Release Notes: Junos® OS Release 13.2X50 for the EX Series and QFX Series

3 October 2013
Revision 2

Contents

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, J Series, M Series, MX Series, PTX Series, QFabric, QFX Series, SRX Series, and T Series.

These release notes accompany Junos OS Release 13.2X50 for the EX 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 EX Series Switches

These release notes accompany Junos OS Release 13.2X50 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 13.2X50 for the EX Series.

Hardware

- **EX4300 switch hardware**—Juniper Networks EX4300 Ethernet Switches provide connectivity for high-density environments and scalability for growing networks. These switches can be deployed wherever you need high density of Gigabit Ethernet ports (24 to 48 ports) or redundancy. Typically, EX4300 switches are used in large branch offices, campus wiring closets, and data centers, where they can be positioned as the top device in a rack to provide connectivity for all the devices in the rack and provide options for optimized airflow (hot aisle or cold aisle).

  The EX4300 switches are available in 24-port and 48-port models, with or without PoE+, and with four additional 40-gigabit quad small form-factor pluggable (QSFP+) uplink ports. All models provide for an optional 4-port 10-gigabit SFP+ uplink module that accepts 1-gigabit or 10-gigabit SFP or SFP+ transceivers.

  The following optics are supported on EX4300 switches:

  - EX-SFP-10GE-SR
  - EX-SFP-10GE-LR
  - EX-SFP-10GE-LRM
  - EX-SFP-1GE-SX
  - EX-SFP-1GE-LX
  - EX-SFP-10GE-USR
  - EX-SFP-10GE-ER
  - EX-QSFP-40GE-SR4
  - EX-SFP-10GE-DAC-1M
  - EX-SFP-10GE-DAC-3M
  - EX-SFP-10GE-DAC-5M
  - EX-SFP-10GE-DAC-7M
  - EX-QSFP-40GE-DAC-50CM
  - EX-QSFP-40GE-DAC-1M
  - EX-QSFP-40GE-DAC-3M

  [See EX4300 Hardware Documentation.]

- **SFP+ Media Access Control Security (MACsec) uplink module**—The SFP+ MACsec uplink module is now available. You can install the SFP+ MACsec uplink module in EX4200 switches to provide the switch with MACsec-capable ports. You can configure the SFP+ MACsec module to support up to four MACsec-capable 1-gigabit small

form-factor pluggable (SFP) transceivers or up to two MACsec-capable 10-gigabit small form-factor pluggable (SFP+) transceivers. [See Uplink Modules in EX4200 Switches.]

- **48-port SFP line card for EX6200 switches**—The 48-port SFP line card for EX6200 switches (EX6200-48F) comprises:
  - 48 100-gigabit or 1000-gigabit small form-factor (SFP) ports
  - Line card LEDs
  - Network port LEDs

  The line cards in EX6200 switches combine a Packet Forwarding Engine and Ethernet interfaces on a single assembly. They are field-replaceable units (FRUs) that can be installed in the line card slots on the front of the switch chassis. The line cards are hot-insertable and hot-removable.

  The following optics are supported on EX6200-48F line cards:
  - EX-SFP-GE10KT13R15
  - EX-SFP-GE10KT15R13
  - EX-SFP-GE10KT13R14
  - EX-SFP-GE10KT14R13
  - EX-SFP-GE40KT13R15
  - EX-SFP-GE40KT15R13
  - EX-SFP-FE20KT13R15
  - EX-SFP-FE20KT15R13

  [See 48-Port SFP Line Card in an EX6200 Switch.]

Authentication and Access Control

- **Access control on EX4300 switches**—Support for controlling access to your network through an EX Series switch by using several different authentication methods, by configuring 802.1X, MAC RADIUS, or a captive portal. You now enable the **authentication-whitelist** statement at the **[edit switching-options]** hierarchy level instead of at the **[edit ethernet-switching-options]** hierarchy level. [See Access Control on EX4300 Switches.]
- **Captive portal authentication for Layer 3 interfaces on EX3300 switches**—Captive portal is now supported on EX3300 switches. [See Understanding Authentication on EX Series Switches.]

Bridging and Learning

- **Layer 2 bridging and control protocol features on EX4300 switches**—Support for the following Layer 2 bridging and control protocol features:

- Layer 2 access and trunk interfaces. Configure access or trunk mode at the **[edit interfaces *interface-name* unit 0 family ethernet-switching interface-mode]** hierarchy level.
- A VLAN specified with a VLAN ID, or a VLAN specified with multiple ranges of VLAN IDs. Configure VLANs at the **[edit vlans]** hierarchy level.
- 802.1Q VLAN trunking. Note that only EtherType 0x8100 is supported for identifying VLAN-tagged packets. Configure the interface mode as trunk with the **set interfaces *interface-name* unit 0 family ethernet-switching interface-mode trunk** command.
- Integrated routing and bridging (IRB). This feature is similar to routed VLAN interfaces (RVIs), which are supported on EX Series switches that are not running ELS. Configure IRB interfaces at the **[edit vlans]** and **[edit interfaces]** hierarchy levels. Determine the current state of IRB interfaces by entering the **show interfaces irb** command.
- The configuration of up to 4093 VLANs. Configure VLANs at the **[edit vlans]** hierarchy level.
- The ability to enable or disable MAC address learning. By default, MAC address learning is enabled globally, per interface, and per VLAN. Disable MAC address learning globally at the **[edit protocols l2-learning]** and **[edit switch-options no-mac-learning]** hierarchy levels, per interface at the **[edit switch-options interface *interface-name* no-mac-learning]** hierarchy level, and per VLAN at the **[edit vlans *vlan-name* switch-options no-mac-learning]** hierarchy level.
- The ability to set a global MAC address aging time. Configure this feature at the **[edit protocols l2-learning global-mac-table-aging-time]** hierarchy level.
- The ability to enable or disable a MAC limit per interface and per VLAN, and to set an action to take on the next packet the interface or VLAN receives after the limit is reached. Configure a MAC limit per interface at the **[edit switch-options interface *interface-name* interface-mac-limit *limit* packet-action *action*]** hierarchy level. For a VLAN, you can do the following:
  - Configure a limit for the number of MAC addresses that can be learned by a VLAN at the **[edit vlans *vlan-name* switch-options mac-table-size packet-action *action*]** hierarchy level.
  - Configure a MAC limit for all interfaces in the VLAN at the **[edit vlans *vlan-name* switch-options interface-mac-limit *limit* packet-action *action*]** hierarchy level.
  - Configure a MAC limit for a specified interface in the VLAN at the **[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit* packet-action *action*]** hierarchy level.
- Static MAC address. The Ethernet switching table, also known as the forwarding table, specifies the known locations of VLAN nodes and the addresses of devices within those nodes. There are two ways to populate the Ethernet switching table on a switch. The easiest method is to let the switch update the table with MAC addresses. The second way is to manually insert addresses into the table. You can do this to reduce flooding and speed up the switch's automatic learning process. To configure a device to have a static MAC address, issue the **static-mac *mac-address***

statement at the **[edit vlans *vlan-name* switch-options interface *interface-name*]** hierarchy level.

- MAC move limit. When MAC move limiting is configured, MAC address movements are tracked by the switch, and if a MAC address changes more than the configured number of times within 1 second, the changes to MAC addresses are dropped, logged, or ignored, or the interface is shut down. Configure a limit for the number of times that a MAC address can move, and specify the action to be taken by issuing the **mac-move-limit *limit* packet-action *action*** statement at the **[edit vlans *vlan-name* switch-options]** hierarchy level.
- MAC address filtering, which is supported on a single Ethernet interface but not on aggregated Ethernet interfaces.
- Unknown unicast forwarding.
- Multiple VLAN Registration Protocol (MVRP), which is used to manage the addition, deletion, and renaming of VLANs. Configure MVRP at the **[edit protocols mvrp]** hierarchy level.

[See Ethernet Switching on EX4300 Switches.]

Class of Service

- **CoS support on EX4300 switches**—Support for CoS features, including ingress classification, re-marking (egress rewrite), scheduling, queuing, buffer management, and ingress policers, Virtual Chassis support for CoS, and CoS for host outbound traffic. Additional supported features include:
  - Support for applying rewrite rules on a port (interface).
  - Support for configuring the percentage of excess traffic for each queue. A new configuration statement, **excess-rate**, is now available at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.
  - Support for the WRED congestion management mechanism. A new configuration statement, **drop-probability percentage**, is available at the **[edit class-of-service drop-profiles *profile-name*]** hierarchy level to support the WRED configuration.
  - Support for four loss priorities: high, low, medium-high, and medium-low. The medium-low and medium-high loss priorities are new priorities for EX4300 switches.
  - Support for scheduling strict-high priority queues in round-robin scheduling and taking the queue number into consideration.
  - Support for default forwarding classes for multidestination traffic; the classes are:
    - multicast-af
    - multicast-be
    - multicast-ef
    - multicast-nc

    The default bandwidth distribution rate is:

- Best effort: 75 percent
- Network control: 5 percent
- Multicast best effort: 15 percent
- Multicast network control: 5 percent

[See Class of Service for EX4300 Switches.]

High Availability (HA) and Resiliency

- **Nonstop bridging on EX4300 Virtual Chassis**—You can configure NSB to provide resilience for Layer 2 protocol sessions on an EX4300 Virtual Chassis. Nonstop bridging operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because all session information is already synchronized to the backup Routing Engine. Traffic disruption for the NSB-supported Layer 2 protocol is minimal or nonexistent as a result of the switchover. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the NSB-supported Layer 2 protocol sessions on the switch. [See Understanding Nonstop Bridging on EX Series Switches.]
- **Nonstop active routing on EX4300 Virtual Chassis**—You can configure NSR on an EX4300 Virtual Chassis to enable the transparent switchover of the Routing Engines in the event that one of the Routing Engines goes down. NSR provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without requiring restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and each can take over for the other. The switchover is transparent to neighbor routing devices, which do not detect that a change has occurred. Enable NSR when neighbor routing devices are not configured to support graceful restart of protocols or when you want to ensure graceful restart of protocols for which graceful restart is not supported (such as PIM). [See Understanding Nonstop Active Routing on EX Series Switches.]
- **High availability (HA) feature support added on EX3300 switches**—Support for these HA features is now provided on EX3300 switches:
  - GRES on LACP
  - GRES for Layer 3 protocols
  - GRES for port security (DHCP snooping, DAI, and IP source guard)
  - Graceful protocol restart for Layer 3

  [See High Availability Features for EX Series Switches Overview.]

Infrastructure

- **Controlled software image on EX Series switches**—A controlled version of Junos OS that supports Media Access Control Security (MACsec) is now available for EX Series switches. On prior Junos OS releases, you could only install the domestic version of Junos OS on all EX Series switches. Two versions of a Junos OS image—a controlled

version that supports Media Access Control Security (MACsec) and a domestic version that does not support MACsec—are now available for EX Series switches. You can download both versions of the Junos OS image from the Software Download Center if you are located in a geography where you are allowed to download the controlled version.

All EX Series switches are shipped with a domestic version of Junos OS.

You can identify whether a software package is the controlled or domestic version of Junos OS by viewing the package name:

- A software package for a controlled version of Junos OS is named using the following format: **package-name-m.nZx.y-controlled-signed.tgz**.
- A software package for a domestic version of Junos OS is named using the following format: **package-name-m.nZx.y-domestic-signed.tgz**.

If you plan on enabling Media Access Control Security (MACsec), download the controlled version of your Junos OS release. The controlled version of a Junos OS release contains all features and functionality available in the domestic version of the Junos OS release while also supporting MACsec.

If you are unsure which version of Junos OS is running on your switch, enter the **show version** command. If the "JUNOS Crypto Software Suite" description appears in the output, you are running the controlled version of Junos OS.

The controlled version of Junos OS contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS is also subject to controls imposed under the laws of other countries.

If you have questions about acquiring the controlled version of Junos OS in your country, contact the Juniper Networks Trade Compliance group at **compliance_helpdesk@juniper.net**.

[See also "Migration, Upgrade, and Downgrade Instructions" on page 30 .]

- **Uniform Enhanced Layer 2 Software CLI configuration statements and operational commands**—ELS provides a uniform CLI for configuring and monitoring Layer 2 features on EX Series switches that support ELS, such as EX9200 switches and EX4300 switches, and on MX Series routers in LAN mode (MX-ELM). With ELS, for example, you can configure a VLAN and other Layer 2 features on an EX9200 switch, an EX4300 switch, *or* an MX-ELM router by using the same configuration commands. [See Getting Started with Enhanced Layer 2 Software.]

