Release Notes: Junos® OS Release 17.1R2 for the ACX Series, EX Series, MX Series and PTX Series, QFX Series, and Junos Fusion

3 September 2020

Contents

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 17.1R2 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

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2 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 SNMP 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

Changes in Behavior and Syntax | 16

Changes in Behavior and Syntax

IN THIS SECTION

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 17.1R2 for the ACX Series.

Interfaces and Chassis

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

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.

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-InterfaceIpaddress'. Junos OS now follows the latest format.

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 | 10

Known Behavior | 18

Known Issues | 19

Resolved Issues | 21

Documentation Updates | 22

Migration, Upgrade, and Downgrade Instructions | 23

Product Compatibility | 24

Known Behavior

There are no known limitations in Junos OS Release 17.1R2 for the ACX Series.

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

SEE ALSO

New and Changed Features | 10

Changes in Behavior and Syntax | 16

Known Issues | 19

Resolved Issues | 21

Documentation Updates | 22

Migration, Upgrade, and Downgrade Instructions | 23

Product Compatibility | 24

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R2 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.

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

Generic Routing Encapsulation

- Traffic loss is seen after restarting the **chassis-control** when 64 **gr-** logical interfaces are configured. This occurs when you restart the Packet Forwarding Engine (PFE) and when there are multiple **gr-** logical interfaces configured. The traffic automatically resumes once all the ARP entries for the traffic are learned. PR1228216

Firewall

- On the ACX5000 line of routers, if you apply firewall filter to an interface using **input-list** at the [**edit interfaces *interface-name* unit *unit-name* family ethernet-switching filter**] hierarchy level, then commit does not happen. PR1037604

Layer 2 Features

- On the ACX5000 line of routers, when you issue the **show ethernet-switching table summary vlan-name** CLI command, an **l2ald.core.0.gz** core is generated. PR1042995
- When interface flaps or process restarts occurs, the interface configured for RSTP with root protection may not transit to DESG state. There is no workaround available.PR1223137

MPLS

- The link protection does not work properly when auto bandwidth is configured on the ACX5000 line of routers. After the interface disable has been deleted, the backup will remain active for 90 seconds. The auto-adjustment of bandwidth does not happen at the first instance when the auto-adjustment timer expires and the bandwidth is adjusted only at the second instance when the timer expires. PR1233761

SNMP

- ACX Series routers do not have control board and when you issue the **show snmp mib walk jnxOperatingState** CLI command, the parameter always shows online.

  The following is an example of the **show snmp mib walk jnxOperatingState** CLI command output:

```
user@host> show snmp mib walk jnxOperatingState
jnxOperatingState.1.1.0.0 = 2
jnxOperatingState.6.1.0.0 = 2
jnxOperatingState.7.1.0.0 = 2
jnxOperatingState.9.1.0.0 = 2
jnxOperatingState.12.0.0.0 = 2
```

PR1191995

Timing and Synchronization

- 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
- When interface flaps or process restarts occurs, the interface configured for RSTP with root protection may not transit to DESG state. There is no workaround available.PR1223137

SEE ALSO

New and Changed Features | **10**

Changes in Behavior and Syntax | **16**

Known Behavior | **18**

Resolved Issues | **21**

Documentation Updates | **22**

Migration, Upgrade, and Downgrade Instructions | **23**

Product Compatibility | **24**

Resolved Issues

IN THIS SECTION

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: 17.1R2

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 | 10
- Changes in Behavior and Syntax | 16
- Known Behavior | 18
- Known Issues | 19
- Documentation Updates | 22
- Migration, Upgrade, and Downgrade Instructions | 23
- Product Compatibility | 24

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R2 for the ACX Series documentation.

SEE ALSO

- New and Changed Features | 10
- Changes in Behavior and Syntax | 16
- Known Behavior | 18
- Known Issues | 19
- Resolved Issues | 21
- Migration, Upgrade, and Downgrade Instructions | 23
- Product Compatibility | 24

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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 | 10
- Changes in Behavior and Syntax | 16
- Known Behavior | 18
- Known Issues | 19

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2: EX4300, EX4600, and EX9200.

NOTE: In Junos OS Release 17.1R2, 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 17.1R2 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 multidestination 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

IN THIS SECTION

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 17.1R2 for the EX Series.

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-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-InterfaceIpaddress'. Junos OS now follows the latest format.

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.

SEE ALSO

- New and Changed Features | 25
- Known Behavior | 32
- Known Issues | 33
- Resolved Issues | 36
- Documentation Updates | 41
- Migration, Upgrade, and Downgrade Instructions | 42
- Product Compatibility | 43

Known Behavior

There are no known limitations for the EX Series switches in Junos OS Release 17.1R2.

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 | 25
- Changes in Behavior and Syntax | 30
- Known Issues | 33
- Resolved Issues | 36
- Documentation Updates | 41
- Migration, Upgrade, and Downgrade Instructions | 42
- Product Compatibility | 43

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R2 for the EX Series.

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

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

- 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

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. PR1123764
- On an EX4300 or a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. PR1125155
- In a rare scenario, GRES might not reach the ready state and might fail to start, because the Routing Engine does not receive the state ack message from the Packet Forwarding Engine after performing

GRES. This is a timing issue. It might also stop Routing Engine resource releasing and then cause resource exhausting. Reboot the system if this problem occurs. 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

Interfaces and Chassis

- On EX Series platforms with a Junos OS release 15.1R1 or later, LLDP PDU gets dropped on the FXP interface. PR1188342
- 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 to avoid this issue, when configuring PoE interfaces, use the set poe interface all configuration command instead of configuring specific interfaces individually. To recover connections after seeing this issue, disable and reenable the ports affected by the issue. PR1203880

Junos Fusion Enterprise

- On a Junos Fusion Enterprise, Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) fast start does not work. PR1171899
- Loss of connectivity of the link connecting the standalone box might lead to conversion failure from Junos OS to SNOS. PR1232798
- On a Junos Fusion Enterprise, in order to use a non-default port as a clustering port in a clustering port policy, the policy must include at least one port that is a default uplink/clustering port for that platform. PR1241808
- On a Junos Fusion Enterprise, the satellite device might not come online when the systems is converted from cluster to non-cluster mode without accompanying topology changes. PR1251790

Network Management and Monitoring

- On EX9200 switches, analyzer configurations with analyzer input and output stanzas 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

