

Release Notes: Junos[®] OS Release 13.2X51-D25 for the EX Series and QFX Series

Release 13.2X51-D25
29 July 2014
Revision 1

Contents

Junos OS Release Notes for EX Series Switches	4
New and Changed Features	4
Hardware	4
Ethernet Switching	5
Software Installation and Upgrade	6
Changes in Behavior and Syntax	6
Interfaces	6
Known Behavior	7
Authentication and Access Control	8
High Availability	8
Interfaces and Chassis	8
Traffic Management	8
Known Issues	10
Authentication and Access Control	10
Bridging and Learning	11
Class of Service	11
High Availability	11
Infrastructure	11
Interfaces and Chassis	12
J-Web	12
Layer 2 Protocols	13
Multicast Protocols	13
Network Management and Monitoring	13
Platform and Infrastructure	13
Port Security	13
Routing Policy and Firewall Filters	14
Software Installation and Upgrade	14

Virtual Chassis	15
Virtual Chassis Fabric	15
Resolved Issues	16
Issues Resolved in Junos OS Release 13.2X51-D25	16
Issues Resolved in Junos OS Release 13.2X51-D21	17
Issues Resolved in Release 13.2X51-D20	17
Issues Resolved in Release 13.2X51-D15	20
Issues Resolved in Release 13.2X50-D15	24
Documentation Updates	24
Migration, Upgrade, and Downgrade Instructions	25
Upgrade and Downgrade Support Policy for Junos OS Releases	25
Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations	25
Upgrading to a Controlled Version of Junos OS	26
Product Compatibility	26
Hardware Compatibility	27
Junos OS Release Notes for the QFX Series	27
New and Changed Features	28
Hardware	28
Ethernet Switching	29
Software Installation and Upgrade	30
Virtual Chassis	30
Changes in Behavior and Syntax	30
Interfaces	30
IPv6	31
Software Upgrade	31
System Management	31
Known Behavior	32
Multiprotocol Label Switching (MPLS)	32
Network Management and Monitoring	32
Platform and Infrastructure	32
Traffic Management	32
Known Issues	35
Class of Service (CoS)	35
High Availability (HA) and Resiliency	36
Interfaces and Chassis	36
Platform and Infrastructure	36
Routing Protocols	38
Software Installation and Upgrade	38
Storage and Fibre Channel	38
Resolved Issues	39
Issues Resolved in Junos OS Release 13.2X51-D25	39
Issues Resolved in Junos OS Release 13.2X51-D21	39
Issues Resolved in Junos OS Release 13.2X51-D20	40
Issues Resolved in Junos OS Release 13.2X51-D15	40
Documentation Updates	42
System Management	42

Migration, Upgrade, and Downgrade Instructions	42
Upgrading Software on QFX5100 Standalone Switches	43
Performing an In-Service Software Upgrade (ISSU)	44
Preparing the Switch for Software Installation	44
Upgrading the Software Using ISSU	45
Product Compatibility	46
Hardware Compatibility	47
Third-Party Components	47
Finding More Information	47
Documentation Feedback	47
Requesting Technical Support	48
Self-Help Online Tools and Resources	48
Opening a Case with JTAC	49
Revision History	49

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 13.2X51 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 on page 4](#)
- [Changes in Behavior and Syntax on page 6](#)
- [Known Behavior on page 7](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions on page 25](#)
- [Product Compatibility on page 26](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 13.2X51-D25 for the EX Series.

- [Hardware on page 4](#)
- [Ethernet Switching on page 5](#)
- [Software Installation and Upgrade on page 6](#)

Hardware

- **EX4600-40F switch**—Starting with Junos OS Release 13.2X51-D25, the Juniper Networks EX4600-40F switch is supported. The EX4600 is a compact 1 U model that provides wire-speed packet performance, very low latency, and a rich set of Layer 2 and Layer 3 features. In addition to a high-throughput Packet Forwarding Engine, the performance of the control plane running on the EX4600 switch is enhanced by the 1.5GHz dual-core Intel CPU with 8 gigabits of memory and 32 gigabits of solid-state-drive (SSD) storage. The EX4600-40F is a 10-Gigabit Ethernet small form-factor pluggable plus (SFP+) top-of-rack switch with 24 fixed SFP+ access ports and 4 quad SFP+ (QSFP+) high-speed links.

The EX4600-40F switch can be used as:

- A standalone switch.
- A master, backup, or “linecard” member in a Virtual Chassis with EX4600 switches or with EX4300 switches. An EX4600 Virtual Chassis allows you to interconnect a mixture of up to 10 switches into one logical device and manage the device as a single chassis. An EX4600 Virtual Chassis is cabled in a ring topology.



BEST PRACTICE: Use EX4600 switches as the master and backup members in a mixed EX4600 and EX4300 Virtual Chassis, because the EX4600 does not support Power over Ethernet (PoE), 802.1X, or Link Layer Discovery Protocol (LLDP).

Eligible Virtual Chassis members are:

- EX4600 as a master, backup, or linecard
- EX4300 as a linecard
- **Support for the EX4600-EM-8S expansion module (EX4600-40F switch)**—Starting with Junos OS Release 13.2X51-D25, the EX4600-EM-8S expansion module is supported. It provides eight additional 10-Gigabit Ethernet SFP+ ports. The EX4600-EM-8S expansion module can coexist with the QFX-EM-4Q expansion module in the same EX4600 switch. Expansion modules can be hot-inserted or hot-removed.

Ethernet Switching

- **CoS Explicit congestion notification (ECN) (EX4600)**—Enables end-to-end congestion notification between two endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

[See [Understanding CoS Explicit Congestion Notification](#).]

- **Q-in-Q tunneling enhancement (EX4300, EX4600)**—With previous releases of Junos OS 13.2X51, you cannot create a regular VLAN on an interface if you have created an S-VLAN or C-VLAN on that interface for Q-in-Q tunneling. This means that you cannot create an integrated routing and bridging (IRB) interface on that interface because regular VLANs are a required part of IRB configuration. With Junos OS 13.2X51-D25, you can create a regular VLAN on a trunk interface that has an S-VLAN, which means that you can also create an IRB interface on the trunk. You still cannot create a regular VLAN on an access interface that has a C-VLAN. [See [Understanding Q-in-Q Tunneling](#).]



NOTE: A regular VLAN and an S-VLAN on the same trunk interface cannot share the same VLAN ID.

Software Installation and Upgrade

- **Junos OS Release 13.2X51-D25 image (QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches)**—Junos OS Release 13.2X51-D25 is a superset of Release 13.2X51-D20 and is provided for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches. Release 13.2X51-D25 is the recommended release for those switches.

Related Documentation

- [Changes in Behavior and Syntax on page 6](#)
- [Known Behavior on page 7](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions on page 25](#)
- [Product Compatibility on page 26](#)

Changes in Behavior and Syntax

This section lists the changes in default behavior and syntax in Junos OS Release 13.2X51 for EX Series switches.

Interfaces

- **Link aggregation group (LAG) bundle and Equal-Cost Multipath (ECMP) next-hop hashing configuration (EX4300 switches)**—You can now configure the fields that the hashing algorithm uses to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled. For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device. Configuring the fields used by the hashing algorithm helps users manage traffic flows when a switch is using LAG bundles or ECMP, and is especially helpful in scenarios when most of the traffic is similar. You configure the hashing algorithm at the **[edit forwarding-options enhanced-hash-key]** hierarchy level. [See [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic](#).]

Related Documentation

- *New and Changed Features*
- *Known Behavior*
- *Known Issues*
- *Resolved Issues*
- *Documentation Updates*
- *Migration, Upgrade, and Downgrade Instructions*
- *Product Compatibility*

Known Behavior

This section lists the limitations in Junos OS Release 13.2X51-D25 for the EX Series.

