Release Notes: Junos® OS Release 17.1R1 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 17.1R1 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 17.1R1 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the features and enhancements in Junos OS Release 17.1R1 for ACX Series Universal Metro Routers.

Application Level Gateways (ALGs)

- **Support for Application Level Gateways (ALGs) for NAT processing (ACX500)**—Starting with Junos OS Release 17.1R1, ACX500 routers supports 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 supports 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 supports 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.]

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 supports 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 supports 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 supports 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 supports 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 supports 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 supports 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 supports 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 supports 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

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 17.1R1 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.

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.

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 17.1, 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 17.1, 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 | 11
- Known Behavior | 19
- Known Issues | 20
- Resolved Issues | 22
- Documentation Updates | 22
- Migration, Upgrade, and Downgrade Instructions | 23
- Product Compatibility | 24

Known Behavior

There are no known limitations in Junos OS Release 17.1R1 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 | 11
- Changes in Behavior and Syntax | 17
- Known Issues | 20
- Resolved Issues | 22
- 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 17.1R1 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 towards the service engine, the resource scale requirements at the service engine cannot be met and service processor may reboot. A workaround is to 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. A workaround is to 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

SEE ALSO

New and Changed Features | **11**

Changes in Behavior and Syntax | **17**

Known Behavior | **19**

Resolved Issues | **22**

Documentation Updates | **22**

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

Product Compatibility | **24**

Resolved Issues

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

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

SEE ALSO

- New and Changed Features | 11
- Changes in Behavior and Syntax | 17
- Known Behavior | 19
- Known Issues | 20
- Documentation Updates | 22
- Migration, Upgrade, and Downgrade Instructions | 23
- Product Compatibility | 24

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R1 for the ACX Series documentation.

SEE ALSO

- New and Changed Features | 11
- Changes in Behavior and Syntax | 17
- Known Behavior | 19
- Known Issues | 20
- Resolved Issues | 22
- 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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 14.2 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

SEE ALSO

- New and Changed Features | 11
- Changes in Behavior and Syntax | 17
- Known Behavior | 19
- Known Issues | 20

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 17.1R1 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 17.1R1 for the EX Series.

> **NOTE:** The following EX Series switches are supported in Release 17.1R1: EX4300, EX4600, and EX9200.

> **NOTE:** In Junos OS Release 17.1R1, 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.

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. QFX5100 switches support 802.1X features including 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 Defining CoS Multidestination (Multicast, Broadcast, DLF) BA 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.

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 17.1R1 for the EX Series.

High Availability 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 Junos OS Release 17.1R1.

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.

Services Applications

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

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

  [See Device Discovery Overview.]

System Management

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

User Interface and Configuration

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

directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1, the generated output files are placed in the **/var/tmp** directory.

Virtual Chassis

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

SEE ALSO

New and Changed Features | 25

Known Behavior | 32

Known Issues | 33

Resolved Issues | 36

Documentation Updates | 39

Migration, Upgrade, and Downgrade Instructions | 39

Product Compatibility | 41

Known Behavior

There are no known limitations for the EX Series switches in Junos OS Release 17.1R1.

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 | 39

Migration, Upgrade, and Downgrade Instructions | 39

Product Compatibility | 41

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R1 for the EX Series.

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

High Availability (HA) and Resiliency

- During a nonstop software upgrade (NSSU) on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. PR1123764
- On an EX4300 Virtual Chassis, when you perform an NSSU, there might be more than 5 seconds of traffic loss for multicast traffic. PR1125155
- Sometimes there could be continuous traffic drop after performing the GRES. As a workaround, reset the new backup member of Virtual Chassis after GRES. PR1189592
- On an EX9200 switch, when you perform an in-service software upgrade (ISSU) while protocol sessions are active, the protocols might go down and come back up again, which can cause traffic loss for up to 40 seconds.PR1247937

Infrastructure

Interfaces and Chassis

- On EX2300 and EX3400 switches, IPV6 neighborship is not created on the IRB interface PR1198482

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 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
- On EX4300 switches, starting in Junos OS Release 15.1R3, a pfex_junos core file might be generated when you add or delete a native VLAN configuration with **flexible-vlan-tagging**. PR1089483
- On an EX9200 Virtual Chassis with JDHCP_Relay_LSYS configurations, the Virtual Chassis linecard members might go up and down after you reboot the switch. PR1108402
- On EX4300-VC platforms, if a Q-in-Q S-VLAN interface with MC-LAG is configured, when the backup EX4300 switch is acting as master, you might lose connection to the management IP address through the interface. As a result, a drop in management traffic might occur. PR1131755
- With nonstop active routing (NSR) configured, for RIP and RIPng, the protocol state information might not be replicated correctly on the backup Routing Engine. PR1149740
- On EX3400 and EX2300 switches, adaptive sFlow sampling might not get triggered when sending packets. PR1150644
- On EX9200 switches, after an ISSU is performed, storm control takes effect only after you delete the storm control configuration and then re-create it. PR1151346
- Unicast reverse-path forwarding is not supported on the EX2300 switch. PR1151632
- On EX2300 switches, when a large number of packets reach the CPU because of a firewall reject action configuration, interface flapping occurs. PR1156553
- 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
- Switch login sometimes fails, with **LOGIN_PAM_ERROR** after reboot due to file sytem full errors. PR1171120

- On EX2300 and EX3400 switches, data traffic is forwarded when the supplicant mode is changed from single to single-secure mode or multiple mode, even when the supplicant is not authenticated. PR1175835
- On EX3400, MF classification does not work on the IRB interface. PR1176253
- On EX2300 and EX3400 switches, Internet Group Management Protocol (IGMP) reports sent to a port are not flooded to the port of the multicast router interface. IGMP general query packets inbound to the multicast router interface port are not flooded to ports connected to a host that is in the same VLAN. PR1177296
- On EX3400 switches, the**request system software add** command takes more time to complete, depending on the number of VC members. PR1178337
- EX2300 switches do not support the Energy Efficient Ethernet feature. PR1178790
- Output for the **show chassis hardware** and **show chassis routing-engine** commands shows the master FPC's SKU for both Routing Engines. PR1178803
- On EX3400 switches, alarms are not generated when fans are removed from the system. PR1179485
- On EX2300 switches, for CoS, medium-low packet loss priority (PLP) is not supported in firewall configurations. PR1180586
- On EX9200 switches, periodic packet management (PPM) core files might be generated following a commit. This happens only on a large-scale setup, when the logical interface number of PFE exceeds 64. PR1187104
- On EX4300 and EX4600 switches and QFX Series switches with VSTP enabled for multiple VLANs and participated in a VSTP topology, when the BPDU packets are received on the Packet Forwarding Engine from other switches, the switch will send BPDU packets to the Routing Engine for further VSTP computing. However, in rare cases, the switch might not send VSTP packets for all VLANs to the Routing Engine. For example, for an uncertain VLAN, BPDU packets do not reach the Routing Engine, even though VSTP is enabled for that VLAN. As a result, the VLAN considers itself to be the root bridge. Because the VLAN advertises itself as the root bridge and sends BPDUs to other VSTP switches, other switches might block the related port. This result might not follow the network design. PR1187499
- On EX Series switches running Junos OS Release 15.1 or later, LLDP PDU gets dropped on the fxp interface. PR1188342
- On EX2300 and EX3400 switches, the **hash-mode** option is not available for the **forwarding-options enhanced-hash-key** command. PR1188866
- On EX3400 switches, MACsec is not supported for uplink ports (xe/ge-*/2/*). PR1189042
- On EX3400 switches, CoS rewrite does not work on the IRB interface. PR1190361
- On EX3400 switches, after you configure PSU redundancy to N+N, it does not revert back to N+0 mode. PR1191731

- On EX2300, EX3400, EX4300, EX4600, and QFX5100 Series switches in a Virtual Chassis configuration, IPv6 multicast packets might not be flooded in a VLAN if IGMP snooping is enabled and the ingress interface is on a different FPC than the egress interface. PR1205416
- On EX9200 switches in a virtual chassis configuration, when DHCP-relay is configured on default and non-default routing instances Virtual Chassis, instability is seen after switch reboot. PR1211648

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. This issue is fixed in Junos OS Release 14.1X53-D40 and late releases.PR1093295

SEE ALSO

New and Changed Features | 25

Changes in Behavior and Syntax | 30

Known Behavior | 32

Resolved Issues | 36

Documentation Updates | 39

Migration, Upgrade, and Downgrade Instructions | 39

Product Compatibility | 41

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: 17.1R1

Authentication and Access Control

- A dot1xd core file might observed during a CoA update for juniper-switching-filter in D55 for 34xx/23xx platform. PR1219538

Infrastructure

- On EX4300 Series switches, when you configure firewall filter on a loopback(lo0) interface to accept BGP flow and another term with discard action, and the receiving host-inbound traffic with a designated TCP port 179 to the Routing Engine, existing BGP sessions might go down. PR1090033
- On an EX4300 switch or EX4300 Virtual Chassis that has a generic routing encapsulation (GRE) tunnel configured on an integrated routing and bridging interface (IRB), the associated GRE statistical counters might not be updated after the GRE interface is deactivated and then reactivated. PR1183521
- On EX4300 switches with DHCP relay configured, DHCP return packets, (such as DHCPREPLY and DHCPOFFER), that are received across a GRE tunnel might not be forwarded to clients, which can impact DHCP services. PR1226868

Interfaces and Chassis

- The interface fxp0 might flap upon some specific commit, this may impact the normal work of out-of-band management.PR1213171
- On an EX9208 switch, when an "interface process" DCD restarts, the switch might disable MC-LAG member links, resulting in traffic loss. PR1229001

Platform and Infrastructure

- On an EX4300, if you install a firewall filter with filter-based forwarding rules to multiple bind points, it might exhaust the available TCAM. In this case, the filter is deleted from all the bind points. PR1214151
- On EX4300, eBGP packets with ttl=1 and non-eBGP packet with ttl=1, whether destined for the device or even transit traffic, go to the same queue. In the event of a heavy inflow of non-eBGP ttl=1 packets, occasionally valid eBGP packets would get dropped causing eBGP to flap. In Junos OS Release 14.1X53-D40 and later releases, eBGP packets with ttl=1 go to a different queue. PR1215863
- When **set vlans interface all** configuration is used on an EX4300 the device control process (dcd) might generate a core file because this is not a supported configuration option using L2NG syntax. PR1221803
- On EX4300 series switches, if L3 interface receives a frame with the CFI/DEI bit set to 1, then this frame would be dropped and not processed further. PR1237945
- In Junos OS Release 16.1, due to the factory-default file that gets activated post zeroize, EX4300 can contain more interfaces to cater to a 10-member Virtual Chassis default configuration, even if the interfaces are not physically there or if there is only a standalone device. PR1238848
- On EX4300 switches, problems with connectivity might arise on 100M interfaces set to full duplex and half duplex or on 10M interfaces set to full duplex or half duplex. The links appear, but connectivity to

end devices might not work. The port does not transmit packets even though port statistics show packets as transmitted. PR1249170

- On EX3400, binding does not occur on some of the IPV6 clients if two DHCPv6 relays are present with VRRP between them. PR1189333
- When a firewall filter has tcp-reset applied to the IRB interface, the action does not work properly. PR1219953
- In single supplicant mode, the **show captive-portal** command might not display the client, even though the client gets authenticated successfully. This issue is seen in Junos OS Release 17.1R1 only. PR1240259
- On EX3400 switches in a Virtual Chassis configuration, the RA guard enabled interface stays in trusted state even after the **mark-interface trusted** configuration is deleted. PR1242937
- On EX3400 switches in a Virtual Chassis configuration, executing **request access-security router-advertisement-guard-block interface** and **restart dhcp-service** commands triggers the jdhcpd process to generate a core file.PR1243147
- On EX3400 switches in a Virtual Chassis configuration, RA guard does not discard unauthorized packets.PR1244666

Virtual Chassis

- On EX4300 switches, a message such as **/kernel: %KERN-5: tcp_timer_keep: Dropping socket connection due to keepalive timer expiration** might be seen repeatedly. There is no service impact from the condition that causes the message (a Packet Forwarding Engine timeout trying to connect to a daemon that is not active). PR1209847

SEE ALSO

New and Changed Features | **25**

Changes in Behavior and Syntax | **30**

Known Behavior | **32**

Known Issues | **33**

Documentation Updates | **39**

Migration, Upgrade, and Downgrade Instructions | **39**

Product Compatibility | **41**

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 17.1R1 for the EX Series switches documentation.

Virtual Chassis

- Starting with Junos OS Release 17.1R1, EX9200 Virtual Chassis is no longer supported, and the EX9200 Virtual Chassis documentation has been archived. See EX Series Documentation Archives. For deployments with EX9200 switches, we recommend planning or moving to MC-LAG or Junos Fusion Enterprise architectures instead of using a Virtual Chassis.