  The Web-based ELS Translator tool is available for registered customers to help them become familiar with the ELS CLI and to quickly translate existing EX Series switch–based CLI configurations into ELS CLI configurations. [See ELS Translator.]

- **Storm control on EX4300 switches**—Support for storm control that enables the switch to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level—is exceeded, thereby preventing packets from proliferating and degrading the LAN.

Storm control is enabled by default on all Layer 2 interfaces. You can modify the storm-control configuration with a two-step process:

1. Configure a storm-control profile at the **[edit forwarding-options]** hierarchy level.
2. Bind the storm-control profile to a specific logical interface or to a group of logical interfaces. The group can include a range of interfaces or all interfaces on the switch.

[See Understanding Storm Control on EX Series Switches.]

- **IPv4 and IPv6 unicast routes on EX4300 switches**—Support for IPv4 and IPv6 unicast routes, which along with IPv4 and IPv6 multicast routes, are stored in a flexible forwarding table. This forwarding table functions as a dynamic storage resource. There are no configurable parameters for this forwarding table. [See Understanding Bridging and VLANs on EX Series Switches.]

Interfaces and Chassis

- **Q-in-Q tunneling on EX4300 switches**—Support for Q-in-Q tunneling on EX4300 switches. Q-in-Q tunneling 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. The Juniper Networks Junos operating system (Junos OS) implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard. The configuration of Q-in-Q tunneling is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

  Configure Q-in-Q tunneling using one of the following methods to map customer VLANs (C-VLANs) to service-provider-defined service VLANs (S-VLANs):

  - All-in-one bundling—–Map packets from all C-VLAN interfaces on a switch to an S-VLAN. Configure this at:
    - The **[edit interfaces *interface-name*]** hierarchy level with **flexible-vlan-tagging** and **encapsulation extended-vlan-bridge** statements
    - The **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level with **input-vlan-map push**, **output-vlan-map pop**, **vlan-id-list**, and **vlan-id** statements
    - The **[edit vlans *vlan-name* interface *interface-name.logical-unit-number*]** hierarchy level
  - Many-to-many bundling—Map packets from multiple C-VLANs to multiple S-VLANs. Configure this at:
    - The **[edit interfaces *interface-name*]** hierarchy level with **flexible-vlan-tagging** and **encapsulation extended-vlan-bridge** statements
    - The **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level with **input-vlan-map push**, **output-vlan-map pop**, **vlan-id-list**, and **vlan-id** statements
    - The **[edit vlans *vlan-name* interface *interface-name.logical-unit-number*]** hierarchy level

- Specific interface mapping with VLAN ID translation—Map packets from a specified C-VLAN to a specified S-VLAN. In addition, while the packets are transmitted to and from the S-VLAN, you can specify that the 802.1Q C-VLAN tag be removed and replaced with the S-VLAN tag or vice versa. Configure this at:
  - The **[edit interfaces *interface-name*]** hierarchy level with **flexible-vlan-tagging** and **encapsulation extended-vlan-bridge** statements
  - The **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level with **input-vlan-map swap**, **output-vlan-map swap**, **vlan-id-list**, and **vlan-id** statements
  - The **[edit vlans *vlan-name* interface *interface-name.logical-unit-number*]** hierarchy level

  [See "Documentation Updates" on page 29.]

- **Generic routing encapsulation on EX4500, EX4550, and EX6200 switches and EX4500 and EX4550 Virtual Chassis**—These switches and Virtual Chassis now support GRE, a tunneling protocol to transport packets over a network. You can use GRE tunneling services to encapsulate any network layer protocol over any other network layer protocol. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first encapsulates the payload packet in a GRE packet and then encapsulates the resulting GRE packet in a delivery protocol. A switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect noncontiguous networks and to provide options for networks that contain protocols with limited hop counts. [See Understanding Generic Routing Encapsulation.]

- **Link aggregation on EX4300 switches**—Support for the ability to group together one or more Ethernet interfaces to form a LAG, which a MAC client can treat as a single link. The link aggregation feature has the following benefits:
  - Increased bandwidth
  - Incremental bandwidth
  - Graceful degradation as failure occurs
  - Increased availability

  The EX4300 switch supports up to 112 LAGs, and each LAG can have up to 8 members. Aggregated Ethernet interfaces are configured at the **[edit interfaces ae*x* aggregated-ether-options]** hierarchy level.

  [See Understanding Aggregated Ethernet Interfaces and LACP.]

- **RTG support**—Support for redundant trunk groups (RTGs), which provide redundancy in cases of link or line card failures. RTGs are configured at the **[edit switch-options redundant-trunk-group group *group-name*]** hierarchy level. [See Understanding Redundant Trunk Links on EX Series Switches.]

IPv6

- **IPv6 support added on EX3300 switches**—Support for these IPv6 features is now provided on EX3300 switches:
  - IPv6 routing IPv6 filter-based forwarding
  - IPv6 routing virtual routing and forwarding (VRF) support for IPv6 filter-based forwarding
  - IPv6 routing VRF support for IPv6 multicast

  [See Understanding Virtual Routing Instances on EX Series Switches and Understanding Filter-Based Forwarding for EX Series Switches

Junos XML API and Scripting

- **Packaging Python scripts on EX4300 switches and EX4300 Virtual Chassis**—Python is available on EX4300 switches and EX4300 Virtual Chassis running Junos OS when the following package is installed:
  - jinstall-ex-4300

  The Python interpreter only runs scripts that have been installed from signed packages created with the Junos SDK. For more information, see "Using Python on Junos" and "Building a Junos SDK Package Containing Python Scripts" in the *Junos SDK Developer Guide*.

J-Web

- **The J-Web interface supports EX4300 switches.** [See J-Web User Interface for EX Series Switches Overview.]

Multicast

- **IGMP snooping on EX4300 switches**—Support for Internet Group Management Protocol (IGMP) snooping. IGMP snooping constrains the flooding of IPv4 multicast traffic on VLANs on the switch. When IGMP snooping is enabled, the switch examines IGMP messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. Based on what it learns, the switch then forwards multicast traffic only to those interfaces that are connected to interested receivers instead of flooding the traffic to all interfaces. IGMP snooping is configured at the **[edit protocols igmp-snooping]** hierarchy level.

  The EX4300 switch supports IGMPv1, IGMPv2, and IGMPv3.

  In a pure Layer 2 environment, traffic is not forwarded between VLANs. For a host in one VLAN to receive the multicast traffic from the source in another VLAN, you must create an IRB interface in each of the VLANs to permit routing of the multicast traffic between the VLANs. In addition, you must enable Protocol Independent Multicast (PIM) on the switch to perform the multicast routing.

The EX4300 switch also supports the routing of multicast traffic between multicast routing domains. You must enable PIM and Multicast Source Discovery Protocol (MSDP) on the switch for the switch to perform multicast routing.

[See IGMP Snooping on EX Series Switches Overview.]

- **Multicast virtual routing and forwarding with IPv4 for IGMP snooping on EX3300 switches**—Support for multicast VRF for IGMP snooping is now provided on EX3300 switches. [See Understanding Virtual Routing Instances on EX Series Switches .]

Network Management and Monitoring

- **sFlow monitoring technology and LLDP-MED on EX4300 switches**—Support for:
  - sFlow technology, a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology 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. The sFlow technology is configured at the **[edit protocols sflow]** hierarchy level. [See Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch.]
  - Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the switch to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently, exchanging information related to VoIP VLAN, endpoint location, and power capabilities between the switch and the IP telephone. LLDP-MED is configured at the **[edit protocols lldp-med]** hierarchy level. [See Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches.]
- **RFC 4293 EX Series switch support**—Mandatory MIBs (system-wide statistics) are supported on EX Series switches. [See Juniper Networks Enterprise-Specific MIBs.]
- **Native analyzer support on EX4300 switches**—Support for native analyzers and remote port-mirroring capabilities. A native analyzer configuration contains both an input stanza and an output stanza in this analyzer hierarchy for mirroring packets. In remote port mirroring, the mirrored traffic is flooded into a remote mirroring VLAN that can be specifically created for the purpose of receiving mirrored traffic. On EX4300 switches, the analyzer configuration is available at the **[edit forwarding-options]** hierarchy level. [See Understanding Port Mirroring and Analyzers on EX4300 Switches.]

Port Security

- **Media Access Control Security (MACsec) support on EX4200, EX4300, and EX4550 switches**—MACsec is now supported on EX4200 switches, EX4300 switches, and EX4550 switches. EX4200 switches must have the SFP+ MACsec uplink module to support this feature. MACsec is an 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. You can use it in combination with other security protocols to provide end-to-end network security. MACsec is standardized in IEEE 802.1AE.

NOTE: On EX4200, EX4300, and EX4550 Virtual Chassis, you cannot use NSSU to upgrade from the controlled version (MACsec) of the software image to a domestic version of the software image.

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

[See also "Changes in Behavior and Syntax" on page 17 and "Migration, Upgrade, and Downgrade Instructions" on page 30 for more information about MACsec.]

- **Port security on EX4300 switches**—Support for the following port security features:
  - DHCP snooping
  - Dynamic ARP inspection (DAI)
  - IP source guard
  - DHCP option 82
  - Static IP address
  - Trusted DHCP interface

  You can enable or disable these features at the [**edit vlans forwarding-options dhcp-security]** hierarchy level. If you enable DAI or IP source guard, DHCP snooping is also enabled implicitly for all VLANs. Also, if you configure a group of access interfaces at the **[edit vlans]** hierarchy level, the configuration implicitly enables DHCP snooping.

  - **DAI**—DAI protects switches against ARP spoofing. DAI inspects 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 cache poisoning. To enable DAI on a VLAN, issue the **arp-inspection** statement at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level.
  - **DHCP option 82**—You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. To enable DHCP option 82, issue the **option-82** statement at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level.
  - **IP source guard**—You can use the IP source guard access port security feature to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, it discards the packet. To enable IP source guard on a VLAN, issue the **ip-source-guard** statement at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level.
  - **Static IP**—You can add static (fixed) IP addresses and bind them to fixed MAC addresses in the DHCP snooping database. To configure a static IP/MAC binding in the DHCP snooping database, you must first create a group of access interfaces at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level. To

configure a specific interface within the group to have a static IP address that is bound to a fixed MAC address, issue the **group *group-name* interface *interface-name* static-ip *ip-address mac-address*** statement at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level.

- **Trusted DHCP server interface**—You can configure any interface on a switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases. By default, all access interfaces are untrusted, and all trunk interfaces are trusted. However, you can override the default setting for access interfaces by configuring a group of access interfaces within a VLAN, specifying an interface to belong to that group, and then configuring the group as trusted. To configure an untrusted access interface as a trusted interface for a DHCP server, issue the **overrides trusted** statement at the **[edit vlans *vlan-name* forwarding-options dhcp-security group *group-name*]** hierarchy level.

  To monitor and manipulate the DHCP snooping database, ARP inspection statistics, and IP source guard statistics, issue the following operational commands:

  - **show dhcp-security binding**
  - **clear dhcp-security binding**
  - **show dhcp-security arp inspection statistics**
  - **clear arp**
  - **clear interfaces statistics all**

  [See Port Security on EX4300 Switches.]

Routing Policy and Firewall Filters

- **Firewall filters on EX4300 switches**—Support for firewall filters that can be configured on ports, VLANs, and Layer 3 interfaces. You configure firewall filters at the **[edit firewall]** hierarchy level. [See Routing Policy and Packet Filtering for EX4300 Switches.]

Spanning-Tree Protocols

- **Spanning-tree protocols on EX4300 switches**—Support for spanning-tree protocol versions that enable you to create a loop-free topology in Layer 2 networks:
  - The original Spanning Tree Protocol (STP), which is defined in the IEEE 802.1D 1998 specification.
  - Rapid Spanning Tree Protocol (RSTP), which was originally defined in the IEEE 802.1w draft specification. RSTP provides faster reconvergence time than the original STP by identifying certain links as point to point and by using protocol handshake messages rather than fixed timeouts.
  - Multiple Spanning Tree Protocol (MSTP), which was originally defined in the IEEE 802.1s draft specification and later incorporated into the IEEE 802.1Q-2003 specification. MSTP provides the capability to logically divide a Layer 2 network into

regions. Every region has a unique identifier and can contain multiple instances of spanning trees.