Authentication and Access Control

- On EX4300 switches, after you clear the MAC addresses from an Ethernet-switching table, the MAC RADIUS authentication sessions is not cleared from the authentication table if the traffic is continuous. [PR833888](#)

High Availability

- On EX4300 Virtual Chassis, the configuration database might get stuck in the Synchronizing state and the Virtual Chassis might not be able to do a switchover after multiple Routing Engine switchovers or mastership changes. As a workaround, issue either the **commit synchronize** command or the **commit synchronize force** command on the master Routing Engine. This is a known software limitation. [PR965661](#)

Interfaces and Chassis

- For aggregated Ethernet interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands do not include broadcast packet information. Also, for aggregated Ethernet interfaces, the SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are not supported. The counter values are always 0. [This is a known software limitation.]

Traffic Management

- The packet replication engine on EX4600 switches replicates flood traffic at line rate. (Flood traffic is unicast traffic for which the switch has not learned the destination address and other traffic such as broadcast or multicast traffic that must be replicated on all of the member interfaces of a VLAN). The engine can replicate two packets per clock cycle.

The number of packets replicated per second depends on the clock speed of the replication engine. A faster clock speed means a faster clock cycle. A faster clock cycle means more packets replicated per second. If the number of flood packets to replicate exceeds the engine capacity, flood packets are dropped.

All of the switch interfaces share the replication engine. When the number of flood packets to replicate exceeds the engine capacity, all excess unlearned unicast and multdestination traffic is dropped on all VLANs, not just on the VLAN that is oversubscribing the replication engine. [Table 1 on page 9](#) shows the clock speed and the maximum number of packets the engine can replicate per second on an EX4600 switch.

Table 1: EX4600 Switch Clock Speed and Replication Engine Capacity

Clock Speed	Maximum Number of Packets Replicated Per Second
635 MHz	1.27 Gigapackets

Packet size affects how many packets must be replicated. The absolute number of packets-per-second to replicate decreases as packet size increases (the larger the packet size, the fewer the number of packets-per-second to process).

The calculation to determine the number of packets being flooded in a VLAN is:

$$(\text{interface-speed} / \text{packet-size}) * \text{number-of-packet-replications} * \text{number-of-ports-flooding-traffic} = \text{number-of-packets-per-second}$$

For example, if a VLAN has ten 10-Gigabit Ethernet member interfaces, then every flood packet must be replicated nine times (the packet is not flooded out of the ingress interface). If a flood of 64-byte packets occurs on the VLAN, and every interface is flooding traffic to the other interfaces in the VLAN, the calculation to determine the number of packets the engine must replicate per second at line rate is:

$$(10\text{-Gbps} / (64 * 8)) * 9 * 10 = 1.758 \text{ Gigapackets}$$

In this example, *10-Gbps* is the interface speed, *(64 * 8)* is the packet size in bits (64 bytes * 8 bits), *9* is the number of times the engine must replicate each packet, and *10* is the number of ports flooding at line rate. In this case, the requirement to replicate *1.758 Gigapackets-per-second* exceeds the capacity of the replication engine, so some flood traffic is dropped.

For another example, if all interfaces in the same VLAN flood 128-byte packets instead of 64-byte packets:

$$(10\text{-Gbps} / (128 * 8)) * 9 * 10 = 0.879 \text{ Gigapackets}$$

In this case, the engine can replicate all of the packets at line rate because the total number of packets-per-second is well within the replication engine capacity.

Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 6](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions on page 25](#)
- [Product Compatibility on page 26](#)

Known Issues

The following issues are outstanding in Junos OS Release 13.2X51-D25. The identifier following the description is the tracking number in our bug database.

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>.

- [Authentication and Access Control](#)
- [Bridging and Learning](#)
- [Class of Service](#)
- [High Availability](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [J-Web](#)
- [Layer 2 Protocols](#)
- [Multicast Protocols](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Port Security](#)
- [Routing Policy and Firewall Filters](#)
- [Software Installation and Upgrade](#)
- [Virtual Chassis](#)
- [Virtual Chassis Fabric](#)

[Authentication and Access Control](#)

- On EX4300 switches, when you enter the **set protocols dot1x authenticator interface all** command, a commit warning might appear. [PR892082](#)
- On EX4300 switches, if you change the supplicant mode on an interface, the interface does not fall back to 802.1X authentication from captive-portal authentication. [PR920134](#)

- On EX4300 switches, if a VoIP phone is attached to the switch, the phone might not receive an IP address from the DHCP server until the phone is unplugged and plugged back in. [PR985856](#)

Bridging and Learning

- On EX4300 switches, under the [edit vlans *vlan-name* switch-options] hierarchy level, a value for the **mac-table-size** option might commit without committing **interface-mac-limit**, and the **mac-table-size** option then would not work properly—that is, the number of MAC table entries might exceed the **mac-table-size** value. [PR977984](#)
- On EX4300 switches that are enabled for 802.1X authentication and configured in multi-supplicant mode with guest VLAN, if a client is authenticated via guest VLAN and later requests authentication through EAP for the same MAC address, the client might not be immediately authenticated. [PR1009117](#)
- On EX4300 switches that are enabled for 802.1X authentication and configured in multisuppliant mode with guest VLAN and server-reject VLAN, if a client is authenticated through the guest VLAN method and later tries to get authenticated via EAP with incorrect credentials for the same MAC address, the client might be reauthenticated on the guest VLAN again. [PR1009134](#)
- On EX4300 switches that are enabled for 802.1X authentication and configured in either single-secure or multisuppliant mode, previous session information might be retained after the client has logged out. [PR1010602](#)

Class of Service

- On EX4300 switches, using IEEE 802.1p rewrite rules to set CoS code-point bits in outbound packets might not work properly when both IEEE 802.1p and DSCP rewrite rules are configured on a Layer 3 subinterface. [PR914889](#)

High Availability

- On EX4300 switches, rebooting the master (FPC1) might cause VRRP to flap. Also, there might be an STP loop for a short period. [PR857822](#)

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, 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 a different VLAN at runtime. [PR817456](#)
- 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, if you configure more than 512 VSTP instances, the switch might create a core file. [PR848278](#)
- On EX4300 switches, proxy ARP is not working after Layer 3 routes are changed. [PR889003](#)
- When a 40-gigabit link between an EX4300 switch and an EX4550 switch is connected with a DAC cable, the link does not come up if auto-negotiation is set on the EX4300. As a workaround, disable auto-negotiation on the EX4300 switch using the **set interfaces *interface-name* ether-options no-auto-negotiation** command. [PR935197](#)
- On EX4300 switches, an active interface participating in MVRP might not register and declare the VLANs that are included under **vlan-id-list** in a VLAN range. [PR950081](#)
- On EX4300 switches, MAC entries might be deleted from the Ethernet-switching table after you change the **interface-mac-limit packet-action** from **drop-and-log** to **drop**. [PR951001](#)
- On EX4300 switches, Ethernet ring protection (ERP) switching time does not happen within 50 ms; ERP data traffic loss occurs for approximately 145 ms. [PR968262](#)

Interfaces and Chassis

- On EX4300 switches, for Layer 3 logical interfaces, the traffic statistics for output packets displayed by the **show interfaces** command are incorrect. [PR824894](#)
- On EX4300 switches, VLAN MAC limit with drop action does not work. [PR911753](#)
- On EX4300 switches, when there is a limit on the number of MAC addresses that can be learned on an aggregated Ethernet interface, and the action configured on the interface is to shut down after reaching the MAC limit, the aggregated Ethernet interface might not shut down. [PR933168](#)

J-Web

- 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](#)
- On an EX4300 switch, if you use the J-Web user interface to request support information for all members at the same time, the switch might not be able to retrieve the information. As a workaround, request support information for each member one at a time. [PR911551](#)
- On EX4300 switches, the structured data format for system log messages is not supported in the J-Web interface. If system log messages are configured to be written in structured data, the event logs in J-Web will not be populated, and you will not be able to view them using **Monitor > Events and Alarms > View Events**. As a workaround, use the **show log** operational mode command for viewing structured-data format files. [PR959505](#)
- On an EX4300 switch, when you use the **Configure > Interface** menu in the J-Web user interface, you cannot add a VLAN range to any interface. As a workaround, configure