SEE ALSO

New and Changed Features | 25

Changes in Behavior and Syntax | 30

Known Behavior | 32

Known Issues | 33

Resolved Issues | 36

Migration, Upgrade, and Downgrade Instructions | 39

Product Compatibility | 41

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 14.2 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

SEE ALSO

- New and Changed Features | 25
- Changes in Behavior and Syntax | 30
- Known Behavior | 32
- Known Issues | 33
- Resolved Issues | 36
- Documentation Updates | 39
- Product Compatibility | 41

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

- New and Changed Features | 25
- Changes in Behavior and Syntax | 30
- Known Behavior | 32
- Known Issues | 33
- Resolved Issues | 36
- Documentation Updates | 39
- Migration, Upgrade, and Downgrade Instructions | 39

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R1 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 17.1R1 for Junos Fusion Enterprise.

> **NOTE:** For more information about the Junos Fusion Enterprise features, see the Junos Fusion Enterprise User Guide.

Hardware

- **Satellite device support (EX2300)**—Starting with Junos OS Release 17.1R1, you can configure EX2300 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].

- **Satellite device support (EX3400)**—Starting with Junos OS Release 17.1R1, you can configure 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 17.1R1 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 | 42

Known Behavior | 48

Known Issues | 50

Resolved Issues | 52

Documentation Updates | 53

Migration, Upgrade, and Downgrade Instructions | 53

Product Compatibility | 62

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 17.1R1 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

- In a Junos Fusion Enterprise topology with dual aggregation devices, firewall statistics are not synchronized across the aggregation devices. PR1105612
- 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
- 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

SEE ALSO

New and Changed Features | 42

Changes in Behavior and Syntax | 48

Known Issues | 50

Resolved Issues | 52

Documentation Updates | 53

Migration, Upgrade, and Downgrade Instructions | 53

Product Compatibility | 62

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R1 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, 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, 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

- Loss of connectivity of the link connecting the Standalone box might lead to conversion failure from Junos OS to satellite software (SNOS). 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, 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

SEE ALSO

New and Changed Features | 42

Changes in Behavior and Syntax | 48

Known Behavior | 48

Resolved Issues | 52

Documentation Updates | 53

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.

Junos Fusion Enterprise

- On Junos Fusion Enterprise, Power over Ethernet (PoE) telemetries do not work. PR1112953
- On a Junos Fusion Enterprise, Power over Ethernet (PoE) configuration changes may not be reflected on satellite devices that are not in the online state at the time of the configuration change. 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
- In a Junos Fusion Enterprise that has rebooted a satellite device in a satellite device cluster, traffic can be lost for several seconds after the satellite device returns to an operational state. PR1168820
- In a Junos Fusion Enterprise, the SNMP trap that should be sent for a satellite device reboot event is not sent. PR1182895
- In a Junos Fusion Enterprise, LLDP might stop working if it is reenabled after being manually disabled. PR1188254

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R1 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 17.1R1:

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*/junos-install-ex92xx-x86-64-17.1R1.n.tgz**

All other customers, use the following commands.

`user@host>` **request system software add validate reboot *source*/junos-install-ex92xx-x86-64-17.1R1.n-limited.tgz**

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

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

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

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

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 17.1R1 **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 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 16.1R1 to 17.1R1.

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 17.1R1 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 17.1R1 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 higher.
- The Junos switch must be either set to factory default configuration 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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 14.2 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Downgrading from Release 17.1

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

NOTE: It is not recommended to downgrade the aggregation device from 17.1R1 to 16.1 if there are cluster satellite devices in the setup.

To downgrade a Junos Fusion Enterprise from Junos OS Release 17.1R1 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 16.1R1.

2. Downgrade the aggregation device to 16.1R1. 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 | 42
- Changes in Behavior and Syntax | 48
- Known Behavior | 48
- Known Issues | 50
- Resolved Issues | 52
- Documentation Updates | 53
- Product Compatibility | 62

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

- New and Changed Features | 42
- Changes in Behavior and Syntax | 48
- Known Behavior | 48
- Known Issues | 50
- Resolved Issues | 52
- Documentation Updates | 53
- Migration, Upgrade, and Downgrade Instructions | 53

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R1 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 17.1R1 for Junos Fusion Provider Edge.

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 17.1R1 or later 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 17.1R1 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

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R1 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

- On Junos Fusion Provider Edge, the transit traffic received from the LAG of the extended ports is still forwarded even when the minimum-link condition is not met. To avoid this issue, make sure you have the same number of minimum links on the other end. PR1188482

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.

Junos Fusion

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

  1. The Junos OS switch must use factory-default settings or it must have include **set chassis auto-satellite-conversion** in its configuration.

  2. The package used to do the conversion must be SNOS 3.0, SNOS 1.0R5, SNOS 2.0R2, or later. PR1249877

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R1 for Junos Fusion Provider Edge documentation.

SEE ALSO

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 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.

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

- For 64-bit software:

> **NOTE:** We highly recommendthat you see 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

```
user@host> request system software add validate reboot
source/jinstall64-17.1R1.9-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-17.1R1.9-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

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