- VLAN Spanning Tree Protocol (VSTP), which is compatible with the Per-VLAN Spanning Tree Plus (PVST+) and Rapid-PVST+ protocols supported on certain vendors' routers and switches. VSTP maintains a separate spanning-tree instance for each VLAN. Different VLANs can use different spanning-tree paths.

By default, RSTP is enabled on the switch. You can configure options for RSTP at the **[edit protocols rstp]** hierarchy level. You can enable and configure options for MSTP or VSTP at the **[edit protocols (mstp | vstp)]** hierarchy level, respectively. For compatibility with devices that do not support RSTP or VSTP, include the **force-version stp** statement at the **[edit protocols (rstp | vstp)]** hierarchy level to enable xSTP.

[See Spanning-Tree Protocols for EX4300 Switches.]

Virtual Chassis

- **EX4300 Virtual Chassis support**—EX4300 switches can be interconnected to form a Virtual Chassis. [See Understanding EX4300 Virtual Chassis.]

Related Documentation

- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 18
- Resolved Issues on page 23
- Documentation Updates on page 29
- Migration, Upgrade, and Downgrade Instructions on page 30
- Product Compatibility on page 36

Changes in Behavior and Syntax

There are no changes in default behavior and syntax in Junos OS Release 13.2X50 for EX Series switches.

Related Documentation

- New and Changed Features on page 4
- Known Behavior on page 18
- Known Issues on page 18
- Resolved Issues on page 23
- Documentation Updates on page 29
- Migration, Upgrade, and Downgrade Instructions on page 30
- Product Compatibility on page 36

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 13.2X50 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.

- Infrastructure

Infrastructure

- On EX4300 switches and EX4300 Virtual Chassis, when an Internet Group Management Protocol (IGMP) leave message is sent, the multicast group might not be cleared from the IP multicast (IPMC) table. This is a known software limitation. PR853369
- EX4300 switches might continuously drop packets when they have to forward the packets over ECMP based on integrated routing and bridging (IRB) interfaces, if the next-hop addresses are reached through different underlying Layer 2 logical interfaces. This is a known software limitation. PR858509
- On an EX4300 Virtual Chassis, forwarding might be delayed and the count of flooded packets might not be the same as the count of transmitted packets. This is a known software limitation. PR863481
- On an EX4300 switch, packets with reserved registered multicast IPv4/IPv6 DIP are rate-limited even though storm control is disabled for registered multicast. This is a known software limitation. PR891348

Related Documentation

- New and Changed Features on page 4
- Changes in Behavior and Syntax on page 17
- Known Issues on page 18
- Resolved Issues on page 23
- Documentation Updates on page 29
- Migration, Upgrade, and Downgrade Instructions on page 30
- Product Compatibility on page 36

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 13.2X50 for the EX Series.

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

- Authentication and Access Control
- Bridging and Learning
- Routing Policy and Firewall Filters

Authentication and Access Control

- On EX4300 switches, when scaling up to 10,000 clients, the switch might create a dot1x core file. PR886279
- On EX4300 switches, after you clear the Ethernet switching table, MAC RADIUS authenticated sessions are not cleared if the traffic is continuous. PR833888
- On an EX4300 Virtual Chassis, if a large number of clients have been authenticated and then you issue the **clear dot1x interface** command, the system might not remove all entries from the Ethernet switching table. PR867518
- On an EX4300 Virtual Chassis, when a large number of dynamic VLAN users are authenticated on multiple interfaces, dynamic VLAN associations are not removed even after all authenticated 802.1X sessions have cleared. PR881777
- On EX4300 switches, deactivating captive-portal configuration on an interface might clear 802.1X clients; authenticated sessions on the same interface with a fallback configuration and vice versa. PR878361
- On an EX4300 Virtual Chassis, after you add and then delete a protocol, a few interfaces might drop traffic. As a workaround, restart the l2-learning process or reboot the switch. PR886316
- On EX4300 switches, after a client that has been authenticated on a VoIP VLAN interface sends a logoff message, the VoIP VLAN binding on that interface might be deleted. PR896091

Bridging and Learning

- On EX4300 switches and EX4300 Virtual Chassis, after you enable unknown unicast forwarding, the MAC table might not forward packets, and then rebooting the system might cause a Layer 2 address learning daemon (l2ald) core file to be created. PR865485.

Routing Policy and Firewall Filters

- On EX4300 switches, the **from interface *interface-name*** match condition is not supported on egress firewall filters. PR817979

Infrastructure

- On EX4300 switches, if you create more than one Ethernet ring protection (ERP) instance on the same interface, traffic on that interface might be lost. PR815700
- On EX4300 switches, if you create an Ethernet ring protection (ERP) instance with a specified control VLAN, and then create a data VLAN for the same ERP instance, traffic might be lost. PR816517
- On EX4300 switches, Ethernet ring protection (ERP) fails if the control VLAN is replaced with another VLAN at runtime. PR817456
- On an EX4300 Virtual Chassis, the device control process (dcd) creates core files when mastership is switched over to another Routing Engine. PR818726
- On EX4300 switches and EX4300 Virtual Chassis, ICMP ping packets with sizes exceeding the configured MTU on the interface might be allowed to pass, whereas data packets with sizes exceeding the configured MTU on the interface, for example, up to MTU+12, are allowed to egress. PR832358
- On EX4300 switches, interfaces are not marked as m-router interfaces when they are connected to a multicast router that is not an IGMP querier. PR832877
- On EX4300 switches and EX4300 Virtual Chassis, if you configure more than 512 VSTP instances, the switch might create a core file. PR848278
- On EX4300 switches, if you change VLAN configurations by using the **load override *config-file*** commands, the system might create a core file.PR849522
- On EX4300 switches and EX4300 Virtual Chassis, when Virtual Router Redundancy Protocol (VRRP) is enabled on a logical interface on which proxy ARP is enabled, the switch might reply with the VRRP's MAC address instead of the integrated routing and bridging (IRB) MAC address. PR849664
- On an EX4300 Virtual Chassis, after you enable and then disable persistent learning on an interface, the number of entries in the Ethernet switching table might be twice the actual count. PR854673
- On EX4300 switches and EX4300 Virtual Chassis, the system might create a multicast snooping process (mcsnoopd) core file. PR859009
- On EX4300 switches and EX4300 Virtual Chassis, unicast packets that are routed on a unicast reverse path forwarding (RPF)-enabled equal-cost multipath (ECMP) path

might be dropped when one of the other ECMP paths through which traffic is not getting routed goes down. PR860426

- On an EX4300 switch, when you issue the **set protocols dot1x authenticator interface all** command, a commit error appears. PR892082
- On EX4300 switches, if an access interface is configured in both a data VLAN and a VoIP VLAN, then if IP source guard is enabled on the data VLAN, traffic on the VoIP VLAN might be affected. As a workaround, enable IP source guard on both the data VLAN and the VoIP VLAN. PR898192
- On EX4200 switches, the 1-Gbps copper tranceivers (SPF-T) on uplink modules might not work properly. PR872404
- On a mixed EX4200 and EX4550 Virtual Chassis, multicast traffic in a GRE tunnel from the EX4550 member switch to an MX Series router might be dropped. PR920309

Interfaces and Chassis

- On an EX4300 or QFX Series Virtual Chassis, the **show virtual-chassis vc-port diagnostics optics** operational command is not supported. PR855574

J-Web

- If you have accessed the J-Web interface using an HTTPS connection through the Microsoft Internet Explorer Web browser, you might not be able to download and save reports from some pages on the Monitor, Maintain, and Troubleshoot tabs. Some affected pages are at these locations:
  - Maintain > Files > Log Files > Download
  - Maintain > Config Management > History
  - Maintain > Customer Support > Support Information > Generate Report
  - Troubleshoot > Troubleshoot Port > Generate Report
  - Monitor > Events and Alarms > View Events > Generate Report
  - Monitor > Routing > Route Information > Generate Report

  As a workaround, use the Mozilla Firefox Web browser to download and save reports using an HTTPS connection. PR566581
- On EX2200-C switches, if you have changed the media type and committed the change, the Ports configuration page (Configure > Interfaces > Ports) might not list the uplink port. PR742847
- After you remove or reboot a Virtual Chassis member (either the backup or a member in the linecard role), when you click other members in the J-Web interface, the Chassis Viewer for those members might not expand, and the dashboard might log the following error: **stackImg is null or not an object**. As a workaround, manually refresh the dashboard. PR771415
- On EX8200 Virtual Chassis, while using the Virtual Chassis Wizard in the J-Web interface in the Mozilla Firefox version 3.x browser, if you have selected more than six port pairs

from the same member to convert from VCPs to network ports, the wizard might display incorrect port conversion status. Also, if you double-click **Next** after deleting an active member in the Members page, the J-Web interface might stop responding. PR796584

- On EX4300 switches, when you commit a configuration using EZSetup, if the laptop becomes disconnected, the J-Web interface reports that the commit operation was successful regardless of whether the commit operation actually succeeded. PR866976

Multicast Protocols

- On an EX4300 switch, when you configure a multicast route, multicast traffic might not go out the egress interface, and the multicast route is not installed in the Packet Forwarding Engine. PR894175
- On EX4300 switches, do not issue the **show igmp snooping membership | match Groups** command if you have a large number (1000+) of groups, because processing uses high CPU. As a workaround, to see a specific group for an interface or all groups for an interface, issue the **show igmp snooping membership** command with filters such as **group** or **interface**. PR914908

Platform and Infrastructure

- On EX4300 switches, Ethernet ring protection (ERP) fails if the control VLAN is replaced with a different VLAN at runtime. PR817456

Port Security

- On EX4200 and EX4550 switches, MACsec dynamic keys are not supported on 1-Gigabit Ethernet fiber interfaces. PR901344
- On EX4300 switches, when you enable MACSec dynamically on a Layer 3 physical interface, the STP state of the port in hardware is set incorrectly to “blocking” and traffic is dropped. As a workaround, delete the family inet/inet6 configuration on the port and reconfigure it. PR912123

Routing Policy and Firewall Filters

- On EX4300 switches, the **from interface *interface-name*** match condition is not supported on egress firewall filters. PR817979

Spanning-Tree Protocols

- When configuring xSTP on EX4300 switches, you *must add all the interfaces* in the applied VLANs in configurations. For MSTP, configure all interfaces in all VLANs at the **[edit protocols mstp interface]** hierarchy level. PR860226
- On EX4300 Virtual Chassis, VSTP convergence might take more time than expected. PR915914

- On EX4300 Virtual Chassis, if the system has a VSTP configuration with a large number (253) of VLANs, CPU idle might remain at 0 percent and an l2cpd core file might be created. PR919835

Related Documentation

- New and Changed Features on page 4
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Resolved Issues on page 23
- Documentation Updates on page 29
- Migration, Upgrade, and Downgrade Instructions on page 30
- Product Compatibility on page 36

Resolved Issues

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.