VLAN ranges in the CLI by including the **vlan-id-list** statement at the **[edit vlans vlan-name]** hierarchy level and applying the VLAN to the desired interface. [PR987059](#)

Layer 2 Protocols

- On EX4300 switches, the MSTI identifier range for MSTP is limited to 1--64. It should be 1--4094. [PR846878](#)
- 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](#)
- EX4300 switches might not switch packets across a VSTP-enabled interface and a redundant trunk group interface that belong to the same VLAN. [PR877467](#)
- On EX4300 switches, despite an administrative link being down, child members of an aggregated Ethernet group that is part of a multicast downstream IRB VLAN might be programmed into a multicast route index in the PFE. This situation results in the failure of multicast replication of packets for some VLANs. [PR880769](#)

Multicast Protocols

- On EX4300 switches, if MC data packets that fail an RPF check are received on a nonshared tree, the packets might be trapped on the Routing Engine at a high rate, resulting in poor PIM convergence. [PR911649](#)
- 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](#)

Network Management and Monitoring

- On EX4300 switches, there might be a difference of several milliseconds in the results for a two-way delay measurement using an SLA iterator profile and manual on-demand. [PR831541](#)
- On EX4300 switches, the **adjacencies** option is not available in the **show ethernet oam connectivity-fault-management** command. [PR848776](#)
- On EX4300 switches, two-way Ethernet frame delay measurement (OAM CFM) does not work in centralized mode. [PR960168](#)

Platform and Infrastructure

- On EX4600 switches, if you replace the EX4600-EM-8F module with the QFX-EM-4Q module or the reverse, all the interfaces on the switch will be brought down. [PR1003291](#)

Port Security

- 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 EX4300 switches, DHCPv6 snooping does not work when the client uses the DHCPv6 Rapid Commit Option. [PR941953](#)

Routing Policy and Firewall Filters

- On EX4300 switches, the **from interface** *interface-name* match condition is not supported on egress firewall filters. [PR817979](#)
- On EX4300 switches, in an egress router-based firewall filter, IPv6 Layer 4 headers (**icmp-type**) might not work. [PR912483](#)
- On EX4300 Virtual Chassis, packets that are generated in the CPU and egress out of a non-master FPC port might be subjected to an egress port-based firewall filter and be egress filtered, while packets that egress on a master FPC port might not be egress filtered. [PR923659](#)
- On EX4300 switches, the following match conditions configured in IPv6 egress router-based firewall filters and applied to the me0 or vme0 interface do not work: **source-address**, **destination-address**, **source-prefix-list**, and **destination-prefix-list**. [PR934196](#)
- On EX4300 switches, the following actions do not work when they are configured in ingress router-based firewall filters for IPv4 or IPv6 and applied to me0 or vme0 interfaces: **port-mirror** and **port-mirror-instance**. [PR935140](#)
- On EX4300 switches, the following actions do not work when you configure them in ingress router-based firewall filters for IPv4 or IPv6 and apply them to me0 or vme0 interfaces: **forwarding-class** and **loss-priority**. [PR935485](#)
- On EX4300 switches, the ingress router-based firewall filter action **three-color-policer** might not take effect for packets received on me0 and vme interfaces. [PR935859](#)
- On EX4300 switches, when a firewall filter is configured with the action **vlan**, traffic is not forwarded to the specified VLAN. [PR951798](#)
- On EX4300 switches, if you restart a firewall filter process (dswd) and then reboot the switch, the filter might stop working. [PR952306](#)
- If you apply a firewall filter that filters multicast traffic to a loopback interface on an EX4300 switch, the filter might also be applied to multicast traffic that is transiting the switch. [PR999000](#)
- On an EX4300 switch, if you apply a firewall filter containing the **reject** action to the loopback interface and perform an upgrade from Junos OS Release 13.2X50-D15 to Release 13.2X51-D21.1, the **reject** action might fail and accept the traffic. As a workaround, include the **discard** action in place of the **reject** action in the firewall filter. [PR1012014](#)

Software Installation and Upgrade

- NSSU to Junos OS Release 13.2X51-D25 is not supported for EX4300. [PR1010403](#)

- In an EX4300 Virtual Chassis, if you attempt to upgrade using NSSU from Junos OS Release 13.2X51-D20.3 to Release 13.2-X51-D25, you might see traffic outages for long durations and multiple core files might be created. [PR1012124](#)

Virtual Chassis

- 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 Virtual Chassis, if you renumber the members, a pfex process core file might be created. [PR954351](#)
- On EX4300 Virtual Chassis, when line-rate Layer 2 multicast traffic is sent on 10-gigabit uplink modules, a Virtual Chassis split might occur. [PR969005](#)
- On EX4300 Virtual Chassis, after a linecard member splits from the Virtual Chassis, the Virtual Chassis ports (VCPs) on other Virtual Chassis members that connected them to the split member might not go down. [PR977199](#)
- On EX4300 Virtual Chassis, traffic might be duplicated for 5–6 seconds after you renumber a Virtual Chassis member. [PR978115](#)
- On EX4300 Virtual Chassis, Virtual Chassis functionality might be broken after you reboot the master—for example, members might become disconnected, resulting in traffic duplication, or members might be split from the Virtual Chassis even though you had not set them to split. [PR979295](#)
- On an EX4600 switch in a Virtual Chassis, if lossless traffic is passing over a 10-Gbps SFP+ link configured as a Virtual Chassis port (VCP), traffic on the link might be dropped when the link is congested. As a workaround, configure 40-Gbps QSFP+ links into VCPs whenever possible if you are creating a Virtual Chassis that is supporting lossless traffic. This issue is seen on 10-Gbps SFP+ links only. [PR1006974](#)
- On EX4300 Virtual Chassis, the output from the **show chassis hardware** command might not display 40-gigabit transceiver information. [PR1012649](#)

Virtual Chassis Fabric

- In a mixed-mode Virtual Chassis Fabric (VCF) that contains EX4300 switches as leaf nodes, egress port mirroring does not work on the EX4300 switch interfaces. [PR987829](#)
- On a mixed-mode Virtual Chassis Fabric (VCF), if you reboot the VCF, Layer 2 unknown traffic might not be load balanced proportionally across all the links in the LAG bundle. [PR987841](#)
- On a Virtual Chassis Fabric (VCF), if you enable OSPF and a member reboots, the routing protocol might transition to the *init* state on all members and then return to the *full* state when the rebooting member rejoins the VCF. [PR987828](#)
- In a mixed-mode Virtual Chassis Fabric, the PPS field in the **show interfaces** command output for aggregated Ethernet subinterfaces might show double the packet count, and the **monitor interfaces** command output for AE subinterfaces might not show packet statistics. [PR1001539](#)

Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 6](#)
- [Known Behavior on page 7](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions on page 25](#)
- [Product Compatibility on page 26](#)

Resolved Issues

This section lists the issues fixed in Junos OS Release 13.2X51 for EX Series.

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

- [Issues Resolved in Junos OS Release 13.2X51-D25 on page 16](#)
- [Issues Resolved in Junos OS Release 13.2X51-D21 on page 17](#)
- [Issues Resolved in Release 13.2X51-D20 on page 17](#)
- [Issues Resolved in Release 13.2X51-D15 on page 20](#)
- [Issues Resolved in Release 13.2X50-D15 on page 24](#)

Issues Resolved in Junos OS Release 13.2X51-D25

The following issues have been resolved since Junos OS Release 13.2X51-D21. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- When using 802.1X multiple supplicant mode on an EX4300 switch, client authentication requests might be dropped when a client is already logged into the authentication server. The authentication requests are dropped until the client that is logged into the server is logged off of the server. [PR987115](#)

High Availability

- On EX4300 switches, VRRP on an IRB logical interface stops working if another IRB logical interface's VRRP transitions from backup to master to backup. [PR933735](#)