```
user@host> request system software add validate reboot
source/jinstall64-17.1R1.9-export-signed.tgz
```

- For 32-bit software:

  `user@host>` **request system software add validate reboot *source*/jinstall-17.1R1.9-export-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*** (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 17.1R1 **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 17.1R1, 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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 14.2 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Downgrading from Release 17.1

To downgrade from Release 17.1 to another supported release, follow the procedure for upgrading, but replace the 17.1 **jinstall** package with one that corresponds to the appropriate release.

> **NOTE:** You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

SEE ALSO

- New and Changed Features | 64
- Changes in Behavior and Syntax | 65
- Known Behavior | 65
- Known Issues | 66
- Resolved Issues | 67
- Documentation Updates | 68
- Product Compatibility | 77

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

New and Changed Features | 64
Changes in Behavior and Syntax | 65
Known Behavior | 65
Known Issues | 66
Resolved Issues | 67
Documentation Updates | 68
Migration, Upgrade, and Downgrade Instructions | 68

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

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R1 for the MX Series routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 17.1R1 for the MX Series routers.

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 Release17.1R1, 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 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)**—In Junos OS release 17.1R1, support is provided for the 1-port 100-Gigabit dense wavelength division multiplexing (DWDM) optical transport network (OTN) MIC (MIC3-100G-DWDM) with CFP2 analog

coherent optical (CFP2-ACO) pluggable optics on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG). The 100-Gigabit Ethernet DWDM OTN MIC supports the following features:

- Transparent transport of 100-Gigabit Ethernet signals with optical channel transport unit, OTU4 (V) framing
- Dual-polarization quadrature phase shift keying (DP-QPSK) modulation with coherent receiver and soft-decision forward error correction (SD-FEC) for long-haul and metro applications
- International Telecommunication Union (ITU)-standard OTN performance monitoring and alarm management
- Extensive optical, digital signal processing (DSP), and bit error ratio (BER) performance monitoring statistics for the optical link

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

Class of Service (CoS)

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

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

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

EVPNs

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

- MPLS-reachability failure between the local PE device and a remote PE device

[See EVPN Multihoming Overview]

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

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

  [See Overview of VPWS Service with EVPN Signaling Mechanisms]

- **Support for inter-data center connectivity over pure Layer 3 network with EVPN (MX Series routers with MPCs)**—Starting in Junos OS Release 17.1R1, the control plane EVPN Type-5 supports IP prefix for inter-subnet connectivity across data centers. The data packet is sent as the L2 Ethernet frame encapsulated in the VXLAN header over the IP network across the data centers to reach the tenant through the connectivity provided by the EVPN Type-5 IP prefix route.

  [See EVPN Type-5 Route with VXLAN encapsulation for EVPN/VXLAN]

- **Support for LACP in EVPN active-active multihoming (MX Series routers with MPCs)**—Starting with Junos OS Release 17.1R1, an extra level of redundancy can be achieved in an Ethernet VPN (EVPN) active-active multihoming network by configuring the Link Aggregation Control Protocol (LACP) on both the endpoints of the link between the multihomed customer edge (CE) and provider edge (PE) devices. The link aggregation group (LAG) interface of the multihomed CE-PE link can either be in the active or in the standby state. The interface state is monitored and operated by LACP to ensure fast convergence on isolation of a multihomed PE device from the core.

  When there is a core failure, a traffic black hole can occur at the isolated PE device. With the support for LACP on the CE-PE link, at the time of core isolation, the CE-facing interface of the multihomed PE device is set to the standby state, thereby blocking data traffic transmission from and toward the multihomed CE device. After the core recovers from the failure, the interface state is switched back from standby to active.

  To configure LACP in 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**—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 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**—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**—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**—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 17.1 introduces abstract hops, which are user-defined router clusters or groups that can be sequenced and used for setting up a label-switched path (LSP), similar to real-hop constraints.

  The router groups are created using constituent lists that include constituent attributes, which 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.]

  [See 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 17.1, 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.
  - CoS features such as shaping applied to an agent circuit identifier (ACI) interface set and its members.

  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 17.1, you can deploy vMX routers on x86 servers. FreeBSD 10 is the underlying OS for Junos OS for vMX.

  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.

Routing Policy and Firewall Filters

- **Support for force premium firewall filter added for the Bridge, CCC, and VPLS families (MX Series)**—Starting in Junos OS Release 17.1R1, you can include the **force-premium** option at the **[edit firewall family *(bridge | ccc | vpls)* filter *filter-name* term *term-name*]** hierarchy level to ensure that traffic matching the firewall filter term is treated as premium traffic by a subsequent hierarchical policer, regardless of its forwarding class.

  > **NOTE:** The **force-premium** filter option is supported only on MPCs.

  [See force-premium (Firewall Filter Action).]

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 17.1, 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, that 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 17.1, 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 17.1, 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 | 102

Known Behavior | 109

Known Issues | 111

Resolved Issues | 122

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 17.1R1 for MX Series routers.

Interfaces and Chassis

- **Support for maximum queues configuration on MPC7E, MPC8E, and MPC9E (MX Series)**—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.]

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**—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']/**.

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**—Starting with JUNOS OS Release 17.1R1, the termination of single, multiple, and dual-tagged service delimited vlans are transported over a single ethernet-ccc pseudowire using ps virtual port devices. This feature provides scaled layer-3 service application at the pseudowire head-end termination appliance. This behavior is as an extension and evolution for ethernet pseudowire that is described in RFC 4448.

Network Management and Monitoring

- **SNMP syslog messages changed (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 SNMP 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.

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**—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.

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.]

Security

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

cdn	**scccn**
hello	**sccrq**

iccn	**stopccn**
icrq	**unclassified**

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

[See protocols (DDoS).]

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

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

  See Configuring How Flow Detection Operates Globally and Configuring How Traffic in a Culprit Flow Is Controlled Globally.

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.]

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 17.1, 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 17.1R1 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

- **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

- **rpd process may 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 may 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.

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.

SEE ALSO

New and Changed Features | **78**

Changes in Behavior and Syntax | **102**

Known Issues | **111**

Resolved Issues | **122**

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 config commit with following configurations: **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

- 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. That is the architecture. Since there is no route change, there is no route flash, so route_record is blissfully unaware. In order to change this, we would need to decide that we want a route flash for this case. Currently, for direct and local routes / nexthops, these are "don't care" in rpd, as far as route updates go. We just update our nexthop info, without marking for any other notifications. To change this, we would need to find the correct place to decide we need to flash the route, and at the same time, make sure we don't do any harm to anything else. That is what I am currently working on finding. A complication for the solution is a change that was done for PR 1002287, where if the NOTINSTALL flag is set, do not send the update to srrd. That flag is set for direct and local routes. Incidently, this is day-one operation. If the interface is up at startup, it should all work correctly. Why is the pfe depending on rpd / srrd to get the info for sampling when it is already there in the forwarding table ? +++ FIB table can provide OIF/GW only. SRC_MASK, DST_MASK, SRC_AS and DST_AS are not available in PFE FIB Table. So SRRD connection is required. Listening to both SRRD and FIB table, and consolidating information will complicate implementation. Scanning entire FIB Table just for the few such routes will have performance impact and will complicate present implementation. This is day 1 implementation for SRRD/Sampled. Workarounds: ++++++++++++ There are two possible workarounds a) A workaround would be to have the far end interface up when the DUT interface is brought up. In the case where that is not happening, a recovery would be to disable the DUT interface, then enable it again. At that point, everything should be initially brought up in the state we are looking for. b) enable nexthop-learning knob. Please refer to the documentation on the working of this knob before enabling. PR1224105

- Firewall filter family "any" with shared-bandwidth-policer on MC-AE interface does not reconfigure bandwidth or carve-up policer when standby becomes active after A/S switchover, it drops all packets. PR1232607

- "dfwinfo: tvptest:dfwlib_owner_create tvp driven policer_byte_count support 0" message is seen after show firewall. This behavior is 16.1 specific and cosmetic issue. << sample config >> set interfaces ge-0/0/0 unit 0 family inet filter input test_filter set interfaces ge-0/0/0 unit 0 family inet address 100.100.100.1/24 set firewall family inet filter test_filter term policer then policer policer_test set firewall policer policer_test if-exceeding bandwidth-limit 100m set firewall policer policer_test if-exceeding burst-size-limit 125k set firewall policer policer_test then loss-priority low PR1248134
- FreeBSD 10.x based Junos OS is not supported on 32-bit Routing Engines from Junos OS release 17.1R1. PR1252662

General Routing

- In an IPsec load-balancing environment using MS-MPC cards, the ICMP request and ICMP reply can go through two different IPsec tunnels due to asymmetric routing; that is, ICMP request goes through one PIC, and ICMP reply goes through another PIC. Because of this, the ICMP reply will get dropped and never reach the other side of the IPsec tunnel. PR1059940
- ICMP echo_reply traffic with applications like IPSec will not work with the MS-MIC and MS-MPC cards in a asymmetric traffic environment since these cards employ a stateful firewall by default. The packet will be dropped at the Stateful Firewall since it sees an ICMP Reply that has not 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(e.g. an ae interface), mirrored packets may get dropped. PR1134523
- This is an intermittent issue. Assuming that AE is configured with the bypass-queuing-chip configuration statement. Now followup configuration changes are such that removing child link(s) from AE bundle, 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
- During SIB yanking (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
- Starting with JunOS 15.1F5-S2 15.1F6 16.2R1 17.1R1 on vMX series, it introduces new CLI command "set chassis fpc X performance-mode num-of-ucode-workers Y" to support dedicated users for control and multicast traffic. This will avoid unicast traffic to be hashed to users doing ucode processing. As usual vMX X86 can be configured to run in lite-mode or performance-mode. With this new CLI option, users are allowed to configure number of ucode workers to process multicast and control traffic on separate worker cores. Intention of this command to separate flow cache and non-flow cache traffic, but as part of this fix, only control and multicast traffic will be separated from remaining traffic. In future, we could move other non-flow cache traffic to this dedicated ucode workers. Whenever there is change

in num-of-ucode-workers, RIOT will be rebooted and first Y workers will process control and multicast traffic and remaining workers will process flow cache traffic. PR1178811

- 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. As it is seen below, 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 at a few times 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
- When IPv4 firewall filter have 2625/32 destination in prefix-list , filter attached to subscriber interface is found broken. PR1184543
- AMS redundant interfaces not listed under possible-completions of operational commands. PR1185710
- On MX platform with Junos release equals to or lager than 15.1, LLDP PDU gets dropped on FXP interface. PR1188342
- On MX series with NAT service configured on AMS interfaces. after rebooting FPC/PIC, the NAT pool split between AMS members is incorrect, there are overlapping IP pools and sometimes missing pools, causing NAT not working correctly. PR1190461
- 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 error 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. Without this software improvement, such error can be cleared by rebooting the FPC. PR1200503
- when ppm deviation exceeds 10 ppm, do not display off-frequency if the clock source is still being locked. Display as 'in-use#' instead. This indicates that it is still locked to the source, although the clock has 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. With this fix, allow stable and persistent dynamic tunnel even for non-forwarding routes. PR1202926
- 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 firewall filter related events result in filter structure getting deleted and re-created again. PR1205325
- 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. That causes a circular dependency as 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
- There is no ISSU from 15.1 and older releases to 16.2R1. PR1222540
- The problem of tunnel stream getting mis configured for LT interfaces is due to internal programming and the same has been corrected to evaluate multiple lt interfaces for FPC and PIC slot combination. PR1223087
- with qmon sensor, when you issue an operational clear command, such as clear interfaces statistics all, the counters at the telemetry Junos Telemetry Interface server are not reset. Hence qmon sensor stats at Junos Telemetry Interface server won't match with the CLI/VTY commands output, after the "clear interfaces statistics" commands. PR1226948
- Change of behavior of reflexive keyword. If the "reflexive" keyword is configured in a COS rule, then the COS-service-plugin will store COS-VALUES [DSCP, Forwarding class] received in the forward flow and apply the same COS-VALUES [DSCP, Forwarding class] to packets going back in reverse flow. PR1227021
- Continuously increasing normal discard count in 'show pfe statistics traffic' without any user traffic due to an internal control traffic which is expected to be dropped silently is unexpectedly being counted as 'normal discard'. There's no impact on user traffic with this issue. PR1227162
- A wrong PE is being attached to an ESI when the router receives two copies of the same AD/ESI route (e.g. one through eBGP and another one received from an iBGP neighbor). This will causes partial traffic blackhaule and stale MAC entries. You can confirm the issue by checking the members of the ESI: labroot@MX2> 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 DEP router. The IPSEC Sas are distributed on the DEP router. After that the neighbor state between the ospf peers also move to full but after that it doesn't stay in that state consistently and it changes to various states like init, 2-way, ex-start to full again.

Because of that I see the data traffic between the routers getting dropped. So tunnel distribution with protocol traffic in not supported PR1232277

- Changing virtual switch type from IRB type to regular bridge, interfaces under openflow protocol got all removed. Openflow daemon failed 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 would still send 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 might be seen, and causing rpd crash. PR1238333
- On MX with Junos Telemetry Interface deployed, streaming telemetry stops being sent for Message Queue Telemetry Transport (mqttd) memory reaches to max of 4G. PR1238803
- For ANCP subscribers in Idle state the previously reported speed in ANCP Port UP message is not applied. PR1242992
