantized congestion notification (QFX 10000 Series)**—Starting with Junos OS Release 17.1R1, QFX 10000 Series switches support data center bridging quantized congestion notification, which is a congestion management mechanism that sends a congestion notification message through the network to the ultimate source of the congestion, stopping congestion at its source.

[See [Understanding DCB Features and Requirements](#).]

- **New show interfaces command for virtual output queues (QFX 10000 Series)**—Starting with Junos OS Release 17.1R1, QFX 10008 Series switches support the **show interfaces voq interface-name** command, which enables you to view statistics for virtual output queues.

[See [show interfaces voq](#).]

- **Support for data center bridging standards (QFX 10000 Series)**—Starting with Junos OS Release 17.1R1, QFX 10008 Series switches support three data center bridging standards:

- Priority-based flow control (PFC) allows you to select traffic flows within a link and pause them, so that the output queues associated with the flows do not overflow and drop packets.
- Enhanced transmission selection (ETS), also called CoS hierarchical port scheduling, is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues (forwarding classes) and to groups of queues (forwarding class sets).
- Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks.

[See [Understanding DCB Features and Requirements](#).]

- **Support for data center bridging standards (QFX 5100 Series)**—Starting with Junos OS Release 17.1R1, class of service (CoS) features can be configured on OVSDB-managed VXLAN interfaces on QFX5100 switches. An OVSDB-managed VXLAN interface uses an OVSDB controller to create and manage the VXLAN interfaces and tunnels.

[See [Understanding CoS on OVSDB-Managed VXLAN Interfaces](#).]

Dynamic Host Configuration Protocol

- **Virtual-router aware DHCP server/DHCP relay agent (QFX10008)**—Starting with Junos OS Release 17.1R1, QFX10000 switches can be configured to act as a DHCP server or DHCP relay agent for IPv4 and IPv6. If you have virtual router instances on the switch, the DHCP implementation can work with them. This feature was previously supported in an “X” release of Junos OS.

[See [DHCP and BOOTP Relay Overview](#).]

High Availability (HA) and Resiliency

- **Support for high availability features (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, the following features are supported:
 - Graceful Routing Engine switchover (GRES)—Enables a switch with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails.
To configure GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level and the **synchronize** statement at the **[edit system commit]** hierarchy level.
 - Nonstop active routing (NSR)—Uses the same infrastructure as GRES to preserve interface and kernel information. NSR also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine.
To configure NSR, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level.
 - Nonstop bridging (NSB)—Uses the same infrastructure as GRES to preserve interface and kernel information. NSB also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.
To configure NSB, include the **nonstop-bridging** statement at the **[edit protocols layer2-control]** hierarchy level.

These features were previously supported in an “X” release of Junos OS.

Infrastructure

- **Support for Secure Boot (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.1R1, a significant system security enhancement, Secure Boot, has been introduced. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement Secure Boot.

This feature was previously supported in an “X” release of Junos OS.

Interfaces and Chassis

- **LACP hold-up timer configuration and initialization delay timer support on LAG interfaces (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, 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. 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 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 interface-name aggregated-ether-options lacp]** hierarchy level.

You can configure an initialization delay timer value on link aggregation group (LAG) interfaces. When a standby multichassis aggregated Ethernet (MC-AE) interface reboots to come up in active-active MC-AE mode, the Link Aggregation Control Protocol (LACP) protocol comes up faster than the Layer 3 protocols. As soon as LACP comes up, the interface is UP and starts receiving traffic from the neighboring interfaces. In absence of the routing information, the traffic received on the interface is dropped, causing traffic loss. The initialization delay timer, when configured, delays the MC-AE node from coming UP for a specified amount of time. This gives the Layer 3 protocols time to converge on the interface and prevent traffic loss. To configure the initialization delay timer on an MC-AE interface, use the **init-delay-timer** statement at the **[edit interfaces ae interface-name aggregated-ether-options mc-ae]** hierarchy level.

These features were previously supported in an “X” release of Junos OS.

[See [Configuring LACP Hold-UP Timer to Prevent Link Flapping on LAG Interfaces](#) and [mc-ae](#).]

- **Support for 10-Gigabit Ethernet on QFX10000-30C line card (QFX10008 and QFX10016)**—Starting with Junos OS Release 17.1R1, QFX10008 and QFX10016 switches support 10-Gigabit Ethernet interfaces in addition to 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces on the QFX10000-30C line card.

When a particular provider edge (PE) is working in mode A to support 10-Gigabit Ethernet, ports 6, 7, 16, 17, 26, and 27 at the PE0 to PE5 level are non-operational. However, once the PE goes into mode

A, these ports can operate at 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet as well.

Depending on the optics that are plugged in, the interface works in 40-Gigabit Ethernet or 100-Gigabit Ethernet speed. For 10-Gigabit Ethernet, you must configure the port using the channelization command. Because there is no port-groups option for the 100-Gigabit Ethernet line card, you must use individual port channelization commands.

In 30C line card, by default FPC comes up in Mode D. when you channelize first port in any PE, whole FPC restarts and corresponding PE comes up in Mode A. Further channelization in that PE does not restart the FPC. But if you channelize some other ports in other PE, then the whole FPC restarts again. If you undo the channelization of all ports in any PE, then FPC gets restarted and corresponding PE comes up in Mode D which is the default mode. [See [QFX10000-30C Line Card](#).]

NOTE: If any mode changes (A to D or D to A) occur at the PE, you must perform a cold reboot on the Packet Forwarding Engine.

- **Support for multichassis link aggregation groups (MC-LAG) (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can use MC-LAG to enable a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without Running STP.

On one end of an MC-LAG is an MC-LAG client that has one or more physical links in a LAG. This client does not need to detect the MC-LAG. On the other side of the MC-LAG are two MC-LAG switches. Each of these switches has one or more physical links connected to a single client. The switches coordinate with each other to ensure that data traffic is forwarded properly.

This feature was previously supported in an “X” release of Junos OS.

[See [Multichassis Link Aggregation Features, Terms, and Best Practices](#).]

- **Support for link aggregation (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can use multiple network cables and ports in parallel to increase link speed and redundancy.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Aggregated Ethernet Interfaces and LACP](#).]

- **LAG local minimum links per Virtual Chassis or VCF member (QFX5100 switches)**—Starting with Junos OS Release 17.1R1, you can use the local minimum links feature to help avoid traffic loss due to asymmetric bandwidth on link aggregation group (LAG) forwarding paths through a Virtual Chassis or Virtual Chassis Fabric (VCF) member switch when one or more LAG member links local to that chassis have failed.

When this feature is enabled, if a user-configured percentage of local LAG member links has failed on a chassis, all remaining local LAG member links on the chassis are forced down, and LAG traffic is redistributed only through LAG member links on *other* chassis.

To enable local minimum links for an aggregated Ethernet interface (aex), set the **local-minimum-links-threshold** configuration statement with a threshold value that represents the percentage of local member links that must be up on a chassis for any local LAG member links on that chassis to continue to be active in the aggregated Ethernet bundle. Otherwise, all remaining LAG member links on that chassis are also forced down. The feature responds dynamically to bring local LAG member links up or down if you change the configured threshold, or when the status or configuration of LAG member links changes. Note that forced-down links also influence the minimum links count for the LAG as a whole, which can bring down the LAG, so enable this feature only in configurations where LAG traffic is carefully monitored and controlled.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Local Minimum Links](#).]

- **Support for Micro BFD over child links of AE or LAG bundle (cross-functional Packet Forwarding Engine/kernel/rpd) (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, this feature provides a Layer 3 BFD liveness detection mechanism for child links of the Ethernet LAG interface. In scenarios in which you do not have a point-to-point link, and a Layer 1 device fails at one end of the link, Micro BFD detects failures faster than traditional LACP. Micro BFD sessions are independent of each other despite having a single client that manages the LAG interface. Micro BFD is not supported on pure Layer 2 interfaces.

To enable failure detection for aggregated Ethernet interfaces, include the **bfd-liveness-detection** statement at the **[edit interfaces aex aggregated-ether-options bfd-liveness-detection]** hierarchy level.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Independent Micro BFD Sessions for LAG](#).]

- **PVLAN and Q-in-Q on the same interface (QFX5100 Switches)** —Starting with Junos OS Release 17.1R1, you can configure a private VLAN and Q-in-Q tunneling on the same Ethernet port. To configure both PVLAN and Q-in-Q on the same physical interface, you must configure flexible Ethernet services to support dual methods of configuring logical interfaces. Q-in-Q requires a service provider configuration method, and PVLAN requires an enterprise configuration method.

This feature was previously supported in an “X” release of Junos OS.

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

- **Support for configuration synchronization for MC-LAG (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, Multichassis Link Aggregation group (MC-LAG) configuration synchronization enables you to easily propagate, synchronize, and commit configurations from one MC-LAG peer to another. You can log into any one of the MC-LAG peers to manage both MC-LAG peers, thus having a single point of management. You can also use configuration groups to simplify the configuration process. You can create one configuration group for the local MC-LAG peer, one for the remote MC-LAG peer, and one for the global configuration, which is essentially a configuration that is common to both MC-LAG peers.

In addition, you can create conditional groups to specify when a configuration is synchronized with another MC-LAG peer. You can enable the **peers-synchronize** statement at the **[edit system commit]** hierarchy to synchronize the configurations and commits across the MC-LAG peers by default. NETCONF over SSH provides a secure connection between the MC-LAG peers, and Secure Copy Protocol (SCP) copies the configurations securely between them.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding MC-LAG Configuration Synchronization](#).]

- **Support for configuration consistency check for MC-LAG (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, Multichassis Link Aggregation group (MC-LAG) configuration consistency check alerts you of both severe and moderate configuration inconsistencies across MC-LAG peers. The configuration consistency check feature checks MC-LAG configuration parameters, such as chassis ID, session establishment time, and so on, on each peer and notifies you of any errors, so you can fix the inconsistencies. Configuration inconsistencies are categorized as severe or moderate. If there is a severe inconsistency, the MC-LAG interface is prevented from coming up. Once you have corrected the inconsistency, the system will bring up the interface. If there is a moderate inconsistency, you are notified of the error and can then fix the inconsistency. After you fix any inconsistency, you must commit the changes to take effect.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Multichassis Link Aggregation Group Configuration Consistency Check](#).]

- **Configuration support to improve MC-LAG Layer 2 and Layer 3 convergence (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can configure multichassis link aggregation (MC-LAG) interfaces to improve Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet link goes down or comes up in a bridge domain. To use this feature, ensure that the Inter-Chassis Link (ICL) is configured on an aggregated Ethernet interface. For Layer 2 convergence, configure the **enhanced-convergence** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. For Layer 3 convergence, configure the **enhanced-convergence** statement on an integrated routing and bridging (IRB) interface at the **[edit interfaces irb unit unit-number]** hierarchy level.

This feature was previously supported in an “X” release of Junos OS.

[See [enhanced-convergence](#).]

- **Channelizing 40-Gigabit Ethernet QSFP+ ports (QFX10008 switch)**—This feature enables you to channelize four 10-Gigabit Ethernet interfaces from the 40-Gigabit Ethernet QSFP+ interfaces. Channelization is supported on fiber break-out cable using standard structured cabling techniques.

NOTE: This feature is not supported on the QFX10000-30C line card.

By default, the 40-Gigabit Ethernet QSFP+ interfaces are named **et-fpc/pic/port**. The resulting 10-Gigabit Ethernet interfaces appear in the following format: **xe-fpc/pic/port:channel**, where channel can be a

value of 0 through 3. To channelize a 40-Gigabit Ethernet QSFP+ interface into four 10-Gigabit Ethernet interfaces, include the **10g** statement at the **[edit chassis fpc fpc-slot pic pic-slot (port port-number | port-range port-range-low port-range-high) channel-speed]** hierarchy level. To revert the 10-Gigabit Ethernet channels to a full 40-Gigabit Ethernet interface, remove the **10g** statement from the same hierarchy level.