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 a EX4300-VC platform, if a Q-in-Q S-VLAN interface with MC-LAG is configured, 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 will be dropped. PR1131755
- On EX4300 Series switches, certain multicast traffic might impact the network, for example, cause OSPF to flap. Issues might occur when multicast packets use the same interface queue as certain network protocol packets (for example, OSPF, RIP, PIM, and VRRP). PR1244351

Port Security

- When LACP is configured together with MACsec, the links in the bundle might not all work. Rebooting the switch might solve the problematic links, but could also create the same issue on other child interfaces. PR1093295
- On a dot1x-enabled interface, sometimes when you log in, log off, and then log in within a short interval (within subseconds), the logical interface plus the bridge domain or VLAN remain in a pending state, and you will not be able to access the network. As a workaround, restart the l2-learning process to recover the port/interface from the problematic state. PR1230073

Security

- On EX4300 switches, when storm-control or storm-control-profiles with action-shutdown is configured, if the storm-triggered traffic is control traffic such as LACP, the physical interface might be put into an STP blocking state rather than turned down, so valid control traffic might be trapped at the control plane and unrelated interfaces might be set down as an LACP timeout. PR1130099

Virtual Chassis

- When the linecard role FPC is removed and rejoined to the Virtual Chassis immediately, the LAG interface on the master/backup would not be reprogrammed in the rejoined FPC. PR1255302

SEE ALSO

New and Changed Features | 25

Changes in Behavior and Syntax | 30

Known Behavior | 32

Resolved Issues | 36

Documentation Updates | 41

Migration, Upgrade, and Downgrade Instructions | 42

Product Compatibility | 43

Resolved Issues

IN THIS SECTION

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: 17.1R2

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: SDPD core files might be seen while converting an EX2300 or EX3400 cluster from Junos OS to SNOS. PR1239915
- 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
- On QFX10000 switches with MC-LAG configured, CDP packets are looping to the other QFX10000 devices in the MC-LAG. PR1237227

MPLS

- Virtual Chassis/Virtual Chassis Fabric-l2ckt: FXPC core file is seen when deactivating core interface on MPLS l2ckt 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
- On EX3400 some of the IPV6 clients do not get bind if two dhcpv6 relays are present with VRRP between them. PR1189333
- FF reject tcp-reset does not work on IRB interface. PR1219953
- On a Junos Fusion Enterprise: SDPD core files might be seen while converting an EX2300 or EX3400 cluster from Junos OS to SNOS. PR1239915
- Issue with the **show** command occurs in single supplicant mode captive portal. PR1240259
- On EX3400 Virtual Chassis, RA guard-enabled Interface stays in Trusted mode even after the **mark-interface trusted** statement is deleted. PR1242937
- On EX3400 Virtual Chassis, executing **request access-security router-advertisement-guard-block interface** and **restart dhcp-service** commands triggers the jdhcpd to generate a core file. PR1243147
- On EX3400 Virtual Chassis, RA guard Policy discard does not discard the packet matching with policy-option. PR1244666
- 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 | 25

Changes in Behavior and Syntax | 30

Known Behavior | 32

Known Issues | 33

Documentation Updates | 41

Migration, Upgrade, and Downgrade Instructions | 42

Product Compatibility | 43

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R2 for the EX Series switches documentation.

SEE ALSO

New and Changed Features | 25

Changes in Behavior and Syntax | 30

Known Behavior | 32

Known Issues | 33

Resolved Issues | 36

Migration, Upgrade, and Downgrade Instructions | 42

Product Compatibility | 43

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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://www.juniper.net/support/eol/junos.html.

SEE ALSO

- New and Changed Features | 25
- Changes in Behavior and Syntax | 30
- Known Behavior | 32
- Known Issues | 33
- Resolved Issues | 36

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2 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

IN THIS SECTION

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 17.1R2 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 | 44

Known Behavior | 50

Known Issues | 53

Resolved Issues | 55

Documentation Updates | 56

Migration, Upgrade, and Downgrade Instructions | 56

Product Compatibility | 65

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R2 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- On a Junos Fusion Enterprise, 'show ethernet-switching table' takes a few minutes to show entries when an extended port receives with MAC count set to 150K. PR1117567
- 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, 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\n1\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
- 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 | 44

Changes in Behavior and Syntax | 49

Known Issues | 53

Resolved Issues | 55

Documentation Updates | 56

Migration, Upgrade, and Downgrade Instructions | 56

Product Compatibility | 65

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R2 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
- 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, 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
- 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

SEE ALSO

New and Changed Features | 44

Changes in Behavior and Syntax | 49

Known Behavior | 50

Resolved Issues | 55

Documentation Updates | 56

Migration, Upgrade, and Downgrade Instructions | 56

Product Compatibility | 65

Resolved Issues

IN THIS SECTION

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: 17.1R2

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 telemetrics 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

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R2 for Junos Fusion Enterprise documentation.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

Basic Procedure for Upgrading Junos OS on an Aggregation Device

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

> **NOTE:** Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:
>
> ```
> user@host> request system snapshot
> ```
>
> The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the Junos OS Administration Library.

To download and install Junos OS Release 17.1R2:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list on the right of the page.

5. Select the **Software** tab.

6. Select the software package for the release.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

10. Install the new **junos-install** package on the aggregation device.

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

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

`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 Release 17.1R2 **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%` **vty 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 17.1R2 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 the satellite device from the Junos Fusion topology.

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

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads

2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos Fusion** from the menu and select the switch platform series and model for your satellite device.

4. Select the software image for your platform. For satellite device software requirements, see Understanding Junos Fusion Enterprise Software and Hardware Requirements.

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

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

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

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

   ```
   [edit]
   ```

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

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

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

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

8. Commit the configuration.

   To commit the configuration to both Routing Engines:

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

   To commit the configuration to a single Routing Engine:

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

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

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

   For example, to install a software package stored in the **/var/tmp** directory on the aggregation device onto an EX4300 switch acting as the satellite device using FPC slot 102:

   ```
   [edit]
   user@aggregation-device> request chassis satellite install
   /var/tmp/jinstall-ex-4300-14.1X53-D35.3-domestic-signed.tgz fpc-slot 102
   ```

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

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

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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 | 44
- Changes in Behavior and Syntax | 49
- Known Behavior | 50
- Known Issues | 53
- Resolved Issues | 55
- Documentation Updates | 56
- Product Compatibility | 65

Product Compatibility

IN THIS SECTION

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

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2 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

IN THIS SECTION

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 17.1R2 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

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 17.1R2 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 | 67
- Changes in Behavior and Syntax | 68
- Known Issues | 69
- Resolved Issues | 70
- Documentation Updates | 71
- Migration, Upgrade, and Downgrade Instructions | 72
- Product Compatibility | 80

Known Issues

IN THIS SECTION

- Junos Fusion | 70

This section lists the known issues in hardware and software in Junos OS Release 17.1R2 for Junos Fusion Provider Edge.