- Issues Resolved in Release 13.2X50-D10 on page 23
- Issues Resolved at Release 13.2X50-D15 on page 25

Issues Resolved in Release 13.2X50-D10

The following issues have been resolved since Junos OS Release 13.2X50-D10. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On EX Series switches, the LLDP-MED media endpoint class is shown as invalid. This problem is just a display issue—there is no functional impact. PR840915

Bridging and Learning

- For all EX Series switches except EX8200 switches, hash index collisions were causing problems with the learning of MAC addresses in the forwarding database (FDB). You can now increase the maximum number of searchable hash indexes in increments of 4, from 4 to a maximum of 32 entries, using the CLI command **set ethernet-switching-options max-lookup-length**. PR842439
- On an EX4200 switch configured for VLAN translation, Windows NetBIOS traffic might not be translated. PR791131
- If an EX Series switch has a redundant trunk group (RTG) link, a MAC Refresh message might be sent on a new active link of the RTG when RTG failover occurs. The switch sends the RTG MAC Refresh message with a VLAN tag even though RTGs are configured on access ports. PR853911

High Availability (HA) and Resiliency

- On EX8200 switches, multiple rpd process core files might be created on the backup Routing Engine after a nonstop software upgrade (NSSU) has been performed while multicast traffic is on the switch. PR841848

Infrastructure

- If you reboot the switch with the RVI disabled, then even if you reenable the RVI, RVI traffic is not routed in the Packet Forwarding Engine; the traffic is trapped to the CPU and is policed by the rate limit in the Packet Forwarding Engine. PR838581
- On EX8200 switches, the **commit synchronize** command might fail with the error message **error: could not open configuration database (juniper.data+)**. PR844315
- On EX Series switches, if you configure a physical interface's MTU with a large value and you do not reconfigure the family inet MTU, OSPF packets that are larger than 1900 bytes might be dropped when they reach the internal logical interface. All communication traffic between Routing Engines and between FPCs passes through the internal logical interface. The OSPF neighbor does not receive the OSPF transmissions and ends the OSPF session. The switch displays the error message **bmeb_rx failed**. PR843583

Layer 2 Features

- On EX Series switches, the Cisco Discovery Protocol and the VLAN Trunking Protocol do not work through L2PT. PR842852

Network Management and Monitoring

- On EX Series switches, the Q-BRIDGE-MIB OID 1.3.6.1.2.1.17.7 reports the VLAN internal index instead of the VLAN ID. PR850299
- On EX Series switches, a configured OAM threshold value might get reset when the chassis is rebooted. PR829649
- The SNMP query or walk on ipNetToMediaPhysAddress does not match the **show arp** command output. PR850051
- An SNMP MIB walk might show unwanted data for newly added objects such as jnxVirtualChassisPortInPkts or jnxVirtualChassisPortInOctets. PR791848

Routing Policy and Firewall Filters

- On EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6210 switches, a firewall filter with family set to **ethernet-switching** and configured for IPv4 blocks specific transit IPv6 traffic if the ether_type match condition in the filter is not explicitly set to ipv4. As a workaround, set ether_type to ipv4 in the filter. PR843336

Software Installation and Upgrade

- The **unlink** option in the **request system software add** *package-name* **unlink** command does not work on EX Series switches. PR739795

Virtual Chassis

- In a mixed EX4200 and EX4500 Virtual Chassis, link aggregation might cause a PFEM core in some member switches. PR846498
- On EX4200 Virtual Chassis, **CHASSISD_SNMP_TRAP6: SNMP trap generated: Fan/Blower Removed** messages might be generated periodically, even when member switches cited in the messages are not present in the Virtual Chassis. PR858565
- On EX2200 Virtual Chassis, when there are multiple equal-cost paths, the **show virtual-chassis vc-path source-interface** *interface-name* **destination-interface** *interface-name* command displays the first discovered shortest path, even though traffic might be flowing through an alternate path. PR829752

Issues Resolved at Release 13.2X50-D15

Authentication and Access Control

- On EX4300 switches, moving MAC-RADIUS clients across VLANs might not work as expected. PR840104
- On EX4300 switches, when scaling up to 10,000 clients, the switch might create a dot1x core file. PR886279
- On EX4300 switches, previously learned MAC addresses might not be flushed from the Ethernet-switching table in these cases:
  - If you change the supplicant mode of an 802.1X-enabled interface from single to single-secure
  - If you change the supplicant mode of a captive-portal-enabled interface
  - If you change a captive-portal-enabled interface to an 802.1X-enabled interface

  PR848336, PR850690, PR852890
- On an EX4300 Virtual Chassis, clients might be authenticated even when the **no-mac-learning** statement is configured. PR862724
- On an EX4300 Virtual Chassis, after you issue the **restart l2-learning** command, an interface authenticated on the guest VLAN might have other port VLAN associations. PR866927

- On an EX4300 Virtual Chassis, a MAC-RADIUS user and an 802.1X user with the same MAC might be authenticated on two 802.1X-enabled interfaces. PR867959
- On an EX4300 Virtual Chassis, after you disable and enable a dynamic VLAN user-authenticated interface, dynamic VLAN association with 802.1X might not work as expected. PR867994
- On EX4300 switches, if a PC is connected to the switch through a captive-portal-enabled interface and you move that interface from one VLAN to another VLAN, the PC might not be able to obtain an IP address through DHCP. PR868140
- On an EX4300 Virtual Chassis, network policy Type Length Value (TLV) messages are not sent out after an 802.1X (dot1x) configuration is deleted from a VoIP-enabled interface. PR868840
- On an EX4300 Virtual Chassis, after clients have been authenticated and then you reboot the l2-learning process, the VLAN on which the clients were originally authenticated might be flooded before authentication. PR877854
- On EX4300 switches, with 802.1X and MAC RADIUS enabled on an interface and successful MAC-RADIUS authentication on that interface with assignment to a dynamic VLAN, if you disable the interface and then reenable it, the MAC-RADIUS client is authenticated, but the MAC address of the client is not learned in the Ethernet switching table. PR880077
- On EX4300 switches, after you have enabled 802.1X on an interface and a client has been authenticated through bypass authentication, if you then clear the authenticated 802.1X sessions on that interface, the MAC address for the client is not removed from the Ethernet switching table. PR880804
- On an EX4300 Virtual Chassis, the Layer 2 address learning (l2ald) process does not learn the media access control (MAC) address with a dynamic VoIP VLAN. PR895188
- On EX4300 switches, if a client is authenticated in single-supplicant mode on an interface that is 802.1X-enabled, that interface is assigned dynamically to a VoIP VLAN, and then that client logs off, the VoIP VLAN binding for that interface might be dropped. PR897684

Bridging and Learning

- On an EX4300 Virtual Chassis, when **interface-mac-limit** is configured, the switch might not be able to differentiate the same port number on two different Virtual Chassis members while the switch is counting MACs. For example, if ge-0/0/0 and ge-1/0/0 are configured with an **interface-mac-limit** value of 2, and if only ge-0/0/0 is configured with the packet action drop, if ge-1/0/0 learns two MACs, then traffic from ge-0/0/0 is dropped. PR900283

Class of Service

- On EX4300 switches, you can configure rewrite rules only at the port level and not at the sub-interface level. PR806487
- On EX Series switches, EXP CoS classification does not occur if you delete EXP CoS classifiers and then add them back. PR848273

- On EX4300 switches, after you configure a fixed classifier on a Layer 2 interface, then delete the classifier, and then execute a configuration rollback operation, the classifier does not take effect on the egress Layer 2 interface. PR892473

Firewall Filters

- On EX4300 switches, Layer 3 and Layer 4 match conditions are not supported on egress port-based firewall filters or egress VLAN-based firewall filters. PR816888

Infrastructure

- On an EX4300 Virtual Chassis, MAC learning and ARP resolution might fail among interfaces in a VLAN that are connected to the backup when VSTP is enabled on some VLANs and not on others. As a workaround, bring the affected interfaces down and then up again. PR822708
- On EX4300 switches, traffic might continue to be routed on an integrated routing and bridging (IRB) interface even after the interface is disabled. PR831484
- On an EX4300 Virtual Chassis, an SNMP MIB walk for the jnxAnalyzerMIB does not produce any output. PR840835
- On EX4300 switches and EX4300 Virtual Chassis, maintenance intermediate point (MIP) functionality might not work properly. PR842780
- On an EX4300 Virtual Chassis, configuring and then removing the configuration for IGMP snooping multiple times might create a VMcore file. PR853988
- On an EX4300 Virtual Chassis that is sending Layer 3 multicast traffic, the system might create a periodic packet management process (ppmd) core file. PR854630
- On an EX4300 Virtual Chassis, if you issue the **clear spanning-tree statistics interface** command, the system might create a Layer 2 Control Protocol process (l2cpd) core file. PR854972
- On an EX4300 Virtual Chassis, if you issue the **clear ethernet-switching recovery-timeout** command, the system displays the error message **error: timeout communicating with l2-learning daemon**. PR856639
- On EX4300 switches, when IGMP snooping is enabled and a new interface is added to a VLAN, packet loss occurs on the interface configured as the multicast receiver in the VLAN. PR858197
- On EX4300 switches, if you issue the command **clear interfaces statistics all** or **clear interfaces statistics *interface-name***, the ARP inspection statistics are not cleared. PR858517
- On EX4300 switches and EX4300 Virtual Chassis, after a graceful Routing Engine switchover (GRES), spanning-tree protocol statistics might not be distributed. PR858542
- In some conditions, a kernel panic occurs on an EX4300 switch, and the device stays in db> mode. PR860338
- On an EX4300 Virtual Chassis, the **action-shutdown** configuration for storm control does not take effect, even though storm control is working as configured and the traffic

storm is contained. When traffic exceeds the configured limit, the interface does not shut down; instead, the traffic is rate-limited. PR865700

- On EX4300 switches and EX4300 Virtual Chassis, when you restart the routing process, the periodic packet management process (ppmd) might create a core file. PR866173
- On an EX4300 Virtual Chassis, OAM connectivity fault management (CFM) sessions might not form on LAG member links after a system reboot. PR867590
- On EX4300 switches, the uplink failure detection (UFD) process creates core files when changes are made to the UFD configuration groups. PR869121
- On an EX4300 Virtual Chassis, if you reboot the master switch or run a graceful Routing Engine switchover (GRES), the periodic packet management process (ppmd) might create a core file. PR887431
- In a mixed mode Virtual Chassis having EX4200 and EX4500 switches, loading the Junos OS image using the **request system software add set** command fails, and the following error message is displayed: **invalid use of multiple package names or files**.

  As a workaround, install the Junos OS image on each member in the Virtual Chassis, one member at a time, using the command **request system software add *package* member *member-id***. PR874239
- On an EX4300 Virtual Chassis, when you reboot the master switch, the master switch regains mastership, but the chassis control subsystem might not run. PR891832
- On an EX4300 switch, when you remove and then reinsert a 10-gigabit DAC cable into an uplink port, a pfex core file might be created. PR893483
- On EX4300 switches, storm control default rate limiting might not occur on 40-gigabit ports and on 10-gigabit ports.PR904937
- On EX Series switches, if an interface is a member of an interface range and is not also configured directly at the **[edit interfaces *interface-name*]** hierarchy level, the MTU value configured with the interface range might not be deleted when you delete the interface range configuration. PR859160

Network Management and Monitoring

- On EX4200 and EX4500 switches, adaptive sampling is triggered on interfaces configured for sFlow monitoring technology even though the sampling rate is less than 300 pps. PR840858

Port Security

- On EX4300 switches, an unsuccessful DHCP request might result in a stale snooping entry in the DHCP snooping table. PR896751
- On EX4300 switches, if persistent DHCP snooping is configured, the snooping entries are written to the configured file. If you then reboot the switch, the DHCP snoop entries are retrieved in the control plane. The entries might not be present in the kernel and the Packet Forwarding Engine, and that fact affects DAI and IP source guard if those features are configured. PR900238

Software Upgrade and Installation

- On EX3300 switches, an SFID core file might be created when you change the software on downstream switches. PR856455

Related Documentation

- New and Changed Features on page 4
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 18
- Documentation Updates on page 29
- Migration, Upgrade, and Downgrade Instructions on page 30
- Product Compatibility on page 36

Documentation Updates

This section lists the errata and changes in Junos OS Release 13.2X50 for EX Series switches documentation.

- Changes to Junos OS for EX Series Switches Documentation on page 29
- Errata on page 29

Changes to Junos OS for EX Series Switches Documentation

No changes have been made to the documentation for Junos OS Release 13.2X50 for EX Series switches since it was published.

Errata

This section lists outstanding issues with the published documentation for Junos OS Release 13.2X50 for EX Series switches.

- Q-in-Q tunneling is now supported on EX4300 switches. The documentation does not contain details at this time. See "New and Changed Features" on page 4 for a description of the Q-in-Q tunneling feature and configuration options.

Related Documentation

- New and Changed Features on page 4
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 18
- Resolved Issues on page 23
- Migration, Upgrade, and Downgrade Instructions on page 30
- Product Compatibility on page 36

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies 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.

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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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

For information on software installation and upgrade, see the Installation and Upgrade Guide.

Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations

If you are upgrading to Junos OS Release 12.1R2 or later from Release 12.1R1 or earlier, ensure that any VSTP configurations on the switch meet the following guidelines. If the VSTP configurations do not meet these guidelines and you run the upgrade, the upgrade fails, and you have to connect the console, change the invalid VSTP configurations, and

commit the changed configurations through the console. Guidelines for VSTP configurations are:

- If you have specified physical interfaces for VSTP-configured VLANs, ensure that those interfaces are members of the VLANs specified in the VSTP configuration. If the VSTP configuration specifies **vlan all**, then the interfaces configured at the **[edit protocols vstp vlan all]** hierarchy level must be members of all VLANs.
- If the interfaces are not members of the VLANs in the VSTP configurations but are already added to the VSTP configurations, remove them from those configurations, add them to the VLANs, and then add them back to the VSTP configurations.