There are 100-Gigabit Ethernet ports that work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet but are recognized as 40-Gigabit Ethernet by default. You cannot channelize the 100-Gigabit Ethernet ports when they are operating as 100-Gigabit Ethernet interfaces. The 40-Gigabit Ethernet ports can operate independently or be channelized into four 10-Gigabit Ethernet ports as part of a port range. Ports cannot be channelized individually. Only the first and fourth port in each 6XQSFP cage is available to channelize as part of a port range. In a port range, the ports are bundled with the next two consecutive ports. For example, if you want to channelize ports 0 through 2, you channelize port 0 only. If you try to channelize a port that is not supported, you receive an error message when you commit the configuration. Auto-channelization is not supported on any ports.

When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports.

This feature was previously supported in an “X” release of Junos OS.

[See [Channelizing Interfaces](#).]

IP Tunneling

- **Generic Routing Encapsulation support (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.1R1, you can configure GRE tunnels. GRE provides a private, secure path for transporting packets through a public network by encapsulating (or tunneling) the packets. GRE tunneling is accomplished through tunnel endpoints that encapsulate or de-encapsulate traffic. To configure a GRE tunnel interface, include the **gre-fpc/pic/port** set of statements at the **[edit interfaces]** hierarchy level.

This feature was previously supported only on the QFX10002 switch.

[See [Configuring Generic Routing Encapsulation Tunneling](#).]

IPv4

- **IPv4 address conservation method for hosting providers (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can configure a static route on an IRB interface with or without pinning to a specific underlying interface, thereby conserving the usage of IP address space.

Configure the interface on the router with an address from the reserved IPv4 prefix for shared address space (RFC 6598) and by using static routes pointed at the interface. IANA has recorded the allocation of an IPv4 /10 for use as shared address space. The shared address space address range is 100.64.0.0/10.

This way, the interface in the router is allocated an IP address from the shared address space, so it is not consuming publicly routable address space, and connectivity is handled with static routes on an interface. The interface in the server is configured with a publicly routable address, but the router interfaces are not. Network and broadcast addresses are consumed out of the shared address space rather than the publicly routable address space.

[See [IPv4 Address Conservation Method for Hosting Providers](#).]

Layer 2 Features

- **Support for Layer 2 Features (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, the following features are supported:
 - VLAN support—Enables you to divide one physical broadcast domain into multiple virtual domains.
 - LLDP—Enables a switch to advertise its identity and capabilities on a LAN, as well as receive information about other network devices.
 - Q-in-Q tunneling support—Allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.
 - Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP)—Provide Layer 2 loop prevention.

These features were previously supported in an “X” release of Junos OS.

[See [Overview of Layer 2 Networking](#).]

- **NNI and UNI on same interface (QFX5100 switches)**—Starting with Junos OS Release 17.1R1, this feature allows you to configure the same interface as a network-to-network interface (NNI) and a user-network interface (UNI) when you use Q-in-Q tunneling. This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Q-in-Q Tunneling](#).]

- **Q-in-Q tunneling support (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.1R1, this feature allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer’s 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag. This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Q-in-Q Tunneling](#).]

- **Support for IRB interfaces on Q-in-Q VLANs (QFX5100 switches and QFX5100 Virtual Chassis)**—Starting with Junos OS Release 17.1R1, integrated routing and bridging (IRB) interfaces are supported on Q-in-Q VLANs—you can configure the IRB interface on the same interface as one used by an S-VLAN, and you can use the same VLAN ID for both the VLAN used by the IRB interface and for the VLAN used as an S-VLAN. This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Q-in-Q Tunneling](#).]

- **Dual VLAN tag translation (QFX5100 switches and QFX5100 Virtual Chassis)**—Starting with Junos OS Release 17.1R1, you can use the dual VLAN tag translation (also known as dual VLAN tag rewrite) feature to deploy switches in service-provider domains, allowing dual-tagged, single-tagged, and untagged VLAN packets to come into or exit from the switch. Operations added for dual VLAN tag translation are swap-push, swap-swap, and pop-push. This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Q-in-Q Tunneling](#).]

Layer 3 Features

- **Support for Layer 3 unicast features (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, the following layer 3 features for unicast IPv4 and IPv6 traffic are supported on QFX10000 switches:
 - Neighbor Discovery Protocol (IPv6 only)
 - Virtual Routers
 - OSPF
 - IS-IS
 - BGP
 - VRRP

This feature set was previously supported in an “X” release of Junos OS.

[See [IPv6 Neighbor Discovery Overview](#).]

Management

- **Support for adding non-native YANG modules to the Junos OS schema (QFX Series)**—Starting in Junos OS Release 17.1R1, you can load custom YANG models on devices running Junos OS to add data models that are not natively supported by Junos OS but can be supported by translation. Doing this enables you to extend the configuration hierarchies and operational commands with data models that are customized for your operations. The ability to add data models to a device is also beneficial when you want to create device-agnostic and vendor-neutral data models that enable the same configuration or RPC to be used on different devices from one or more vendors. You can load custom YANG modules by using the **request system yang add** operational command.

[See [Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS](#).]

Multicast

- **Layer 2 and layer 3 multicast support (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, IGMP, including versions 1, 2, and 3, IGMP snooping, PIM-SM and PIM-SSM are supported. You can also configure IGMP, IGMP snooping and PIM in virtual-router instances. MSDP is also supported. Configure IGMP at the **[edit protocols igmp]** hierarchy level. Configure IGMP snooping at **[edit protocols igmp-snooping]** hierarchy level. Configure PIM at the **[edit protocols pim]** hierarchy level. Configure MSDP at the **[edit protocols msdp]** hierarchy level.

This feature set was previously supported in an “X” release of Junos OS.

[See [Multicast Overview](#).]

MPLS

- **Path Computation Element Protocol (QFX10000 switch)**—Starting in Junos OS Release 17.1R1, QFX10000 switch supports the Path Computation Element Protocol (PCEP). A Path Computation Element (PCE) is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE. PCEP enables communications between a PCC and a PCE, or between two PCEs (defined in RFC 5440).

[See [PCEP Overview](#).]

- **Static label-switched path with resolve next hop (QFX5100 switches)**—Starting in Junos OS Release 17.1R1, you can configure a static label-switched path (LSP) to be resolved to a next hop that is not directly connected. This feature provides simplicity and scalability to your configuration, because you are no longer required to configure multiple, directly connected next hops if you have multiple links.

This feature was previously supported in an “X” release of Junos OS.

[See [MPLS Stitching for Virtual Machine Connection](#).]

- **MPLS support (QFX5100 switches)**—Starting in Junos OS Release 17.1R1, MPLS is supported on the QFX10008 and QFX10016 switches. MPLS provides both label edge router (LER) and label switch router (LSR) and provides the following capabilities:
 - Support for both MPLS major protocols, LDP and RSVP
 - IS-IS interior gateway protocol (IGP) traffic engineering
 - Class of service (CoS)
 - Object access method, including ping, traceroute, and Bidirectional Forwarding Detection (BFD).
 - Fast reroute (FRR), a component of MPLS local protection (both one-to-one local protection and many-to-one local protection are supported).
 - Loop-free alternate (LFA)
 - SixPE devices
 - Layer 3 VPNs for both IPv4 and IPv6
 - LDP tunneling over RSVP

This feature was previously supported in an “X” release of Junos OS.

[See [MPLS Overview for Switches.](#)]

- **Support for IRB interfaces over MPLS (QFX5100 switches)**—Starting in Junos OS Release 17.1R1, you can configure integrated routing and bridging (IRB) interfaces over an MPLS network. An IRB is a logical Layer 3 VLAN interface used to route traffic between VLANs. An IRB interface functions as a logical switch on which you can configure a Layer 3 logical interface for each VLAN. The switch relies on its Layer 3 capabilities to provide this basic routing between VLANs.

This feature was previously supported in an “X” release of Junos OS.

[See [Example: Configuring IRB Interfaces on QFX5100 Switches over an MPLS Core Network.](#)]

- **Support for MPLS automatic bandwidth allocation and dynamic label switched path (LSP) count sizing (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and rely on this feature to dynamically adjust the bandwidth allocation based on current traffic patterns. Dynamic LSP count sizing provides an ingress router with the capability of acquiring as much network bandwidth as possible by creating parallel LSPs dynamically. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

This feature was previously supported in an “X” release of Junos OS.

[See [Configuring Automatic Bandwidth Allocation for LSPs.](#)]

- **Support for MPLS filters (QFX10000 switches)**—Starting in Junos OS Release 17.1R1, you can configure firewall filters to filter MPLS traffic. To use an MPLS firewall filter, you must first configure the filter and then apply it to an interface that you have configured for forwarding MPLS traffic. You can also configure

a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached.

This feature was previously supported in an “X” release of Junos OS.

[See [Configuring MPLS Firewall Filters and Policers](#).]

- **BGP link state distribution (QFX Series and QFX10000)**—Starting with Junos OS Release 17.1R1, you can deploy a mechanism to distribute topology information across multiple areas and autonomous systems (ASs) by extending the BGP protocols to carry link state information. Previously, this information was acquired using an IGP. Using BGP provides a policy-controlled and scalable means of distributing the multi-area and multi-AS topology information. This information is used for computing paths for MPLS LSPs spanning multiple domains, such as inter-area TE LSP. This information also enables external path computing entities.

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

- **Ethernet-over-MPLS L2 circuit (QFX10000 switches)**—Starting in Junos OS Release 17.1R1, you can configure a Layer 2 circuit to create a point-to-point Layer 2 connection using MPLS on the service provider's network. Ethernet-over-MPLS allows sending Layer 2 (L2) Ethernet frames transparently over MPLS. Ethernet-over-MPLS uses a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. It encapsulates Ethernet protocol data units (PDUs) inside MPLS packets and forwards the packets, using label stacking, across the MPLS network. To enable a Layer 2 circuit, include the **l2circuit** statement at the **[edit protocols mpls labeled-switched-path lsp-name]** hierarchy level.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Ethernet-over-MPLS \(L2 Circuit\)](#).]

Network Management and Monitoring

- **Support for hrProcessorTable object (QFX Series)**—Starting in Junos OS Release 17.1R1, support is provided for the **hrProcessorTable** object (object id: 1.3.6.1.2.1.25.3.3) described in the RFC2790, *Host Resources MIB*. The **hrProcessorTable** object provides the load statistics information per CPU for multi-core devices.

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

- **IEEE 802.3ah (QFX10002, QFX10008, QFX10016)**—QFX Series switches support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning.
- **Port mirroring support (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.1R1, port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring

and predicting traffic patterns, correlating events, and so on. This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Port Mirroring](#).]

- **sFlow technology support (QFX10008/QFX10016 switches)**—Starting in Junos OS Release 17.1R1, the QFX10008 and QFX10016 switches support monitoring technology for high-speed switched or routed networks. You can configure sFlow technology to monitor traffic continuously at wire speed on all interfaces simultaneously. sFlow technology also collects samples of network packets, providing you with visibility into network traffic information. You configure sFlow monitoring at the **[edit protocols sflow]** hierarchy level. sFlow operational commands include **show sflow** and **clear sflow collector statistics**.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding How to Use sFlow Technology for Network Monitoring on a Switch](#).]

Port Security

- **Support for MAC limiting and MAC move limiting on OVSDB-managed interfaces (QFX5100 switches)**—Starting in Junos OS Release 17.1R1, you can configure MAC limiting and MAC move limiting on interfaces managed by a Contrail controller through the Open vSwitch Database (OVSDB) management protocol. MAC limiting protects against flooding of the Ethernet switching table. MAC move limiting detects MAC movement and MAC spoofing on access interfaces.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding MAC Limiting and MAC Move Limiting for Port Security](#).]

Routing Policy and Firewall Filters

- **IPv4 filter-based GRE tunneling (QFX10000 switches)**—Starting in Junos OS Release 17.1R1, QFX10000 switches support filter-based generic routing encapsulation (GRE) tunneling across IPv4 networks. GRE tunneling is performed by tunnel endpoints that encapsulate or de-encapsulate traffic. With filter-based GRE tunneling, you can use a firewall filter to de-encapsulate traffic over an IPv4 network. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. This provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation.