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

Junos Fusion

- Junos Fusion transit traffic fails between two VLANS on the same extended port. As a workaround, configure extended port with LAG.PR1264900

SEE ALSO

Resolved Issues

IN THIS SECTION

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: 17.1R2

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 | 67

Changes in Behavior and Syntax | 68

Known Behavior | 69

Known Issues | 69

Documentation Updates | 71

Migration, Upgrade, and Downgrade Instructions | 72

Product Compatibility | 80

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R2 for Junos Fusion Provider Edge documentation.

SEE ALSO

New and Changed Features | 67

Changes in Behavior and Syntax | 68

Known Behavior | 69

Known Issues | 69

Resolved Issues | 70

Migration, Upgrade, and Downgrade Instructions | 72

Product Compatibility | 80

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

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

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos Fusion** to 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 17.1R2 **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://www.juniper.net/support/downloads

2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Select **By Technology > Junos Platform > Junos Fusion** from the drop-down list and select the switch platform series and model for your satellite device.

4. Select the Junos OS Release 14.1X53-D30 software image for your platform.

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

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

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

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

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

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

   ```
   [edit]
   ```

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

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

9. Commit the configuration.

   To commit the configuration to both Routing Engines:

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

   Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

    ```
    [edit]
    user@aggregation-device> request chassis satellite install
    /var/tmp/install-media-pxe-qfx-5-14.1X53-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 17.1R2, 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

Product Compatibility

IN THIS SECTION

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

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

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2 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-ACOand 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-1OC192-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 MLPS 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-instances*instance-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

- jnxClksyncIfIndex
- 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 (MXSeries 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 softwire statistics, and clear services inline softwire 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 reasssembly 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 translater (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 Ssrvices 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

IN THIS SECTION

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 17.1R2 for MX Series routers.

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.

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>
```

LDP

- **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.

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.]

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-InterfaceIpaddress'. 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 <url href=https://www.juniper.net/documentation/en_US/junos/topics/reference/command-summary/show-system-statistics-mpls.html>show system statistics mpls </url>]

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 SMNPv3 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.]

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.

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).]

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.

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
        }
    }
```

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.

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

Known Behavior

IN THIS SECTION

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R2 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.

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.

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.

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).

Subscriber Management and Services

- If a graceful Routing Engine switchover (GRES) is triggered by an operational mode command, the state of aggregated services interfaces (ASIs) are not preserved. For example:

  ```
  request interface <switchover | revert> asi-interface
  ```

  However, if GRES is triggered by a CLI commit or FPC restart or crash, the backup Routing Engine updates the ASI state. For example:

  ```
  set interface si-x/y/z disable
  commit
  ```

  Or:

  ```
  request chassis fpc restart
  ```

SEE ALSO

New and Changed Features | 82

Changes in Behavior and Syntax | 106

Known Issues | 120

Resolved Issues | 130

Documentation Updates | 147

Migration, Upgrade, and Downgrade Instructions | 148

Product Compatibility | 156

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R2 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.

Forwarding and Sampling

- Firewall module (daemon dfwd) on Routing Engine always leaks some memory upon configuration commit with following statements: **set routing-options forwarding-table export qos3** , **set policy-options policy-statement <policy-name> term 1 from source-address-filter <ip-address>**, and **set policy-options policy-statement <policy-name> term 1 then forwarding-class <forwarding-class>**. PR1157714
- When a policing filter is applied to an active LSP carrying traffic, the LSP resignals and drops traffic for approximately 2 seconds. It can take up to 30 seconds for the LSP to come up under the following conditions: (1) Creation of the policing filter and application of the same to the LSP through 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 is followed by a commit.PR1160669

- Inline JFlow (MXVC): NextHop Address/OIF being reported by IPv6 template on MXVC setup is correct—Root Cause of the Problem: ++++++++++++++++++++++++++ As per the investigation from RPD : we have is 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 nexthop in rpd local storage. route and nexthop for the interface are taken care of in the kernel. There is no route change in rpd. route_record depends on route flash to find out about updates. Since there is no route change, there is no route flash, so route_record is unaware. PR1224105
- Firewall filter family "any" with shared-bandwidth-policer on MC-AE interface does not reconfigure bandwidth or carve up the policer when standby becomes active after A/S switchover; it drops all packets. PR1232607
- After executing **show firewall** command, "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 config. << 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 lowPR1248134
- FreeBSD 10.x based Junos OS is not supported on 32-bit Routing Engines in Junos OS Release 17.1R1. PR1252662

General Routing

- ICMP echo_reply traffic with applications like IPSec will not work with the MS-MIC and MS-MPC cards in an asymmetric traffic environment since these cards employ a stateful firewall by default. The packet will be dropped at the stateful firewall because it acknowledges an ICMP Reply that has no matching session. PR1072180
- Show evpn vpws-instance SID NNN is not supported. PR1122695
- In a situation where both mirrored interface and mirrored destination are on MPC card and mirror destination interface is a unilist next-hop(for example,an ae interface), mirrored packets may get dropped. PR1134523
- Queue bps rates is more than expected when AE child reconfigured with per-unit-scheduler. This is an intermittent issue. Assuming that Aggregated Ethernet is configured with the bypass-queuing-chip configuration statement. Now , followup configuration changes are such that removing child link(s) from AE bundle, and configuring per-unit-scheduler on the removed child link(s) in a single commit causes intermittent issues with per-unit-scheduler configuration updates to cosd and the Packet Forwarding Engine. Hence, dedicated scheduler nodes might not be created for all units or IFLs. PR1162006