- ANCP neighbors going down after commit in case any ANCP related configuration was changed. PR1243164
- On MX2000 MPC6E, EOAM LFM adjacency flaps when 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 with CLI "show vpls mac-table", though all subscribers have traffic. Thus it is considered a cosmetic issue. PR1257605
- Issue: Upon restart na-grpcd overwrites the previous log file. Analysis: na-grpcd opens log file in write mode due to which whenever na-grpcd restart, the files were overwritten. Fix: Fix is to open the log file in append mode. PR1258484
- 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 each such events, all fabric traffic is queued for this PFE reporting the error and causes very high amount of fabric drops. PR1265385
- For very large tables or very slow consumers, the BgpRouteMonitorRegister() may not send the operation BgpRouteMonitorOper.END_OF_RIBS after all of the initial updates are sent. PR1265427

- Database status may remain as not ready after several SO PR1271306
- With MPC7E, MPC8E or MPC9E, when a 40GE or 100GE port is configured under an Aggregated Ethernet bundle, some received packets are incorrectly dropped with "DA reject" reason under the "show interface extensive" output on the corresponding physical interface. This is due to a misconfiguration on the Aggregated Ethernet MAC address under the Packet Forwarding Engine. PR1274073

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 config : "set system ports console log-out-on-disconnect" when set, logs the user out from the console and closes the console connection . If the config "set system syslog console any warning" is used along with the earlier config and when 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. This issue can be worked around by removing the later config, "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 IPv6 Link Local address is not assigned. PR1063404
- During configuration changes and reuse of Virtual IP on an interface as a interface address; It is required to delete the configuration do a commit and then add the interface address configuration in the following commit. PR1191371
- On MX Series IPV6 neighbor-ship is not created on IRB interface PR1198482
- 1. Delay Measurement support for 5-port 100G DWDM PIC and 5-port 100G DWDM MIC is *ONE TIME Delay Measurement*. This means that Customer intending 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 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 customer wants to verify this, 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 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

- t3 interfaces configured with "compatibility-mode digital-link" may fail to come up due incorrect subrate. To verify, 'show interfaces t3-0/0/0 extensive' will display the subrate. DSU configuration: Compatibility mode: Digital Link, Scrambling: Enabled, Subrate: 4195338 Kbps <<<< // expected result: DSU configuration: Compatibility mode: Digital Link, Scrambling: Enabled, Subrate: Disabled PR1238395
- 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

Layer 2 Features

- On routers running Junos OS with Routing Engine GRES enabled, if vpls is configured with a dynamic-profile association, some traffic loss would be observed when Routing Engine switches from master to standby. This is due to a change in underlying database that handles the dynamic-profile sessions which causes the vpls connection to destroy and re-create after a Routing Engine switchover. PR1220171

Layer 2 Ethernet Services

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

MPLS

- In BGP prefix-independent convergence (PIC) edge scenario, when the ingress route (the primary route) fails, due to the fact that LDP may fail to send the session down event to PFE correctly, the PFE may still use the primary path to forward traffic until (in some cases, 3- 5 seconds for 30k prefixes) the global convergence is completed by the interior gateway protocol (IGP). In addition, the issue may also be seen when the "delay-delete" knob is configured, in this scenario, the session down event may get sent to the PFE correctly, however, due to local reversion, the primary path may also be chosen as forwarding path when it is deleted. PR1097642
- When graceful Routing Engine switchover (GRES) is done between the master and backup Routing Engines of different memory capabilities (such that one has only enough memory to run routing protocol process (rpd) in 32-bit mode while the other is capable of 64-bit mode, which could be caused by using Junos OS Release 13.3 onwards with the configuration statement "auto-64-bit" configured, or, using Junos OS Release 15.1 onwards even without the configuration statement), rpd might crash on the new master Routing Engine. As a workaround, this issue could be avoided by the CLI command "set system processes routing force-32-bit". PR1141728

- In MVPN scenario, if active primary path goes down, then PLR(Point of Local Repair) needs to send Label Withdraw for old path and new Label Mapping for new path to the new upstream neighbor. In this case, LDP P2MP path may stay in "Inactive" state for 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
- On MX-Series and PTX-Series platform, the LDP may fail to install LDP route in inet.3 table if ISIS is configured with source-packet-routing and ldp-tunneling is enable, which might cause the LDP to fail to install routes when L-ISIS routes are present. PR1248336
- A new configuration "protocols mpls traffic-engineering bgp-igp-both-ribs" in the routing-instance is required to make COC work. PR1252043
- During MBB, Next-Hop will change in PFE RSVP route doesn't request for a Next-Hop ACK before changing the route pointing to new NH When scale is high, traffic loss can be seen up to 1 second PR1264089
- When a container LSP has >10 member LSPs only the first 10 LSP will be shown in the **show mpls container-lsp name <lsp-name> statistics** output. PR1267774

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 below 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 may not use the selective acknowledgement (SACK) TCP extensions. PR1024798
- In configurations with IRB interfaces, during times of interface deletion, such as 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 re-converging on the expected state. PR1054798
- On MX platform, parity memory errors might happen in pre-classifier engines within a MPC. Packets will be silently discarded as such errors are not reported and makes 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
- CoS error messages might appear when nonexistent path for database file is configured for CoS, these messages do not affect any service and traffic. PR1158127
- The delegated BFD session over AE interface failed to come up after FEB switchover with FEB redundancy group (1:1 and 1:N) PR1169018

- 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
- In the earlier Junos releases < 17.2, MGD with extend-db feature supports 2.5G size of database (maximum) on 64 bit platforms which is a bug fixed through this PR. After this PR the max config db size supported with extend-db feature is 1.5G on i386 platforms (on both 32 bit and 64 bit) PR1228629
- From 13.3, the SRX cluster needs running auditd on both node. But on MX-VC Bm and TXP all LCC also add auditd. Because LCC and VC-BM don't have route for accounting server, so it will generate 1813 unreachable infor. test@router-re0> 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
- If there are more than 500 AS numbers in the AS path the routing process could restart. PR461329
- 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 large number of configured VPNs, routes changing in the midst of a bgp path-selection configuration change can sometimes lead to an rpd core. This core has been seen with the removal of the "always-compare-med" option. PR1213131
- RPD leaks memory with the topology and configuration attached to this PR. However it's been confirmed that adding/deleting static flowspec routes in isolation doesn't cause any memory leak. Exact Configuration which causes the leak is still unknown. PR1213959
- PIM NSR Design :- - With GRES+ NSR enabled, master Routing Engine replicate kernel states and protocol states on backup Routing Engine - Both kernel state (ifstates) and protocol state replication are independent processes. - ksyncd takes care of ifstates replication - RPD infra takes care of replication (mirror) connection between two Routing Engine - And NSR supported protocols have their own mechanism to replicate their database using mirror connection - As per PIM/MVPN NSR design, on backup Routing Engine, it walks through replication database (RDB) with ?consume & delete? action i.e. once a PIM/MVPN states is processed on backup, associated RDB is deleted - If ?kernel replication? is restarted, which lead to interface delete/add on backup Routing Engine. PIM states on backup goes out of sync . That’s a caveat. - ?kernel replication? restart lead to interface delete/add on backup Routing Engine only - PIM/MVPN does not have RDB on backup Routing Engine, so - On interface delete, it delete the relevant PIM state - Once interface is added by kernel, PIM has no state to consume - No change on master Routing Engine to re-initiate the protocol replication - 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

Services Applications

- On MX series with L2TP configured, for some reason 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 current interaction scheme between ANCP and L2TP daemons and can last up to 5 seconds. In a production network scenario this delay shouldn't be visible as L2TP daemon checks for state updates each time when there is an L2TP packet that has to be sent or received. PR1234674
- In case if l2tp subscriber has static pp0 interface on LAC side, LCP re-negitiation is configured on LNS side and CPE has been changed, it can cause an issue with successful negotiation of PPP session between LNS and CPE PR1235554

Subscriber Access Management

- On MX Series routers with subscriber management feature enabled, after GRES switchover "show network-access aaa statistics radius" CLI command display only zeros and "clear network-access aaa statistics radius" doesn't clear statistics as it should. It's a cosmetic issue and communication with Radius server is working fine, the only impact is that affected CLI commands do not work as expected. PR1208735
- Subscribers 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

SEE ALSO

New and Changed Features | **78**

Changes in Behavior and Syntax | **102**

Known Behavior | **109**

Resolved Issues | **122**

Documentation Updates | **133**

Migration, Upgrade, and Downgrade Instructions | **134**

Product Compatibility | **143**

Resolved Issues

IN THIS SECTION

- Resolved Issues for 17.1R1 | **123**

This section lists the issues fixed in the Junos OS 17.1R1 Release 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.

Resolved Issues for 17.1R1

Class of Service (CoS)

- When the "chained-composite-next-hop" is enabled for Layer 3 VPN routes, MPLS CoS rewrite rules attached to the core-facing interface for "protocol mpls-inet-both-non-vpn" are applied not only to non-VPN traffic (which is the correct behavior) but also to Layer 3 VPN traffic. That is, both MPLS and IP headers in Layer 3 VPN traffic receive CoS rewrite. PR1062648
- The CLI command **show interfaces queue <if-name>** has three display options: 1. **show interfaces queue <if-name>** Displays queued/transmitted/dropped packets/bytes for all IFD children. 2. **show interfaces queue <if-name> aggregate** Displays queued/transmitted/dropped packets/bytes for all IFD children except for IFD RTP traffic 3. **show interfaces queue <if-name> remaining** Displays queued/transmitted/dropped packets/bytes for IFD RTP traffic only.

  Note that unlike queued/transmitted/dropped counters, queues depth values cannot be aggregated. With changes done in this PR, the following is true for queues depth values: 1. **show interfaces queue <if-name>** Displays queues depth values for RTP queues 2. **show interfaces queue <if-name> aggregate** Displays queues depth values for RTP queues 3. **show interfaces queue <if-name> remaining** Displays queues depth values for RTP queues. The above logic is the same for physical interfaces, interface-sets and for logical interfaces units. PR1226558
- Following error log message will be seen with Hierarchical CoS configuration. Dec 27 11:08:02.293 mand-re0 fpc1 cos_check_temporal_buffer_status: IFD ge-1/2/1 IFL 358: **Delay buffer computation incorrect. If hierarchical scheduler is configured for an IFD and if guaranteed rate is not set for an IFL under this IFD, then the temporal buffer configured.** The display of error message is valid when guaranteed rate is 0, but it is not valid when guaranteed rate is disabled. PR1238719

Forwarding and Sampling

- If a two-color policer is configured on MX Series MPCs/MICs, more traffic than the limited traffic might be passed when packets size is less than 128 bytes. PR1207810
- Bandwidth-percent policer does not work on ps interface, which will result in commit error. PR1225977
- On MX Series and EX Series device with the "ipv4-flow-table-size" or "ipv6-flow-table-size" configuration statement, if sampling instance is not defined under chassis hierarchy (sampling instance is not associated to FPC), after rebooting the device, the "ipv4-flow-table-size" or "ipv6-flow-table-size" does not propagate to FPC. PR1234905
- J-Flow version 9 cannot get TCP flag information from IPv6 fragment packets. However, it can get other information like src and dst ports information. It can get sampling information partially from TCP header in IPv6 fragment packets. PR1239817

General Routing

- On MX Series platforms with MS-MPC/MS-MIC in use, due to some reason if the NAT session is freed/removed but without removing timer wheel entry, then it might cause MS-MPC/MS-MIC crash. It is a timing issue where just before invoking the timer wheel callback the NAT session extension got freed/removed. PR1117662
- On MX Series router with services PIC (MS-DPC/MS-MPC/MS-MIC), the ICMP time exceeded error packet is not generated on an IPsec router on the decap side. PR1163472
- This is a display issue and does not affect functionality of the power. The **show chassis power** and **show chassis environment pem** commands have been amended.PR1177536
- Fragmented ALG control traffic is not supported on the MS-MPC or MS-MIC. PR1182910
- During unified in-service software upgrade the following log messages might be displayed: **SFP: pointer Null, sfp_set_present**. This might trigger a flap in the interfaces on MX Series routers while upgrading using unified ISSU. PR1200045
- When PPPoE subscribers log in or log out of the device, the, SNMP link up/down trap will be generated by the system if "no-trap" is configured in the corresponding dynamic profile. PR1204949
- Problem -- In case of local source and with asm MoFRR enabled, the default MDT traffic loops back to the originating router on the MoFRR backup interface, thereby causing continous IIF_mismatches. MoFRR behavior after fix –- with the current MoFRR code, since the source is local, SPT BIT is set by default, hence an (S,G,rpt) PRUNE is sent out of MoFRR active interface. But an (S,G,rpt) PRUNE is not sent out of MoFRR backup interface (Missing Code). With the new fix, (S,G,rpt) PRUNE is sent over the MoFRR backup path also (if there is already an (S,G,rpt Prune) going out of the MoFRR active Path) in order to avoid IIF_Mismatches. PR1206121
- On MX Series platforms, if any inline feature is configured(for example, inline BFD, CFM, or PPP), the FPC might crash and core files are generated. PR1210060
- MS-MPC/MS-MIC might crash when large fragmented (larger than 2048 bytes) traffic hits the pinhole opened by an ALG. PR1214134
- The AMS interface is configured in warm-standby mode and when failover occurs, a percentage of the traffic might fail to get NAT. The issue is after the failover the internal mappings driving traffic back to the service PIC might fail. PR1216030
- When VPLS instances are configured for the first time or when a system with VPLS instances is rebooted, RPD will consume high CPU usage (100 percent) for 10-20 minutes. The installation of other routes might be deferred and traffic will be lost. Many other RPD services might also slow down or be unavailable. PR1216332
- Due to a software issue, replacing an MQ FPC (MPC Type1, 2, MPC 3D 16x10GE) with an XM one (MPC Type 3,4,5 6. 2E-NG, 3E-NG) might cause all other MQ-based cards to report "FI Cell underflow at the state stage". Packets will be dropped. PR1219444
- On MX Series Virtual Chassis, partial or complete traffic loss for streams via AE interfaces might be observed in certain scenarios. For example, if VCP port configurations were deleted and then reconfigured,

then two consecutive global GRES switchovers were performed and an MPC hosting AE child links was reloaded, traffic loss would be observed after the MPC boots up because of incorrect programming of the AE interface on its Packet Forwarding Engine. PR1220934

- PPPoE/DHCP subscribers fail to bind due to ProcessPADIFailedUiflNotActive/SML_CLIENT_DELETE_SDB_ADD_FAILED errors. It is seen during inflight tests. PR1221690
- After Junos OS Release 15.1, the behavior of storage devices enumeration in kernel level has been changed. Device enumeration in earlier releases of Junos OS (before 15.1) show CF and Disk as ad0 and ad1 respectively. Device enumeration after Junos OS Release 15.1 show CF and Disk as ad1 and ad0 instead in the result of **show chassis hardware**. This might be inconsistent for other result of output, such as **show system boot-messages** and **show log messages**. PR1222330
- On setup with IRB configuration statement and non-enhanced-IPmode, when actions result in the underlying AE interface of IRB going down, the backup Routing Engine might experience panic and hence reboot. The panic is caused by the inability to allocate the next-hop index that the master Routing Engine has requested. Because the panic and reboot occurs backup Routing Engine, routing, forwarding, or any other functionality will not be affected. Some examples of triggers are - continuous child link flaps of AE or back-to-back commits of different IRB configuration statements or activating/deactivating bridge family on an underlying interface. PR1222582