[See [Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100, QFX10000, or OCX Switch](#).]

Routing Protocols

- **Support for BGP flow routes for traffic filtering (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can propagate flow routes as part of BGP through flow-specification network-layer reachability information (NLRI) messages. Flow routes provide traffic filtering and rate-limiting capabilities much like firewall filters. Propagating flow routes as part of BGP enables you to propagate filters against denial-of-service (DOS) attacks dynamically across autonomous systems. Include the **flow route name** set of statements at the **[edit routing-options]** hierarchy level.

[See [Example: Enabling BGP to Carry Flow-Specification Routes](#).]

- **Support for advertising multiple paths in BGP (QFX5100 switches and QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can configure BGP to advertise multiple paths to the same destination, instead of advertising only the active path. The potential benefits of advertising multiple paths for BGP include fault tolerance, load balancing, and maintenance. Include the **add-path** set of statements at the **[edit protocols bgp group group-name family family-type]** hierarchy level.

[See [add-path](#).]

- **Enhancement to ECMP next-hop groups (QFX5100 switches)**—Starting with Junos OS Release 17.1R1, equal-cost multipath (ECMP) next hops are allocated dynamically. A dynamic, rather than fixed, allocation of ECMP next hops, or paths, effectively increases the number of ECMP groups available for route resolution. For example, if the maximum number of ECMP next hops is set to 16, a dynamic allocation

means that as many 1,000 ECMP groups are supported. To configure the maximum limit for ECMP next hops, include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level.

This feature was previously introduced in an "X" release of Junos OS.

[See [Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing.](#)]

- **Support for BGP Monitoring Protocol (BMP) Version 3 (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can configure BMP, which sends BGP route information from the switch to a monitoring application, or station, on a separate device. To deploy BMP in your network, you need to configure BMP on each switch and at least one BMP monitoring station. Only version 3 is supported. To configure BMP, include the **bmp** set of statements at the **[edit routing-options]** hierarchy level. To configure a BMP monitoring station, include the **station-address ip-address** and the **station-port number** statements at the **[edit routing-options bmp]** hierarchy level.

This feature was previously introduced in an "X" release of Junos OS.

[See [Configuring BGP Monitoring Protocol Version 3.](#)]

Security

- **Firewall filter support (QFX10008/QFX10016 switches)**—Starting in Junos OS Release 17.1R1, you can define firewall filters on the switch that defines whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces.

This feature was previously supported in an "X" release of Junos OS.

[See [Overview of Firewall Filters.](#)]

- **Policing support (QFX10008/QFX10016 switches)**—Starting in Junos OS Release 17.1R1, you can use policing to apply limits to traffic flow and to set consequences for packets that exceed those limits. A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.

This feature was previously supported in an "X" release of Junos OS.

[See [Overview of Policers.](#)]

- **Support for policers on OVSDB-managed interfaces (QFX5100 switches)**—Starting in Junos OS Release 17.1R1, you can configure two-rate three-color markers (policers) on interfaces managed by a Contrail controller through the Open vSwitch Database (OVSDB) management protocol.

This feature was previously supported in an "X" release of Junos OS.

[See [Understanding Policers on OVSDB-Managed Interfaces.](#)]

- **Support for firewall filters on OVSDB-managed interfaces (QFX5100 switches)**—Starting in Junos OS Release 17.1R1, you can configure firewall filters on interfaces managed by a Contrail controller through

the Open vSwitch Database (OVSDB) management protocol. Firewall filters enable you to control packets transiting a device to a network destination as well as packets destined for and sent by a device.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Firewall Filters on OVSDB-Managed Interface.](#)]

Software Defined Networking

- **Support for EVPN-VXLAN (QFX5100 and QFX10000 switches)**—Traditionally, data centers use Layer 2 technologies such as STP and multi-chassis link aggregation groups (MC-LAGs) for compute and storage connectivity. As the design of data centers shifts to scale-out, service-oriented multi-tenant networks, a new data center architecture emerges that allows decoupling of an underlay network from the tenant overlay network with VXLAN. Starting with Junos OS Release 17.1R1, you can use a Layer 3 IP-based underlay coupled with an EVPN-VXLAN overlay to deploy larger networks than those possible with traditional Layer 2 Ethernet-based architectures. With an EVPN-VXLAN overlay, endpoints (servers or virtual machines) can be placed anywhere in the network and remain connected to the same logical Layer 2 network.

This feature was previously supported in an “X” release of Junos OS.

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

- **Support for LACP in EVPN active-active multihoming (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, 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.](#)]

- **OVSDB schema updates (QFX5100, QFX5100VC)**—Starting with Junos OS Release 17.1R1, the Open vSwitch Database (OVSDB) schema (for physical devices) implemented on QFX5100 switches is version 1.3.0. In addition, this schema now supports the multicast MACs local table.

This feature was previously supported in an “X” release of Junos OS.

[See [OVSDB Schema for Physical Devices](#).]

- **Class-of-service support for OVSDB-managed VXLAN interfaces (QFX5100 switches)**—Starting with Junos OS Release 17.1R1, class-of-service (CoS) features can be configured on OVSDB-managed VXLAN interfaces on QFX5100 switches. An OVSDB-managed VXLAN interface uses an OVSDB controller to create and manage the VXLAN interfaces and tunnel. T

his feature was previously supported in an “X” release of Junos OS.

[See [Understanding CoS on OVSDB-Managed VXLAN Interfaces](#).]

- **Support for ping and traceroute with VXLANs (QFX5100 switches)**—Starting with Junos OS Release 17.1R1, you can use ping and traceroute to troubleshoot the physical underlay that supports a VXLAN overlay.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Overlay ping and traceroute Packet Support](#).]

- **PIM NSR support for VXLAN (QFX5100 Virtual Chassis)**—Starting in Junos OS Release 17.1R1, the QFX5100 Virtual Chassis supports Protocol Independent Multicast (PIM) nonstop active routing (NSR) for Virtual Extensible LANs (VXLANs).

The Layer 2 address learning daemon (l2ald) passes VXLAN parameters (VXLAN multicast group addresses and the source interface for a VXLAN tunnel **vtep-source-interface**) to the routing protocol process on the master Routing Engine. The routing protocol process forms PIM joins with the multicast routes through the pseudo-VXLAN interface based on these configuration details.

Because the l2ald daemon does not run on the backup Routing Engine, the configured parameters are not available to the routing protocol process in the backup Routing Engine when NSR is enabled. The PIM NSR mirroring mechanism provides the VXLAN configuration details to the backup Routing Engine, which enables creation of the required states. The routing protocol process matches the multicast routes on the backup Routing Engine with PIM states, which maintains the multicast routes in the Forwarding state.

[See [PIM NSR Support for VXLAN Overview](#).]

Software Installation and Upgrade

- **Support for FreeBSD 10 kernel for Junos OS (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, FreeBSD is the underlying OS that enables SMP for Junos OS, rather than the FreeBSD 6.1 version that is used in some older Juniper Networks devices. If you compare the switch to devices that run the older kernel, you will notice that some system commands display different output and a few other commands are deprecated.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Junos OS with Upgraded FreeBSD.](#)]

System Management

- **Support for Precision Time Protocol (PTP) transparent clock (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, PTP synchronizes clocks throughout a packet-switched network. With a transparent clock, the PTP packets are updated with residence time as the packets pass through the switch. There is no master/slave designation. End-to-end transparent clocks are supported. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that the timestamps are sent in one packet. In a two-step process, estimated timestamps are sent in one packet, and additional packets contain updated timestamps. In addition, User UDP over IPv4 and IPv6, and unicast and multicast transparent clock are supported.

You can configure the transparent clock at the **[edit protocols ptp]** hierarchy.

This feature was previously supported in an “X” release of Junos OS.

[See [Understanding Transparent Clocks in Precision Time Protocol.](#)]

- **Support for reporting FATAL and MAJOR FAULT information (QFX10000 switches)**—Starting in Junos OS Release 17.1R1, FATAL and MAJOR errors are reported in the output of the **show chassis fpc errors** command.

This feature was previously supported in an “X” release of Junos OS.

VPNs

- **Support for carrier-of-carriers Layer 3 VPNs (QFX10000 switches)**—Starting in Junos OS 17.1R1, this feature is supported for customers who want to provide VPN service. Layer 3 VPNs based on BGP MPLS are used by service providers to provide VPN services to end-user customers, enabling these customers to use the MPLS backbone network to connect their multiple sites seamlessly. Include the **labeled-unicast** statement in the configuration for the IBGP session to the carrier-of-carriers customer’s CE device and include the **family-inet-vpn** statement in the configuration for the IBGP session to the carrier-of-carriers PE device on the other side of the network.

[See [Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service.](#)]

- **IPv6 Layer 3 VPNs (QFX5100 and QFX10000 switches)**—You can now configure switch interfaces in a Layer 3 VPN to carry IPv6 traffic. This feature, commonly referred to as 6VPE, allows for the transport of IPv6 traffic across an MPLS-enabled IPv4 backbone to provide VPN service for IPv6 customers.

This feature was previously supported in an “X” release of Junos OS.

[See [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks.](#)]

SEE ALSO

[Changes in Behavior and Syntax | 296](#)

[Known Behavior | 302](#)

[Known Issues | 304](#)

[Resolved Issues | 309](#)

[Documentation Updates | 326](#)

[Migration, Upgrade, and Downgrade Instructions | 326](#)

[Product Compatibility | 339](#)

Changes in Behavior and Syntax

IN THIS SECTION

- [Class of Service | 297](#)
- [General Routing | 297](#)
- [MPLS | 297](#)
- [Network Management and Monitoring | 297](#)
- [Security | 300](#)
- [Services Applications | 300](#)
- [Software Defined Networking | 300](#)
- [Software Installation and Upgrade | 300](#)
- [System Management | 300](#)
- [User Interface and Configuration | 300](#)
- [Virtual Chassis | 301](#)
- [VPNs | 301](#)

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 17.1R3 for the QFX Series.

Class of Service

- When you configure the **transmit-rate** statement, you must also configure the **guaranteed-rate** statement under **traffic-control-profiles**. If you commit the configuration of the **transmit-rate** statement without configuring **guaranteed-rate**, a warning message is displayed and the default scheduler map is applied.

General Routing

- **Support for deletion of static routes when the BFD session goes down (QFX Series)**—Starting with Junos OS Release 17.1R3, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

MPLS

- **Representation for OSPF designated router node**—Up until version -10 of the Internet Engineering Task Force (IETF) BGP-LS draft, the OSPF designated router node representation was ambiguous. One could represent designated router nodes as 'AdvertisingRouterId-InterfaceIpAddress' or 'InterfaceIpAddress-1'. Junos OS used to follow the 'InterfaceIpAddress-1' format. Starting with version -11 of the IETF BGP-LS draft, the representation for OSPF designated router node must be 'AdvertisingRouterId-Interfaceaddress'. Junos OS now follows the latest format.

Network Management and Monitoring

- **SNMP syslog messages changed (QFX Series)**—Starting in Junos OS Release 17.1R1, two misleading SNMP syslog messages have been rewritten to accurately describe the event:
 - OLD—AgentX master agent failed to respond to ping. Attempting to re-register
 - NEW—AgentX master agent failed to respond to ping, triggering cleanup!
 - OLD—NET-SNMP version %s AgentX subagent connected
 - NEW—NET-SNMP version %s AgentX subagent Open-Sent!

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

- **Cloud Analytics Engine disabled in Junos OS by default (QFX Series)**—Starting in Junos OS Release 17.1R1 and later, Cloud Analytics Engine network analytics probe processing is disabled by default in Junos OS. Probe processing is enabled automatically when you configure any supported Cloud Analytics Engine configuration statement in the **[edit system services cloud-analytics]** configuration statement hierarchy. In Junos OS Release 16.1R3 and earlier, Cloud Analytics Engine Junos OS functionality is enabled by default, and no configuration steps are required for the Junos OS to process and respond to probes.