This issue is being tracked by PR/736488 in our bug database.

Upgrading from Junos OS Release 10.4R3 or Later

This section contains the procedure for upgrading from Junos OS Release 10.4R3 or later to Junos OS Release 12.2 or later. You can use this procedure to upgrade Junos OS on a standalone EX Series switch with a single Routing Engine and to upgrade all members of a Virtual Chassis or a single member of a Virtual Chassis.

To upgrade Junos OS on an EX6200 or EX8200 switch with dual Routing Engines, see Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure).

On switches with dual Routing Engines or on Virtual Chassis, you might also be able to use nonstop software upgrade (NSSU) to upgrade Junos OS. See "Upgrading EX Series Switches Using NSSU" on page 33 for more information.

To upgrade Junos OS on a switch with a single Routing Engine or on a Virtual Chassis:

1. Download the software package as described in Downloading Software Packages from Juniper Networks.
2. (Optional) Back up the current software configuration to a second storage option. See the *Junos OS Installation and Upgrade Guide* for instructions.
3. (Optional) Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.

   This step is optional because you can also upgrade Junos OS using a software image that is stored at a remote location.
4. Install the new software package on the switch:

   ```
   user@switch> request system software add package
   ```

   Replace ***package*** with one of the following paths:

   - **/var/tmp/*package*.tgz**—For a software package in a local directory on the switch
   - **ftp://*hostname/pathname/package*.tgz** or **http://*hostname/pathname/package*.tgz**—For a software package on a remote server

   package.tgz is the name of the package; for example, **jinstall-ex-4200-11.4R1.8-domestic-signed.tgz**.

To install software packages on all switches in a mixed EX4200 and EX4500 Virtual Chassis, use the **set** option to specify both the EX4200 package and the EX4500 package:

```
user@switch> request system software add set [package package]
```

To install the software package on only one member of a Virtual Chassis, include the **member** option:

```
user@switch> request system software add package member member-id
```

Other members of the Virtual Chassis are not affected. To install the software on all members of the Virtual Chassis, do not include the **member** option.

NOTE: To abort the installation, do not reboot your device. Instead, finish the installation and then issue the **request system software delete** *package*.tgz command, where *package*.tgz is the name of the package; for example, **jinstall-ex-8200-11.4R1.8-domestic-signed.tgz**. This is the last chance to stop the installation.

5. Reboot the switch to start the new software:

   ```
   user@switch> request system reboot
   ```

   To reboot only a single member in a Virtual Chassis, include the **member** option:

   ```
   user@switch> request system reboot member
   ```

6. After the reboot has finished, log in and verify that the new version of the software is properly installed:

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

7. Once you have verified that the new Junos OS version is working properly, copy the version to the alternate slice to ensure that if the system automatically boots from the backup partition, it uses the same Junos OS version:

   ```
   user@switch> request system snapshot slice alternate
   ```

   To update the alternate root partitions on all members of a Virtual Chassis, include the **all-members** option:

   ```
   user@switch> request system snapshot slice alternate all-members
   ```

Upgrading from Junos OS Release 10.4R2 or Earlier

To upgrade to Junos OS Release 13.2 from Junos OS Release 10.4R2 or earlier, first upgrade to Junos OS Release 11.4 by following the instructions in the Release 11.4 release notes. See *Upgrading from Junos OS Release 10.4R2 or Earlier* or *Upgrading from Junos OS Release 10.4R3 or Later* in the Junos OS 11.4 Release Notes.

Upgrading from Junos OS Release 12.3 or Earlier

If you are upgrading from Junos OS Release 12.3 or earlier to Release 13.2 or later, notice that the display output for the **show version** operational command has changed. In Junos OS Release 13.2, several of the software packages installed on the switch that were displayed in the **show version** output in Junos OS Release 12.3 or earlier no longer appear.

Instead, these software packages have been combined into fewer software packages, with some of the new packages having new names, such as JUNOS EX Software Suite.

Upgrading to a Controlled Version of Junos OS

Starting in Junos OS Release 13.2X50-D15, two versions of a Junos OS image—a controlled version that supports Media Access Control Security (MACsec) and a domestic version that does not support MACsec—are available for EX Series switches. In previous Junos OS releases for EX Series switches, the domestic version of Junos OS was the only available Junos OS. If you want to enable Media Access Control Security (MACsec), you must install the controlled version of Junos OS in your switch.

If you are upgrading your switch between the domestic version of Junos OS and the controlled version of Junos OS, keep the following issues in mind:

- You can use NSSU to upgrade or downgrade from a domestic version of Junos OS to a controlled version of Junos OS. You cannot use NSSU to upgrade or downgrade from a controlled version of Junos OS to a domestic version of Junos OS, however.
- In a Virtual Chassis, all member switches must be running the same release of Junos OS. If you connect member switches that are running domestic and controlled versions of the same Junos OS release, the switches do successfully join together in a Virtual Chassis. To support MACsec, however, all member switches in the Virtual Chassis must be running the *controlled* version of Junos OS.

The upgrade or downgrade procedure from a domestic version of Junos OS to a controlled version of Junos OS is, otherwise, identical to any other Junos OS upgrade. See Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure) or Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure).

[See also "New and Changed Features" on page 4 for more information about MACsec and about the controlled version.]

Upgrading EX Series Switches Using NSSU

You can use nonstop software upgrade (NSSU) to upgrade Junos OS releases on standalone EX6200 and EX8200 switches with dual Routing Engines and on EX3300, EX4200, EX4500, and EX8200 Virtual Chassis. For instructions on how to perform an upgrade using NSSU, see:

- Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, or Mixed EX4200 and EX4500 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)
- Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)
- Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)

Table 1 on page 34 details the switch platforms on which NSSU is supported and the required Junos OS release.

Table 1: Platform and Junos OS Upgrade Support for NSSU

Switch Platform	Upgrade from Junos OS Release *x.x*	Upgrade to Junos OS Release 12.3
EX3300 Virtual Chassis	Releases earlier than 12.1R2	Not supported
	12.1R2 or later	Supported
	12.2R1 or later	Supported
EX4200 Virtual Chassis, EX4500 Virtual Chassis, and mixed EX4200 and EX4500 Virtual Chassis	Releases earlier than 12.1R1	Not supported
	12.1R1 or later	Supported
	12.2R1 or later	Supported
EX6200 standalone switch	Releases earlier than 12.1R2	Not supported
	12.1R2 or later	Supported
	12.2R1 or later	Supported
EX8200 standalone switch	10.4R1 or 10.4R2	Not supported
	10.4R3 or later	Supported
	11.1R1 or later	Supported
	11.2R1 or later	Supported
	11.3R1 or later	Supported
	11.4R1 or later	Supported
	12.R1 or later	Supported
	12.2R1 or later	Supported

Table 1: Platform and Junos OS Upgrade Support for NSSU *(continued)*

Switch Platform	Upgrade from Junos OS Release *x.x*	Upgrade to Junos OS Release 12.3
EX8200 Virtual Chassis	10.4R1 or later	Not supported
	11.1R1, 11.1R2, or 11.1R3	Not recommended
	11.1R4 or later	Supported
	11.2R1 or later	Supported
	11.3R1 or later	Supported
	11.4R1 or later	Supported
	12.1R1 or later	Supported
	12.2R1 or later	Supported

NOTE: On EX4200, EX4300, and EX4550 Virtual Chassis, you can use NSSU to upgrade from a domestic version of Junos OS to a controlled (MACsec) version of Junos OS. You cannot, however, use NSSU to upgrade from the controlled version of Junos OS to a domestic version of Junos OS.

NOTE: **NSSU upgrades to Junos OS Release 12.3R2**—You can use the following releases for NSSU upgrades to Release 12.3R2:

- EX8200 standalone switches and EX8200 Virtual Chassis—Releases 11.4R7 or later and 12.1R5 or later
- EX3300 Virtual Chassis—Releases 12.2R3 or later and 12.3R1 or later
- EX4500 and EX4200 Virtual Chassis—Release 12.1R5 or later
- EX4550 Virtual Chassis—Releases 12.2R3 or later and 12.3R1 or later
- EX6200 switches—Releases 12.2R3 or later and 12.3R1 or later

On an EX8200 Virtual Chassis, an NSSU operation can be performed only if you have configured the XRE200 External Routing Engine member ID to be 8 or 9.

NOTE: Do not use nonstop software upgrade (NSSU) to upgrade the software on an EX8200 switch from Junos OS Release 10.4 if you have configured the IGMP, MLD, or PIM protocols on the switch. If you attempt to use NSSU, your switch might be left in a nonfunctional state from which it is difficult to recover. If you have these multicast protocols configured, upgrade the software on the EX8200 switch from Junos OS Release 10.4 by following the instructions in Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure). This issue does not apply to upgrades from Junos OS Release 11.1 or later.

NOTE: If you are using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 10.4 or Junos OS Release 11.1 and sFlow technology is enabled, disable sFlow technology before you perform the upgrade using NSSU. After the upgrade is complete, you can reenable sFlow technology. If you do not disable sFlow technology before you perform the upgrade with NSSU, sFlow technology does not work properly. This issue does not affect upgrades from Junos OS Release 11.2 or later.

NOTE: If you are using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 11.1 and NetBIOS snooping is enabled, disable NetBIOS snooping before you perform the upgrade using NSSU. After the upgrade is complete, you can reenable NetBIOS snooping. If you do not disable NetBIOS snooping before you perform the upgrade with NSSU, NetBIOS snooping does not work properly. This issue does not affect upgrades from Junos OS Release 11.2 or later.

Related Documentation

Product 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 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: http://pathfinder.juniper.net/feature-explorer/

Related Documentation

- New and Changed Features on page 4
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 18
- Resolved Issues on page 23
- Documentation Updates on page 29
- Migration, Upgrade, and Downgrade Instructions on page 30

Junos OS Release Notes for the QFX Series

These release notes accompany Junos OS Release 13.2X50 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

- New and Changed Features on page 37
- Changes in Behavior and Syntax on page 42
- Known Behavior on page 44
- Known Issues on page 48
- Migration, Upgrade, and Downgrade Instructions on page 53
- Product Compatibility on page 59

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 13.2X50-D15 for the QFX Series.

- Virtual Chassis on page 39
- Junos OS Software on page 39
- Interfaces and Chassis on page 41

- Network Management and Monitoring on page 42
- **Junos XML and API Scripting**

Virtual Chassis

- **Virtual Chassis support (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—QFX3500 and QFX3600 devices configured as standalone switches can now be interconnected to form a QFX Series Virtual Chassis. You can configure up to 10 total member switches—the 10 total member switches can be any combination of QFX3500 and QFX3600 series switches—into a Virtual Chassis and manage the interconnected switches as a single chassis. The advantages of connecting multiple switches into a Virtual Chassis include better-managed bandwidth at a network layer, simplified configuration and maintenance because multiple devices can be managed as a single device, increased fault tolerance and high availability (HA) because a Virtual Chassis can remain active and network traffic can be redirected to other member switches when a single member switch fails, and a flatter, simplified Layer 2 network topology that minimizes or eliminates the need for some first hop redundancy protocols such as Virtual Router Redundancy Protocol (VRRP) and loop prevention protocols such as Spanning Tree Protocol (STP).

  You interconnect QFX Series devices into a Virtual Chassis by cabling two QFX3500 or QFX3600 devices together using the 40-Gigabit Ethernet QSFP+ uplink ports or the fixed SFP+ 10-Gigabit Ethernet ports and issuing the **request virtual-chassis vc-port set** command to configure the connecting interfaces into Virtual Chassis ports (VCPs).

  [See Understanding QFX Series Virtual Chassis.]

  

  NOTE: You cannot designate a 10-Gigabit Ethernet port channelized from a 40-Gigabit Ethernet QSFP+ port as a VCP.

Junos OS Software

- **Junos OS Enhanced Layer 2 Software (ELS) configuration statements and operational commands (QFX3500 and QFX3600 switches)**—Provides a uniform CLI for configuring and monitoring Layer 2 features on newer model EX Series, MX Series, and QFX Series devices that support ELS, such as the EX4300 and EX9200 switches and MX Series routers in LAN mode (MX-ELM).

  The Web-based *ELS Translator tool* is available for registered customers to help them become familiar with the ELS CLI and to translate existing QFX Series switch-based CLI configurations into ELS CLI configurations quickly.