- In enhanced subscriber management environment (**set system services subscriber-management enable**) when the **remove-when-no-subscribers** configuration statement is configured in auto-configure stanza, when last subscriber logs out (which is triggering the dynamic VLAN IFL removal) and a new subscriber logs in before the IFL is set to inactive, the dynamic profile deletion might fail. Subsequent subscriber login failure is also seen. PR1222829
- In MX Series Virtual Chassis with subscriber management environment, the bbe-smgd process might leak memory in the backup Routing Engine when running continuous subscriber login logout loop tests. Memory utilization increases with each login logout loop untill it reaches 809 MB and it does not increase beyond that. PR1223625
- On MX Series platforms, executing command of **show chassis ucode-rebalance** without a special FPC slot number, it might cause chassisd to crash. PR1227445
- Flowstat reply has incorrect DL type. For example, for the following flow rule, the flowstat reply shows DL type as 0xcc88 instead of 0x88cc. PR1228383
- The dynamic-profile service filter matches the traffic that is not defined in prefix-list applied to the filter. This causes the filter to not work as expected or even match all the traffics. PR1230997
- ICMP Identifier not translated back to expected value during traceroute for TTL exceeded packets on NAT using Multiservice MPC. This occurs for ICMP ID >255 and causes all hops (except first and last) appearing as "*". PR1231868
- IPsec tunnels anchored on service-set are not cleared when ms interface inside IFL is disabled through CLI command. PR1232276

- Some Packet Forwarding Engine statistics counters do not work in MPC7/8/9. 1. Fabric Input/Output pps counters do not work in **show pfe statistics traffic** 2. Output and Fabric Input/Output counters don't work in **show pfe statistics traffic detail**. PR1232540
- Packet Forwarding Engine statistics input packets pps counter has a large error. Compared to the output pps counter, Input pps counter is very rough. PR1232547
- On XQ based linecard, in rare condition, if offline/online the FPC or link flap, some error messages might be seen. PR1232686
- When you set port-mirror for the MX Series router, LSP ping might fail and IP packets with options will not get mirrored due to unexpected echo reply from DUT: **<---------- echo request ----------> echo reply [R1]----------[DUT]----------[R2] A | ----------> echo reply (unexpected behavior) | mirror** PR1234006
- 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. This is because the code checks if any dynamic profile is configured before enabling **dynamic-profile-options versioning**. If so, it throws a commit error. But there is no need to check when the Routing Engine is in backup state. PR1234453
- MX Series MPC7 and above might receive noise on the FPC console port, and interprets it as valid signals. This might cause login fails on the console port, core-dumps or even reloads. PR1234712
- VLNS(VBNG) - Commit generated a **warning: requires 'l2tp-inline-lns' license** but valid license is installed. PR1235697
- When per-packet load sharing is enabled under aggregated Ethernet (AE) interface, egress traffic over the AE interface might be dropped unexpectedly. PR1235866
- When PIC-based MPLS J-Flow is configured and MPLS packets are being sampled at egress (to be sent to service pic), the sampled packets do not reach service PIC, which results in no MPLS J-Flow flows getting created. PR1236892
- When the interface configured under **router-advertisement** physically comes up for the first time, the rpd might repeatedly send the router-advertisement, which might result in as high as 100 percent Routing Engine CPU usage. PR1237894
- After the number of licenses for the scale-subscriber feature is exceeded, the following endless logs are seen on the backup Routing Engine every 10 seconds: **Dec 12 13:22:41 antelope-re0 license-check[4900]: Routing Engine protocol backup state = 0 Dec 12 13:22:42 antelope-re0 license-check[4900]: Empty license directory copied from the master Dec 12 13:22:51 antelope-re0 license-check[4900]: Routing Engine protocol backup state = 0 Dec 12 13:22:52 antelope-re0 license-check[4900]: Empty license directory copied from the master backup Routing Engine: has all licenses in state permanent master Routing Engine: shows the license with the expiry date.** The log messages disappear after the master switchover. When changing the master back, the above messages will start again. These messages do not appear on the master Routing Engine, which has the expire day set, regardless of the mastership state. PR1238615

- Due to a regression issue, presence of errors and or traps during unified ISSU might result in LU/XL-based MPC crash. PR1239304
- In a BGP-PIC scenario, a change in the IGP topology, (for example, a link failure in the IGP path) causes traffic outage for certain prefixes. The reason for this is that the unilist next hops for these prefixes are in a broken state. PR1239357
- During scaled subscriber setup, the lowest dynamic-profile CoS service rate might be applied to other sessions. PR1241201
- The PTP clock class changes are delayed. When PTP fails and the system goes into holdover, it will be send clock class 6 for the next 10-15 minutes. The same behavior occurs when the system goes from holdover in state "locked". It will be send clock class 248 for the next 10-15 minutes. PR1241211
- Auto route insertion (ARI) IPv6 routes installed for IPSec dynamic endpoints might disappear from the routing-table after performing a graceful Routing Engine switchover (GRES) with nonstop active routing (NSR) enabled. The issue is triggered for IPv6 ARI routes with masks of /98 or longer. PR1242503

High Availability (HA) and Resiliency

- On all platforms, when running unified ISSU, connection might be broken between master Routing Engine and the backup Routing Engine. PR1234196

Infrastructure

- The GNU debugger, gdb, can be exploited in a way that might allow execution of arbitrary unsigned binary applications. PR968335
- During the upgrade, harmless "invalid SMART checksum logs" might be seen. PR1222105

Interfaces and Chassis

- The interface fxp0 might flap upon some specific commit, this may impact the normal work of out-of-band management. PR1213171
- PPPoE tunneled subscriber (L2TP) might get stuck in terminating state if radius sendsÂ Framed-IP-Address and Framed-IP-Netmask via access-accept in LAC. PR1228802
- The configuration change where for a static vlan demux interface the underlying physical interface is changed to a one with a lower bandwidth (for example, from xe to ge) can fail with the following error: **“error: Bandwidth on IFL demux0.7000 cannot be greater than that of its IFD". For example: user@router# show | compare [edit interfaces demux0 unit 7000 demux-options] - underlying-interface xe-0/1/0; + underlying-interface ge-0/3/9; user@router# commit re0: error: Bandwidth on IFL demux0.7000 cannot be greater than that of its IFD error: DCD Configuration check FAILED. error: configuration check-out failed.** PR1232598
- On MX Series platforms acting as broadband network gateway (BNG), in Point-to-Point Protocol (PPP) scenario, when using the Internet Protocol Control Protocol (IPCP) and Internet Protocol version 6 Control Protocol (IPv6CP) for negotiation and IPv6CP is negotiated first, if the router receives an IPCP Configure-Request packet from the client, MX BNG sends the Configure-Request packet, but does not

send the Configure-Ack packet in case it does not receive the Configure-Ack that corresponds to the Configure-Request packet it sent. The behavior does not follow RFC 1661, which requires that both actions Send-Configure-Request (i.e. ConfReq from MX Series to client) and Send-Configure-Ack (i.e. ConfAck from MX Series to client) must be conducted on the router without any significant delay. PR1235261

Junos Fusion Satellite Software

- Junos OS Fusion is not supported with any MX Series platform as Aggregation Device in Junos OS Release 16.2R1. PR1231227

Layer 2 Features

- When MSTP is configured under routing-instance, both the primary and standby VPLS pseudowires are struck in ST state. PR1206106
- On MX Series platforms, if chassis-level configuration is used to offline FPC after detecting major errors, FPC will be offlined. However, if you commit a configuration afger offlining the FPC, it will be brought online back. PR1218304
- MX Series is not including Delegated-IPv6-Prefix in accounting interim. PR1231665
- DHCPv6 renegotiation-lockout configuration statement range has expanded to 4-600 seconds, enabling you to reduce the MX Series BNG wait time for responding to DHCPv6 Solicit retransmissions messages. PR1234009

MPLS

- You can configure both **load-balance-label-capability** and **no-load-balance-label-capability** together. This is incorrect and confusing. PR1126439
- When PCE updates the delegated LSP, **no-install-to-address** configured under the LSP stanza is not honored. PR1169889
- In a scaled environment, when there are many Unicast NHs related to the same transport LSP (for example, same RSVP or LDP label), MPLS traffic statistics collection might take too much CPU time in kernel mode. This can in turn lead to various system-impacting events, like scheduler slips of various processes and losing connections between the backup Routing Engine and FPCs. PR1214961
- Carrier-over-carrier VPN PE router "protocol mpls" under RI breaks existing "protocol connection." In a multi instance RSVP scenario we support MPLS in the VRF routing-instance but we still do not support Connections protocol inside the VRF routing instance. So, when we are adding any interface under MPLS inside VRF routing-instance, then it should affect the Connections protocol inside the main instance. When we were adding the CE facing interface under MPLS in VRF instance it was deleting the Patricia which was having CCC information as we do not have CCC information inside the VRF instance. So to resolve this issue we have added a check that before acting on the Connections protocol we should

check whether the instance passed is master instance or not and if it is not the master instance we do not trigger the functionality related to CCC. PR1222570

- The **rsvp-lsp-enh-lp-upstream-status** command takes a lot of time to synchronize the backup Routing Engine on egress. PR1242324

Multicast

- RPD creates an indirect next hop when a multicast route (S,G) needs to be installed when listeners show their interest to S,G traffic. Kernel would then creates a composite NH. In this case, this appears to be P2MP MCNH which gets created. When any member interface is not a Packet Forwarding Engine specific interface (for example, Vt, LSI, IRB or any other pseudointerfaces), the kernel throws this message indicating that FMBB cannot be supported. These messages are harmless and do not have any impact. PR1230465

Platform and Infrastructure

- Junos OS: key attribute that is emitted in the XML format of configuration will not be emitted in the JSON format of configuration. PR1195928
- Blank firewall logs for IPv6 packets with next-header hop-by-hop are fixed in Junos OS Release 14.1R9.PR1201864
- In large-scale configurations or environments with high rates of churn, the FPC ASIC memory will become "fragmented" over time. It is possible in an extreme case that memory of a particular size will become exhausted and due to the fragmentation, the available memory will not fulfill the pending allocation. PR1216300
- Next hop used for Routing Engine-generated TCP traffic might differ from the one used for Routing Engine-generated non-TCP traffic if the prefix not subjected to **then load-balanced per-packet** action and is pointing to an indirect nexthop resolved via unilist nexthop (ECMP). Before the fix for PR1193697 this leads to non-TCP traffic generated from Routing Engine taking one unicast next-hop while TCP traffic generated from the Routing Engine is load-balanced across different next hops. After the fix for PR1193697 this behavior might lead to non-TCP host outbound traffic taking one unicast next-hop, while TCP host outbound traffic takes another. PR1229409
- Firewall filter index mapping gets incorrect after Routing Engine switchover, because the contents of **/var/etc/filters/filter-define.conf** are incorrectly changed after Routing Engine switchover. PR1230954
- Incoming interface index cannot be used as a loadbalancing input factor under family multiservice if the traffic payload is non-Ethernet frame. PR1232943
- The scale-subscriber license count might increase to an invalid license state with L2TP/LTS clients. This is due to the L2TPD daemon not going through proper state transition on L2TP/LTS clients logout. Hence, license count was not getting updated. The fix will ensure license count is updated on logout regardless of whether the daemon goes through proper state transition or not. PR1233298
- Login for flow-tap DTCP-over-SSH service fails when SSH key-based authentication is configured for the flow-tap user. When such a login attempt is performed, depending on the Junos OS version, ssh-relay

process might crash or the following log message might be printed: dfcd[21043]: **DFCD_DTCP_USER_NOT_AUTHORIZED: Unauthorized user ft-user tried to log in for flow-tap service. The following configuration will cause the login for the flow-tap-dtcp service to fail: system { login { class ft-class { permissions [flow-tap flow-tap-control flow-tap-operation]; } user ft-user { uid 2012; class ft-class; authentication { ssh-rsa " ssh-rsa "}}** PR1234464

- MX2010/2020 cannot sample multicast traffic when this multicast is copied to multiple interfaces. This behavior seems MX2010/MX2020 specific. **Sample Topology MC traffic(ipv4 or ipv6) ----------> xe-0/0/0+------+xe-0/0/1 +------+ [Tester]----------|MX2020|----------|MX2020| | |----------|LSYS | +------+et-1/0/0 +------+ | | sampling| | V | [collector] FCP0 : MPC6E FPC1 : MPC9E** PR1237164
- On MX Series platform with MPC5/MPC7/MPC8/MPC9, when a low value of temporal buffer-size (for example,10,000) is configured, the threshold in the drop rule in the Packet Forwarding Engine differs from expectation. PR1240756

Routing Protocol

- In a dual Routing Engines scenario with NSR and PIM configuration, when the backup Routing Engine is handling mirror updates about PIM received from the master Routing Engine, it will delete the PIM session information from its database. But due to a software defect, a leak of 2 memory blocks (8 or 16 byte leaks) will occur for every PIM leave. If the memory is exhausted, the rpd process might crash on the backup Routing Engine. There is no impact seen on the master Routing Engine when the rpd generates a core file on backup. PR1155778
- In Junos OS Release 16.1R2, when BGP add-path is configured and the same prefix is received from multiple peers with different source AS, depending on the order that the prefix advertisements are received, the rpd might crash. PR1223651
- On all platforms, if MPLS goes down due to link flap or FPC reboot or restart, rpd might generate a core file. 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. Refer to JSA10778 for more information.PR1229868
- Memory redzone check fails when attempting to free reference memory block. Redzone check fails because an address that is not the beginning of a memory block is passed. For this particular case, because of the incorrect beginning address for session reference, when you try to pick up session thread, you actually get the rts_entry. The pointer to the rts_parms happens to be the same value as the session_id, so rpd considers it as a match on processing a delete.When rpd tries to free the session reference block, it crashes. PR1232742
- When a rib-group is configured with a nonexisting routing-instance, after deleting rib-group and deactivating static flow route, a stale route is present in inetflow.0 rib. It might affect traffic forwarding. PR1236636

- When there are different LSPs toward the same egress endpoint and they are up and advertised in IS-IS or IS-IS TE shortcuts is configured, the active route is expected to use the LSPs as ECMP next hops in inet.0. If in addition, rsvp load-balance bandwidth is configured, it would be expected that traffic is load-balanced taking into a consideration the LSP's bandwidth. The latter was not happening and the traffic was load-balanced equally across all ECMP LSPs, which should not have been the case. PR1237531
- When MX Series router is running protocol BGP and policy configuration is modified, an assertion condition might be hit where the routing protocol daemon(rpd) generates a corefile. PR1239990
- In BGP configuration, if the static rt-constrain feature is configured but family route-target is not present on any BGP configuration, rpd might generate a core file. This is due to cleanup code attempting to free state that was not created since family route-target was not configured. PR1247625

Services Applications

- In an L2TP scenario, when the LNS is flooded by a high rate of L2TP messages from LAC, the CPU on the Routing Engine might become too busy to bring up new sessions. PR990081
- When using NAT on the MX Series, the FTP ALG 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
- When loading or rolling back a configuration that removes a service-set and changes where the MS interfaces are assigned, traffic may be blackholed to a series of the existing service-sets. PR1223302
- When the stateful firewall flows time out repeatedly, there can be memory leaks on the MS-DPC PIC. It can eventually exhaust the memory and result in failure of new flow creation. PR1242556

Subscriber Access Management

- In a subscriber management environment with two or more RADIUS servers connected to an MX Series router, syslog is not generated when the RADIUS server is marked dead. **Primary Server Aug 13 20:16:58 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.247 set to DEAD (profile radius) Aug 13 20:17:58 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.247 set to ALIVE (profile radius) Aug 13 20:19:36 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.247 set to ALIVE (profile radius) Secondary Server Aug 13 20:17:17 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to UNREACHABLE (profile radius) Aug 13 20:17:47 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to ALIVE (profile radius) Aug 13 20:18:04 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to UNREACHABLE (profile radius) Aug 13 20:18:34 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to ALIVE (profile radius) Aug 13 20:18:51 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to UNREACHABLE (profile radius) Aug 13 20:19:21 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to ALIVE (profile radius) Aug 13 20:19:41 ERX-MX240-2-RE1 authd[1976]: AUTHD_RADIUS_SERVER_STATUS_CHANGE: Status of radius server 10.219.48.248 set to UNREACHABLE (profile radius)** PR1207904