[See [Configuring Cloud Analytics Engine on Devices.](#)]

- **Update to SNMP support of apply-path statement (QFX Series)**—In Junos OS Release 17.1R2, SNMP implementation for the apply-path configuration statement supports only two lists:

- **apply-path "policy-options prefix-list <list-name> <*>"**

This configuration has been supported from day 1.

- **apply-path "access radius-server <*>"**

This configuration is supported as of this release.

- **Juniper MIBs Loading Errors Fixed (QFX Series)**—In Junos OS Release 17.1R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:

- jnx-gen-set.mib
- jnx-ifotn.mib
- jnx-optics.mib

[See [MIB Explorer.](#)]

- **Change in default log level setting (QFX Series)**—In Junos OS Release, 17.1R3, 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 that come from a nondefault routing instance and a nondefault logical system (QFX Series)**—In Junos OS Release 17.1, a new option, **context-oid**, for the **trap-options** statement enables 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.](#)]

- **Need to reconfigure SNMPv3 configuration after upgrade (QFX Series)**—In Junos OS Release 17.1R2, you might need to reconfigure SNMPv3 after upgrading from an earlier release to this release. This is necessary only if you are using SNMPv3 and if the engine ID is based on the MAC address because the engine ID is changed. In releases before Junos OS Release 17.1R3, you have to reconfigure SNMPv3

every time after a reboot. This problem is fixed. If you upgrade, you must still reconfigure the SNMPv3, but only once—if you have already reconfigured SNMPv3 in an earlier release, you do not need to reconfigure SNMPv3 again. To reconfigure the SNMPv3, use the **delete snmp v3** command, commit, and then reconfigure SNMPv3 parameters. Platforms affected are QFX5100, QFX10000, QFX10008, and QFX10016.

[See [Configuring the Local Engine ID.](#)]

Security

- **Syslog or log action on firewall drops packets (QFX5000 switches)**—Starting in 17.1R3, 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.

Services Applications

- **Device discovery with device-initiated connection (QFX Series)**—Starting in Junos OS Release 17.1R1 and later, when you configure statements and options under the **[system services ssh]** hierarchy and commit the configuration, make sure that the system reaches a stable state before you commit any **outbound-ssh** configurations.

You use the device discovery feature in the Devices workspace to add devices to Junos Space Network Management Platform. By default, Junos Space manages devices by initiating and maintaining a connection to the device.

[See [Device Discovery Overview](#).]

Software Defined Networking

- On QFX10000 switches running Junos OS Release 17.1R3 or later, the local preference setting for an Ethernet VPN (EVPN) pure type-5 route is inherited by IP routes that are derived from the EVPN type-5 route. Further, when selecting an IP route for incoming traffic, the QFX10000 switches consider the local preference of the route. A benefit of the QFX10000 switches including local preference in their route selection criteria is that you can set up a policy to manipulate the local preference, thereby controlling which route the switch selects.

Software Installation and Upgrade

- **In-service software upgrade (QFX5100 switches)**—Unified ISSU is not supported from earlier Junos OS releases to Junos OS Release 17.1R1.

System Management

- **Peers option not supported in batch configuration mode**— Starting in Junos OS Release 17.1R1, the **peers** option at the **[edit system commit]** hierarchy level is not supported in batch configuration mode.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (QFX**

Series)—Starting in Junos OS Release 17.1R1, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 17.1R1, integers in JSON configuration data were treated as strings and enclosed in quotation marks.

- **Changes to the `show system schema module juniper-command` output directory (QFX Series)**—Starting in Junos OS Release 17.1, when you issue the `show system schema module juniper-command` operational command in the Junos OS CLI, the device places the generated output files in the current working directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1, the generated output files are placed in the `/var/tmp` directory.

Virtual Chassis

- **Adaptive load balancing (ALB) feature (Virtual Chassis Fabric)**—Starting in Junos OS Release 17.1R3, 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 statement before initiating the upgrade.

See [Understanding Traffic Flow Through a Virtual Chassis Fabric](#) and [fabric-load-balance](#).

VPNs

- **Enhancements to output for `show route` and `show evpn ip-prefix-database` commands**—The `show route` command now displays a **Multipath** field for EVPN pure type-5 routes. This field shows the path selected by the routing protocol process. For the `show evpn ip-prefix-database extensive` command, the **IP Route Status** field is now displayed in the **Remote Advertisements** section. Previously, this field was displayed in the **Prefix** section. Also, the **inactive/active** field for each advertisement has been removed.

SEE ALSO

[New and Changed Features | 276](#)

[Known Behavior | 302](#)

[Known Issues | 304](#)

[Resolved Issues | 309](#)

[Documentation Updates | 326](#)

[Migration, Upgrade, and Downgrade Instructions | 326](#)

[Product Compatibility | 339](#)

Known Behavior

IN THIS SECTION

- [EVPN | 302](#)
- [General Routing | 303](#)
- [High Availability and Resiliency | 303](#)
- [Layer 2 Features | 304](#)
- [MPLS | 304](#)
- [Routing Protocols | 304](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R3 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.

EVPN

- A PE device running EVPN IRB with an IGP configured during VRF associated with the EVPN instance is unable to establish an IGP adjacency with a CE device attached to a remote PE device. The IGP instance running in the VRF on the PE device might be able to discover the IGP instance running on the remote CE device through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE device. [PR977945](#)
- On QFX10000 switches configured as type-5 route peers, when only peer 1 advertises routes, that peer might not install the decapsulated next-hop (NH) route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any type-5 route. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. [PR1191092](#)
- When you activate and deactivate route target per bridge domain in EVPN, the rpd process might crash, resulting in traffic loss. [PR1244956](#)

General Routing

- On QFX5100 switches, Zero Touch Provisioning might take some time to complete because TFTP might take a long time to fetch required data. [PR980530](#)
- On a QFX10002 switch, insert a small form-factor pluggable (SFP) transceiver on the management interface (em1). After a system reboot, if you replace the SFP transceiver with a copper SFP transceiver, the management interface might not work properly with speed 10m/100m. [PR1075097](#)
- On QFX Series switches, nonstop software upgrade (NSSU) cannot be used to upgrade from a Junos OS Release 14.1X53 image to a Junos OS Release 15.1 or later image. [PR1087893](#)
- On a fully loaded QFX10008 chassis, line cards might take as long as 15 minutes to become operational after startup. [PR1124967](#)
- On a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than 5 seconds of traffic loss for multicast traffic. [PR1125155](#)
- With a multihop BFD, traffic loss of around 5 to 10 seconds is observed when the intermediate interface is shut down. After 5 to 10 seconds, traffic recovers and no action is needed. [PR1150695](#)
- On disabling and reenabling a 1-Gigabit Ethernet port on a 60-port 10-Gigabit Ethernet line card in both QFX10008 and QFX100016 systems, **pechip_cmerror_set_error:3113: Level: Major, cmerror_code: 0x21060e (id=1550), recover_err: 0 (counter: 0), fh_msg: 0x0** messages are logged. No functionality impact is observed. [PR1238269](#)
- When software is upgraded to Junos OS Release 17.1R1 from a earlier release of Junos OS on QFX5100, the host platform is upgraded. As a result, unified ISSU from earlier releases to Junos OS Release 17.1R1 on these platforms is not supported. [PR1257220](#)

High Availability and Resiliency

- **Unified ISSU incompatibility with VPLS dynamic profiles (QFX Series)**—Using unified ISSU to upgrade from an earlier Junos OS Release to Junos OS Release 17.1R1 does not work if VPLS dynamic profiles are configured and enhanced subscriber management is not configured.
[See [ISSU System Requirements](#).]
- During a nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a traffic loop or loss might occur if the Junos OS Release that you are upgrading from and the Junos OS Release that you are upgrading to use different internal message formats. [PR1123764](#)

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging is enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. [PR1075230](#)

MPLS

- On QFX5100 switches with Layer 2 circuit configured on the PE switches, enabling VLAN bridge encapsulation on a CE device 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 GRES on QFX10000 switches, some IPv6 groups might experience momentary traffic loss. This issue occurs when a IPv6 traffic is running with multiple paths to the source, and the `join-load-balance` statement for PIM is also configured. [PR1208583](#)

SEE ALSO

New and Changed Features 276
Changes in Behavior and Syntax 296
Known Issues 304
Resolved Issues 309
Documentation Updates 326
Migration, Upgrade, and Downgrade Instructions 326
Product Compatibility 339

Known Issues

IN THIS SECTION

- [EVPN | 305](#)
- [General Routing | 306](#)

- MPLS | 307
- Network Management and Monitoring | 307
- Platform and Infrastructure | 307
- Routing Protocols | 308

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.1R3.

EVPN

- On QFX10000 switches, when you upgrade to Junos OS Release 15.1X53-D60 from Release 15.1X53-D33, traffic over route type-5 on the tunnel ingress node might drop if you have the forwarding-table **no-indirect-next-hop** statements configured at the **[edit routing-options]** hierarchy level. As a workaround, delete the configuration **routing-options forwarding-table no-indirect-next-hop** before you perform an upgrade. This configuration is not needed when route type-5 is configured. [PR1187482](#)
- On QFX10000 switches configured as type-5 route peers, when only peer 1 advertises routes, that peer might not install the decapsulated next-hop (NH) route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any type-5 route. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. [PR1191092](#)
- Error message **JPRDS_DLT_ALPHA KHT** shows as failed, but the entries in the hardware are programmed correctly. This might cause confusion between a working and a nonworking condition. [PR1258933](#)
- In an EVPN-VXLAN scenario, a previously learned MAC address from a remote Ethernet segment Identifier (ESI) cannot be changed to local even if it is connected directly. The MAC address of the host might remain as learned from ESI instead of the local interface until the MAC address is aged out. [PR1303202](#)
- The rpd generates unreproducible core file with scaling EVPN-VXLAN configuration on QFX10000 platforms because of the memory depletion on the EVPN MAC route entries queue for L2ALD. L2ALD closed the IPC connection that caused rpd-cumulated EVPN MAC route entries in the queue and ends up running out of memory. [PR1339979](#)

General Routing

- On QFX10002 switches, the **request system snapshot** command does not work. [PR1048182](#)
- On a QFX10002 switch, when a new interface is added to an existing link aggregation group (LAG) interface, which acts as an input analyzer interface, the traffic sent to the added interface might not be mirrored. [PR1057527](#)
- While using SSH to log in to a VNF the error message **Unrecognized command** is seen. This error has no impact on the functionality. [PR1108785](#)
- After sending leave and rejoin in a few seconds, L3 multicast traffic does not converge up to 100 percent and a few traffic drops are seen continuously. This behavior is seen when scaling beyond 2000 VLANs or 2000 IRBs with VLAN replication in the system. [PR1135045](#)
- L3 multicast traffic does not converge to 100 percent and a few continuous drops are observed after bringing an interface down and back up again or while an FPC comes online after FPC restart. This behavior is seen when scaling beyond 2000 VLANs or 2000 IRBs with VLAN replication configured. [PR1161485](#)
- When per-packet load balancing is removed or deleted, the next-hop index might change. [PR1198092](#)
- On QFX10002 platforms, some random ports, using 100-gigabit Lumentum optics, might not come up after a reboot. This is a timing issue because of failures during optics read on some ports. As a workaround, when you encounter this issue, remove and reinsert the optics, which might bring up the ports. [PR1227029](#)
- On QFX10008 switches, the IPv6 packets/bytes counter shows higher values than the total packets/bytes of the interface if LAG child members belong to the same PE device. As a workaround, if you monitor IPv6 statistics over the LAG, choose LAG child members across PE devices. [PR1232388](#)
- On QFX10000 line switches, sFlow monitoring technology output might display a negative number of samples after a long run. As a workaround, issue the **clear sflow collector** command to show or reset the count. [PR1244080](#)
- QFX10000 platforms do not support a discontinuous mask within source-address or destination-address of a firewall filter. When the user commits a firewall filter with a discontinuous mask prefix (for example, x.x.x.x/255.255.0.240) on QFX10000 platforms, the commit is successful but the filter does not take effect (the firewall compilation returns an error because discontinuous IP address mask is not supported and the filter is not programmed in hardware). [PR1267498](#)
- In QFX Series devices, if the em0 interface is unplugged, **Management Ethernet Links Down Alarms** flap is observed. [PR1271325](#)
- Every load override and rollback operation increases the refcount by 1 and after it reaches the maximum value of it (65,535), the mgd crash will be observed and the session will get killed. When mgd crashes, the active lock might remain, preventing any further commits. [PR1313158](#)
- The management process (mgd) might panic after modifying aggregated Ethernet interface members under the "ethernet-switching vlan" stanza. After mgd panic, your remote session is terminated as a result. [PR1325736](#)