  

  NOTE: If you upgrade a QFX3500 or QFX3600 switch to Junos OS Release 13.2X50-D15 or later from Junos OS Release 12.3X50 or earlier, you need to convert your configuration from the original style Junos OS CLI to the ELS CLI format. If you do not convert your configuration, the upgrade fails.

  To upgrade your switch from a version of Junos OS that does not support ELS to a version of Junos OS that supports ELS, we recommend performing the following procedure.

  

  NOTE: Because this procedure can cause service outages, we recommend that you avoid performing this procedure on switches carrying traffic in a production network.

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

     

     NOTE: Only perform this procedure from the console port. You can lose connectivity to your device if you perform this procedure from a management port or any other interface.

  2. Copy your entire existing configuration into a text file. Save the file to a remote location or USB drive.
  3. Retain the portion of your existing configuration related to management network connectivity (such as **[edit system]**). Delete all other top-level configuration hierarchy levels (such as **[edit interfaces]**, **[edit protocols]**, and **[edit vlans]**). Issue a **commit** operation to remove the deleted configuration hierarchy levels.
  4. Perform the software upgrade and reboot your device to complete the upgrade. Maintain your console port connection during the reboot.
  5. Using a web browser, navigate to the ELS Translator Tool . Follow the instructions on the page to convert your saved configuration file to the new ELS CLI format.

6. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator Tool and load it in to your switch.

7. Issue a commit operation to activate the translated configuration.

[See Getting Started with Enhanced Layer 2 Software and ELS Translator Tool.]

Interfaces and Chassis

- **New method for channelizing 40-Gigabit Ethernet QSFP+ interfaces (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—Enables you to configure four 10-Gigabit Ethernet interfaces from a 40-Gigabit Ethernet QSFP+ interface. By default, a 40-Gigabit Ethernet QSFP+ interface is 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. [See Channelizing Interfaces on QFX3500 and QFX3600 Standalone Switches.]

  NOTE: When you configure 10-Gigabit Ethernet channelization for a 40-Gigabit Ethernet QSFP+ interface, the Packet Forwarding Engine restarts and a mastership switchover occurs.

- **64th port available (QFX3500 switches)**—Enables you to configure port **xe-0/1/0** (on QSFP+ port Q0), which was previously unavailable. To make this port available for use, issue the **request chassis port-mode extended** command at the root level. To disable this port, issue the **request chassis port-mode standard** command. After enabling or disabling port **xe-0/1/0**, issue a commit operation and reboot the system to allow the change to take effect.

Network Management and Monitoring

- **Network analytics feature (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—Provides visibility into the performance and behavior of the data center infrastructure by enabling high-frequency traffic statistics collection and microburst monitoring. You use network analytics to monitor queue statistics and traffic statistics, which include queue depth, latency, and traffic information, at user configurable intervals. The network analytics reports can help identify problem areas in your network traffic and applications so that you can adjust resources as needed. You configure network analytics at the **[edit services analytics]** hierarchy level. Supported platforms are the QFX3500 and QFX3600 standalone switches, and the QFX3500 and QFX3600 Virtual Chassis. [See Network Analytics Overview.]

Junos XML and API Scripting

- **Packaging Python Scripts**—Python is available on devices running Junos OS when the following packages are installed:
  - jinstall-ex-4300
  - jinstall-qfx
  - jinstall-dc-re

  The Python interpreter only runs scripts that have been installed from signed packages created with the Junos SDK. For more information, see "Using Python on Junos" and "Building a Junos SDK Package Containing Python Scripts" in the *Junos SDK Developer Guide*.

Related Documentation

- Changes in Behavior and Syntax in Junos OS Release 13.2X50 for the QFX Series on page 42
- Known Behavior in Junos OS Release 13.2X50 for the QFX Series on page 44
- Known Issues in Junos OS Release 13.2X50 for the QFX Series on page 48
- Migration, Upgrade, and Downgrade Instructions for Junos OS Release 13.2X50 for the QFX Series on page 53
- Product Compatibility on page 59

Changes in Behavior and Syntax

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 13.2X50-D15 for the QFX Series.

- Interfaces and Chassis
- Platform and Infrastructure

Interfaces and Chassis

- **Consistent CLI for channelizing 40-Gigabit Ethernet QSFP+ interfaces (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—In previous software releases, you would channelize a 40-Gigabit Ethernet QSFP+ port on a QFX3600 switch into four 10-Gigabit Ethernet channels by including the **xe** statement at the **[edit chassis fpc *fpc-slot* pic *pic-slot*]** hierarchy level. To make the implementation consistent for both the QFX3500 and QFX3600 switches, you now channelize a 40-Gigabit Ethernet port into four 10-Gigabit Ethernet channels by including 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.

  40-Gigabit Ethernet QSFP+ interfaces in previous releases were called **xle-** interfaces, but are now called **et-** interfaces.

Platform and Infrastructure

- **Alarms and traps are generated when the primary slice fails and the system boots from the alternate slice (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—If the active slice fails and the system boots from the alternate slice, a banner is displayed before the login prompt: **WARNING: THIS DEVICE HAS BOOTED FROM THE BACKUP JUNOS IMAGE. It is possible that the primary copy of JUNOS failed to boot up properly, and so this device has booted from the backup copy. Please re-install JUNOS to recover the primary copy in case it has been corrupted.** Additionally, an alarm is generated if you disable the automatic snapshot function.
- **Read-only directories in the root partition (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—All of the directories that reside in the "/" partition are read only.
- **Automatic snapshot with manual overrride (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—Allows you to take a snapshot of the active slice and copy it to the alternate slice on the boot media by issuing the **request system snapshot slice alternate** command.
- **Extra PIC slots are displayed in the CLI (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—Extra PIC slots are displayed in the CLI to accommodate members in different Virtual Chassis platforms that contain more PIC slots (such as the EX4300 switch). The range of available PIC slots is 0 to the maximum number of PICs available on a Virtual Chassis platform. The **request chassis pic fpc-slot *slot-number* pic-slot *slot-number*** command and **[edit chassis fpc *fpc-slot* pic *pic-slot*]** hierarchy level display the PIC slot range as **0..2** to accommodate the PIC slot range on other Virtual Chassis platforms, even though PIC slot 2 is unused on the standalone QFX3500 switch (which only uses PIC slots 0 and 1), and PIC slots 1 and 2 are unused on the standalone QFX3600 switch (which only uses slot 0).

Related Documentation

- New and Changed Features in Junos OS Release 13.2X50 for the QFX Series on page 37
- Known Behavior in Junos OS Release 13.2X50 for the QFX Series on page 44
- Known Issues in Junos OS Release 13.2X50 for the QFX Series on page 48

- Migration, Upgrade, and Downgrade Instructions for Junos OS Release 13.2X50 for the QFX Series on page 53
- Product Compatibility on page 59

Known Behavior

This section contains the known behavior, system maximums, and limitations in Junos OS Release 13.2X50-D15 for the QFX Series.

- Interfaces and Chassis
- Network Management and Monitoring
- Traffic Management
- User Interface and Configuration
- Virtual Chassis

Interfaces and Chassis

- **Virtual Chassis port (VCP) limitation (QFX3500 and QFX3600 switches running Enhanced Layer 2 Software)**—When creating Virtual Chassis ports for system management, you cannot designate a 10-Gigabit Ethernet port channelized from a 40-Gigabit Ethernet QSFP+ port as a VCP. Instead, use fixed 10-Gigabit Ethernet ports in the QFX3500 switch or nonchannelized 40-Gigabit Ethernet ports in either the QFX3500 or QFX3600 switch.
- After you insert a QSFP+ transceiver in a QFX3500 device, the status LED does not light green for up to 15 seconds.
- Filter-based forwarding is not supported with IPv6 on a QFX Series Virtual Chassis.
- IP directed broadcast is not supported on a QFX Series Virtual Chassis.

Network Management and Monitoring

- If a QFX3500 switch drops traffic because of an ingress firewall filter, the switch does not generate an sFlow monitoring technology flow sample packet that contains this dropped traffic.

Traffic Management

- **Access interface CoS support**—CoS on Virtual Chassis access interfaces is the same as CoS on QFX Series access interfaces with the exception of shared buffer settings. All of the documentation for QFX Series CoS on access interfaces applies to Virtual Chassis access interfaces.

Virtual Chassis access interfaces support the following CoS features:

- Forwarding classes—The default forwarding classes, queue mapping, and packet drop attributes are the same as on QFX Series access interfaces:

Default Forwarding Class	Default Queue Mapping	Default Packet Drop Attribute
best-effort (be)	0	drop
fcoe	3	no-loss
no-loss	4	no-loss
network-control (nc)	7	drop
mcast	8	drop

- Packet classification—Classifier default settings and configuration are the same as on QFX Series access interfaces. Support for behavior aggregate, multifield, multidestination, and fixed classifiers is the same as on QFX Series access interfaces.
- Enhanced transmission selection (ETS)—This data center bridging (DCB) feature that supports hierarchical scheduling has the same defaults and user configuration as on QFX Series access interfaces, including forwarding class set (priority group) and traffic control profile configuration.
- Priority-based flow control (PFC)—This DCB feature that supports lossless transport has the same defaults and user configuration as on QFX Series access interfaces, including support for six lossless priorities (forwarding classes).
- Ethernet PAUSE—Same defaults and configuration as on QFX Series access interfaces.
- Queue scheduling—Same defaults, configuration, and scheduler-to-forwarding-class mapping as on QFX Series access interfaces. Queue scheduling is a subset of hierarchical scheduling.
- Priority group (forwarding class set) scheduling—Same defaults and configuration as on QFX Series access interfaces. Priority group scheduling is a subset of hierarchical scheduling.
- Tail-drop profiles—Same defaults and configuration as on QFX Series access interfaces.
- Code-point aliases—Same defaults and configuration as on QFX Series access interfaces.
- Rewrite rules—As on the QFX Series access interfaces, there are no default rewrite rules applied to egress traffic.
- Host outbound traffic—Same defaults and configuration as on QFX Series access interfaces.

The default shared buffer settings and shared buffer configuration are also the same as on QFX Series access interfaces, except that the shared buffer configuration is global and applies to all access ports on all members of the Virtual Chassis. You cannot configure different shared buffer settings for different Virtual Chassis members.

- **Similarities in CoS support on VCP interfaces and QFabric system Node device fabric interfaces**—VCP interfaces support full hierarchical scheduling (ETS). ETS includes:
  - Creating forwarding class sets (priority groups) and mapping forwarding classes to forwarding class sets.
  - Scheduling for individual output queues. The scheduler defaults and configuration are the same as the scheduler on access interfaces.
  - Scheduling for priority groups (forwarding class sets) using a traffic control profile. The defaults and configuration are the same as on access interfaces.
  - No other CoS features are supported on VCP interfaces.

NOTE: You cannot attach classifiers, congestion notification profiles, or rewrite rules to VCP interfaces. Also, you cannot configure buffer settings on VCP interfaces. Similar to QFabric system Node device fabric interfaces, you can only attach forwarding class sets and traffic control profiles to VCP interfaces.

The behavior of lossless traffic across 40-Gigabit VCP interfaces is the same as the behavior of lossless traffic across QFabric system Node device fabric ports. Flow control for lossless forwarding classes (priorities) is enabled automatically. The system dynamically calculates buffer headroom that is allocated from the global lossless headroom buffer for the lossless forwarding classes on each 40-Gigabit VCP interface. If there is not enough global headroom buffer space to support the number of lossless flows on a 40-Gigabit VCP interface, the system generates a syslog message.

NOTE: After you configure lossless transport on a Virtual Chassis, check the syslog messages to ensure that there is sufficient buffer space to support the configuration.

NOTE: If you break out a 40-Gigabit VCP interface into 10-Gigabit VCP interfaces, lossless transport is not supported on the 10-Gigabit VCP interfaces. Lossless transport is supported only on 40-Gigabit VCP interfaces.

- **Differences in CoS support on VCP interfaces and QFabric system Node device fabric interfaces**—Although most of the CoS behavior on VCP interfaces is similar to CoS behavior on QFabric system Node device fabric ports, there are some important differences:

  - Hierarchical scheduling (queue and priority group scheduling)—On QFabric system Node device fabric interfaces, you can apply a different hierarchical scheduler (traffic control profile) to different priority groups (forwarding class sets) on different interfaces. However, on VCP interfaces, the schedulers you apply to priority groups are global to all VCP interfaces. One hierarchical scheduler controls scheduling for a priority group on all VCP interfaces.