- In Junos OS Releases 14.1X50-D110.9+ and 14.1X50-D110-J2, 3GPPP -SGSN -MCC -MNC svp with value “999999” will be sent in all CCR-GY requests. In later releases a hidded unreleased configuration statement is added to configure both the value of this avp and whether it will be sent (**set access ocs partition <name> sgsn-mcc-mnc <string>**).PR1233847

- On MX Series with subscriber management, the DHCPv6 solicit packets with IA_PD option from the subscriber are ignored if the DHCPv6 server does not have a prefix to allocate to this subscriber. This behavior is incorrect. According to the RFC standard, the DHCPv6 server should reply to such packets using special Status Code: NoPrefixAvail (6), which should be included in Advertise/Reply in case if no delegated prefix is available. PR1234042

- On MX Series router with dual Routing Engines, after router GRES, if you add a trace options filter before the GRES is ready, the authd process might crash. PR1234395

User Interface and Configuration

- Some configuration objects are not properly handled by dexp. The objects affected can be classified as “leaf-list” and container with presence in YANG Data Modeling Language (RFC 6020) terminology. Leaf-list can include a list of values (for example, a list of group names or a list of policy-statement names) but cannot have any other configuration statement nested below it. For example, - **protocols bgp group X apply-groups - protocols bgp group X import - routing-options forwarding-table export 'Container with presence'** does not convey any additional information and serves just as a present/not_present flag, for example, - **chassis fpc 7 pic 0 tunnel-services**. The annotations for such statements were not processed correctly by dexp, were not written to juniper.conf, and were shown in 'show | compare' output right after entering configuration mode. PR1245187

VPNs

- In an NG MVPN scenario with asm-override-ssm configuration statement for source-specific multicast (SSM) group, if you issue the **clear pim join** command on the source PE router downstream interfaces get pruned, causing the multicast flow to stop. PR1232623
- When the l2circuit neighbor is changed in NSR enabled, an rpd core file is observed on the backup Routing Engine. PR1241801

SEE ALSO

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 17.1R1 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.

SEE ALSO

New and Changed Features | **78**

Changes in Behavior and Syntax | **102**

Known Behavior | **109**

Known Issues | **111**

Resolved Issues | **122**

Migration, Upgrade, and Downgrade Instructions | **134**

Product Compatibility | **143**

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 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.1 remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see Understanding Junos OS with Upgraded FreeBSD.

NOTE: In Junos OS Release 15.1, Junos OS (FreeBSD 10.x) is not available to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to use the existing Junos OS (FreeBSD 6.1).

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.1-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

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 Guideand Upgrading Junos OS with Upgraded FreeBSD.

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.

Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

NOTE: This section does not apply to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to refer to the next section.

To download and install from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x):

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.

- For 32-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-32-17.1R1.9-signed.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-17.1R1.9-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*** (available only for Canada and U.S. version)

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.1) 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 6.1) 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 6.1) 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.

Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)

Products impacted: MX80, and MX104.

NOTE: Customers in Belarus, Kazakhstan, and Russia must use the following procedure for all MX Series routers running Junos OS Release 17.1.

To download and install from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1):

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.

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

  `user@host>` **request system software add validate reboot *source*/jinstall-17.1R1.9-domestic-signed.tgz**

- All other customers, use the following command:

  `user@host>` **request system software add validate reboot *source*/jinstall-17.1R1.9-export-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*** (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.

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 14.2 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

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.

Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1

In releases prior to Junos OS Release 10.1, the draft-rosen multicast VPN feature implements the unicast lo0.x address configured within that instance as the source address used to establish PIM neighbors and create the multicast tunnel. In this mode, the multicast VPN loopback address is used for reverse path forwarding (RPF) route resolution to create the reverse path tree (RPT), or multicast tunnel. The multicast VPN loopback address is also used as the source address in outgoing PIM control messages.

In Junos OS Release 10.1 and later, you can use the router's main instance loopback (lo0.0) address (rather than the multicast VPN loopback address) to establish the PIM state for the multicast VPN. We strongly recommend that you perform the following procedure when upgrading to Junos OS Release 10.1 if your draft-rosen multicast VPN network includes both Juniper Network routers and other vendors' routers functioning as provider edge (PE) routers. Doing so preserves multicast VPN connectivity throughout the upgrade process.

Because Junos OS Release 10.1 supports using the router's main instance loopback (lo0.0) address, it is no longer necessary for the multicast VPN loopback address to match the main instance loopback address lo0.0 to maintain interoperability.

> **NOTE:** You might want to maintain a multicast VPN instance lo0.x address to use for protocol peering (such as IBGP sessions), or as a stable router identifier, or to support the PIM bootstrap server function within the VPN instance.

Complete the following steps when upgrading routers in your draft-rosen multicast VPN network to Junos OS Release 10.1 if you want to configure the router's main instance loopback address for draft-rosen multicast VPN:

1. Upgrade all M7i and M10i routers to Junos OS Release 10.1 before you configure the loopback address for draft-rosen multicast VPN.

   > **NOTE:** Do not configure the new feature until all the M7i and M10i routers in the network have been upgraded to Junos OS Release 10.1.

2. After you have upgraded all routers, configure each router's main instance loopback address as the source address for multicast interfaces.

   Include the **default-vpn-source interface-name *loopback-interface-name*]** statement at the **[edit protocols pim]** hierarchy level.

3. After you have configured the router's main loopback address on each PE router, delete the multicast VPN loopback address (lo0.x) from all routers.

We also recommend that you remove the multicast VPN loopback address from all PE routers from other vendors. In Junos OS releases prior to 10.1, to ensure interoperability with other vendors' routers in a draft-rosen multicast VPN network, you had to perform additional configuration. Remove that configuration from both the Juniper Networks routers and the other vendors' routers. This configuration should be on Juniper Networks routers and on the other vendors' routers where you configured the lo0.mvpn address in each VRF instance as the same address as the main loopback (lo0.0) address.

This configuration is not required when you upgrade to Junos OS Release 10.1 and use the main loopback address as the source address for multicast interfaces.

NOTE: To maintain a loopback address for a specific instance, configure a loopback address value that does not match the main instance address (lo0.0).

For more information about configuring the draft-rosen multicast VPN feature, see the *Multicast Protocols User Guide for Routing Devices*.

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 | **78**
- Changes in Behavior and Syntax | **102**
- Known Behavior | **109**
- Known Issues | **111**
- Resolved Issues | **122**
- Documentation Updates | **133**
- Product Compatibility | **143**

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 | 78
- Changes in Behavior and Syntax | 102
- Known Behavior | 109
- Known Issues | 111
- Resolved Issues | 122
- Documentation Updates | 133
- Migration, Upgrade, and Downgrade Instructions | 134

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R1 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 17.1R1 for the PTX Series.

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.
- **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

  (For installation requirements, see the PTX5000 Packet Transport Router Hardware Guide.)

  > **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.

- **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)**—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.

- 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**—The 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.]

Interfaces and Chassis

- **Aggregated Ethernet Statistics Enhancements (PTX Series Routers)**—Starting with 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 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. 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 those 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)**—Junos OS Release 17.1 introduces abstract hops, which are user-defined router clusters or groups that can be sequenced and used for setting up a label-switched path (LSP), similar to real-hop constraints.

  The router groups are created using constituent lists that include constituent attributes, which 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.

  [See Multiprotocol BGP MVPNs Overview.]

  > **NOTE:** Next-generation MVPN is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

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.]

  [See 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

- **Secure Boot (PTX3000)**—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.

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 17.1R1 for the PTX Series.

General Routing

- ECMP nexthops 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']/**.

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.

Routing Protocols

- **Change in default behavior of router capability (PTX Series)**—In Junos OS Releases 15.1F7, 16.1R4, 16.1X65, and 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.

Services Applications

- **Device discovery with device-initiated connection (PTX 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 (PTX Series)**—Starting in Junos OS Release 17.1R1, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 17.1R1, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the show system schema module juniper-command output directory (PTX Series)**—Starting in Junos OS Release 17.1, when you issue the **show system schema module juniper-command** operational command in the Junos OS CLI, the device places the generated output files in the current working directory, which defaults to the user's home directory. Prior to Junos OS Release 17.1, the generated output files are placed in the **/var/tmp** directory.

SEE ALSO

New and Changed Features | **144**

Known Behavior | **158**

Known Issues | **159**

Resolved Issues | **161**

Documentation Updates | **162**

Migration, Upgrade, and Downgrade Instructions | **163**

Product Compatibility | **167**

Known Behavior

There are no known limitations in Junos OS Release 17.1R1 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.

SEE ALSO

New and Changed Features | **144**

Changes in Behavior and Syntax | **155**

Known Issues | **159**

Resolved Issues | **161**

Documentation Updates | **162**

Migration, Upgrade, and Downgrade Instructions | **163**

Product Compatibility | 167

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 17.1R1 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 might 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 the PTX Series router 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 can cause a large clock phase shift in the 100GbE-LR4 CFP modules, and this phase shift can cause output optical pulse waveform distortion on the 100GbE-LR4 interfaces. Hence, it results in interface flap. PR1130403
- On PTX Series platforms with FPC3, the octets of IPv4 source and destination addresses in the firewall log are listed in reverse; this might affect troubleshooting. The IPv6 log works fine. This is a minor issue, and there is no other service impact. PR1141495
- On the PTX5000, after plugging a QSFP28 PIC (15x100GE/15x40GE/60x10GE QSFP+ PIC) into FPC-P1, 100 percent FPC CPU usage might be seen. PR1158640
- While upgrading from Junos OS Release 15.1F based images to Junos OS Release 16.x or later images or downgrading from 16.x or later images to 15.1F based images, if the **validate** option is enabled, chassisd might crash and upgrade/downgrade will fail. This issue should not be seen if both base and target images are from Junos OS Release 15.1F or Junos OS Release 16.x or later. PR1171652

- When running interfaces on QSFP28 PIC on PTX Series platform in 40G or 100G mode, some of the interfaces on the QSFP28 PIC might not come up after a system reboot. This issue does not impact interfaces running in 10G mode. PR1176641
- 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 second seconds CMERROR infra raises an alarm and appropriate action is taken. PR1197865
- 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
- The Junos Traffic Vision client might notice slower arrival of packets for the configured sensors. PR1243810

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 Customer intending 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 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 customer wants to verify this, 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 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

Platform and Infrastructure

- In a very rare scenario, during TAC accounting configuration change, the auditd process crashes because of a race condition between auditd and its sigalarm handler. PR1191527
- On PTX Series platforms with **chassis network-services enhanced-mode** configured, the default policy junos-ptx-series-default is not loaded correctly in case of some configuring operations, As a result, some BGP routes are not installed in the forwarding table as expected. As a workaround to avoid this issue, reboot the router after any configuring operations on network services. PR1204827

Routing Protocols

- If IPv6 default route is configured, after you issue an offline/online FPC command, the IPv6 default route might not be seen in the isis route table, and this might cause IPv6 traffic loss. PR1159482

User Interface and Configuration

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

SEE ALSO

New and Changed Features | **144**

Changes in Behavior and Syntax | **155**

Known Behavior | **158**

Resolved Issues | **161**

Documentation Updates | **162**

Migration, Upgrade, and Downgrade Instructions | **163**

Product Compatibility | **167**

Resolved Issues

IN THIS SECTION

- Resolved Issues: 17.1R1 | **162**

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: 17.1R1

Class of Service (CoS)

- Following error log message will be seen with Hierarchical CoS configuration: **Delay buffer computation incorrect.^M If hierarchical scheduler is configured for an IFD and if guaranteed rate is not set for an IFL under this IFD, then the temporal buffer configured** The display of this error message is valid when guaranteed rate is 0, but it is not valid when guaranteed rate is disabled. PR1238719

General Routing

- Attempting to connect to FPC via cty can result in Routing Engine switchover PR1235761

Interfaces and Chassis

- Power default settings for 5-port 100G DWDM PIC are missing. These power-related error messages do not affect any PIC functionality. PR1184415

Routing Protocols

- On all platforms, if MPLS goes down due to link flap or FPC reboot or restart, rpd might generate a core file. PR1228388

SEE ALSO

New and Changed Features | **144**

Changes in Behavior and Syntax | **155**

Known Behavior | **158**

Known Issues | **159**

Documentation Updates | **162**

Migration, Upgrade, and Downgrade Instructions | **163**

Product Compatibility | **167**

Documentation Updates

There are no errata or changes in Junos OS Release 17.1R1 documentation for PTX Series.

SEE ALSO

New and Changed Features | **144**

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.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Upgrading a Router with Redundant Routing Engines

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

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

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

Basic Procedure for Upgrading to Release 17.1

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 17.1R1:

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

    > **NOTE:** After you install a Junos OS Release 17.1R1 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the request system software add validate command and specify the jinstall package that corresponds to the previously installed software.

    The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is for a different release. Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router

displays the login prompt. The loading process might take 5 to 10 minutes. Rebooting occurs only if the upgrade is successful.

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