- Chef for Junos supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev-resources. The netdev-resource developed for interface configuration has a limitation to configure XE interface. Netdev-interface resource assumes that 'speed' is a configurable parameter which is supported on a GE interface but not on an XE interface. Hence netdev-interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages chef-11.10.4_1.1.*.tgz chef-11.10.4_2.0_*.tgz in all platforms {i386/x86-32/powerpc}. PR1181475
- EVPN VPWS convergence and association with traffic loss is tied to the type of redundancy and the route exchange via BGP. In A/A this traffic loss is low due to 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
- With NAT translation-type as napt-44, a few sessions are getting stuck upon deactivating/activating service-set or corresponding applications with traffic running. The same symptom is seen upon deactivating/activating service-set with traffic running and with 'deterministic-napt44' translation type as well. PR1183193
- AMS redundant interfaces are not listed under possible-completions of operational commands. PR1185710
- On MX Series platforms with Junos OS Release 15.1R1 or later, LLDP PDU gets dropped on the FXP interface. PR1188342
- As described in RFC7130, 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 Up state. PR1192161
- GUMEM errors for the same address may continually be logged if a parity errors occurs in a locked location in GUMEM. Since GUMEM utilizes ECC memory, any error is self-correcting and has no impact to router's operation. In a rare case, such parity error may appear repeatedly at a specific location. As a workaround, such errors 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 as 'in-use#' instead. This indicates that it is still locked to the source, although the clock has a considerably large ppm deviation. PR1202327
- A dynamic tunnel gets timed out every 15 mins by default, and then re-tries to create another tunnel. This happens if the route obtained from IGP is non-forwarding. PR1202926
- MPC might crash after firewall filter configuration is changed and all interfaces/protocols are flapped. The issue is due to access to a stale or invalid pointer which caused a particular check based on the pointer structure field to unpredictably fail, resulting in the assert later in the code. The issue happened

when a sequence of events related to firewall filters resulted in filter structure getting deleted and re-created again. PR1205325

- The ptp master streams on IP and Ethernet not supported simultaneously. PR1217427
- The /etc/passwd file is created in the process of the first commit when a pristine jinstall image is used to boot for the first time. If event-options is configured, the system will try to read the configuration from the available event scripts, which requires privileges obtained from the /etc/passwd file. This causes a circular dependency because the commit will not pass if the configuration includes event-options the first time a pristine image boots up (which is the case of an upgrade performed with virsh create). PR1220671
- The problem of tunnel stream getting misconfigured for LT interfaces is due to internal programming and to evaluate multiple lt interfaces for FPC and PIC slot combination. PR1223087
- Wth qmon sensor, when you issue an operational clear command, such as clear interfaces statistics all, the counters at the telemetry jvision server are not reset. Hence qmon sensor stats at jvision server will not match the CLI/VTY commands output, after the **clear interfaces statistics** commands. PR1226948
- Continuously increasing normal discard count in 'show pfe statistics traffic' occurs without any user traffic. This occurs because internal control traffic that is expected to be dropped silently is unexpectedly being counted as 'normal discard'. There is no impact on user traffic with this issue. PR1227162
- An incorrect PE is being 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 partial traffic blackhaule and 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: 00:13:78:00:00:00:00:00:00:01 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
- OSPF is used as routing protocol between the clients and DEP router with TD configured. The OSPF protocol traffic brings the IPsec up on spokes and the DEP router. The IPsec SAs are distributed on the DEP router. The neighbor state between the OSPF peers moves to full but it does not stay in that state. States change 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
- When changing virtual switch type is changed from IRB type to regular bridge, interfaces under openflow protocol are all removed. Openflow daemon fails to program any flows. PR1234141
- To distinguish between flow and kernel IFL 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 "local-gateway" for the IPSec VPN service is not actually assigned to any interface in UP state (not present a local/direct route in the routing-table), the system still sends ISAKMP packets for IKE exchange. As a source address for these packets, an address of the outgoing interface would be selected. PR1238112

- On MX Series with rpd in "ASYNC" mode, if the distributed IGMP is configured, rpd core file might be seen, causing rpd crash. PR1238333
- For ANCP subscribers in Idle state, the previously reported speed in ANCP Port UP message is not applied. PR1242992
- ANCP neighbors go down after commit when any ANCP related configuration is changed. PR1243164
- After connecting 1k L2BSA subscriber and running 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 Ctrl-C can not help to terminate the CLI. In some cases entering Ctrl+C causes ANCPD to crash. PR1250996
- On MX2000 MPC6E, EOAM LFM adjacency flaps when an unrelated MIC accommodated in the same MPC6E slot is onlined with configuring OAM pdu-interval 100 ms and pdu-threshold 3. PR1253102
- VPLS MAC table is not being populated properly when checked the CLI M command **show vpls mac-table**", though all subscribers have traffic. Thus this is considered a cosmetic issue. PR1257605
- Due to transient hardware error conditions only syslog events XMCHIP(x) FI: Cell underflow at the state stage - Stream 0, Count 65535 are reported, which is a sign of fabric stream wedge. Additional traffic flow register pointers are validated and if stalled, a new CMERROR alarm is raised: "XMCHIP(x) FI: Cell underflow errors with reorder engine pointers stalled - Stream 0, late_cell_value 65535, max_rdr_ptr 0x6a9, reorder_ptr 0x2ae." PR1264656
- Due to transient Hardware events, fabric stream may report 'CPQ1: Queue unrderrun indication - Queue <q#>' in continuous occurrence. For such events, all fabric traffic is queued for this Packet Forwarding Engine reporting the error and causes a very high amount of fabric drops. PR1265385
- The MTU configuration option for vt/mt/pd/pe interfaces will be removed after the fix of this PR because the MTU on these interfaces is already set to unlimited, so there is no need for configuring MTU on these interfaces. PR1277600

High Availability (HA) and Resiliency

- In a rare scenario, GRES might not reach the ready state and might fail to start, because the Routing Engine does not receive the state ack message from the Packet Forwarding Engine after performing GRES. This is a timing issue. It might also stop Routing Engine resource releasing and then cause resource exhausting. Reboot the system if this problem occurs. PR1236882

Infrastructure

- The configuration statement: "set system ports console log-out-on-disconnect" logs the user out from the console and closes the console connection . If "set system syslog console any warning" is used along with the previously mentioned statement and there is no active telnet connection to the console, the daemons try to open the console and hang as they wait for a "serial connect", which is received only by doing a telnet to the console. As a workaround, remove the second statement, "set system syslog console any warning", which solves the issue. PR1230657

Interfaces and Chassis

- After changing the MTU on the IFD, on the static vlan demux interface above the IFD the IPv6 Link Local address is not assigned. PR1063404
- During configuration changes and reuse of Virtual IP on an interface as a interface address; you must delete the configuration do a commit and then add the interface address configuration in another commit. PR1191371
- IPV6 neighborship is not created on the IRB interface. PR1198482
- 1. Delay Measurement support for 5-port 100G DWDM PIC and 5-port 100G DWDM MIC is *ONE TIME Delay Measurement*.If you intend to measure Delay 2 points should ensure that 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 Routing Engine CLI. 2. Remote-loop-enable should be configured first on remote end. Next, start-measurement should be configured. 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 highest, SDFEC in the interim, and GFEC being least 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
- In some rare situations Ethernet Connectivity Fault Management Daemon (cfmd) might crash when committing a configuration where CFM filter refers to a firewall policy. When hitting this issue, all CFM enabled interfaces are down. PR1246822