Interfaces and Chassis

- On EX4300 switches, setting the inet MTU on a VLAN-tagged aggregated Ethernet interface might cause routing of frames that are larger than the inet MTU. [PR910933](#)
- On EX4300 switches, Power over Ethernet (PoE) might stop working on all ports if the PoE port from the EX4300 switch is connected to another PoE-enabled port on an EX4200 switch through an RJ-45/straight cable. The problem is due to the large voltage difference between the two power supplies, such as between that of the EX4300 (PoE+, which is 57V) and that of the EX4200 (standard PoE, which is 44V), which in

turn generates significant negative current and turns ports off to protect the power sourcing equipment (PSE). As a workaround, use crossover cables and disable PoE on the EX4200. [PR976551](#)

- On an EX4300 switch, aggregated Ethernet interfaces do not display statistics for logical interfaces. [PR984998](#)
- On EX4300 fiber-based switches, if you remove an SFP-T transceiver from a 4X10-Gigabit Ethernet port, the LED lights might continue to glow. [PR987007](#)
- On EX4300 Virtual Chassis, the LACP Rx counter in the **show lacp statistics interfaces** operational mode command does not increment when it receives LACP hello packets. [PR988068](#)

Layer 2 Protocols

- On EX4300 Virtual Chassis, an IS-IS adjacency does not come up over aggregated Ethernet interfaces. [PR988234](#)

Virtual Chassis

- On EX4300 Virtual Chassis, if you issue the **show virtual-chassis vc-port statistics extensive** command, you might see **Undersized packets** and **Runts** error counts incrementing slowly on 40-gigabit Virtual Chassis ports (VCPs). [PR952196](#)

Issues Resolved in Junos OS Release 13.2X51-D21

The following issues have been resolved since Junos OS Release 13.2X51-D20. The identifier following the description is the tracking number in our bug database.

Interfaces

- On EX Series and QFX Series switches, if you change an interface description, the new description might not appear in the output of the **show lldp neighbors** operational mode command. [PR913792](#)

Issues Resolved in Release 13.2X51-D20

The following issues have been resolved since Junos OS Release 13.2X51-D15. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On an EX4300 Virtual Chassis, if a large number of clients are 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, 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](#)

- On EX4300 switches, if you restart the firewall process, dynamic filter counters might be created for all authenticated hosts even though only one host has dynamic filter configurations. [PR955305](#)
- On EX4300 switches, in an 802.1X configuration with multiple-supplicant mode, if you clear the Ethernet-switching table, traffic might not be forwarded on a dynamic VLAN. [PR959323](#)

Infrastructure

- On EX4300 switches, the `jdhcpd` process might create a core file if you remove a DHCP server or DHCP relay configuration. [PR961684](#)
- On EX4300 switches, when you boot the switch, the `check_configured_tpids: ge-X/X/X: number of configured tpids exceeds the limit(4)` system log message might be displayed. No functionality is affected. [PR966061](#)
- On EX4300 switches, if you issue the `request system power-off` command, an `Unrecognized command` error message appears. [PR968269](#)
- If an EX4300 Virtual Chassis is zeroized and rebooted continuously, the default Virtual Chassis ports (VCPs) might not be created after boot up. As a workaround, delete the VCPs by issuing the `request virtual-chassis vc-port delete pic-slot pic-slot port port` operational-mode CLI command, and then reset them by issuing the `request virtual-chassis vc-port set pic-slot pic-slot port port` command. See the PR for the detailed workaround. [PR975234](#)
- On EX4300 switches, when the routing instance/VLAN/IBFD is deleted on which the client is present, the lease timer is not cleared properly and a JDHCPD core file might be created. [PR1012556](#)

Interfaces and Chassis

- On EX4300 switches, in an RTG, if the member access interfaces are converted to trunk interfaces and are then converted back to access interfaces, the interfaces might lose their association with the VLAN. [PR951336](#)
- On EX4300 switches, the `show interfaces interface-name media` command shows the speed as 1000 Mbps instead of 100 Mbps for SFP-FX interfaces. [PR967119](#)

J-Web

- On EX4300 switches, when you run EZsetup from the J-Web interface, the commit configuration might fail with a timeout error the first time you try to commit the configuration. As a workaround, disconnect the laptop from the switch and then reconnect it, and then use the EZsetup Wizard again. [PR858819](#)

Layer 2 Protocols

- 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](#)

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](#)

Port Security

- 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, policers applied to an egress VLAN-based firewall filter do not work. [PR912027](#)
- On EX4300 switches, in an egress port-based firewall filter, the match condition **learn-vlan-id** might not work. [PR912191](#)
- On EX4300 switches, filter-based forwarding does not work for routes dynamically inserted through routing protocols. [PR913558](#)
- On EX4300 switches, in an egress VLAN-based firewall filter, the IPv4 match condition **interface** might not work. [PR918271](#)
- On EX4300 Virtual Chassis, if an 802.1X client is authenticated and the server fails, that client might not be re-authenticated, even if the server fail fallback action is configured as **use-cache**, and that client will go into the **Held** state, as shown in the output for the **show dot1x interface** command. [PR952144](#)

Spanning-Tree Protocols

- On EX4300 switches, STP BPDUs might be dropped on trunk interfaces. [PR978646](#)

Virtual Chassis

- On EX4300 Virtual Chassis, if you change the topology from a ring topology to a linear topology and then reboot the Virtual Chassis, one of the members might fail to join the Virtual Chassis. [PR953677](#)
- On EX4300 Virtual Chassis, after a reboot, 10-gigabit Virtual Chassis ports (VCPs) might be shown as **Absent** in **show virtual-chassis vc-port** command output. [PR959732](#)

Issues Resolved in Release 13.2X51-D15

The following issues have been resolved since Junos OS Release 13.2X50-D15. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On EX4300 switches, the output for the **show lldp neighbors interface interface-name** command displays type-length-value (TLV) information in alphanumeric codes, which is difficult to understand. [PR882143](#)
- On EX4300 switches with a voice VLAN configured, if there is a configuration change that triggers the Link Layer Discovery Protocol (LLDP) to parse the configuration (for example, an interface description) or an interface addition or deletion, a memory leak might result in the Layer 2 Control Protocol daemon (l2cpd). The l2cpd process creates a core file. [PR948718](#)

Bridging and Learning

- On EX4300 switches, in an MVRP scenario, if you exchange the VLAN IDs of two different VLANs or configure a VLAN with a VLAN ID that you previously configured and deleted on a different VLAN, the Layer 2 learning daemon (l2ald) might create a core file. [PR885891](#)
- On EX4300 Virtual Chassis, configuring static ARP on an IRB interface might result in an error at commit if the l2-interface is configured with all members in trunk mode. [PR915932](#)
- On EX4300 switches, configuring the switch using the **ezsetup** script creates vlan.0 instead of irb.0. [PR934439](#)
- On EX4300 switches, when a timer counter wraps around from its maximum count back to zero, any MAC entries that are inactive for a duration of more than 1 second might be deleted and relearned. [PR954625](#)

Hardware

- On EX4300-24T switches, you might see the following message periodically even though the temperature is within threshold limits: **chassisd[1022]: CHASSISD_SNMP_TRAP6: SNMP trap generated: Over Temperature!**
(jnxContentsContainerIndex 9, jnxContentsL1Index 1, jnxContentsL2Index 0,

jnxContentsL3Index 0, jnxContentsDescr Routing Engine 0, jnxOperatingState/Temp 53) . [PR948901](#)

- On EX4300 switches, you might see this message displayed continuously: **fpc 0: power budget received, old = 490 new = 490**. The message should only be displayed if there is a mismatch between the old power and the new power, but if the mismatch check is missing, the message is displayed continuously. The message is harmless; it has no impact. [PR957480](#)

High Availability

- RFC5798 (*Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*) states that a VRRP backup router should not send router advertisement (RA) messages. EX4300 switches send a copy of an RA message under the following conditions: 1). When VRRP is initially activated. 2). When the master switches over to backup, the old master keeps sending RA messages. [PR833436](#)