```
user@host> request system software add validate reboot
source/jinstall-17.1R1.9-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-17.1R1.9-export-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

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

NOTE: After you install a Junos OS Release 17.1 **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.

SEE ALSO

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

- New and Changed Features | **144**
- Changes in Behavior and Syntax | **155**
- Known Behavior | **158**
- Known Issues | **159**
- Resolved Issues | **161**
- Documentation Updates | **162**
- Migration, Upgrade, and Downgrade Instructions | **163**

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 17.1R1 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the new features for the QFX Series switches in Junos OS Release 17.1R1.

NOTE: The following QFX Series platforms are supported in Release 17.1R1: QFX5100, QFX10002, QFX10008, and QFX10016.

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 100-Gbps and 40-Gbps. Note that as of 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 CoS-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.

- **New Show Interfaces Command for Virtual Output Queues (QFX 10000 Series)**—Starting with Junos OS Release 17.1R1, QFX 10008 Series switches support the new command **show interfaces voq *interface-name***, that enables you to view statistics for virtual output queues.
- **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.
  - CoS Hierarchical Port Scheduling (ETS) 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.
- **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.

Dynamic Host Configuration Protocol

- **Virtual-router aware DHCP server/DHCP relay agent (QFX10008)**—Starting with Junos OSRelease 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 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 (aex), set the **local-minimum-links-threshold** configuration statement with a threshold value that represents the percentage of local member links that must be up on a chassis for any local LAG member links on that chassis to continue to be active in the aggregated Ethernet bundle. Otherwise, all remaining LAG member links on that chassis are also forced down. The feature responds dynamically to bring local LAG member links up or down if you change the configured threshold, or when the status or configuration of LAG member links changes. Note that forced-down links also influence the minimum links count for the LAG as a whole, which can bring down the LAG, so enable this feature only in configurations where LAG traffic is carefully monitored and controlled.

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

[See Understanding Local Minimum Links.]

- **Support for Micro BFD over child links of AE or LAG bundle (cross-functional Packet Forwarding Engine/kernel/rpd) (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, this feature provides a Layer 3 BFD liveness detection mechanism for child links of the Ethernet LAG interface. In scenarios in which you do not have a point-to-point link, and a Layer 1 device fails at one end of the link, Micro BFD detects failures faster than traditional LACP. Micro BFD sessions are independent of each other despite having a single client that manages the LAG interface. Micro BFD is not supported on pure Layer 2 interfaces.

  To enable failure detection for aggregated Ethernet interfaces, include the **bfd-liveness-detection** statement at the **[edit interfaces aex aggregated-ether-options bfd-liveness-detection]** hierarchy level.

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

  [See Understanding Independent Micro BFD Sessions for LAG.]

- **PVLAN and Q-in-Q on the same interface (QFX5100 Switches)** —Starting with Junos OS Release 17.1R1, you can configure a private VLAN and Q-in-Q tunneling on the same Ethernet port. To configure both PVLAN and Q-in-Q on the same physical interface, you must configure flexible Ethernet services to support dual methods of configuring logical interfaces. Q-in-Q requires a service provider configuration method, and PVLAN requires an enterprise configuration method.

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

  [See Understanding Flexible Ethernet Services Encapsulation on Switches.]

- **Support for configuration synchronization for MC-LAG (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, Multichassis Link Aggregation group (MC-LAG) configuration synchronization enables you to easily propagate, synchronize, and commit configurations from one MC-LAG peer to another. You can log into any one of the MC-LAG peers to manage both MC-LAG peers, thus having a single

point of management. You can also use configuration groups to simplify the configuration process. You can create one configuration group for the local MC-LAG peer, one for the remote MC-LAG peer, and one for the global configuration, which is essentially a configuration that is common to both MC-LAG peers.

In addition, you can create conditional groups to specify when a configuration is synchronized with another MC-LAG peer. You can enable the **peers-synchronize** statement at the **[edit system commit]** hierarchy to synchronize the configurations and commits across the MC-LAG peers by default. NETCONF over SSH provides a secure connection between the MC-LAG peers, and Secure Copy Protocol (SCP) copies the configurations securely between them.

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

[See Understanding MC-LAG Configuration Synchronization.]

- **Support for configuration consistency check for MC-LAG (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, Multichassis Link Aggregation group (MC-LAG) configuration consistency check alerts you of both severe and moderate configuration inconsistencies across MC-LAG peers. The configuration consistency check feature checks MC-LAG configuration parameters, such as chassis ID, session establishment time, and so on, on each peer and notifies you of any errors, so you can fix the inconsistencies. Configuration inconsistencies are categorized as severe or moderate. If there is a severe inconsistency, the MC-LAG interface is prevented from coming up. Once you have corrected the inconsistency, the system will bring up the interface. If there is a moderate inconsistency, you are notified of the error and can then fix the inconsistency. After you fix any inconsistency, you must commit the changes to take effect.

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

  [See Understanding Multichassis Link Aggregation Group Configuration Consistency Check.]

- **Configuration support to improve MC-LAG Layer 2 and Layer 3 convergence (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can configure multichassis link aggregation (MC-LAG) interfaces to improve Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet link goes down or comes up in a bridge domain. To use this feature, ensure that the Inter-Chassis Link (ICL) is configured on an aggregated Ethernet interface. For Layer 2 convergence, configure the **enhanced-convergence** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. For Layer 3 convergence, configure the **enhanced-convergence** statement on an integrated routing and bridging (IRB) interface at the **[edit interfaces irb unit unit-number]** hierarchy level.

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

  [See enhanced-convergence.]

- **Channelizing 40-Gigabit Ethernet QSFP+ ports (QFX10008 switch)**—This feature enables you to channelize four 10-Gigabit Ethernet interfaces from the 40-Gigabit Ethernet QSFP+ interfaces. Channelization is supported on fiber break-out cable using standard structured cabling techniques.

NOTE: This feature is not supported on the QFX10000-30C line card.

By default, the 40-Gigabit Ethernet QSFP+ interfaces are named **et-fpc/pic/port**. The resulting 10-Gigabit Ethernet interfaces appear in the following format: **xe-fpc/pic/port:channel**, where channel can be a value of 0 through 3. To channelize a 40-Gigabit Ethernet QSFP+ interface into four 10-Gigabit Ethernet interfaces, include the **10g** statement at the **[edit chassis fpc fpc-slot pic pic-slot (port port-number | port-range port-range-low port-range-high) channel-speed]** hierarchy level. To revert the 10-Gigabit Ethernet channels to a full 40-Gigabit Ethernet interface, remove the **10g** statement from the same hierarchy level.

There are 100-Gigabit Ethernet ports that work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet but are recognized as 40-Gigabit Ethernet by default. You cannot channelize the 100-Gigabit Ethernet ports when they are operating as 100-Gigabit Ethernet interfaces. The 40-Gigabit Ethernet ports can operate independently or be channelized into four 10-Gigabit Ethernet ports as part of a port range. Ports cannot be channelized individually. Only the first and fourth port in each 6XQSFP cage is available to channelize as part of a port range. In a port range, the ports are bundled with the next two consecutive ports. For example, if you want to channelize ports 0 through 2, you channelize port 0 only. If you try to channelize a port that is not supported, you receive an error message when you commit the configuration. Auto-channelization is not supported on any ports.

When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports.

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

[See Channelizing Interfaces.]

IP Tunneling

- **Generic Routing Encapsulation support (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 17.1R1, you can configure GRE tunnels. GRE provides a private, secure path for transporting packets through a public network by encapsulating (or tunneling) the packets. GRE tunneling is accomplished through tunnel endpoints that encapsulate or de-encapsulate traffic. To configure a GRE tunnel interface, include the **gre-*fpc*/*pic*/*port*** set of statements at the **[edit interfaces]** hierarchy level.

  This feature was previously supported only on the QFX10002 switch.

  [See Configuring Generic Routing Encapsulation Tunneling.]

IPv4

- **IPv4 address conservation method for hosting providers (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, you can configure a static route on an IRB interface with or without pinning to a specific underlying interface, thereby conserving the usage of IP address space.

  Configure the interface on the router with an address from the reserved IPv4 prefix for shared address space (RFC 6598) and by using static routes pointed at the interface. IANA has recorded the allocation of an IPv4 /10 for use as shared address space. The shared address space address range is 100.64.0.0/10.

  This way, the interface in the router is allocated an IP address from the shared address space, so it is not consuming publicly routable address space, and connectivity is handled with static routes on an interface. The interface in the server is configured with a publicly routable address, but the router interfaces are not. Network and broadcast addresses are consumed out of the shared address space rather than the publicly routable address space.

  [See IPv4 Address Conservation Method for Hosting Providers.]

Layer 2 Features

- **Support for Layer 2 Features (QFX10000 switches)**—Starting with Junos OS Release 17.1R1, the following features are supported:
  - VLAN support—VLANs 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.

  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 with 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's Carrier (CSC) Layer 3 VPNs (QFX10000 switches)**—Staring with Junos OS 17.1R1, CSC 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 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 (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.]

- **IPv6 Layer 3 VPNs (QFX5100 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 17.1R1 for the QFX Series.

Multiprotocol Label Switching (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.

Network Management and Monitoring

- **Cloud Analytics Engine disabled in Junos OS by default (QFX Series)**—In Junos OS Release 17.1R1 and later, Cloud Analytics Engine network analytics probe processing is disabled by default in the 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 prior releases through Junos OS Release 16.1R3, Cloud Analytics Engine Junos 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)**—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.]

Software Installation and Upgrade

- **In-Service Software Upgrade (QFX5100 switches)**—Starting with Junos OS Release 17.1R1, you cannot perform an unified ISSU from a Junos OS Release earlier than 17.1R1 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 | **169**

Known Behavior | **193**

Known Issues | **194**

Resolved Issues | **199**

Documentation Updates | **201**

Migration, Upgrade, and Downgrade Instructions | **202**

Product Compatibility | **209**

Known Behavior

There are no known limitations for the QFX Series switches in Junos OS Release 17.1R1.

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 | **169**

Changes in Behavior and Syntax | **190**

Known Issues | **194**

Resolved Issues | **199**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 17.1R1.

Interfaces and Chassis

- On QFX10000 switches, in a multichassis link-aggregation group (MC-LAG) configuration, the all option at the [edit protocols igmp-snooping vlan] hierarchy level does not work. As a workaround, enable IGMP snooping on a per-VLAN basis on each of the MC-LAG peers. PR1180494
- It might take more time for traffic to converge when we add or delete members of MC-AE lag as compared to earlier releases. PR1199306

Layer 2 Features

- On QFX5100 VC interfaces on which the flexible-vlan-tagging statement is specified, STP, RSTP, MSTP, and VSTP are not supported. PR1075230
- On QFX Series switches except QFX10000/QFX5200, when transmitting large packet which is more than MTU configured and not be fragmented on the IRB interface, ICMP error packet about type3 with code4 cannot be generated. The large packets are getting silently dropped. PR1089445

- On EX4600/QFX Series switches, after add and delete the fifth logical interface, the first 4 AE subinterfaces might be down and lose connectivity. PR1171488
- On QFX10k, commit error 'Too many VLAN-IDs on interface' will be seen when adding over 1025 sub-interfaces on lag interface. PR1186556
- When same VLAN TAG Id is configured on the NNI and UNI interface belonging to same Bridge Domain, the traffic on the NNI is egressing with Single Tag instead of Dual Tag. PR1192760

MPLS

- With 100 or more Layer 2 circuit configurations in standby mode on a QFX5100 switch, the Layer 2 circuits might go down after issuing the restart routing operational mode command. PR1169575
- Layer 2 circuits on QFX5100 switches might not come up if 100 or more Layer 2 circuit connections are configured in no-standby mode. PR1169659
- For 2 label PUSH cases, both labels are consuming entries in the same label table. This might result in instabilities of MPLS tunnels and packets drop when add/delete routes. Correct behavior should be that tunnel label goes in one table and VRF label should go in another table. PR1185550
- On QFX10000 Series switches, when the L2 CCC MPLS frammes MTU is higher than egress MPLS interface MTU, the frames are still forwarded instead of dropped. PR1190025
- After deactivating interfaces on a QFX5100 switch that is configured as a primary neighbor of a provider edge router, the backup Layer 2 circuit might not get activated as expected. PR1198191
- When changing the "routing-options forwarding-table chained-composite-next-hop" configuration while there are active MPLS LSPs, then LSP traffic loss can be observed at a later time.PR1243088

Platform and Infrastructure

- On QFX10002 switches, the command **request system snapshot** does not work. PR1048182
- On a QFX10002 switch, when a new interface is added to an existing link aggregation group (LAG) interface which acts as an input analyzer interface, traffic sent to the added interface might not mirrored. PR1057527
- On a QFX10002 switch inserting a small form pluggable (SFP) on the management interface (em1). After a system reboot, replace SFP with a copper SFP, the management interface might not work properly with speed 10m/100m. PR1075097
- On a fully loaded QFX10008 chassis, line cards might take as many as 15 minutes to become operational after startup. PR1124967
- With Multi-Hop BFD, traffic loss of around 5 to 10 sec is observed when intermediate interface is shutdown. After 5-10 secs, traffic recovers and no action is needed. PR1150695