- In a VPLS multihoming scenario, the CFM packets are forwarded over the standby PE link, resulting in duplicate packets or loop between the active and standby link PR1253542
- Junos OS upgrade involving releases 14.2R5 (and above in 14.2 maintenance releases) and 16.1 and above mainline releases with CFM configuration can cause CFMD core after the upgrade. This is due to the old version of /var/db/cfm.db. PR1281073

Layer 2 Ethernet Services

- After changing the underlying IFD for a static vlan demux interface, the NAS-Port-ID is still formed based on the previous IFD. PR1255377

Layer 2 Features

- On routers running Junos OS with RE 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 is due to a change in the underlying database that handles the dynamic-profile sessions . As a result, it causes the vpls connection is destroyed and re-created after a Routing Engine switchover. PR1220171

MPLS

- When graceful Routing Engine switchover (GRES) is done between master and backup Routing Engines of different memory capabilities. For example, this issue can occur when one Routing Engine has only enough memory to run rpd in 32-bit mode while the other is capable of 64-bit mode. This scenario can occur when using Junos OS Release 13.3 and later with the statement "auto-64-bit" configured or when using Junos OS Release 15.1 or later (even without the configuration statement). As a workaround, use the statement "set system processes routing force-32-bit" to avoid the issue. PR1141728
- In MVPN scenario, if the active primary path goes down, then PLR (Point of Local Repair) needs to send Label Withdraw for the old path and new Label Mapping for the new path to the new upstream neighbor. In this case, LDP P2MP path may stay in "Inactive" state for an indefinite time if an LSR receives a Label Release, immediately followed by a Label Mapping for the same P2MP LSP from the downstream neighbor. PR1170847
- A new configuration **protocols mpls traffic-engineering bgp-igp-both-ribs** in the routing-instance is required to make cOC work. PR1252043
- The throughput measurement may be inaccurate when doing performance measurement on an MPLS label-switched path. PR1274822

Network Management and Monitoring

- Symptom: "MIB2D_RTSLIB_READ_FAILURE: rtslib_iflm_snmp_pointchange" syslog message during config restore. Cause: mib-process daemon sends to the requests to kernel to update snmp ifIndex for the interfaces that it is learning. If this interface was already deleted from kernel, the syslog message is seen. This interface learning by the mib-process daemon will happen later, once the kernel sends the ADD notification for these interfaces. There is no system impact. PR1279488

Platform and Infrastructure

- FPC reports the following errors and the FPC is not able to connect any subscriber: "Pkt Xfer:** WEDGE DETECTED IN PFE 0 TOE host packet transfer: %PFE-0: reason code 0x1". Also, the MQ FI 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
- When TCP authentication is enabled on a TCP session, the TCP session might not use the selective acknowledgement (SACK) TCP extensions. PR1024798
- In configurations with IRB interfaces, during times of interface deletion, (for example, an FPC reboot) the Packet Forwarding Engine may 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 platform, parity memory errors might happen in pre-classifier engines within an MPC. Packets will be silently discarded because such errors are not reported and are therefore harder to diagnose. The correct behavior is for CM-ERRORs, such as syslogs messages and alarms, to 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 and traffic. PR1158127
- In a very rare scenario, during TAC accounting configuration change, auditd daemon crashes due to a race condition between auditd and its sigalarm handler. PR1191527
- Several files are copied between Routing Engines during 'ffp synchronize' phase of the commit (e.g. /var/etc/mobile_aaa_ne.id, /var/etc/mobile_aaa_radius.id, etc). These files are copied even if there was no corresponding change in the configuration thus unnecessarily increasing commit time. PR1210986
- Starting from Junos OS 13.3, the SRX Series cluster 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 route for the accounting server, the following message is generated: 1813 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

- On rare occasions during the route add/delete/change operation, the kernel might encounter a crash with the panic string "rn_clone_unwire no ifclone parent". PR1253362

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
- On MX Series router, when a instance type is changed from VPLS to EVPN, and in the same commit an interface is added to the EVPN instance, the newly added EVPN interface might not be able to come up. PR1016797
- For devices populated with a master and backup Routing Engines (RE) and configured for nonstop active routing (NSR) and Protocol Independent Multicast (PIM) configuration, the routing protocol process (RPD) may crash on the backup Routing Engine due to a memory leak. This leak occurs when the backup Routing Engine handling mirror updates about PIM received from the master Routing Engine deletes information about a PIM session from its database. But due to a software defect, a leak of 2 memory blocks (8 or 16 bytes) may occur for every PIM leave. If the memory is exhausted, the rpd may crash on the backup Routing Engine. There is no impact seen on the master Routing Engine when the rpd crashes on the backup Routing Engine. Use the **show system processes extensive** command to check the memory. PR1155778
- The VRF related routes which are leaked to the global inet.0 table and advertised by the access routers are not being advertised to global inet.0 table on the core. PR1200883
- In the context of a large number of configured VPNs, routes changing in the midst of a bgp path-selection configuration change can sometimes lead to an rpd core files. This core file has been seen with the removal of the "always-compare-med" option. PR1213131
- RPD leaks memory with the topology and configuration. However, adding/deleting static flowspec routes in isolation does not cause any memory leak. The exact configuration that causes the leak is currently unknown. PR1213959
- PIM NSR Design : With GRES+ NSR enabled, the master Routing Engine (RE) replicates kernel states and protocol states on backup RE. Both kernel state (ifstates) and protocol state replication are independent processes. The ksyncd takes care of ifstates replication. RPD infra takes care of replication (mirror) connection between the two Routing Enginess. NSR supported protocols have their own mechanism to replicate their database using mirror connection. As per PIM/MVPN NSR design, the

backup RE, it walks through the replication database (RDB) with consume and delete action. That is once a PIM/MVPN states is processed on the backup RE, associated RDB is deleted. If kernel replication is restarted, it can lead to interface deletions and additions only on the backup RE. PIM states the backup goes out of sync. - ?kernel replication? restart lead to interface delete/add on Backup-RE only - PIM/MVPN does not have RDB on the backup RE, so on interface delete, it deletes the relevant PIM state..Once an interface is added by kernel, PIM has no state to consume. No change occurs on the master Routing Engine to reinitiate the protocol. replicationThis .PIM/MVPN out-of-sync issue can be seen with following events :- Manually "restart kernel-replication" - PIM out of sync - ksyncd cored & restarted - PIM out of sync - ksyncd restarted as workaround of kernel replication issues- PIM out of sync. PR1224155