Infrastructure and Chassis

- On EX4300 switches, if you configure a DHCP local server using the **dhcp-local-server** configuration statement and commit the configuration, the CLI displays **dfw_check_filter(): /kernel: 1 bad_data_offset errors**. [PR827762](#)
- On EX4300 switches, for DHCP relay, when clients connected to an integrated routing and bridging (IRB) interface send a rebind, the relay bindings remain in the rebinding state. [PR828147](#)
- On EX4300 switches, entering the **request system zeroize media** command does not remove all configuration information on the Routing Engines and reset all key values as expected. [PR834158](#)
- In a DHCP scenario that includes an integrated routing and bridging (IRB) interface on an EX4300 switch, if any event causes VLAN membership deletion and addition in the VLAN corresponding to the IRB interface (for example, if the VLAN membership is deleted and then added back manually), the DHCP service might stop working. As a workaround, restart the **jdhcpd** process. [PR868140](#)
- On EX4300 switches, the password database might get corrupted after you issue the **request system zeroize** command or upgrade the software, thereby preventing you from logging into the switch. [PR872067](#)
- On EX4300 switches with fiber ports, after you reboot the switch, a warning message might appear: **Interrupt storm detected on "irq41:" throttling interrupt source**. [PR896126](#)
- On EX4300 switches, Packet Forwarding Engine log messages might contain typographical errors. [PR922445](#)
- If an EX4300 switch is configured for DHCP relay, if an IRB interface walks through a Layer 2 trunk interface and the corresponding DHCP relay is configured in a routing instance, if you deactivate or activate (or delete or add) a hierarchy that contains a DHCP relay-related configuration, DHCP relay might not work as expected. [PR935155](#)

- On EX4300 switches with Bidirectional Forwarding Detection (BFD) enabled, memory might be leaked in the Packet Forwarding Engine manager (pfex) process due to high numbers of BFD flaps, and a pfex process core file might be generated. [PR951637](#)
- On EX4300 switches, generating a core file using the gcore process and then executing the `ls -l` command in the shell might create a kernel core file. [PR955067](#)
- On network devices, such as EX Series switches, the kernel memory type session might leak memory due to a memory-free issue in the system. After memory consumption exceeds the limit, the device does not forward traffic. This issue is platform-independent. [PR944385](#)
- On EX Series switches, if the Network Time Protocol (NTP) server is not a stratum 1 server, the NTP synchronization process cannot be completed. You can confirm that the process was not completed by looking at the output of the `show ntp status` command. [PR944510](#)
- On an EX4300 switch, if an IPv6 address has nonzero values starting from the fifth octet in the prefix, then the route might not be programmed correctly in hardware, resulting in traffic loss. For example, you might see the issue for the address 5001:1:1::1/64 but not for 5001::1/64. [PR948569](#)
- On EX4300 Virtual Chassis, you might see a memory utilization increase on the Routing Engine due to continuous MIB polling, as shown in the `show chassis routing-engine` command output. [PR951639](#)

Interfaces

- On EX4300 switches, changing the interface type configuration might cause the device configuration daemon (dcd), which is responsible for configuring interface types, to generate a core file. When this issue occurs, the interface will flap and service might be interrupted. [PR917311](#)
- On EX4300 switches with Power over Ethernet (PoE), when some ports are connected to a PoE device and others are used as network ports, if the PoE firmware version that you see when you issue the `show chassis firmware detail` CLI command is greater than 2.3.9 (for example, 2.4.8), PoE might stop working on some of the ports. [PR941205](#)
- On EX4300 switches, if you configure an interface with a speed of 100 Mbps and then set the interface to **disable**, the interface might not actually be disabled, and traffic from peer devices might be lost. [PR943779](#)
- On EX4300 Virtual Chassis that have PoE connections to another vendor's pre-standard PoE phones, the phone power might drop off. [PR956176](#)
- On EX4300 switches, response packets for ARP and ICMP might not be sent out on the egress interface after you enable/disable that interface. [PR956638](#)
- On EX4300 switches, the file `/var/log wttmp` is not rotated once a month or every 10 MB. [PR964118](#)

Multicast Protocols

- On EX4300 switches, when you configure multicast and there is a high rate of PIM or IGMP join and leave messages, the control plane and data plane might have inconsistent entries and some multicast traffic might not be forwarded properly. [PR894175](#)
- On EX4300 switches, if the multicast cache is flushed out repeatedly using the **clear multicast forwarding-cache** command, a routing protocol daemon (rpd) core file might be created. [PR894522](#)
- On EX4300 switches, in a scaled multicast scenario with more than 200 multicast streams, changes in multicast routes might not be completed, which can result in multicast traffic loss. [PR924167](#)
- When IGMP snooping is enabled on EX4300 switches, IPv6 neighbor discovery packets might be dropped. [PR957108](#)

Network Management and Monitoring

- When an EX4300 switch receives packets addressed to a missing destination route, those packets might clog the queue to the Routing Engine, causing other packets (such as SSH and telnet packets) to be dropped. [PR942114](#)
- On EX4300 switches, when the SNMP Management Information Base II daemon (mib2d) polls system statistics from the kernel, a memory leak might occur in system buffers (mbuf), which might cause packets (such as ARP packets) to drop. [PR953664](#)

Routing Policy and Firewall Filters

- EX4300 switches might not create ternary content addressable memory (TCAM) entries for firewall filters that are configured with both a policer and forwarding-class/loss-priority actions in a filter term. [PR939777](#)
- On EX4300 switches, when you configure a firewall filter with an action set to discard and log and apply it to the physical loopback interface, lo0, to protect the Routing Engine from a denial-of-service (DoS) attack, the log action might trap DoS packets to the Routing Engine. A large DoS attack might clog the queue and cause some packets (such as SSH packets) to be dropped from the queue. [PR956807](#)

Routing Protocols

- On EX4300 Virtual Chassis, IPv6 OSPFv3 sessions might get stuck in the init state for all other IRB interfaces on the aggregation switch when one of the LAGs between the top of rack switch and the aggregation switch is disabled. [PR953989](#)

Virtual Chassis

- On EX4300 switches, when a 40-gigabit Virtual Chassis port (VCP) is converted to a network port, the port might not come up, and it cannot be converted to a VCP. [PR906948](#)
- On EX4300 Virtual Chassis with link aggregation groups (LAGs) configured over Virtual Chassis ports (VCPs), if one member link of a LAG is located on the master Routing

Engine, traffic loss might occur for approximately 60 seconds during reboot of the master Routing Engine. [PR907300](#)

- On EX4300 Virtual Chassis, when two different Layer 3 subinterfaces have the same port numbers and are assigned the same VLAN ID, neighbors might become unreachable over the uplink module. [PR922445](#)
- On EX4300 Virtual Chassis, when a LAG is configured with the master switch as one of the LAG members, if the master is rebooted, traffic through the LAG interface might be transmitted through the link that resides on the master as it is going down, resulting in traffic loss. [PR928905](#)
- EX4300 Virtual Chassis failover might cause approximately 30 seconds of packet loss. [PR939154](#)
- On EX4300 Virtual Chassis, a graceful Routing Engine switchover (GRES) might reset the Power over Ethernet (PoE) interfaces of the previous master switch. [PR958085](#)

Issues Resolved in Release 13.2X50-D15

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.

Infrastructure

- 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](#)

Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 6](#)
- [Known Behavior on page 7](#)
- [Known Issues on page 10](#)
- [Documentation Updates on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions on page 25](#)
- [Product Compatibility on page 26](#)

Documentation Updates

There are no errata or changes in Junos OS Release 13.2X51-D25 documentation.

Related Documentation

- *New and Changed Features*
- *Changes in Behavior and Syntax*
- *Known Behavior*
- *Known Issues*
- *Resolved Issues*
- *Migration, Upgrade, and Downgrade Instructions*

- *Product Compatibility*

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 on page 25](#)
- [Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations on page 25](#)
- [Upgrading to a Controlled Version of Junos OS on page 26](#)

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 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 \(CLI Procedure\)](#).

Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 6](#)
- [Known Behavior on page 7](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 24](#)
- [Product Compatibility on page 26](#)

Product Compatibility

- [Hardware Compatibility on page 27](#)

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 6](#)
- [Known Behavior on page 7](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions on page 25](#)

Junos OS Release Notes for the QFX Series

These release notes accompany Junos OS Release 13.2X51-D25 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 28](#)
- [Changes in Behavior and Syntax on page 30](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Resolved Issues on page 39](#)
- [Documentation Updates on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Product Compatibility on page 46](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 13.2X51-D25 for the QFX Series. To view the entire set of software information in PDF format, see the [Complete Software Guide for Junos OS for the QFX Series](#).