- On QFX5100 and QFX10002 switches, 'Rx power low warning set' messages might be logged continuously for channelization ports that are in the DOWN state with snmpwalk running in the background. PR1204988
- Below configurations on 10002 are not supported. Set chassis fpc 0 error fatal action reset set chassis fpc 0 error fatal action offline. If fatal errors are detected in log or in alarm, please run below command to restart the dcpfe OR reboot the system from Junos OS CLI. **request app-engine service restart packet-forwarding-engine**. PR1239325
- On QFX5100 switches in the Open vSwitch Database (OVSDB) scenario with Virtual Extensible LAN (VXLAN) configured, MAC learning might not work well across the interface when it is dynamically changed and bounded on the link from virtual routers between the bare-metal server (BMS) to virtual routers connected locally and also between BMS to virtual routers connected through the spine.PR1115546
- On QFX5100 switches, OVSDB traffic might be dropped after Layer 2 learning is restarted. PR1177012
- On QFX10000 switches, when you upgrade to Junos OS Release 17.1R1 from Release 15.1X53-D33, traffic over route type-5 on the tunnel ingress node might drop if you have the forwarding-table **no-indirect-next-hop** statements configured at the **[edit routing-options]** hierarchy level. As a workaround, delete the configuration **routing-options forwarding-table no-indirect-next-hop** before you perform an upgrade. This configuration is not needed when route type-5 is configured. PR1187482
- On QFX Series switches, a Packet Forwarding Engine or device-control process (dcd) restart might result in traffic loss. PR1188120
- On QFX5100 Virtual Chassis, the DHCP snooping database might be cleared if you change the configuration of the LACP mode from fast to slow. PR1191404
- On QFX10000 switches, a slow increase in CRC error is seen when using Juniper Networks branded 100G LR4 optics. PR1208159
- On QFX10000 switches with enhanced MC-LAG IRB next hops, member links of the aggregate 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 Packet Forwarding Engine uKernel. Those failure logs do not impact traffic or performance on the switch, and they are harmless. PR1221831
- On QFX5100 VC, traffic loss might be observed while adding or deleting trunk members from local minimum links. PR1226488
- On QFX Series switches, in rare cases, the Link Up/Down notification from the Packet Forwarding Engine to the Routing Engine might require additional time. The Packet Forwarding Engine side interface and the remote device interface show Admin Up and Link Up, but the CLI might show the interfaces in Admin Down and Link Down. This issue might take approximately 30 seconds to resolve. PR1227947

- On QFX10008 switches, the IPv6 packets/bytes counter shows higher values than the total packets/bytes of the interface if LAG child members belong to the same PE device. As a workaround, if you monitor IPv6 statistics over the LAG, choose LAG child members across PE devices. PR1232388
- On QFX10008 and QFX10016 switches with 60X10G line cards, when you disable or enable the 1g port on the line card, the following error messages appear in the log: **pechip_cmerror_set_error:3113: Level: Major, cmerror_code: 0x21060e (id=1550), recover_err: 0 (counter: 0), fh_msg: 0x0**. However, no impact on functionality is observed. PR1238269

Routing Protocols

- L3 multicast traffic does not converge completely upto 100% and few drops will be seen continuously, after doing interface unshut that is bringing up ports after bringing it down or while fpc comes online after doing fpc restart. This behavior will be seen while we scale beyond 2k vlan's or 2k irb's with vlan replication in system. PR1161485
- On QFX5100 switches, when parity errors occur on interfaces, they might affect the memory management unit (MMU) memories. MMU counters can be corrupted, the interface buffers might be stuck, and there might be interface flaps and traffic loss on the affected ports. As a workaround (restoration only), reboot the system. PR1169700
- On EX4600/QFX Series switches, after native-vlan-id is configured and rolled back on a vlan-tagged sub-interface, ARP might not be resolved and traffic forwarding can be affected. PR1184985
- On QFX10K 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 QFX10K switches, traffic drop will be seen with ISISv6 traffic during convergence in following two scenarios: 1. While doing port unshutdown that is bringing up the ports, after bringing it down. 2. While FPC comes online after doing FPC restart. The above behavior is seen while flapping one of the ISISv6 sessions. PR1190180
- On QFX5100 switches with MPLS and LDP enabled, for packets with incoming labels that must perform a PHP (penultimate hop popping) operation on the QFX5100 switch, occasionally the packets are not processed and are dropped. PR1190437
- On QFX5100/EX4600 series switches, if a routing loop is created , the TTL of the packet does not reduce to 0 and eventually the packet is not dropped. PR1196354
- On QFX5100/QFX5200/EX4600 series switches, if disable IRB interface then reboot the switch. After the switch is rebooted then, enable IRB interface, after that IRB interface might not be reachable. PR1196380
- On QFX10002 switch that functions as a peer in a multicast group, multicast traffic entering a Layer 3 VLAN-tagged interface might be inadvertently dropped. PR1198502

- On Junos OS based platforms, software does not support route with next hop type other than 'unicast' and 'resolve' leaking from default routing table (inet.0 or inet6.0) to VRF. And, when next-hop of a route need to resolve from different routing-instance, this actually a route leaking, if the next hop type other than 'unicast' and 'resolve', the route leak will fail, the leak route cannot be installed in VRF. PR1210620
- In a Layer 3 VPN, if IRB is used between the penultimate hop and the PE node, if checking VRF connectivity using PE to PE ping, then pinging to the PE loopback address or interface IP address from the remote PE does not work. PR1211462
- On QFX10002, QFX10008, and QFX5100-24Q platforms, the **show interfaces <interface-name> | match pps** command does not display any output and might cause the process to hang during FIB route installation. PR1250328

Software Installation and Upgrade

On a QFX10000 switch running Junos OS Release 15.1X53-D62, the configuration includes the following command: **set protocols evpn vni-options vni *id* vrf-target export target:*community-name***. After upgrading the switch software from Junos OS Release 15.1X53-D62 to Junos OS Release 17.1R1, the **vrf-target export target:*community-name*** configuration statement is inadvertently omitted, which means that the configuration now includes the following command: **set protocols evpn vni-options vni *id***. To work around this issue, you must add the missing configuration statement. PR1243105

Virtual Chassis

- On EX/QFX-VC platforms, if new members are not zeroized prior to being added to VC, after that one of new members splits from VC, in this case whenever running "commit" or "commit check", it might be seen that the commit is hanging for a long time then report time out error on that FPC which splits from VC. PR1211753

SEE ALSO

New and Changed Features | **169**

Changes in Behavior and Syntax | **190**

Known Behavior | **193**

Resolved Issues | **199**

Documentation Updates | **201**

Migration, Upgrade, and Downgrade Instructions | **202**

Product Compatibility | **209**

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 for 17.1R1

High Availability (HA) and Resiliency

- On QFX5100 and EX4600 switches, during a non-stop software upgrade (NSSU), if an aggregated Ethernet (AE) interface is configured with multiple sub-interfaces across multiple Flexible Port Concentrators (FPC's), the AE interface might go down. PR1227522
- If you are performing a topology-independent in-service software upgrade (TISSU) from one version of Junos OS Release 14.1X53 to another on a QFX5100 switch, and the network analytics feature [edit service analytics] is configured, the upgrade might not succeed. In addition, the fxpc process might stop working, and you might notice that a core file is generated. PR1234945

Interfaces and Chassis

- On QFX10008 and QFX10016 switches, an error message such as expr_cos_rw_nh_qix_get @ 150: Unable to get chip num for ill:994 on mc-ae status-control active node might be displayed after an ARP request is sent. These messages are only for information and have no functional impact on the operation of the switches. PR1228080
- CDP packets with destination address 01:00:0c:cc:cc:cc looping with MC-LAG on 10K. PR1237227

Junos Fusion Satellite Software

- The following conditions must be met before a Junos OS switch can be converted to a satellite device when the action is initiated from the aggregation device: 1. The Junos OS switch must be in factory default settings OR it must have included **set chassis auto-satellite-conversion** in its configuration 2. The package used to do the conversion must be one where the individual images are signed (SNOS 3.0, SNOS 1.0R5, SNOS 2.0R2). PR1249877

Layer 2 Features

- On QFX10000 switches, the kernel might fail to allocate IFBD tokens, with the error message 'IFBD hw token couldn't be allocated for <interface>, even though there are enough IFBD tokens, and thus you might be unable to assign some VLANs to the related interfaces. PR1216464
- On QFX5100 switches, an fxpc process might generate a core file. PR1231071
- Mac learning will be very slow when clearing mac address in case of scale mac learning(128k). PR1240114

MPLS

- In MPLS scenario, on EX4600/QFX Series switches with AE interface configured, after change the IGP metric and disable the AE interface, the fxpc crash might be observed because of child nexthop of a UNILIST is pointing to NULL. PR1168150

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

- On QFX10000 platforms, it may experience latency and jitter with ICMP traffic directed towards a local interface of a switch. PR1221053
- On a QFX10002 switches, 40-Gigabit Ethernet ports can take up to 4 seconds to link when using 40-Gigabit optical transceivers. This issue is addressed in the newer release with an option to quickly tune the optics. Use the following hidden CLI command to quickly tune the optical transceivers. **set chassis fpc<fpc-number> fast-tune** PR1219336
- On QFX5200-32C switches running Junos OS Release 15.1X53-D210, a Management Ethernet 1 Link Down major alarm might be raised even though the switch does not use Management Ethernet 1 (em1). PR1228577
- On QFX10002 switches, when you plug in a USB, FRU insertion messages such as **RE0 & CAMGETPASSTHRU ioctl failed cam_lookup_pass: Inappropriate ioctl for device** might be displayed. These are harmless messages and will not be displayed after you have removed the USB. PR1233037
- QFX5100 switches with 48-port or 96-port configurations might incorrectly list the media type of a SFP-T copper module as fiber in the output for the **show interface** command.PR1240681
- On QFX10002 switches, no network ports are detected after the reboot sequence following a code upgrade. PR1247753

Routing Protocols

- On QFX5100 Virtual Chassis, a user-defined filter might not work due to not getting programmed in the Packet Forwarding Engine. PR1175121
- A filter attached to the lo0 interface with terms containing either destination-port-range-optimize or source-port-range-optimize statements will unexpectedly discard traffic. PR1228335
- Reset action in below configuration statement CLI does not work in 10008 **set chassis fpc <slot-no> error fatal action reset**. To reset the fpc, please run below command once fatal errors are detected **reqeuest chassis fpc <slot-no> restart**.PR1233075
- On QFX5100 switches running Junos OS 14.1X53-D30.3, when you apply an IPv6 firewall filter, the system might crash with a Packet Forwarding Engine panic. PR1234729
- On QFX10000 switches, on aggregated Ethernet interfaces with adaptive load balancing enabled, frequent link flaps might result in zero active members in the LAG bundle, causing memory leaks and eventually causing an FPC crash. The FPC restarts automatically after the crash. PR1236046
- On a QFX5100 switch, Gratuitous Address Resolution Protocol (GARP) reply packets are not updating the Address Resolution Protocol (ARP) table. GARP request packets, however, are updating the ARP table as expected. PR1246988

Virtual Chassis

- VCF not communicating properly with backup Spine. PR1141965

SEE ALSO

New and Changed Features | **169**

Changes in Behavior and Syntax | **190**

Known Behavior | **193**

Known Issues | **194**

Documentation Updates | **201**

Migration, Upgrade, and Downgrade Instructions | **202**

Product Compatibility | **209**

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 17.1R1.

SEE ALSO

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-qfx-5-17.1-R3.10-domestic-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 16.1 **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 17.1R1.

> **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-qfx-10-f-flex-16.1R3.10-domestic.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-qfx-10-f-flex-16.1R3.10-domestic.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
```

Performing an In-Service Software Upgrade (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 206
- Upgrading the Software Using ISSU on page 206

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 ISSU

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

To upgrade the switch using unified ISSU:

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

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

     where ***package-name*.tgz** is, for example, **jinstall-132_x51_vjunos.domestic.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-qfx-5-13.2X51-D15.4-domestic ...
     Install jinstall-qfx-5-13.2X51-D15.4-domestic 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. To ensure that the resilient dual-root partitions feature operates correctly, copy 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 | 169

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

Third-Party Components

This product includes third-party components. To obtain a complete list of third-party components, see Copyright and Trademark Information.

Upgrading Using Unified ISSU

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

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

Compliance Advisor

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

Finding More Information

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

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

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

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: http://www2.juniper.net/kb/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/

- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: https://www.juniper.net/cm/

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

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, see https://www.juniper.net/support/requesting-support.html.

Revision History

3 September 2020—Revision 13, Junos OS Release 17.1R1– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

14 February 2019—Revision 12, Junos OS Release 17.1R1– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

22 November 2017—Revision 11, Junos OS Release 17.1R1– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

2 November 2017—Revision 10, Junos OS Release 17.1R1– EX Series, MX Series, PTX Series, QFX Series and Junos Fusion.

18 August 2017—Revision 9, Junos OS Release 17.1R1– 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.