- On rpd crash with ?switchover-on-routing-crash? enabled on box, live vmcores may be seen on both Routing Engines without an impact on the system. PR1267796

Services Applications

- On MX series with L2TP configured, the L2TP packet in ICRQ retransmission message is set to incorrect value, and this causes frequent L2TP session flaps. PR1206542
- On Layer 2 Tunneling Protocol (L2TP) access concentrator (LAC) router where Access Node Control Protocol (ANCP) protocol is used for bandwidth adjustment, L2TP Connect Speed Update Notification (CSUN) message to L2TP network server (LNS) may be sent after a short delay after ANCP Port-Up 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, because the L2TP daemon checks for state updates each time there is an L2TP packet that has to be sent or received. PR1234674
- If the l2tp subscriber has static pp0 interface on the LAC side, LCP renegitiation is configured on the LNS side, and CPE has been changed, an issue with negotiation of PPP session between LNS and CPE can occur. PR1235554
- 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

Subscriber Access Management

- On MX Series routers with subscriber management feature enabled, after GRES switchover results of the **show network-access aaa statistics radius** CLI command display only zeros and **clear network-access**

aaa statistics radius" does not clear statistics as it should. However this is a cosmetic issue and communication with the RADIUS server is working fine; the only impact is that affected CLI commands do not work as expected. PR1208735

- Subscribers get stuck in terminated state during PPPoE login/logout test. PR1262219

User Interface and Configuration

- When persist-groups-inheritance is configured and you issue a rollback, it will be seen that the configuration is not propagated properly after a commit. PR1214743

VPNs

- In NG-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

SEE ALSO

New and Changed Features | **82**

Changes in Behavior and Syntax | **106**

Known Behavior | **117**

Resolved Issues | **130**

Documentation Updates | **147**

Migration, Upgrade, and Downgrade Instructions | **148**

Product Compatibility | **156**

Resolved Issues

IN THIS SECTION

Resolved Issues: 17.1R2 | **131**

Resolved Issues: 17.1R1 | **140**

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: 17.1R2

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 may 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 may be raised incorrectly while an AS-MCC PIC is coming up. PR1036412
- ICMP reply traffic might get dropped on MS-MPC line cards. PR1059940
- 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 may 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 ppman 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-ackhowledgements" 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() may 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

- Vmcores 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 may still be sent after PADT is sent. PR1234027
- t3 interface not coming up due to incorrect subrate. 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

- DHCP-Relay option-82 format may change when an interface is removed from a configuration. PR1253205
- Periodic MLP query mechanism from DPC for unresolved DMAC entries. PR1256555
- An MPC/FPC might go offline during FRU upgrade phase of ISSU. PR1256940
- DHCP client key identifier mismatch due to DHCPv4 Option 82 Suboption 9 change during the release time. PR1257701
- DHCP (V4 or V6) relay - renew from client does not work as expected for asymmetric-lease configured on a static interface when using broadband subscriber service. PR1258415
- Eliminate the impact of DHCPv6 renegotiation lockout timer for DHCP Solicits with rapid commit options. PR1263156

Layer 2 Features

- BUM traffic might be dropped on egress AE interface or equal-cost multipath (ECMP) core-links with "input-vlan-map" and "enhanced-ip" configured. PR1078617
- VPLS unicast traffic loss might be observed when it is passed through LSI interface and the P-facing interface is LAG. PR1240960

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/l2 interface with VXLAN EVPN is getting dropped after l2 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 postion. 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 Chasis environment traffic loss might be observed due to incorrectly programmed Aggregated Ethernet interfaces. PR1220934
- RPD might crash after offilining 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
- DT_LNS: 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 nextheader 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

- MX v44: 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 may 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 l2circuit neighbor in an NSR scenario. PR1241801

SEE ALSO

New and Changed Features | 82

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 17.1R2 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.
- Starting in Junos OS Release 15.1, the *Broadband Subscriber Sessions User Guide* and the CLI Explorer incorrectly included information about the **show extensible-subscriber-services accounting** command. This command is not present in the CLI. Instead, you can use accounting profiles to collect statistics from the Packet Forwarding Engine for Extensible Subscriber Services Manager (ESSM) subscribers. See Flat-File Accounting Overview for information about accounting for ESSM subscribers.

SEE ALSO

New and Changed Features | **82**

Changes in Behavior and Syntax | **106**

Known Behavior | **117**

Known Issues | **120**

Resolved Issues | **130**

Migration, Upgrade, and Downgrade Instructions | **148**

Product Compatibility | **156**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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. For more information, see the Junos OS Administration Library.

Installation and Upgrade Guideand Upgrading Junos OS with Upgraded FreeBSD.

UProcedure 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://www.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

10. Install the new **jinstall** package on the routing platform.

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

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-17.1R2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.1R2.9-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-x86-32-17.1R2.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-17.1R2.9-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://www.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

10. Install the new **jinstall** package on the routing platform.

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

- All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

  `user@host>` **request system software add validate reboot *source*/jinstall-ppc-17.1R2.9-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-17.1R2.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 | **82**
- Changes in Behavior and Syntax | **106**
- Known Behavior | **117**
- Known Issues | **120**
- Resolved Issues | **130**
- Documentation Updates | **147**
- Product Compatibility | **156**

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

- New and Changed Features | 82
- Changes in Behavior and Syntax | 106
- Known Behavior | 117
- Known Issues | 120
- Resolved Issues | 130
- Documentation Updates | 147
- Migration, Upgrade, and Downgrade Instructions | 148

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2 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* gigether-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 Googl e 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

IN THIS SECTION

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 17.1R2 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.

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.]

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 used to follow 'InterfaceIpAddress-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.
- **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

- **Update to SNMP support of apply-path statement (PTX 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.

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

Known Behavior

IN THIS SECTION

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R2 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.