- On QFX5100 platforms with sFlow enabled, when deleting/deactivating the sFlow interface, all other interfaces might go down and fxpc generates a core file. [PR1356868](#)
- When an MC-LAG is configured with **force-up** enabled on MC-LAG nodes, the LACP admin key should not match the key of the access or CE device. [PR1362346](#)
- On QFX10002, QFX10008, and QFX10016 platforms, the IPv6 traffic might be dropped if the IPv6 over IPv4 Generic Routing Encapsulation (GRE) tunnel is configured because when an interface family (IFF) member is removed on the logical interface, the property of the logical interface to learn MAC gets set to NULL which causes ARP failure. [PR1385723](#)
- MPLS configuration changes and topology changes might result in the tunnel initiator clear messages in the syslog. [PR1396014](#)
- In an aggregated interfaces and Spanning Tree Protocol (STP) scenario, the STP does not work when the aggregated interfaces number is "ae1000" or above in QFX5000 and "ae480" or above in other QFX Series platforms. Such interfaces will remain in incorrect STP discarding state and might not forward packets. [PR1403338](#)

MPLS

- Statistics of transit traffic does not increment LSP statistics signaled by RSVP-TE. [PR1362936](#)
- In an MPLS scenario, label-switched path (LSP) "statistic" and "auto-bandwidth" functionality might not take effect with single-hop LSPs on QFX10000 platform. [PR1390445](#)

Network Management and Monitoring

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

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced I2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- On all Junos OS based platforms, the Junos CLI **file copy** command uses **/var/home/<user>** as temporary staging directory for a non-root user, and uses **/var/tmp** for the root user. When a user issues the CLI command **file copy user@x.x.x.x:/dir/ /var/tmp/** to copy a file to the box, and if the file the user is trying to transfer is larger than the temporary staging directory size, the copy might fail. [PR1195599](#)
- When chassis control restart is done with the CoS rewrite rule configured on aggregated Ethernet interface, the **Platform failed to bind rewrite** messages might be seen in syslog. Issue is specific to

aggregated Ethernet interfaces. It is a timing issue that might occur when a logical interface deletion is delayed because of the high scale and when logical interfaces come up again after restart they have different indexes. [PR1315437](#)

Routing Protocols

- On QFX10000 switches, during a Routing Engine switchover, BGP on the IRB interface might flap when the IRB interface and the underlying Layer 2 logical interface (IFL) are configured with different MTU values. [PR1187169](#)
- On QFX10000 line switches, traffic drop is seen with IS-IS version 6 traffic during convergence in either of the following two scenarios: 1. While bringing up the ports after bringing them down. 2. While FPC comes online after doing an FPC restart. This behavior is seen while flapping one of the IS-IS version 6 sessions. [PR1190180](#)
- On QFX5100 and QFX10000 switches, traffic drop might occur in MC-LAG configurations. This occurs when an interchassis link (ICL) interface and then the MC-LAG interface are brought up. The traffic drop occur because the ARP next-hop update is not recognized on the Packet Forwarding Engine. To recover the traffic path over the MC-LAG interfaces, issue the **clear arp** command. As a workaround to avoid the issue, enable ICL interfaces and MC-LAG interfaces at the same time. [PR1236201](#)
- 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](#)
- With multicast traffic enabled, multicast counters statistics creation/deletion fails and the following errors might occur during LAG member enable/disable on QFX51xx devices. The messages do not indicate the traffic impact. However, the multicast statistics will not work when these messages are seen. **Feb 15 07:28:49 switch fpc0 brcm_ipmc_get_multicast_stats:3947 brcm_ipmc_stat_get failure**
Feb 15 07:28:49 switch fpc0 brcm_rt_stats:1906 brcm_ipmc_get_multicast_stats failure err=-7. [PR1392470](#)
- Autonegotiation errors and flush operation failed errors are seen after power cycle of the device. These error messages do not have any functionality impact. **LOG: Err] ifd 153; Ether autonegotiation error (1000) and ch_vchassis_ipc_flush_pipe: flush operation failed for pipe 155333280.** [PR1394866](#)
- The error message **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(),128:l3 nh 6594 uninstall failed** is seen in hardware with mini-PDT-base configurations. There is no functionality impact because of this error message. [PR1407175](#)

SEE ALSO

[New and Changed Features | 276](#)

[Changes in Behavior and Syntax | 296](#)

[Known Behavior | 302](#)

[Resolved Issues | 309](#)

[Documentation Updates | 326](#)

[Migration, Upgrade, and Downgrade Instructions | 326](#)

[Product Compatibility | 339](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 17.1R3 | 309](#)
- [Resolved Issues: 17.1R2 | 321](#)
- [Resolved Issues: 17.1R1 | 323](#)

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

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

Resolved Issues: 17.1R3

Class of Service (CoS)

- On QFX5100, traffic might get dropped when there is more than one forwarding class under **forwarding-class-sets**. [PR1255077](#)
- Storm control might not be programmed correctly in the Packet Forwarding Engine if it is applied with a **port-speed** configuration in a single commit. [PR1255562](#)
- The transmit rate applied with **forwarding-class-set** does not work properly. [PR1277497](#)
- Firewall filter cannot filter packets with DstIP as 224/4 and DST MAC = QFX_intf_mac on loopback interface using a single match condition for source address 224.0.0.0/4. [PR1354377](#)

EVPN

- The `expr_nh_fwd_create_arp_ndp_egress_descr()`, 1237:nh 131650 type Compst, failed to create L2 description failure log message is seen, but there is no impact on traffic or performance. [PR1221831](#)
- The error message `JPRDS_DLT_ALPHA KHT` shows as failed, but the entries in hardware are programmed correctly. [PR1258933](#)
- The fxpc and kernel crash might be observed after adding MTU configuration on QFX5000 Virtual Chassis platform. [PR1283966](#)
- VXLAN license might be invalid if license QFX-ADV-FEATURE-LIC is installed. [PR1288916](#)
- The dynamic routing protocols might not work correctly over the IRB interface in an EVPN-VXLAN scenario with ECMP. [PR1301521](#)
- VXLAN traffic loss is observed after deleting and adding VLANs. [PR1318045](#)
- The remote ARP entry might be incorrect in an EVPN and VXLAN Layer 3 gateway scenario with multihoming mode. [PR1326691](#)
- The MAC movement between remote VTEP and local VTEP might cause traffic to transmit incorrectly in an EVPN-VXLAN scenario. [PR1335431](#)
- In a redundant Layer 3 gateways Ethernet Virtual Private Network (EVPN)/Extensible Local Area Network (VXLAN) scenario on QFX10000 Series switches, when an IP address move occurs (the same IP address, but the media access control is changed), the ARP entry might be deleted from one Layer 3 gateway device, which might cause a few packets to be lost. [PR1336185](#)
- The rpd process generates an unreproducible core file with scaling EVPN-VXLAN configuration on QFX10000 platform because of the memory depletion on EVPN MAC route entries queue for l2ald . l2ald closes the IPC connection that caused the rpd-cumulated EVPN MAC route entries in the queue and ends up running out of memory. [PR1339979](#)
- On QFX5000 and QFX10000 platforms in an EVPN-VXLAN scenario, the VXLAN proxy ARP/NDP suppression might result in incorrect learning of Virtual gateway MAC addresses. [PR1367610](#)

Forwarding and Sampling

- Unexpected messages might be seen in logs. [PR1270686](#)

General Routing

- The 40-Gigabit Ethernet connection between two QFX5100-24Qs might not come up sometimes. [PR1178799](#)
- A major alarm **Host 0 CPU Temperature Hot** is observed. [PR1241744](#)
- An FPC major alarm might be seen with the following error messages **DLU: ilp memory cache error** and **DLU: ilp prot1 detected_imem_even error**. [PR1251154](#)
- MACsec session fails with dot1x generating a core file. [PR1251508](#)

- In QFX5100, the following multicast statistics counter-related following error messages are observed after LAG interface disable/enable: **brcm_rt_ip_mc_ipmc_deinstall**., **brcm_ipmc_route_counter_delete**, **brcm_ipmc_stat_get**, **brcm_ipmc_get_multicast_stats**. [PR1255497](#)
- On QFX Series, license keys entered through the configuration **system license keys** can be lost (not effective anymore) after certain events/changes. [PR1259460](#)
- SFP-T equipped port does not link up properly at booting up when the port has the speed 100m and link-mode full-duplex setting. [PR1262752](#)
- Random interfaces do not come up after a line card is rebooted. [PR1262839](#)
- QFX100002 generated an L2ALD core file for an unknown reason at:
l2ald_mac_process_update_fwd_entry_mask, **l2ald_mclag_update_change_for_learn_mask**, **logging**, **vlogging**, **vlogging_event**. [PR1264432](#)
- In Junos OS environment, after execution of **<rpc> get-configuration-compare-"rollback" rollback-"0"**, the management daemon (MGD) might restart unexpectedly. The MGD restart also causes connections through ssh or console to drop. [PR1271024](#)
- The jdhcpd process might crash and DHCP does not work if scaling prefixes are configured under the **policy-options prefix-list** hierarchy. [PR1272646](#)
- The 40-Gigabit Ethernet interface might flap between QFX5100 and other products. [PR1273861](#)
- When static link protection mode is configured with backup state as down, the primary port is going to down state instead of the secondary port remains up. [PR1276156](#)
- On QFX Series platforms where MC-LAG with IPv6 is supported, the l2ald memory might leak for every IPv6 neighbor discovery (ND) message it receives from a peer MC-LAG and it does not free the memory allocated, causing l2ald memory exhaustion and an l2ald process crash. [PR1277203](#)
- Multicast listener discovery (MLD) messages are seen continuously on QFX5100 when the management ports are connected through a network. [PR1277618](#)
- MAC pause frames might increase when SXE interfaces are erroneously configured. [PR1281123](#)
- In a MACsec scenario, the **show security macsec statistics** command does not show expected results. [PR1283544](#)
- In 802.1X (dot1x) single-supplicant mode, after username and password are configured on interfaces and dot1x supplicants are started, the users are authenticated with the Radius_DataVlan VLAN, but the Ethernet-switching table is not updated for one of the interfaces. [PR1283880](#)
- After upgrading the QFX5100 to Junos OS Release 16.1 or later releases from Junos OS Release 15.1, a commit warning **/boot/ffp.cookie+** might be seen. [PR1283917](#)
- On QFX5100 switches, an aggregated Ethernet interface might flap upon commit if an explicit speed is configured on an aggregated Ethernet member interface. [PR1284495](#)
- BFD sessions might flap when BFD is configured over IRB interfaces. [PR1284743](#)
- Analytics JSON data format is reporting an incorrect value for 'rxbps' counter. [PR1285434](#)