    You attach a scheduler to VCP interfaces using the global identifier (vcp-*) for VCP interfaces. For example, if you want to apply a traffic control profile (which contains both queue and priority group scheduling configuration) named *vcp-fcoe-tcp* to a forwarding class set named *vcp-fcoe-fcset*, you include the following statement in the configuration:

    ```
    [edit]
    user@switch# set class-of-service interfaces vcp-* forwarding-class-set vcp-fcoe-fcset
    output-traffic-control-profile vcp-fcoe-tcp
    ```

    The system applies the hierarchical scheduler *vcp-fcoe-tcp* to the traffic mapped to the priority group *vcp-fcoe-fcset* on all VCP interfaces.

  - You cannot attach classifiers, congestion notification profiles, or rewrite rules to VCP interfaces. Also, you cannot configure buffer settings on VCP interfaces. Similar to QFabric system Node device fabric interfaces, you can only attach forwarding class sets and traffic control profiles to VCP interfaces.

  - Lossless transport is supported only on 40-Gigabit VCP interfaces. If you break out a 40-Gigabit VCP interface into 10–Gigabit VCP interfaces, lossless transport is not supported on the 10-Gigabit VCP interfaces.

- **CPU-generated host outbound traffic**—CPU-generated host outbound traffic is forwarded on the network-control forwarding class, which is mapped to queue 7. If you use the default scheduler, the network-control queue receives a guaranteed minimum bandwidth (transmit rate) of 5 percent of port bandwidth. The guaranteed minimum bandwidth is more than sufficient to ensure lossless transport of host outbound traffic.

  However, if you configure a scheduler, you must ensure that the network-control forwarding class (or whatever forwarding class you configure for host outbound traffic) receives sufficient guaranteed bandwidth to prevent packet loss.

  If you configure a scheduler, we recommend that you configure the network-control queue (or the queue you configure for host outbound traffic if it is not the network-control queue) as a strict-high priority queue. Strict-high priority queues receive the bandwidth required to transmit their entire queues before other queues are served.

NOTE: As with all strict-high priority traffic, if you configure the network-control queue (or any other queue) as a strict-high priority queue, you must also create a separate forwarding class set (priority group) that contains only strict-high priority traffic, and apply the strict-high priority forwarding class set and its traffic control profile (hierarchical scheduler) to the VCP interfaces.

- You cannot apply classifiers and rewrite rules to IRB interfaces because the members of an IRB are VLANs, not interfaces. You can apply classifiers and rewrite rules to Layer 2 logical interfaces and Layer 3 physical interfaces that are members of VLANs that belong to IRB interfaces.

User Interface and Configuration

- On a QFX Series Virtual Chassis, if you configure RIPng, the system might not receive any Update packets from the neighbor.

Virtual Chassis

- On a QFX Series Virtual Chassis, if you configure sFlow and an egress firewall filter, the filter might not block the sFlow samples and they might arrive at the collector.

Related Documentation

- New and Changed Features in Junos OS Release 13.2X50 for the QFX Series on page 37
- Changes in Behavior and Syntax in Junos OS Release 13.2X50 for the QFX Series on page 42
- Known Issues in Junos OS Release 13.2X50 for the QFX Series on page 48
- Migration, Upgrade, and Downgrade Instructions for Junos OS Release 13.2X50 for the QFX Series on page 53
- Product Compatibility on page 59

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 13.2X50-D15 for the QFX Series.

For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at http://www.juniper.net/prsearch .

- Interfaces and Chassis
- Multiprotocol Label Switching (MPLS)
- Platform and Infrastructure
- Routing Protocols

- User Interface and Configuration
- Virtual Chassis

Interfaces and Chassis

- On an EX4300 or QFX Series Virtual Chassis, the **show virtual-chassis vc-port diagnostics optics** operational command is not supported. PR855574
- On a QFX Series Virtual Chassis, if two linecard members are connected by both a 40-Gigabit Ethernet Virtual Chassis port (VCP) link and a 10-Gigabit Ethernet VCP link, it is likely that only the 40-Gigabit Ethernet VCP link is actively used. In this case, if the active 40-Gigabit Ethernet VCP link is administratively deleted, the Virtual Chassis connection is temporarily broken, even though the 10-Gigabit Ethernet link is present. When the system recovers, the Virtual Chassis connection is re-established on the 10-Gigabit Ethernet link. PR868340
- On a QFX Series Virtual Chassis, if you include the **drop** or **drop-and-log** options at the **[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action]** hierarchy level, frames do not flood in the same VLAN as expected. PR913210
- On a QFX Series Virtual Chassis, after you reboot a linecard member of the Virtual Chassis, there might be some traffic loss as the Virtual Chassis path resumes forwarding to the shortest path. PR914064
- On a QFX Series Virtual Chassis, if you initiate a graceful Routing Engine switchover, BFD sessions might take up to 2 minutes to become established. PR915462
- On a QFX Series Virtual Chassis, if the linecard members are connected in a linear topology and if you reboot them members, it is likely that the linecards will not back come online and rejoin the Virtual Chassis. As a workaround, connect the Virtual Chassis linecard members in a ring topology. PR915597
- On a QFX Series Virtual Chassis, if you configure a static IP/MAC address mapping for the **dhcp-security** forwarding option for an IRB interface and issue the **show dhcp-security binding** command, the static IP/MAC address binding entry is not displayed in the DHCP security binding table. PR915877
- On a QFX Series Virtual Chassis, if you configure DHCP security and a lease period of 300 seconds or more, when the DHCP server is down, the DHCP client might be released before the lease time expires. PR916285
- On a QFX Series Virtual Chassis, if you configure persistent learning for DHCP snooping and reboot the system, the learned entries might not remain in the output of the **show dhcp-security binding** command. PR917255
- On a QFX Series Virtual Chassis, if you configure DHCP Relay and apply DHCP security to a trusted port, the system might drop DHCP REPLY messages from the server. PR917263
- On a QFX Series Virtual Chassis, if you delete an active Virtual Chassis link by removing a QSFP+ transceiver or issuing the **request virtual-chassis vc-port delete pic-slot *slot* port *port-number* member *member-name*** command, traffic is impacted even if there is a redundant Virtual Chassis link present. PR917937

- On a QFX Series Virtual Chassis, when you configure filter-based forwarding for IPv6 and apply an ingress filter to the participating interface, the system might generate an error during the commit operation. PR918170
- On a QFX Series Virtual Chassis, if you configure DHCP Relay for the system and a DHCP client sends an INFORM message, the DHCP server might not respond with an acknowledgement (ACK). PR918508
- On a QFX Series Virtual Chassis, if you configure IP directed broadcast on the IRB interface, and change the address prefix or disable and reenable the interface, traffic is not forwarded to the destination. PR920665

Multiprotocol Label Switching (MPLS)

- On a QFX Series Virtual Chassis, if you configure LDP over RSVP and apply an EXP rewrite on a LAG, the rewrite might not work and traffic might not be directed to the desired queue. PR918501

Platform and Infrastructure

- On a QFX Series Virtual Chassis, if the Virtual Chassis dumps core, some important files (such as log messages and configuration files) might be missing from the core output. PR875519
- On a QFX Series Virtual Chassis, if you channelize a 40-Gigabit Ethernet QSFP+ interface into four 10-Gigabit Ethernet ports and try to apply an interface range configuration to one of the channels, the interface range feature does not work. PR890292
- On a QFX Series Virtual Chassis, if you issue an IPv6 **ping** command between MPLS customer edge devices, the ping fails. PR920220

Routing Protocols

- On a QFX Series Virtual Chassis, if you edit a firewall filter, the counters of other firewall filters might reset. PR868450
- On a QFX Series Virtual Chassis, when you explicitly configure an IPv6 firewall filter that discards OSPFv3 packets, the ingress filter might not discard the OSPFv3 packets. PR897786
- On a QFX Series Virtual Chassis, if you apply a firewall filter to the management port, the filter might not work. PR897981
- On a QFX Series Virtual Chassis, if you delete a member of a LAG associated with an IRB interface, the counter for the filter applied to the IRB interface might reset. PR898171
- On a QFX Series Virtual Chassis, if you assign a forwarding class and loss priority to traffic using firewall filters instead of assigning a forwarding class and loss priority using the **[edit class-of-service]** hierarchy level, then the rewrite rules do not work. PR898734
- On a QFX Series Virtual Chassis, if you configure a firewall filter to count OSPF database description (DD) packets, the filter might not work. PR905349

- On a QFX Series Virtual Chassis, when you issue a single commit operation for an interface that has both a VLAN and a firewall filter configured, the VLAN application might not take effect. PR910864
- On a QFX Series Virtual Chassis, if you apply a firewall filter on a virtual management (**vme**) interface, the filter might not work. PR915668

User Interface and Configuration

- On a QFX Series Virtual Chassis, if you configure a routing trunk group (RTG), the configuration might not be programmed into the Packet Forwarding Engine. PR875172
- On a QFX Series Virtual Chassis, if you configure MAC limiting on a logical interface, the system might incorrectly set the packet action to **drop**, resulting in packets being discarded. PR881763
- On a QFX Series Virtual Chassis, if a port is shut down by storm control and you issue the **show interfaces extensive** command, the output might show incorrect information (such as the interface being up, inbound packets per second, and other erroneous details). PR897064
- On a QFX Series Virtual Chassis, if you configure RIPng, the system might not receive any Update packets from the neighbor. PR900868
- On a QFX Series Virtual Chassis, if you issue the **clear dhcp server bindings interface** command or try to perform a server release from the DHCP client, the DHCP bindings are not cleared and remain intact until the lease timer expires. PR901844
- On a QFX Series Virtual Chassis, if you remove a DHCP configuration, the commit operation might fail and generate the message **error: DHCP service may not be de-configured while clients are present. Please clear bindings.** As a workaround, issue the **clear dhcp server** or **clear dhcp relay** commands to clear the bindings of any clients before committing a change that impacts client bindings (such as VLANs or interfaces that contain clients). PR901868
- On a QFX Series Virtual Chassis, if you issue the **traceroute** command for an IPv6 destination, the output does not show details for the first hop. PR907410
- On a QFX Series Virtual Chassis, if you configure two VLANs on an interface, include the **packet-action drop** statement at the **[edit vlans *vlan-name* switch-options mac-table-size]** hierarchy level for only one of the VLANs, and issue the **clear ethernet-switching table** command, traffic might be dropped on both VLANs (rather than just the VLAN configured for this action). In this case, the traffic loss is seen on both VLANs until the switch learns the source MAC address. PR915432
- On a QFX Series Virtual Chassis, if traffic is flowing across an aggregated Ethernet interface, the output of the **show interfaces ae*X*** command might display interface statistics that do not increment properly. PR915565
- On a QFX Series Virtual Chassis, if you include the **drop-and-log** statement at the **[edit vlans *vlan-name* switch-options mac-move-limit packet-action]** hierarchy level, the system does not drop the packets. PR915942

- On a QFX Series Virtual Chassis, if you enable IGMP snooping on all VLANs while multicast data traffic is active, and you issue the **clear pim join** command, multicast data packet flooding might occur. PR917246
- On a QFX Series Virtual Chassis, if you include the **shutdown** option at the **[edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action]** hierarchy level and issue the **commit** operation, the system generates a commit error. PR918262
- On a QFX Series Virtual Chassis, if you include the **persistent-learning** option at the **[edit switch-options interface *interface-name*]** hierarchy level on a port, each MAC address learned on the port is counted twice. If you also configure MAC limiting on the port, then counting the learned MAC addresses twice can lead to inadvertent MAC limiting because the system learns only half the of the configured MAC limit. PR918482
- On a QFX Series Virtual Chassis, if you delete an IGMP snooping configuration for one or more multicast groups, multicast route entries might be deleted from the Routing Engine and Packet Forwarding Engine, and cause traffic loss for about 40 seconds. PR918543
- On a QFX Series Virtual Chassis, if you configure your system with both the **interface-mac-limit** and **static-mac** switch options, the system does not restrict traffic based on the static MAC address entries as expected. PR918671
- On a QFX Series Virtual Chassis, if you configure VSTP and then disable one of the LAG members, the VSTP port role changes but traffic does not get forwarded through the other active LAG members. PR919457
- On a QFX Series Virtual Chassis, if you issue the **clear dhcp server binding** command and attempt to delete a DHCP server configuration, the commit operation fails. PR921717