- [Hardware on page 28](#)
- [Ethernet Switching on page 29](#)
- [Software Installation and Upgrade on page 30](#)
- [Virtual Chassis on page 30](#)

Hardware

- **QFX5100-48T Switch**—The Juniper Networks QFX5100 line of switches is the next generation of top-of-rack standalone switches. With the availability of the Junos OS 13.2X51-D25 software release, the model QFX5100-48T provides 48 copper access ports and 6 QSFP+ ports. Each 40-Gigabit QSFP+ port can operate either as a native 40-Gigabit port or be channelized into 4 independent 10 Gigabit ports. The 6 QSFP+ ports can be used as either access ports or as uplinks. The QFX5100-48T provides full duplex throughput of 720 Gbps. The QFX5100-48T has a 1 U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

The QFX5100-48T switch can be used as:

- A standalone switch.
- A Node device in a QFabric system.

The QFX5100-48T is supported on both the QFX3000-G and QFX3000-M QFabric systems.

- A master, backup, or linecard in a QFX Virtual Chassis.

A QFX Series Virtual Chassis allows you to interconnect up to ten QFX5100, QFX3500, or QFX3600, switches into one logical device and manage the device as a single chassis using a ring topology.

- A spine device or a leaf device in a Virtual Chassis Fabric (VCF).

VCF uses Virtual Chassis technology to interconnect multiple devices into a single logical device and manage that device as a single logical device inside of a fabric architecture. VCF architecture supports up to twenty total devices in a spine and leaf topology.

- **Support for EX4600-EM-8S Expansion Module (QFX5100-24Q switch)**—Provides eight additional 10-Gigabit Ethernet Small Form-Factor Pluggable (SFP+) ports. With the availability of the 13.2X51-D25 release, the EX4600-EM-8S expansion module can coexist with the QFX-EM-4Q expansion module in the same QFX5100-24Q. Expansion modules can be hot-inserted or hot-removed.

Ethernet Switching

- **CoS Explicit congestion notification (ECN)**—Enables end-to-end congestion notification between two endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

[See [Understanding CoS Explicit Congestion Notification](#).]

- **Q-in-Q tunneling enhancement (QFX5100)**—With previous releases of Junos OS 13.2X51, you cannot create a regular VLAN on an interface if you have created an S-VLAN or C-VLAN on that interface for Q-in-Q tunneling. This means that you cannot create an integrated routing and bridging (IRB) interface on that interface because regular VLANs are a required part of IRB configuration. With Junos OS 13.2X51-D25, you can create a regular VLAN on a trunk interface that has an S-VLAN, which means that you can also create an IRB interface on the trunk. You still cannot create a regular VLAN on an access interface that has a C-VLAN. [See [Understanding Q-in-Q Tunneling](#).]



NOTE: A regular VLAN and S-VLAN on the same trunk interface cannot share the same VLAN ID.

Software Installation and Upgrade

- **Junos OS Release 13.2X51-D25 image (QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches)**—Junos OS Release 13.2X51-D25 is a superset of Release 13.2X51-D20 and is provided for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches. Release 13.2X51-D25 is the recommended release for those switches.

Virtual Chassis

- **Virtual Chassis support (QFX5100 switches)**—QFX5100 switches can now be interconnected to form a Virtual 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 switches can be managed as a single switch, increased fault tolerance and high availability 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 simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP). For all models of the QFX5100 switch, you can configure up to ten switches into a QFX Virtual Chassis. QFX5100 switches, can form a mixed Virtual Chassis that includes QFX3600, QFX3500, and EX4300 switches. [See [Understanding QFX Series Virtual Chassis](#).]

Related Documentation

- [Changes in Behavior and Syntax on page 30](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Resolved Issues on page 39](#)
- [Documentation Updates on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Product Compatibility on page 46](#)

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 13.2X51-D21 or later for the QFX Series.

- [Interfaces](#)
- [IPv6](#)
- [Software Upgrade](#)
- [System Management](#)

Interfaces

- **Increased number of LAG members and groups (QFX3500, QFX3600, and QFX5100 switches)**—Provides support for a total of 64 members per link aggregation group

(LAG) and up to 448 total LAGs per switch. The previous limit was 32 members per LAG and 63 total LAGs. [See [Understanding Aggregated Ethernet Interfaces and LACP](#).]

- **Link aggregation group (LAG) bundle and Equal-Cost Multipath (ECMP) next-hop hashing configuration (QFX5100 switches)**—You can now configure the fields that the hashing algorithm uses to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled. For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device. Configuring the fields used by the hashing algorithm helps users manage traffic flows when a switch is using LAG bundles or ECMP, and is especially helpful in scenarios when most of the traffic is similar. You configure the hashing algorithm at the **[edit forwarding-options enhanced-hash-key]** hierarchy level. [See [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic](#).]

IPv6

- **IPv6 network discovery retry limit (QFX5100 switches)**—Determines the number of times the IPv6 neighbor discovery mechanism will attempt to look for adjacent IPv6 neighbors. You can enable the network discovery retry limit by including the **nd-maxucast-retry** statement at the **[edit system]** hierarchy level. The default value is 3, and the recommended range of retries is 1 to 100.

Software Upgrade

- **Change for the request system software rollback command (QFX5100 switches)**—The **reboot** option has been removed from the **request system software rollback** operational mode command. To reboot the switch after a software rollback, issue the **request system reboot** command as a separate, secondary command.

System Management

- **Zero Touch Provisioning (ZTP) script support (QFX5100 switches using the Enhanced Automation software image)**—Enables you to execute scripts during the ZTP process. ZTP downloads files from the DHCP server based on the information contained in the option `NEW_OP.config-file-name` DHCP vendor option in the `dhcpd.conf` file. ZTP determines that the file is a script instead of a configuration file based on the first line that is included in the file. If the first line of the file consists of the `#!` characters followed by an interpreter path, then ZTP considers the file as a script and executes the script with the mentioned interpreter—for example, `#!/usr/bin/slax`. For the script to run successfully, make sure the script specifies how to fetch and load a valid configuration file.

ZTP supports shell, SLAX, and Python scripts. For SLAX script support, the configuration file must include the `.slax` extension in the configuration file. To execute SLAX scripts, include `?#!/usr/libexec/ui/cscript?` in the first line of the script.

- **BOOTP support (QFX5100 and EX4300 switches)**—Enables BOOTP support for DHCP servers and DHCP relay agents. To enable BOOTP for DHCP servers, include the **bootp-support** statement at the **[edit groups dhcp-server overrides]** hierarchy level. To

enable BOOTP support for DHCP relay agents, include the **bootp-support** statement at the **[edit groups bootp forwarding-options dhcp-relay overrides]** hierarchy level.

- **Update to operational command output (QFX5100 switches using the Enhanced Automation software image)**—When you install the Junos OS automation enhancement software on a QFX5100 switch, the output of the **show version** command displays the phrase **Junos for Automation Enhancement** at the bottom of the output to indicate which software image is running.

Related Documentation

- [New and Changed Features on page 28](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Resolved Issues on page 39](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Documentation Updates on page 42](#)
- [Product Compatibility on page 46](#)

Known Behavior

This section lists the limitations in Junos OS Releases 13.2X51 for the QFX Series.

Multiprotocol Label Switching (MPLS)

- On a QFX5100 switch acting as an MPLS penultimate hop popping (PHP) router, an MPLS label route with a bottom of the stack (BOS) next hop will consume one filter entry and one extra next hop entry. The total number of filter entries used will be the number of next hops in the system with PHP and BOS operation.

Network Management and Monitoring

- If a QFX5100 switch drops traffic because of an ingress firewall filter, the switch does not generate an sFlow technology monitoring flow sample packet that contains this dropped traffic.
- On the QFX5100 switch, the J-Web interface is not supported. As a result, the **web-management** configuration statement in the **[edit system services]** hierarchy level is not available in the CLI.