- The 1-Gigabit copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. [PR1286709](#)
- OVSDB and Openflow have some limitations on 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](#)
- 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](#)
- On QFX5100, an incorrect alarm type might be displayed. [PR1291622](#)
- Traffic might not be received on a 1-Gigabit Ethernet interface if autonegotiation is disabled and speed or duplex is configured on both the QFX Series switch and the peer host. [PR1292275](#)
- On QFX5100 switches with EVPN-VXLAN deployed, broadcast and multicast traffic might not be sent to other switches through VTEP interfaces. [PR1293163](#)
- On QFX5100, the fxpc process generates a core file. [PR1294033](#)
- High heap memory utilization might be seen if multiple SFP-T optics are inserted or **set interface <> link-mode full-duplex** is enabled. [PR1294208](#)
- For ULC-60S-6Q LC on QFX10008, the port becomes unusable after inserting a third-party SFP-T optic. [PR1294394](#)
- The received ARP reply packet whose destination MAC address is the same as the MAC address of the IRB interface might be flooded on the VLAN. [PR1294530](#)
- The 40-Gigabit Ethernet interface might not come up if a specific vendor's DAC cable is used. [PR1296011](#)
- Network analytics process might be incorrect instantiated leading to traffic statistics not being transmitted. When this occurs the 'Sent' value for **show analytics collector** displays as zero and **show analytics traffic-statistics** will be empty: `root@host> show analytics collector Address Port Transport Stream format State Sent 10.10.10.72 50020 udp json n/a 0 10.10.10.167 50020 udp json n/a 0 root@host> show analytics traffic-statistics` CLI issued at 2018-03-26 22:15:56.411671. [PR1297535](#)
- On QFX Series platforms with ZTP environment, the DHCP clients are not getting an IP address with /31 subnet in server configuration. [PR1298234](#)
- Disabled 10-Gigabit Ethernet interfaces might stay up on QFX10000 line switches. [PR1300775](#)
- In QFX10008 and QFX10016, a commit error is seen when mixed speeds are configured. [PR1301923](#)
- The rpd might crash when toggling the **vrf-propagate-ttl** and **no-vrf-propagate-ttl** configuration statements. [PR1302504](#)
- The sFlow records are missing "extendedType ROUTER" fields as well as an outbound interface for traffic that is using BGP multipath. [PR1303236](#)
- When MPLS LSP self-ping is enabled (self-ping is enabled by default), the kernel might panic with an error message **Fatal trap 12: page fault while in kernel mode**. [PR1303798](#)

- Platforms running 32-bit Junos OS might generate an rpd core file when traceoptions are enabled. [PR1305440](#)
- QFX5100 crashes and the fxcp process generates a core file. [PR1306768](#)
- Some error messages can be observed on EVPN-VXLAN setup. [PR1307014](#)
- Run time pps statistics value might show zero for a subinterface of aggregated Ethernet interface. [PR1309485](#)
- Traffic loss might be seen if sending traffic through the 40-Gigabit Ethernet interface. [PR1309613](#)
- The FPC memory might be exhausted with SHEAF leak messages seen in the syslog. [PR1311949](#)
- A traffic loss is observed while performing NSSU. [PR1311977](#)
- On QFX Series standalone switches or their Virtual Chassis with dot1x configured, there will be memory leaks for PNACAUTH in dot1xd. Once the memory block of PNACAUTH used by dot1xd grows to its maximum size, the switch might not process the client's authentication further and results in dot1x clients reauthenticating constantly. The dot1xd process always runs irrespective of configuration, and as part of its initialization it tries connection with authd. If authd is not running, then there is a memory leak in dot1xd. [PR1313578](#)
- Transit traffic over GRE tunnel might hit the CPU and trigger a DDoS violation on the Layer 3 next hop. [PR1315773](#)
- Packet Forwarding Engine might crash after changing analyzer configuration if output includes LAG interface. [PR1316245](#)
- On an Layer 2 next-generation switch platform (QFX5100 and QFX10000), l2cpd might generate core files repeatedly if an interface is connected to a VoIP product with LLDP and LLDP-MED enabled. [PR1317114](#)
- Packets such as TDLS without IP header are looped between the virtual gateway. [PR1318382](#)
- 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 Series Virtual Chassis unified ISSU. [PR1320370](#)
- FPCs go offline because of the error **CHASSISD_IPC_CONNECTION_DROPPED: Dropped IPC connection for FPC**. [PR1321198](#)
- The openflow session cannot be established correctly with controller and interfaces options configured on QFX5100 switches. [PR1323273](#)
- VLAN or VLAN bridge might not be added or deleted if there is an IFBD hardware token limit exhaustion. [PR1325217](#)
- Deleting one VXLAN might cause traffic loop on another VXLAN in a multihoming EVPN-VXLAN scenario with service provider style interface. [PR1327978](#)

- After IP address move, ARP table information is not in synchronization between the two spines. [PR1330663](#)
- The rpd process generates a core file on the new backup Routing Engine at **task_quit**, **task_terminate_timer_callback**, **task_timer_dispatch**, and **task_scheduler** after disabling NSR and GRES. [PR1330750](#)
- The analyzer status might show as down when port mirroring is configured to mirror packets from an aggregated Ethernet member. [PR1338564](#)
- DDoS counters for OSPF might not increase. [PR1339364](#)
- l2ald process generates a core file at `../..../src/junos/usr/sbin/l2ald/l2ald_vxlan_evpn.c:1603` when moving host between two multihop interfaces. [PR1339543](#)
- In ovsdb vxlan network QFX5100 broadcast Layer 2 traffic gets forwarded to the same receiving ports. [PR1342637](#)
- FXPC process might generate a core file when removing VXLAN configuration. [PR1345231](#)
- On any platform that does not clear out `/mfs` when installing a new software release such as EX Series or QFXSeries platform, when upgrading from certain releases to Junos OS Release 18.1R1 the statistics of the pfd process might generate a core file. The issue does not impact the service. [PR1346925](#)
- From Junos OS Releases 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R3, 17.2X75-D90, 17.3R2, 17.4R1, 18.1R1, 18.1X75-D10, 18.2X75-D5, and later releases, QFX5100-48T 10G interface might be autonegotiated at 100-MB speed instead of 10 Gbps after peer device reboot. [PR1347144](#)
- The pfd process consumes 80-90 percent CPU when running subscriber management on PPC-based routers and switches. [PR1351203](#)
- On QFX5000 switches, the Packet Forwarding Engine might drop the ARP reply packets after changing the interface MAC address. [PR1353241](#)
- A major error **PE Error code: 0x2104be** is observed. [PR1354582](#)
- Commit error is observed if the device is downgraded from Junos OS Releases 18.2 or 18.3 to Junos OS Release 17.3R3. On loading the new image, certain stale symlinks from previous image contents need to be removed, which impacts mgd. In this case, the `.slax` script symlinks from `/var/db/scripts/translation` are not getting removed, which causes issues in the initial commit by mgd. The issue is only seen when the previous image was having translation scripts (as part of Junos image) and the new image does not have these translation scripts. [PR1355542](#)
- When rpd reads next hops from the kernel on restart, for INH -> FWD NH{List NH} -> {Chain NH} scenario, the rpd should not create an old-style list next hop for the forwarding next hop. [PR1360354](#)
- On QFX5100VC, the VME interface might be unreachable after link flap of em0 on the master FPC. [PR1362437](#)
- Even if **no-auto-negotiation** is not configured autonegotiation is off default if **gigether-options** autonegotiation is mismatched between the link and its partner, the 1-Gigabit Ethernet interface might stop working. [PR1362977](#)

- On QFX5100 Virtual Chassis/Virtual Chassis Fabric, while doing a unified ISSU from Junos OS Releases 15.1R7 to 16.1R7, the LAG interface might flap. This might result in traffic loss of more than 5 seconds depending on how fast the LAG interface recovers. [PR1365316](#)
- On QFX10000 platforms, the Junos OS boot menu cannot appear because Ctrl+c does not give the menu during the boot process. Root password recovery option might not be available. [PR1365740](#)
- The l2cpd process might crash if MVRP is configured and RSTP is enabled with the statement **interface all**. [PR1365937](#)
- On QFX5000 switches in a Virtual Chassis/Virtual Chassis Fabric scenario, the chassisd might crash after issuing the CLI **show chassis hardware**. This might result in VCP down and traffic drop. [PR1366746](#)
- On QFX5000 switches, when an IS-IS packet is received with DMAC as 09:00:2b:00:00:05 (ISO 9542, all Intermediate System Network Entities Address) and jumbo frame with EtherType as 0x8870 (non-standard, used by Cisco), the packet will be dropped, resulting in failure in the adjacency. [PR1368913](#)
- On QFX5000 Series platforms, performing optics insertion/removal on a port might result in the Packet Forwarding Engine manager CPU spike and eventually microcode failure. [PR1372041](#)
- When VRRP is enabled on an interface, when the interface is disabled and then enabled, the IPv6 routed packet might be transmitted over VRRP virtual IP address. The IPv6 routed packet VRRP state is in the non-master state. When this happens, the peer interface might return to normal later than this interface. At this time, the packets sent out through this interface might be dropped. [PR1372163](#)
- On QFX Series platforms, if RTG redundant trunking group (RTG) is enabled with a large-scale the MAC address, the MAC refresh frame might not be sent out from the new primary link after RTG failover by deactivating the former primary link on the peer side. [PR1372999](#)
- On the QFX5100 platform, the auto-negotiation interface might go down if the peer device supports only 10-MB or 100-MB autonegotiation. [PR1377298](#)
- Debug logs are printed as error logs in **/var/log/** messages. The debug log message, **expr_nh_flapel_check_overwrite: Caller nh_id params** is classified as an error log when it should be LOG_INFO. [PR1377447](#)
- On QFX10000 platforms, the L3VPN traffic might be dropped if one core-facing interface goes down in an L3VPN multipath scenario. [PR1380783](#)
- On QFX5000 platforms, the Packet Forwarding Engine might show **DISCARD next-hop for overlay-bgp-lo0-ip**. [PR1380795](#)
- In an Open vSwitch Database (OVSDb) environment with solid-state drive (SSD) installed on the backup Routing Engine side, the master Routing Engine copies **/var/db/ovsdatabase** to the backup Routing Engine in a very short interval (for example, every 10 seconds), and the backup Routing Engine might write the whole OVSDb file to the SSD card frequently. Therefore, the SSD lifetime might be shortened because of the exceeded amount of read/write. Because of this issue, SSD card failure might be observed. [PR1381888](#)
- QFX Series switches might not be able to establish a complete LACP session ("collecting/distributing") depending on the configuration of the QFX Series interface. If an interface has **native-vlan-id** configured

and that same **native-vlan-id** VLAN is in vlan members list and the VLAN is VxLAN enabled, then QFX Series switches stop processing received LACP PDUs. [PR1382209](#)

- Because of an API introduced in Junos OS Release 18.1R3, a kernel might generate a core file when a configuration change is done. This results from invalid pointer access by the API. [PR1384750](#)
- When DDoS configurations for Virtual Chassis are initialized, DDOS_POL_FLAGS_ASIC is not set. [PR1387508](#)
- The sdk-vmmd might consistently write to the memory. [PR1393044](#)
- 10G copper link flapping might happen during TISSU operation of QFX5100-48T switches. [PR1393628](#)
- The **show chassis fpc** command displays an incorrect amount of available memory on a QFX10000's FPCs. [PR1394978](#)
- If GRES/NSR is enabled on a QFX5100 (single Routing Engine), DHCP subscribers are failing to bind. [PR1396470](#)
- Persistent MAC entries are cleared after system reboot. [PR1400507](#)
- A single Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- PEM alarm for backup FPC will be remained on master FPC though backup FPC is detached from VC. [PR1412429](#)

High Availability (HA) and Resiliency

- Line card reboots after GRES. [PR1286393](#)

Infrastructure

- The **show interface** command is not returning any values and sometimes it gets completely stuck. [PR1250328](#)
- When 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](#)
- VMcore generates a core file because of mbuf leak. [PR1261996](#)
- The QFX5100 switch might be sending a packet with an incorrect destination MAC address in an MPLS PHP scenario. [PR1334929](#)