High Availiablity (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

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R2 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

- PTX 100GbE-LR4 interfaces may flap when the reference clock switches over from "line clock" to "holdover" initiated by offlining the PIC, on which the "line clock" sources reside. When PTX uses "line clock" sources and when it does not have any external clocks from BITS-a or BITS-b, offlining the PIC, which is recovering clock from line, brings "line clock" down and the reference clock is switched from "line clock" to "holdover". This reference clock transition may cause a large clock phase-shift in the 100GbE-LR4 CFP modules, and this phase-shift may cause the output optical pulse waveform distortion on the 100GbE-LR4 interfaces. Hence, it results in interface flap. This issue cannot be fixed by software due to hardware limitation. PR1130403
- While upgrading from 15.1F based images to 16.x+ images or downgrading from 16.x+ images to 15.1F based images, if validate option is enabled, there may be a chassisd crash and upgrade/downgrade will fail. This issue should not be seen if both base and target images are from 15.1F train or 16.x+ train. PR1171652

- For FPC3 on PTX Series platform, in rare scenarios, while restarting FPC, a PIC index mismatch issue might result in FPC crash if it is configured with inline-JFlow. PR1183215
- This is a resiliency feature. If more than 10 FO CRC errors are seen in an interval of 30 seconds, then CMERROR infra raises an alarm and appropriate action is 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
- PTX Series FPC3 might receive noise on the FPC console port and interpret it as valid signals. This might cause login fails on the console port, core files to be generated, or even reloads. PR1224820
- On PTX1000/QFX10002 platform, some random ports, using 100G Lumentum optics, might not come up after a reboot. This is a timing issue because of failures during optics read on some ports. When hitting this issue, remove and insert the optics, which might bring up the ports. PR1227029
- When pulling a SIB out without offline on PTX platform with FPC3, it is possible that traffic may be dropped resulting in an overall reduction in traffic throughput.PR1162977

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 a customer intending to measure Delay 2 points should ensure that the link is up on both sides and then conduct this one-time test. The result value is valid one time once the test is finished. The test result on CLI is not valid after one-time measurement as the old result might show up on Routing Engine CLI. 2. Remote-loop-enable should be configured first on remote end. Only after this start-measurement should be configured. 3. Each time the customer wants to verify this, the test has to be repeated. 4. Processing delays in each mode is different HGFEC [For 5-port 100G DWDM MIC] being highest, SDFEC in the interim and GFEC being least for the same cable length. 5. In summary, any breakage in the transmit/receive path during the Delay Measurement test will hinder the delay measurement. This is true for all FEC modes - GFEC, SDFEC, HGFEC. 6. Currently SNMP walk is not available for Delay Measurement. PR1233917

Platform and Infrastructure

- On a PTX Series platform, parity memory errors might happen in pre classifier engines within an MPC. Packets will be silently discarded as such errors are not reported and make it harder to diagnose. After the change in this PR, CM-ERRORs, such as syslogs and alarms, will be raised when parity memory errors occur. PR1059137
- In a very rare scenario, during TAC accounting configuration change, auditd daemon crashes due to a race condition between auditd and its sigalarm handler. PR1191527
- On PTX Series platform with [chassis network-services enhanced-mode] configured, the default policy junos-ptx-series-default is not loaded correctly during some configuration operations, which causes BGP

routes not to be installed in the forwarding table as expected. To avoid this issue, reboot the router after any configuration operations on network services. PR1204827

User Interface and Configuration

- When persist-groups-inheritance is configured and you issue a rollback, the configuration is not propagated properly after a commit. PR1214743

SEE ALSO

New and Changed Features | **157**

Changes in Behavior and Syntax | **167**

Known Behavior | **171**

Resolved Issues | **174**

Documentation Updates | **177**

Migration, Upgrade, and Downgrade Instructions | **178**

Product Compatibility | **182**

Resolved Issues

IN THIS SECTION

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: 17.1R2

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 gauranteed 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 | **157**

Changes in Behavior and Syntax | **167**

Known Behavior | **171**

Known Issues | **172**

Documentation Updates | **177**

Migration, Upgrade, and Downgrade Instructions | **178**

Product Compatibility | **182**

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R2 documentation for PTX Series.

SEE ALSO

New and Changed Features | **157**

Changes in Behavior and Syntax | **167**

Known Behavior | **171**

Known Issues | **172**

Resolved Issues | **174**

Migration, Upgrade, and Downgrade Instructions | **178**

Product Compatibility | **182**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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 Junos OS Administration Library.

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

To download and install Junos OS Release 17.1R2:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

10. Install the new **jinstall** package on the router.

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

    All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

    `user@host>` **request system software add validate reboot *source*/junos-install-ptx-x86-64-17.1R2.9.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-x86-64-17.1R2.9-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://www.juniper.net/support/eol/junos.html.

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

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

For the detailed procedure, see the Installation and Upgrade Guide.

SEE ALSO

New and Changed Features | **157**

Changes in Behavior and Syntax | **167**

Known Behavior | **171**

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R2 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 http://www.juniper.net/techpubs/software/junos/.

New and Changed Features

IN THIS SECTION

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 17.1R2: QFX5100, QFX10002, QFX10008, and QFX10016.

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 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.

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

  [See Configuring LACP Hold-UP Timer to Prevent Link Flapping on LAG Interfaces.]

  Initialization delay timer feature support on LAG interfaces (QFX10000 switches)—Starting with Junos OS Release 17.1R1, 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.

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

  [See 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 (ae*x*), 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, Multiprotocol Label Switching (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 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 is 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)**—Staring 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

IN THIS SECTION

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 17.1R2 for the QFX Series.

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-InterfaceIpaddress'. Junos OS now follows the latest format.

Network Management and Monitoring

- **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.]

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 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.

SEE ALSO

New and Changed Features | **183**

Known Behavior | **205**

Known Issues | **206**

Resolved Issues | **209**

Documentation Updates | **214**

Migration, Upgrade, and Downgrade Instructions | **215**

Product Compatibility | **228**

Known Behavior

There are no known limitations for the QFX Series switches in Junos OS Release 17.1R2.

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

SEE ALSO

New and Changed Features | **183**

Changes in Behavior and Syntax | **203**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.1R2.