Platform and Infrastructure

- On a QFX5100 switch, when the system mode is in default-mode, the latency for 10-Gigabit Ethernet ports in cut-through mode is the same as it is for store and forward mode. This issue does not apply to 40-Gigabit Ethernet ports.

Traffic Management

- The packet replication engine on QFX5100 switches replicates flood traffic at line rate. (Flood traffic is unicast traffic for which the switch has not learned the destination

address and other traffic such as broadcast or multicast traffic that must be replicated on all of the member interfaces of a VLAN). The engine can replicate two packets per clock cycle.

The number of packets replicated per second depends on the clock speed of the replication engine. A faster clock speed means a faster clock cycle. A faster clock cycle means more packets replicated per second. If the number of flood packets to replicate exceeds the engine capacity, flood packets are dropped.

All of the switch interfaces share the replication engine. When the number of flood packets to replicate exceeds the engine capacity, all excess unlearned unicast and multdestination traffic is dropped on all VLANs, not just on the VLAN that is oversubscribing the replication engine.

Replication engine capacity on QFX5100 platforms differs because the clock speed of the replication engine can differ. [Table 2 on page 33](#) shows the clock speeds and the maximum number of packets the engine can replicate per second on each QFX5100 platform.

Table 2: QFX5100 Platform Clock Speeds and Replication Engine Capacity

QFX5100 Switch Platform	Clock Speed	Maximum Number of Packets Replicated Per Second
QFX5100-48S	635 MHz	1.27 Gigapackets
QFX5100-40T	635 MHz	1.27 Gigapackets
QFX5100-96S	793 MHz	1.58 Gigapackets
QFX5100-24Q	793 MHz	1.58 Gigapackets

Packet size affects how many packets must be replicated. The absolute number of packets-per-second to replicate decreases as packet size increases (the larger the packet size, the fewer the number of packets-per-second to process).

The calculation to determine the number of packets being flooded in a VLAN is:

$$(\text{interface-speed} / \text{packet-size}) * \text{number-of-packet-replications} * \text{number-of-ports-flooding-traffic} = \text{number-of-packets-per-second}$$

For example, if a VLAN has ten 10-Gigabit Ethernet member interfaces, then every flood packet must be replicated nine times (the packet is not flooded out of the ingress interface). If a flood of 64-byte packets occurs on the VLAN, and every interface is flooding traffic to the other interfaces in the VLAN, the calculation to determine the number of packets the engine must replicate per second at line rate is:

$$(10\text{-Gbps} / (64 * 8)) * 9 * 10 = 1.758 \text{ Gigapackets}$$

In this example, *10-Gbps* is the interface speed, *(64 * 8)* is the packet size in bits (64 bytes * 8 bits), *9* is the number of times the engine must replicate each packet, and *10* is the number of ports flooding at line rate. In this case, the requirement to replicate

1.758 Gigapackets-per-second exceeds the capacity of the replication engine, so some flood traffic is dropped.

For another example, if all interfaces in the same VLAN flood 128-byte packets instead of 64-byte packets:

$$(10\text{-Gbps} / (128 * 8)) * 9 * 10 = 0.879 \text{ Gigapackets}$$

In this case, the engine can replicate all of the packets at line rate because the total number of packets-per-second is well within the replication engine capacity.

- On a QFX5100 switch, 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 relevant 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.

Related Documentation

- [New and Changed Features on page 28](#)
- [Changes in Behavior and Syntax on page 30](#)
- [Known Issues on page 35](#)
- [Resolved Issues on page 39](#)
- [Documentation Updates on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Product Compatibility on page 46](#)

Known Issues

The following issues are outstanding in Junos OS Release 13.2X51-D25. The identifier following the description is the tracking number in our bug database.

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>.

- [Class of Service \(CoS\)](#)
- [High Availability \(HA\) and Resiliency](#)
- [Interfaces and Chassis](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)
- [Software Installation and Upgrade](#)
- [Storage and Fibre Channel](#)

Class of Service (CoS)

- On a QFX5100 Virtual Chassis Fabric (VCF), if you configure a shared buffer with a low value (such as 10%), the VCF traffic might be affected and cause a split in the Virtual Chassis Fabric. [PR947585](#)
- On a QFX5100 Virtual Chassis Fabric (VCF), if you configure an ECN enabled queue with a WRED maximum threshold and the traffic exceeds the threshold, the system will continue marking traffic with 100 percent probability until the queue reaches its maximum allowed size rather than dropping all packets. [PR968718](#)
- On a Virtual Chassis Fabric (VCF), the default behavior for Routing Engine generated host traffic has changed slightly:
 - All Routing Engine generated unicast host-traffic behavior remains the same. All packets go out of UC-NC 7 (network-control) by default for both the master and remote Packet Forwarding Engines and the user can change the unicast queue number either through host-outbound or finer-grain classification CLI.
 - Multicast host traffic will be treated like unicast traffic on the master PFE and exit by way of the unicast queue (7 by default).
 - For the remote Packet Forwarding Engine, these packets will exit by way of queue 11. As a result, host-outbound and finer-grain CLI configuration will not take effect for multicast host traffic on the remote PFE. However, there is some bandwidth reserved on queue 11 by default so these packets are treated fairly. [PR977166](#)

High Availability (HA) and Resiliency

- On a QFX5100 switch, if you have a large scale IPv6 multicast topology and perform an in-service software upgrade, some multicast groups might lose traffic. [PR987261](#)

Interfaces and Chassis

- On a QFX5100 Virtual Chassis, if you mistakenly configure a member in fabric mode, this member might not be marked as inactive as expected. [PR927517](#)

Platform and Infrastructure

- On a QFX5100 Virtual Chassis Fabric (VCF), you must configure all members with consistent modes (such as mixed or non-mixed, and fabric or non-fabric). If their modes are not consistent, only one mode will become eligible and the rest become ineligible and inactive in most cases. However, if a new member configured for fabric mixed mode tries to join an existing fabric non-mixed mode VCF, the new mixed-mode member preempts the existing VCF, becomes the master, and causes all the other former VCF members to become inactive linecards. [PR933074](#)
- On QFX Series switches, if you connect a 40G DAC cable between the switch and an EX4300 switch, the link might not become active. As a workaround, disable autonegotiation on the EX4300 and retry the connection. [PR935197](#)
- On a QFX5100 switch or Virtual Chassis Fabric, if you do not provide user input at the prompt during a password recovery procedure, the system might reboot automatically. [PR951370](#)
- On a mixed-mode Virtual Chassis Fabric (VCF), if you enable graceful Routing Engine switchover and make any changes to system mode, slot ID, forwarding mode, VC formation, or the Packet Forwarding Engine restarts, VCP ports might be reset. [PR959867](#)
- On a mixed-mode Virtual Chassis Fabric, if you issue the command to perform a graceful Routing Engine switchover, routing protocols running in distributed mode on spine member links might transition down and up. [PR964974](#)
- On a mixed-mode Virtual Chassis Fabric, during a Routing Engine switchover, you might see up to 1 second of traffic loss and a few seconds of duplicated multicast traffic. [PR964987](#)
- On a QFX5100 switch configured in cut-through mode, if the MTU on a port is less than the size of the traffic, the port might transition down and up continuously. [PR966052](#)
- On a mixed-mode Virtual Chassis Fabric (VCF), if a GRE tunnel interface (gr-) takes a long time to respond (more than 10-15 seconds), an SNMP walk operation for the interface MIB (IF-MIB) might time out. As a workaround, increase the timeout value for the SNMP walk to more than 15 seconds. [PR966983](#)
- On a QFX5100 switch, if you configure firewall-based port mirroring on the switch and delete or modify this configuration, the output interface might remain programmed in the Packet Forwarding Engine. This can cause traffic to be incorrectly forwarded to the output interface. As a workaround, perform the actions separately (for example,

delete and perform the commit operation, then modify the configuration and perform a second commit operation.).[PR968471](#)