Interfaces and Chassis

- Deactivation followed by activation of both aggregated Ethernet and MC-AE interfaces blocks the flow of multicast traffic. [PR1257586](#)
- Multicast data packets are looping in MC-LAG. [PR1281646](#)
- ARP reply drops in MC-LAG scenario. [PR1282349](#)
- Upgrading might encounter commit failure if **redundancy-group-id-list** is not configured under ICCP. [PR1311009](#)

- On QFX5000 platform, if the ICL link is configured on a single interface (such as GE-0/0/0, without LAG) and one member of MC-LAG is down, and both MC-LAG peers are rebooted, packets might drop on the ICL of the MC-LAG peer where MC-LAG is up. [PR1345316](#)
- CVLANs range of 16 might not pass traffic in a Q-in-Q scenario. [PR1345994](#)
- On QFX5000 switches, MC-LAG peer might not send ARP request to the host. [PR1360216](#)
- When l2cpd daemon is restarted, **parse_remove_ifl_from_routing_inst() ERROR : No route inst on et-0/0/16.16386**, errors are seen. [PR1373927](#)

Layer 2 Ethernet Services

- The jdhcpd process generates a core file after making DHCP configuration changes. [PR1324800](#)
- If BOOTP-support is not enabled at the global level, bootstrap protocol (BOOTP) packets might be dropped while receiving them on an interface because there is a defect that the device only checks BOOTP-support at the global level. [PR1373807](#)

Layer 2 Features

- Action-shutdown in storm-control does not bring the physical interface down. [PR1240845](#)
- Interface with **vlan-tagging** and **family ethernet-switching** configuration does not work on QFX10000 platforms. [PR1261915](#)
- Device transmits packets that exceed the interface MTU. [PR1306724](#)
- The **bpdu-block-on-edge** command does not work correctly when **fast-tune** is enabled. [PR1307440](#)
- ARP entry might be learned on STP blocking ports. [PR1324245](#)
- The DHCP discover packets might be looped in MC-LAG and DHCP-relay scenario. [PR1325425](#)
- Interface with **flexible-vlan-tagging** and **family ethernet-switching** does not work on QFX10000. [PR1337311](#)
- On random initialization of QFX5100 the programming of the storm control profile is missed within hardware on random interfaces. This is not visible over the CLI and the configuration still shows as intact. This happens as a result of interface speed not getting properly detected within the hardware. [PR1354889](#)
- On QFX5100, if **native-vlan-id** is configured for the aggregated Ethernet interfaces, after having a reboot, LACP packets might be dropped. [PR1361054](#)
- On QFX5000 switches, IPv6 traffic over VxLAN tunnel does not hash. This might result in some unexpected issue in an ECMP scenario. [PR1368258](#)
- When native-vlan-id is configured for the aggregated Ethernet LACP session to the multihomed server goes down if you have irb.0 configured. This causes incorrect parameters to be pushed to Packet Forwarding Engine causing LACP PDUs to not egress correctly. [PR1369424](#)
- On QFX5000/EX4600 platforms, if changing an interface from Virtual Extensible Local Area Network (VXLAN) to a member of an aggregated Ethernet (AE) interface, the Dynamic Host Configuration Protocol (DHCP) relay would not work and the DHCP client would not get IP addresses normally. [PR1377521](#)

- On EX4300, EX4600, and QFX Series switches (except for QFX10000). In a VLAN service provider style scenario (for example, **flexible-vlan-tagged** under the **[interfaces]** hierarchy), after the egress interface or logical interface is disabled/deactivated/deleted. The switch might continue forwarding Layer 2 traffic on the interface. [PR1379258](#)
- If an aggregated Ethernet interface is configured with LACP, "flexible-vlan-tagging" and "native-vlan-id", then after deleting the "native-vlan-id option", the LACP state will be detached state. [PR1385409](#)
- On QFX Series switches except for QFX10000, in a Virtual Chassis and RTG scenario, if the redundant trunk group (RTG) interface flaps on the Virtual Chassis master, RTG MAC refresh packets are sent out from all the ports that belong to the same VLAN. The MAC refresh packets are used to refresh MAC entries on the peer Layer 2 device connected to the RTG ports. [PR1389695](#)
- A deadlock situation between pfeman thread and Broadcom's linkscan thread causes watchdog trigger and results in generating a dcpfe core file. The issue is seen during the port initialization stage. [PR1398251](#)

MPLS

- In QFX5100, a unified ISSU is not supported with MPLS configuration. [PR1264786](#)
- DHCP clients cannot get IP addresses over BGP-L3VPN. [PR1303442](#)
- LSP stop transferring or passing traffic after MPLS route is changed. [PR1309058](#)
- The rpd might crash on backup Routing Engine because of the memory exhaustion. [PR1328974](#)
- The hot standby for l2circuit does not work on QFX5100. [PR1329720](#)
- In an RSVP scenario, the label-switched path (LSP) might remain UP even if no path is acceptable, because of the constrained shortest path first (CSPF) failure. There are two scenarios which might result in CSPF failure.

Scenario 1 with MBB: Optimization timer fires during make-before-break (MBB).

Scenario 2 without MBB: A link/IGP flap causes CSPF, but it depends on timing. [PR1365653](#)

- On all QFX5000 platforms, if the P/PE router is configured with **no-decrement-ttl**, the rpd sends the NO_PROPAGATE_TTL flag even for the tunnel transit case. [PR1366804](#)

Network Management and Monitoring

- The mib2d syslog messages **MIB2D_RTSLIB_READ_FAILURE: rtllib_iflm_snmp_pointchange** might be seen while removing and restoring configuration. [PR1279488](#)

Platform and Infrastructure

- Dropping the TCP RST packet incorrectly on the Packet Forwarding Engine might cause traffic drop. [PR1269202](#)
- In a Virtual Chassis scenario, when the master member FPC reboots and the interface on which the ARP is learned goes down along with the master FPC, traffic loss might be observed for about 10 seconds. At that time, the ARP entry cannot be learned from the remaining FPC. [PR1283702](#)

- The dexp process might crash after committing **set system commit delta-export**. [PR1284788](#)
- OSPFv3 authentication using IPsec SA does not work if you are using IPsec to authenticate OSPFv3 neighbors on some QFX Series platforms. [PR1301428](#)
- The Virtual Chassis Fabric (VCF) switch is not sending the common technology (Tri speed) with 10Base-T or 100Base-T when negotiating with series devices. Instead it is sending only 1000 negotiation because the QFX Series switch is the master in the Virtual Chassis Fabric. . [PR1311458](#)
- 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](#)
- When a Junos OS image is shipped with translation scripts downgrading to another image, stale symlinks of translation scripts at the time of mgd initialization lead to the device going into amnesiac state. [PR1341650](#)
- Traffic drops occur on the Packet Forwarding Engine as "invalid L2 token" when protocol changes from VPLS to EVPN. [PR1368802](#)
- On QFX Series switches except for QFX10000, pass-through traffic might be dropped while using multiple routes with indirect next hop and load balancing. [PR1376057](#)
- On Virtual Chassis based on QFX5100 switches, the IRB interface associated with aggregated Ethernet interfaces whose member interfaces are only from the master chassis might not turn down when the master chassis is rebooted or halted. [PR1381272](#)

Routing Policy and Firewall Filters

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

Routing Protocols

- The fxpc process might crash and restart when the fxpc process tries to access already-freed-up memory. [PR1271825](#)
- IPv6 packets depending on IPv6 link-local might be lost on channelized interfaces on QFX5100. [PR1283065](#)
- Message **dc-pfe: list_destroy()** is printed on commit. [PR1286209](#)
- GRE tunnel traffic does not switch over to the alternate path if the primary path to the tunnel destination changes. [PR1287249](#)
- FBF with next-ip/next-ip6/next-interface is not working. [PR1289642](#)
- 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](#)
- IPv6 multicast traffic drop occurs in a PIM SSM scenario. [PR1292519](#)
- The mcsnood process generates a 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](#)

- When routes are leaked between routing instances it might be possible for a route to become invalid (reject route) but for this update to not propagate to all routing instances. This issue will eventually lead to the routing table in the Packet Forwarding Engine to become full, which will prevent additional valid routes from being properly installed. [PR1307009](#)
- Packet drop is seen while programming for GRE traffic. [PR1308438](#)
- Some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update. [PR1320723](#)
- The IS-IS Layer 2 hello packets are dropped when they come from a Brocade device. [PR1325436](#)
- The loopbacked IRB interface is not accessible to the remote network. [PR1333019](#)
- QFX Series loopback firewall filter is not able to catch packets with a Martian source address. [PR1343511](#)
- On QFX5100 platforms, if a firewall filter term's action has policer configurations and if the Packet Forwarding Engine command `show filter hw <index>` and `show_terms_brcm` is issued, policer errors are observed. **ERROR (dfw): Unable to create policer ERROR (dfw): brcm_dfw_handle_plcr_cntr_action ()returned error ERROR (dfw): Setting brcm action failed!.** This issue occurs because the following vty show command is issued: `show filter hw <index> show_terms_brcm`. Note that this is an internal Packet Forwarding Engine verification command. Each time the command is executed, a firewall entry in hardware is deleted and this in turn is causing the DFW error logs. [PR1336137](#)
- On QFX5000 Series switches, a high rate of Ethernet pause frames or an ARP packet storm received on the management interface (fxp0) might cause egress interface congestion, resulting in routing protocol packet drops, such as BGP, leading to peering flaps. Refer to <https://kb.juniper.net/JSA10888> for more information. [PR1343597](#)
- On QFX5000 platforms, firewall filter configuration cannot perform packet matching on any IPv6 extension headers. This issue might allow IPv6 packets that should have been blocked to be forwarded. IPv4 packet filtering is unaffected by this vulnerability. See <https://kb.juniper.net/JSA10905> for details. [PR1346052](#)
- On QFX5100 platforms, the device might get into an improper state in which it is unable to correct parity errors in the Packet Forwarding Engine memory. Traffic might silently drop and get discarded for specific destination IPs. [PR1364657](#)
- On QFX5100 switches, when the `switch-options no-arp-trap` statement is configured, the unicast ARP packets that are not destined to the switch-routed interfaces might cause traffic to be transmitted incorrectly or traffic failure because of ARP resolutions failure. [PR1369903](#)
- 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 might not be dropped and then reach the Routing Engine unexpectedly. [PR1379718](#)

- If a QFX5100 device has a host route with equal-cost multipath (ECMP) next hops and receives a better path with a single next hop then the next hop in hardware will not be changed. [PR1387713](#)
- The rpd process might generate a core file when L2VPN is used. [PR1398685](#)

Virtual Chassis

- In a QFX Series Virtual Chassis after first Routing Engine switchover with NSR enabled scenario, BGP adjacency flapped because 20 seconds of traffic loss is observed. Subsequent switchovers work fine. [PR1225829](#)

Resolved Issues: 17.1R2

EVPN

- Route Target per bridge domain for EVPN is not supported. [PR1244956](#)
- On QFX10000/QFX5100 Series with VXLAN/EVPN configured, when multiple IP addresses are configured for VTEP source interface, traffic might be dropped on spines. [PR1248773](#)

Hardware

- QFX10008: After the reboot of 30X100G line card and 36X40G line card with traffic running, a large amount of framing errors are observed. [PR1223330](#)
- QFX10008 and QFX10016: 60x10G ULC 1G mode is not supported in Junos OS Release 17.1R1. [PR1239091](#)
- SFP-T in QFX5100-48S-6Q does not work at 100M full duplex in Junos OS Releases 14.1X53-D35 and later (it works in Junos OS Release 14.1X53-D30). [PR1250453](#)

High Availability (HA) and Resiliency

- ISSU to 17.1R1 from earlier releases is not supported on QFX5100 and EX4600. [PR1255878](#)

Interfaces and Chassis

- Backup links are not carrying traffic when the primary link is disabled on an aggregated interface. [PR1208614](#)
- The traffic might not be transmitted correctly after a logical interface is deleted from one VLAN and added to another VLAN on EX9200, EX4300, QFX Series switches. [PR1228526](#)
- Removal or insertion of a transceiver for a port in a LAG, which is part of scaled VLAN members may cause protocol flap. [PR1229547](#)
- FPC reloads unexpectedly during port speed change from 100G to 40G default. [PR1256267](#)