Hardware

- On a QFX10002 platform, some random ports, when using 100 Gigabit Ethernet Lumentum optics, might not come up after a reboot. This is a timing issue because of failures during optics read on some ports. When hitting this issue, remove and insert the optics, which might bring up the ports. PR1227029

Infrastructure

- When there is a high route churn or when there is a high rate of route updates being pushed to the kernel, the **show interface** command might show a delay or not show all statistics due to route updates being prioritized over statistics messages. PR1250328

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which the flexible-vlan-tagging statement is specified, STP, RSTP, MSTP, and VSTP are not supported. PR1075230

Network Management and Monitoring

- 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, traffic sent to the added interface might not be mirrored. PR1057527

Open vSwitch Database Management Protocol (OVSDB)

- OVSDB (QFX10000 switches)-- Use of OVSDB is not recommended on these switches with Junos OS 17.1R1 and 17.1R2. PR1288227

OpenFlow

- OpenFlow (QFX10000 switches)-- Use of OpenFlow is not recommended on these switches with Junos OS 17.1R1 and 17.1R2. PR1288227

Platform and Infrastructure

- On QFX5100 switches, the amount of time that it takes for Zero Touch Provisioning to complete might be lengthy because TFTP might take a long time to fetch required data. PR980530
- On QFX10002 switches, the "request system snapshot" command does not work. PR1048182

- On a QFX10002 switch, insert a small form-factor pluggable (SFP) on the management interface (em1). After a system reboot, replace the SFP with a copper SFP, the management interface might not work properly with speed 10m/100m. PR1075097
- On a mixed-mode Virtual Chassis Fabric (VCF), 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
- While using SSH to log into a VNF, an error with the message **Unrecognized command** is seen. This has no impact on the functionality. PR1108785
- On a fully loaded QFX10008 chassis, line cards might take as much as 15 minutes to become operational after startup. PR1124967
- With multihop BFD, traffic loss of around 5 to 10 seconds is observed when intermediate interface is shut down. After 5 to 10 seconds, traffic recovers and no action is needed. PR1150695
- The Junos OS CLI **file copy** command uses the "/var/home/<user>" as a temporary staging directory for downloads. This behavior is the same on all Junos OS platforms. When you run file copy as “user1,” then /var/home/user1 is the temporary staging directory. For QFX10000 switches, /var is ~480 MB. So running file copy on a file that is more than 480 MB will fail on QFX10000 switches. PR1195599
- Software enhancements were done to improve link stability of ASIC-Hybrid Memory Cube links. PR1208455
- On disable and re-enable of a 1 GB port on a 60X10GB 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. There was no functionality impact observed. PR1238269

Routing Protocols

- L3 multicast traffic does not converge to 100% 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
- 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 switches, traffic drop is seen with IS-IS version 6 traffic during convergence in either of the following two scenarios: 1. While doing port unshutdown (that is, 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
- During a graceful Routing Engine switchover (GRES) on QFX10000 switches, some IPv6 groups might experience momentary traffic loss. This issue occurs when IPv6 traffic is running with multiple paths to the source, and the join-load-balance statement for PIM is also configured. PR1208583
- On EX4300/EX4600/QFX5100/QFX10000 switches, traffic drops might occur in MC-LAG configurations. This occurs when an interchassis data link (ICL) interface and then the MC-LAG interface are brought

up. The traffic drop occurs 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. To avoid the issue, enable ICL interfaces and MC-LAG interfaces at the same time. PR1236201

- On QFX10000 platforms, during route next-hop churn or earliest deadline first (EDF) job priority changes, memory corruption may happen leading to processing issues and constant packet drop. PR1243724

System Management

- On QFX10000 switches with enhanced MC-LAG IRB next hops, member links of the aggregated underlying Layer 2 interfaces might not be present on all Packet Forwarding Engine instances in a given FPC. Under this condition, during IRB next-hop installation for the Packet Forwarding Engine instance where the underlying Layer 2 interface link is not present, failure logs are generated for the PFE uKernel. Those failure logs do not impact traffic or performance on the switch, and they are harmless. PR1221831

SEE ALSO

New and Changed Features | **183**

Changes in Behavior and Syntax | **203**

Known Behavior | **205**

Resolved Issues | **209**

Documentation Updates | **214**

Migration, Upgrade, and Downgrade Instructions | **215**

Product Compatibility | **228**

Resolved Issues

IN THIS SECTION

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: 17.1R2

EVPN

- Route Target per bridge domain for EVPN is not supported. PR1244956
- On QFX10000/QFX5100/QFX5200 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 loose 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-l2ckt: FXPC core is seen when deactivating core interface on MPLS l2ckt 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 | **183**

Changes in Behavior and Syntax | **203**

Known Behavior | **205**

Known Issues | **206**

Documentation Updates | **214**

Migration, Upgrade, and Downgrade Instructions | **215**

Product Compatibility | **228**

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.1R2.

SEE ALSO

New and Changed Features | **183**

Changes in Behavior and Syntax | **203**

Known Behavior | **205**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to https://www.juniper.net/support/downloads/junos.html.

   The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.

3. Select **17.1** 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-5-17.1R2.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 17.1R2.

> **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-10-17.1R1.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-10-17.1R1.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://www.juniper.net/support.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add *<pathname><source>*** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

   ```
   user@switch> configure
   ```

3. Disable Routing Engine redundancy:

   ```
   user@switch# delete chassis redundancy
   ```

4. Disable nonstop-bridging:

   ```
   user@switch# delete protocols layer2-control nonstop-bridging
   ```

5. Save the configuration change on both Routing Engines:

   ```
   user@switch# commit synchronize
   ```

6. Exit the CLI configuration mode:

   ```
   user@switch# exit
   ```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

   ```
   user@switch> request system software add validate
   /var/tmp/jinstall-host-qfx-10-17.2R1.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)
outing 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 225
- Upgrading the Software Using Unified ISSU on page 225

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

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

  To verify that nonstop active routing is enabled:

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

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

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

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

Upgrading the Software Using 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, ***source*jinstall-host-qfx-5-17.1R2.7-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-5-17.1R2.7-signed.tgz ...
     Install jinstall-host-qfx-5-17.1R2.7-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 | 183

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on Routing and Switching devices, see the High Availability User Guide

For additional information about using ISSU on Security devices, see the Chassis Cluster User Guide for SRX Series Devices

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

Compliance Advisor

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

Finding More Information

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

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

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

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF

format. Find Content Explorer at https://www.juniper.net/documentation/content-applications/content-explorer/.

Requesting Technical Support

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

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://www2.juniper.net/kb/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/

- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: https://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

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

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

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

Revision History

3 September 2020—Revision 11, Junos OS Release 17.1R2– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

14 February 2019—Revision 10, Junos OS Release 17.1R2– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

7 February 2019—Revision 9, Junos OS Release 17.1R2– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

17 May 2018—Revision 8, Junos OS Release 17.1R2– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

29 March 2018—Revision 7, Junos OS Release 17.1R2– 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.