Virtual Chassis

- On a QFX Series Virtual Chassis, if you configure sFlow and an egress firewall filter, the filter might not block the sFlow samples and they might arrive at the collector. PR894386
- On a QFX Series Virtual Chassis, when you delete an interface that is a member of an aggregated Ethernet interface and then configure that interface with an IPv4 address (**family inet**) in one commit operation, the old logical interface (the one the interface used as a member of the aggregated interface) is not deleted. PR905124
- On a QFX Series Virtual Chassis, when you enable IGMPv4 on a interface, IPv6 neighbor solicitation messages might not be forwarded. PR911446
- On a QFX Series Virtual Chassis, unicast reverse-path forwarding (unicast RPF) firewall filter counters are not supported. PR917493
- On a QFX Series Virtual Chassis, if you disable and reenable a LAG interface that receives output from an analyzer, the mirrored traffic does not arrive at the LAG interface. PR920610
- On a QFX Series Virtual Chassis, if the master has a LAG interface and reboots, there might be 50 seconds of traffic loss on one link of the LAG bundle connected to the

master. As a workaround, when you plan to reboot the master switch, disable the LAG interface connected to the master before the reboot.PR923484

Related Documentation

- New and Changed Features in Junos OS Release 13.2X50 for the QFX Series on page 37
- Changes in Behavior and Syntax in Junos OS Release 13.2X50 for the QFX Series on page 42
- Known Behavior in Junos OS Release 13.2X50 for the QFX Series on page 44
- Migration, Upgrade, and Downgrade Instructions for Junos OS Release 13.2X50 for the QFX Series on page 53
- Product Compatibility on page 59

Migration, Upgrade, and Downgrade Instructions

This section discusses the following topics:

Procedure for Upgrading CoS from Junos OS Release 11.1 or Release 11.2 to Release 11.3 or Later

Before you upgrade to Junos OS Release 11.3 or later, you must deactivate the CoS configuration on the QFX3500 switch if the CoS configuration uses the excess-rate option, strict-high or high priority queues, or any of the default multidestination forwarding classes. For full information about this topic, see Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release). A summary of the upgrade steps is included here.

After you upgrade to Junos OS Release 11.3 or later, modify the CoS configuration on the QFX3500 switch to conform to the Junos OS Release 11.3 or later CoS requirements. Then activate the CoS configuration and commit the changes:

1. Deactivate the CoS configuration.

   user@switch# **deactivate class-of-service**

2. Upgrade to Junos OS Release 11.3 or later.
3. Make the following changes to the CoS configuration:
   - Remove the excess-rate option from the CoS configuration if you have used it at the **[edit class-of-service schedulers]** or **[edit class-of-service traffic-control-profiles]** hierarchy level.

- Remove the default multidestination forwarding classes (mcast-be, mcast-af, mcast-ef, and mcast-nc) if you have used them at the **[edit class-of-service schedulers]**, **[edit class-of-service rewrite-rules]**, or **[edit class-of-service classifiers]** hierarchy level. Alternatively, you can change the mapping of the multidestination traffic to use the new default multidestination forwarding class (mcast).

4. If desired, configure strict-high priority queues in accordance with the Junos OS Release 11.3 or later strict-high priority queue rules, and map multidestination traffic to the default multidestination forwarding class (mcast).
5. Activate the CoS configuration.

   user@switch# **activate class-of-service**

6. Commit the CoS configuration.

NOTE: If you have configured the transmit-rate option for any queues at the **[edit class-of-service schedulers]** hierarchy level, if the rate is configured as an exact rate in Mbps, we recommend that you reconfigure the transmit-rate option as a percentage. This is because the scheduler converts exact rates to percentages, and when the exact rate is below 1 Gbps, some granularity may be lost in the conversion. You can avoid this potential issue by specifying the transmit-rate option as a percentage.

CoS Upgrade Requirements from Junos OS Release 12.1 to Junos OS Release 12.2 and Later

Before you upgrade to Junos OS Release 12.2 or later, you might need to edit the class-of-service (CoS) configuration because the way the QFX Series handles lossless forwarding classes has changed starting with Junos OS Release 12.2.

By default, the fcoe and no-loss forwarding classes are mapped to output queue 3 and output queue 4, respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport.

In Junos OS Release 12.1 and earlier, explicitly setting the lossless fcoe and no-loss forwarding classes resulted in the same CoS behavior as using the default configuration. However, in Junos OS Release 12.2 and later, the behavior when you explicitly configure the lossless forwarding classes differs from the behavior when you use the default forwarding classes.

NOTE: The default behavior differs from the explicit configuration behavior even if the explicit configuration is exactly the same as the default configuration.

- If you use the default forwarding class configuration for the lossless queues (the configuration does not include explicit setting of the fcoe or the no-loss forwarding classes), then the fcoe and no-loss queues behave as lossless queues.

If your CoS configuration does not explicitly configure the fcoe and no-loss forwarding classes, you can upgrade from Junos OS Release 12.1 to Junos OS Release 12.2 and later, and the behavior of the two lossless queues remains lossless.

- If your configuration includes statements that explicitly configure the fcoe or the no-loss forwarding class (using the **[set class-of-service forwarding-classes class *class-name* queue-num *queue-number*]** statement), after you upgrade to Junos OS Release 12.2 or later, those queues do *not* receive lossless treatment and behave as lossy (best-effort) queues.

  If your CoS configuration explicitly configures the fcoe and no-loss forwarding classes, to retain the lossless behavior of those queues, you need to remove the explicit configuration for these two forwarding classes from the CoS configuration *before* you upgrade.

  If you upgrade to Junos OS Release 12.2 or later and the fcoe and no-loss forwarding classes are explicitly configured, then those two queues continue to be used, but the traffic is treated as lossy traffic, not lossless traffic. To make the queues for these two forwarding classes lossless, you must delete the explicit forwarding class configuration.

  

  CAUTION: If you explicitly configured the fcoe or the no-loss forwarding class, and you upgrade to Junos OS Release 12.2 or later, the system does not return an upgrade error or a commit error, or generate a syslog message to notify you that these forwarding classes are no longer lossless. Traffic mapped to these forwarding classes is not treated as lossless traffic until you remove the explicit forwarding class configuration.

  Before you upgrade, delete the fcoe and no-loss forwarding classes from the explicit configuration to preserve the lossless behavior of traffic mapped to these forwarding classes.

  - To delete the explicit fcoe forwarding class configuration:

    ```
    [edit]
    user@switch# delete class-of-service forwarding-class class fcoe queue-num 3
    user@switch# commit
    ```

  - To delete the explicit no-loss forwarding class configuration:

    ```
    [edit]
    user@switch# delete class-of-service forwarding-class class no-loss queue-num 4
    user@switch# commit
    ```

  

  NOTE: If you try to delete these forwarding classes and they have not been explicitly configured on the system, the system returns the message **warning: statement not found**. This simply means that there is no explicit configuration to delete and does not change the lossless behavior of the fcoe and no-loss forwarding classes.

  After you delete the explicit configuration for the fcoe and no-loss forwarding classes, traffic mapped to those forwarding classes retains its lossless behavior after the upgrade to Junos OS Release 12.2 or later.

- If you have a Layer 3 classifier attached to a nonzero unit of an interface, we recommend that you deactivate the CoS configuration before you upgrade from Junos OS Release 12.1 to a later Junos OS release. Alternatively, you can delete the following configuration hierarchy for each Layer 3 interface: **set class-of-service interfaces *interface-name* unit *unit-number* classifiers**. Although the CLI in Junos OS Release 12.1 permitted you to assign a Layer 3 classifier to a logical interface, Layer 3 classifiers on logical interfaces are not supported on the QFX Series.

Procedure for Upgrading to an ELS-Based Software package

If you upgrade a QFX3500 or QFX3600 switch to Junos OS Release 13.2X50-D15 or later from Junos OS Release 12.3X50 or earlier, you need to convert your configuration from the original style Junos OS CLI to the Enhanced Layer 2 Software (ELS) CLI format. If you do not convert your configuration, the upgrade fails.

To upgrade your switch from a version of Junos OS that does not support ELS to a version of Junos OS that supports ELS, we recommend performing the following procedure.

NOTE: **Because this procedure can cause service outages, we recommend that you avoid performing this procedure on switches carrying traffic in a production network.**

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

NOTE: **Only perform this procedure from the console port. You can lose connectivity to your device if you perform this procedure from a management port or any other interface.**

2. Copy your entire existing configuration into a text file. Save the file to a remote location or USB drive.
3. Retain the portion of your existing configuration related to management network connectivity (such as **[edit system]**). Delete all other top-level configuration hierarchy levels (such as **[edit interfaces]**, **[edit protocols]**, and **[edit vlans]**). Issue a **commit** operation to remove the deleted configuration hierarchy levels.
4. Perform the software upgrade and reboot your device to complete the upgrade. Maintain your console port connection during the reboot.
5. Using a web browser, navigate to the ELS Translator Tool . Follow the instructions on the page to convert your saved configuration file to the new ELS CLI format.
6. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator Tool and load it in to your switch.
7. Issue a **commit** operation to activate the translated configuration.

Upgrading Software on QFX3500 and QFX3600 Standalone 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 Junos OS Installation and Upgrade Guide and Junos OS Basics in the QFX Series documentation.

NOTE: You cannot upgrade by more than three releases at a time. For example, if your routing device is running Junos OS Release 11.1, you can upgrade to Junos OS Release 11.3 but not to Junos OS Release 12.1. As a workaround, first upgrade to Junos OS Release 11.3 and then upgrade to Junos OS Release 12.1.

NOTE: In some cases, when you downgrade the QFX3500 switch to an earlier software version, the switch might not operate properly. As a workaround, choose one of the following options when downgrading:

1. Issue the **request system software add** command to downgrade to the following or later software versions:
   - Junos OS Release 11.1R5
   - Junos OS Release 11.2R2
   - Junos OS Release 11.3R1
2. Include the **no-validate** option when you issue the **request system software add** command during a downgrade to a software version earlier than the ones listed in option #1.

The download and installation process for Junos OS Release 13.2 is the same as for previous Junos OS releases.

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

1. In a browser, go to http://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 **13.2** 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 Virtual Chassis Install Package for the 13.2 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 validate
source/jinstall-qfx-3-13.2X50-D15.3-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)

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

Adding the **reboot** command reboots the switch after the upgrade is validated and 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 13.2 jinstall package, you cannot issue the request system software rollback command to return to the previously installed software. 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 Extended End-of-Life Software 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. However, you cannot upgrade directly from a non-EEOL release that is more than three releases before or after.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release. For more information on EEOL releases and to review a list of EEOL releases, see Junos Software Release Dates and Milestones.

Related Documentation

- New and Changed Features in Junos OS Release 13.2X50 for the QFX Series on page 37
- Changes in Behavior and Syntax in Junos OS Release 13.2X50 for the QFX Series on page 42
- Known Behavior in Junos OS Release 13.2X50 for the QFX Series on page 44
- Known Issues in Junos OS Release 13.2X50 for the QFX Series on page 48
- Product Compatibility on page 59

Product Compatibility

- Hardware Compatibility on page 59

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 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:
http://pathfinder.juniper.net/feature-explorer/

Related Documentation

- New and Changed Features in Junos OS Release 13.2X50 for the QFX Series on page 37
- Changes in Behavior and Syntax in Junos OS Release 13.2X50 for the QFX Series on page 42
- Known Behavior in Junos OS Release 13.2X50 for the QFX Series on page 44
- Known Issues in Junos OS Release 13.2X50 for the QFX Series on page 48
- Migration, Upgrade, and Downgrade Instructions for Junos OS Release 13.2X50 for the QFX Series on page 53

Third-Party Components

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

For a list of open source attributes for this Junos OS release, see Open Source: Source Files and Attributions.

Finding More Information

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

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: http://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: http://www.juniper.net/techpubs/content-applications/content-explorer/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at http://www.juniper.net/customers/support/downloads/710059.pdf .

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

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool located at 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 http://www.juniper.net/cm/ .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

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

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

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

Revision History

25 September 2013—Revision 1, Junos OS for the EX Series and QFX Series, Release 13.2X50-D15

3 October 2013—Revision 2, Junos OS for the EX Series and QFX Series, Release 13.2X50-D15, added expanded information for upgrading to an ELS-based software package.