Layer 2 Features

- Incorrect statistics might be shown for an AE interface after rebooting a device or clearing interface statistics. [PR1228042](#)
- If RTG and VSTP are configured on the same VLAN, communication doesn't work over RTG interfaces. [PR1230750](#)
- DHCP offer packets (with MPLS header) are getting dropped on ingress of QFX10000 switches; DHCP relay running on VRF. [PR1243936](#)
- QFX10000: IPv6 double tag frame does not pass through QFX10000 switches if a service provider style configuration is used. [PR1254492](#)
- S-Link macs are not moving across MC-LAG chassis on QFX10000 switches. [PR1260316](#)
- The BUM traffic from ESI peer might be transmitted to CE interface after deleting and adding VLAN in a VXLAN/EVPN multihoming scenario. [PR1260533](#)
- QFX5100 does not transfer BPDU packets even though xSTP is disabled. [PR1262847](#)

MPLS

- LSP traffic loss occurs after changing chained-composite-next-hop configuration. [PR1243088](#)

Network Management and Monitoring

- IPv6 packets/bytes counter show higher value than the total packets/bytes of the interface if the LAG child members belong to the same PE device. [PR1232388](#)
- SNMP trap messages about FRU power off might be seen even though the power supply is working fine. [PR1233537](#)
- When the MAC age timer is longer than the ARP age timer, after the ARP timer ages out both MAC and MAC+IP get advertised by all ESI peers regardless of who learns locally. [PR1238718](#)
- Users may lose the sFlow configuration when they upgrade to Junos OS Release 17.1R1 from Junos OS Release 15.1X53-D6x. Also, when they downgrade to Junos OS Release 15.1X53-D6x from Junos OS Release 17.1R1, the downgrade may fail. [PR1240804](#)
- sFlow may show a negative count for a number of samples after a long run. [PR1244080](#)

Platform and Infrastructure

- Protocol flapping and an RE-FPC TCP connection drop are seen on Virtual Chassis setups during image copy using SCP. [PR1213286](#)
- QFX10002: **show chassis fpc** shows the wrong number of slots. [PR1219853](#)
- High latency/jitter might be seen while trying to ping the IP address of a switch. [PR1221053](#)
- The rpd process might crash and restart when a MAC address is learned from a given PE on a different ESI. [PR1247338](#)

- On QFX5100, **show interface** incorrectly displays an interface as **Link-mode: Auto Speed: Auto** even though the interface is configured for, and up at, 100M/Full. [PR1260986](#)
- QFX5100 VCF: Removing force-up causes return-traffic to be dropped by leaf (to spine). [PR1264650](#)
- Description for 40G-AOC cable in **show chassis hardware** shows **UNKNOWN**. [PR1269018](#)

Routing Protocols

- RA packet might not be sent when igmp-snooping is configured for VLAN. [PR1238906](#)
- Layer 3 interface (inet family) is not supported as upstream port in multicast route leaking. [PR1250430](#)
- QFX10008 and QFX10016: While flapping random LAG interfaces with 448 LAG scale, you can see other LAG interfaces getting flapped. [PR1250741](#)
- After running **restart routing** in the master RE, the PIM join states of VXLAN multicast groups in the backup RE are not in sync with the master RE. [PR1255480](#)
- VCF doesn't forward BUM traffic after fabric-tree-root is configured. [PR1257984](#)
- VRRP with MD5 authentication and OSPF3 packets with IPsec do not go the proper host path queue and can cause flapping. [PR1258501](#)
- On a QFX5100 switch, TCP packets with destination IPv6 as link-local address and destination port 179 are dropped in the Packet Forwarding Engine. [PR1267565](#)
- IPv4 traffic drops when changing the member interface of the LAG. [PR1270011](#)

Software Installation and Upgrade

- After upgrading a QFX10000 switch from Junos OS Release 15.1X53-D62 to Junos OS Release 17.1R1, the **vrf-target export target: community-name** configuration statement might be missing from the [protocols evpn vni-options vni] hierarchy level. To work around this issue, you must add the missing configuration statement back to the [protocols evpn vni-options vni] hierarchy level. [PR1243105](#)

Resolved Issues: 17.1R1

High Availability (HA) and Resiliency

- The AE interface might be down after NSSU is done on QFX5100 or EX4600 switches. [PR1227522](#)
- QFX5100 : When network analytics feature is configured, TISSU might fail and cause the generation of fxpc core file. [PR1234945](#)
- ISSU to Junos OS Release 17.1R1 from earlier releases is not supported on QFX5100 and EX4600. [PR1255878](#)

Interfaces and Chassis

- Users may see the error message **expr_cos_rw_nh_qix_get @ 150: Unable to get chip num for ill:994** on mc-ae status-control active node upon sending an ARP request. These messages are for information only and have no functional impact on the operation of QFX10008/QFX10016. [PR1228080](#)

- CDP packets looping with MC-LAG on QFX10000 switches. [PR1237227](#)

Layer 2 Features

- Unable to assign VLAN to an interface after error message **IFBD hw token couldn't be allocated for** is output. [PR1216464](#)
- Incorrect statistics might be shown for an AE interface after rebooting device or clearing interface statistics. [PR1228042](#)
- The fxpc process can generate a core file on QFX5100. [PR1231071](#)
- MAC learning is very slow when clearing MAC addresses in cases of scale MAC learning (128k). [PR1240114](#)
- DHCP offer packets (with MPLS header) are getting dropped on ingress of QFX10000, DHCP relay is running on VRF. [PR1243936](#)

MPLS

- The fxpc crash observed on the switches. [PR1168150](#)
- VC/VCF-I2ckt: FXPC core is seen when deactivating core interface on MPLS I2ckt configuration using IRB interface. [PR1242203](#)

Network Management and Monitoring

- In some cases under heavy logging SD logger messages which report critical events such as daemon restarts are not seen on the aggregator. [PR1239667](#)

Platform and Infrastructure

- Protocol flapping and RE-FPC TCP connection drop seen on VC setups during image copy using scp. [PR1213286](#)
- A high latency/jitter might be seen while trying to ping the IP address of a switch. [PR1221053](#)
- On QFX10000 switches, there is a 4-second delay seen in 40g ports to come up QSFP+-40G-LR4. [PR1219336](#)
- A pfed core file is observed after deleting apply-groups from the configuration. [PR1223847](#)
- The alarm message **Management Ethernet Link Down** might be seen on QFX Series switches. [PR1228577](#)
- On QFX10002 switches, when a USB device is inserted into the switch, field-replaceable unit (FRU) insertion messages such as **RE0 & ?CAMGETPASSTHRU ioctl failed cam_lookup_pass: Inappropriate ioctl for device?** may be displayed. These FRU insertion messages do not affect service and stop after the USB device is removed. [PR1233037](#)
- SNMP trap messages about FRU power off might be seen even though the power supply is working fine. [PR1233537](#)
- The **show interface interface media** command shows the media type for the SFP-T to be fiber. [PR1240681](#)

- The rpd process might crash and restart when a MAC address is learned from a given PE on a different ESI. [PR1247338](#)
- Network ports are not detected on a QFX10002 switch after a reboot. [PR1247753](#)
- On QFX10000 switches, internal comments can be seen in the configuration file after loading the factory default. [PR1248434](#)
- Traffic is dropped on spines in some VXLAN/EVPN scenarios. [PR1248773](#)
- SFP-T in QFX5100-48S-6Q does not work at 100M full duplex in Junos OS Releases 14.1X53-D35 and later (it works in Junos OS Release 14.1X53-D30). [PR1250453](#)

Routing Protocols

- EBGP packets with ttl=1 and non-EBGP packets with ttl=1 go to the same queue. [PR1227314](#)
- The action "reset" is not working for FPC resiliency (fault handling). [PR1233075](#)
- FPC restarts with a dcpfe core. [PR1236046](#)
- Hops through GRE tunnel endpoints are seen in traceroute. [PR1236343](#)
- Packet drop is seen when routing process is restarted, even when graceful restart is configured. [PR1239186](#)
- Kernel crashes in the chassis after FPC reset. [PR1242362](#)
- GARP reply packets are not updating the ARP table. [PR1246988](#)
- Layer 3 interface (inet family) is not supported as upstream port in multicast route leaking. [PR1250430](#)

Virtual Chassis

- VCF not communicating properly with backup spine. [PR1141965](#)

SEE ALSO

[New and Changed Features | 276](#)

[Changes in Behavior and Syntax | 296](#)

[Known Behavior | 302](#)

[Known Issues | 304](#)

[Documentation Updates | 326](#)

[Migration, Upgrade, and Downgrade Instructions | 326](#)

[Product Compatibility | 339](#)

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.1R3.

SEE ALSO

[New and Changed Features | 276](#)

[Changes in Behavior and Syntax | 296](#)

[Known Behavior | 302](#)

[Known Issues | 304](#)

[Resolved Issues | 309](#)

[Migration, Upgrade, and Downgrade Instructions | 326](#)

[Product Compatibility | 339](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 327](#)
- [Installing the Software on QFX10002 Switches | 329](#)
- [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 | 329](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 331](#)
- [Performing a Unified ISSU | 335](#)
- [Preparing the Switch for Software Installation | 336](#)
- [Upgrading the Software Using Unified ISSU | 336](#)

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://support.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **17.1** 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 17.1 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.

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

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

```
user@host> request system software add source/jinstall-host-qfx-5-17.1R3.n-signed.tgz reboot
```

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 17.1 **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 17.1R3.

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-10-17.1R3.n-secure-signed.tgz
reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-17.1R3.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://support.juniper.net/support/>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://support.juniper.net/support/>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-17.2R1.n-secure-signed.tgz
```

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-10-17.2R1.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 336](#)
- [Upgrading the Software Using Unified ISSU on page 336](#)

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, *sourcejinstall-host-qfx-5-17.1R3.7-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-5-17.1R3.7-signed.tgz ...
Install jinstall-host-qfx-5-17.1R3.7-signed completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
```

```

ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

SEE ALSO

[New and Changed Features | 276](#)

Changes in Behavior and Syntax	296
Known Behavior	302
Known Issues	304
Resolved Issues	309
Documentation Updates	326
Product Compatibility	339

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 339

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features	276
Changes in Behavior and Syntax	296
Known Behavior	302
Known Issues	304
Resolved Issues	309

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**—Determine the features supported on MX Series, PTX Series, QFX Series devices. The Juniper Networks Feature Explorer is a Web-based app that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. <https://pathfinder.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/.

Requesting Technical Support

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

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <https://support.juniper.net/support/>
- Search for known bugs: <https://kb.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://support.juniper.net/support/downloads/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://forums.juniper.net>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

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

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

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

Revision History

3 September 2020—Revision 4, Junos OS Release 17.1R3— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

6 June 2019—Revision 3, Junos OS Release 17.1R3— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

25 April 2019—Revision 2, Junos OS Release 17.1R3— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

9 April 2019—Revision 1, Junos OS Release 17.1R3— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

2 November 2017—Revision 6, Junos OS Release 17.1R2— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

5 October 2017—Revision 5, Junos OS Release 17.1R2— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

18 August 2017—Revision 4, Junos OS Release 17.1R2— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

6 July 2017—Revision 3, Junos OS Release 17.1R2— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

29 June 2017—Revision 2, Junos OS Release 17.1R2— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

22 June 2017—Revision 1, Junos OS Release 17.1R2— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

18 May 2017—Revision 8, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

27 April 2017—Revision 7, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

12 April 2017—Revision 6, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

7 April 2017—Revision 5, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

23 March 2017—Revision 4, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

17 March 2017—Revision 3, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

10 March 2017—Revision 2, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

3 March 2017—Revision 1, Junos OS Release 17.1R1— EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

Copyright © 2017, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.