- On a mixed-mode Virtual Chassis Fabric (VCF), OSPF packets are not mirrored by the native analyzer when the output port belongs to another member. [PR969542](#)
- On a QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF), the Virtual Chassis MIB does not include statistics for VCP ports. As a workaround, issue the **show virtual-chassis vc-port statistics extensive** CLI command. [PR972726](#)
- On a QFX5100 Virtual Chassis or Virtual Chassis Fabric, when a packet traverses the fabric, and the ingress and egress ports are on different members, cut-through mode switching latency might become higher than store-and-forward switching latency. [PR977080](#)
- On a QFX5100 Virtual Chassis Fabric (VCF), if you issue the **request support information** command from a console port login session, the device functioning in the master Routing Engine role might become inactive. As a workaround, enter the **request support information** command from a login session that does not use the console port, such as a telnet or SSH session to a management port. [PR978385](#)
- On a QFX Series standalone switch, if you enable FCoE Initialization Protocol (FIP) snooping and perform an in-service software upgrade (ISSU) from Junos OS Release 13.2X51-D15 to 13.2X51-D20, Fibre Channel over Ethernet (FCoE) traffic might be dropped for up to four seconds. [PR981306](#)
- On a QFX5100 switch, if you configure IS-IS and perform an in-service software upgrade (ISSU), there might be approximately 2 seconds of IPv4/IPv6 traffic loss during the em0 handoff. [PR985462](#)
- On a mixed-mode Virtual Chassis Fabric (VCF), if you configure the LACP periodic interval to fast, LACP might transition down and up after a Routing Engine switchover. As a workaround, set the LACP periodic interval to slow. [PR985915](#)
- On a mixed-mode Virtual Chassis Fabric (VCF), if you configure a large number of VLANs and equal-cost multipath (ECMP) or unicast reverse-path-forwarding (uRPF) on an integrated routing and bridging (IRB) interface, the Protocol Independent Multicast (PIM) join rate might be slow on the IRB interface. [PR987302](#)
- On a QFX5100 Virtual Chassis Fabric (VCF), if you use autochannelization to connect a 10-Gigabit Ethernet port from a traffic generator to a channelized 40-Gigabit Ethernet interface on the switch, and then take the traffic generator port offline and bring it back online, the traffic generator port might be removed permanently from the VCF. As a workaround, configure the channelized ports manually on the switch by including the **10g** statement at the **[edit chassis fpc fpc-slot pic pic-slot port port-name channel-speed]** hierarchy level. [PR987321](#)
- In a mixed-mode Virtual Chassis Fabric (VCF) that contains EX4300 switches as leaf nodes, egress port mirroring does not work on the EX4300 switch interfaces. [PR987829](#)
- On a mixed-mode Virtual Chassis Fabric (VCF), if you reboot the VCF, Layer 2 unknown traffic might not be load balanced proportionally across all the links in the LAG bundle. [PR987841](#)

- On a Virtual Chassis Fabric (VCF), if you enable OSPF and a member reboots, the routing protocol might transition to the *init* state on all members and then return to the *full* state when the rebooting member rejoins the VCF. [PR987828](#)
- On a QFX5100 switch in a Virtual Chassis or a Virtual Chassis Fabric (VCF), if lossless traffic is passing over a 10-Gbps SFP+ link configured as a Virtual Chassis port (VCP), traffic on the link might be dropped when the link is congested. As a workaround, configure 40-Gbps QSFP+ links into VCPs whenever possible if you are creating a Virtual Chassis or a VCF that is supporting lossless traffic. This issue is seen on 10-Gbps SFP+ links only. [PR1006974](#)
- In a mixed-mode Virtual Chassis Fabric, the PPS field in the **show interfaces** command output for aggregated Ethernet subinterfaces might show double the packet count, and the **monitor interfaces** command output for AE subinterfaces might not show packet statistics. [PR1001539](#)

Routing Protocols

- On a QFX5100 switch or Virtual Chassis Fabric, the **targeted-broadcast forward-and-send-to-re** option does not work for aggregated Ethernet LAG interfaces. [PR956504](#)
- On a QFX5100 switch with Bidirectional Forwarding Detection (BFD) configured, if the timer of BFD is configured less than 3000ms, when a graceful Routing Engine switchover happens during an in-service software upgrade (ISSU), sometimes the routing protocol might flap, which cause traffic loss of up to 30 seconds. [PR970881](#)
- On QFX5100 switches, if you configure the **next-table** option within a static route configuration for a routing instance, the switch does not forward any packets to any route in the routing table. [PR970895](#)
- On a Virtual Chassis Fabric (VCF), if you power off the switch acting in the master Routing Engine role by using the **request system power-off** command, OSPF adjacencies might transition down and up. As a workaround, physically power off the switch in the master Routing Engine role rather than using the **request system power-off** command. [PR986785](#)

Software Installation and Upgrade

- NSSU to Junos OS Release 13.2X51-D25 is not supported for QFX5100, QFX3600, and QFX3500. [PR1010403](#)

Storage and Fibre Channel

- On a QFX5100 Virtual Chassis, in some cases, the firewall filter installation for a VN2VN FIP snooping session might not happen within 800 milliseconds after receiving the first packet of the transaction. In such cases, the converged network adapter (CNA) will need to re-attempt to establish the transaction. [PR974965](#)

Related Documentation

- [New and Changed Features on page 28](#)
- [Changes in Behavior and Syntax on page 30](#)

- [Known Behavior on page 32](#)
- [Resolved Issues on page 39](#)
- [Documentation Updates on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Product Compatibility on page 46](#)

Resolved Issues

This section lists the issues fixed in the Junos OS Release 13.2X51 for the QFX Series.

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

- [Issues Resolved in Junos OS Release 13.2X51-D25 on page 39](#)
- [Issues Resolved in Junos OS Release 13.2X51-D21 on page 39](#)
- [Issues Resolved in Junos OS Release 13.2X51-D20 on page 40](#)
- [Issues Resolved in Junos OS Release 13.2X51-D15 on page 40](#)

Issues Resolved in Junos OS Release 13.2X51-D25

Platform and Infrastructure

- On a QFX5100 switch, when you delete the OSPF configuration from an interface, the OSPF and ISIS routing protocols during an in-service software upgrade (ISSU) might transition down and up on all other configured interfaces. [PR933536](#)

Issues Resolved in Junos OS Release 13.2X51-D21

This section lists the issues fixed since Junos OS Release 13.2X51-D20 for the QFX Series.

Interfaces and Chassis

- On EX Series and QFX Series switches, if you change an interface description, the new description might not appear in the output of the **show lldp neighbors** operational mode command. [PR913792: This issue has been resolved.](#)

Platform and Infrastructure

- On a QFX3600 switch, if you insert a QSFP+-40G-LR4 module, it might not work even though it operates properly in a QFX3500 switch. [PR987817: This issue has been resolved.](#)
- On a QFX5100 Virtual Chassis Fabric, if you power off the master member and wait 10 minutes, the aggregated Ethernet Layer 3 interface MAC address does not change. As a workaround, change the value of the **mac-persistence-timer** statement at the **[edit virtual-chassis]** hierarchy level. [PR994236: This issue has been resolved.](#)

Routing Policy and Firewall Filters

- On a QFX5100 switch, if you configure a **reject** filter action, it might not work properly and some of the traffic might continue to be forwarded to the destination. [PR992452: This issue has been resolved.](#)

Issues Resolved in Junos OS Release 13.2X51-D20

This section lists the issues fixed since Junos OS Release 13.2X51-D15 for the QFX Series.

High Availability (HA) and Resiliency

- On a QFX5100 switch, when you perform an in-service software upgrade (ISSU), interfaces might go down and up if Link Aggregation Control Protocol (LACP) is configured in fast mode. As a workaround, configure LACP in slow mode, and disable the distributed periodic packet management process (ppmd). [PR965918: This issue has been resolved.](#)

Platform and Infrastructure

- On a QFX5100 switch, the minimum queue statistics polling interval is 10 milliseconds and the minimum traffic statistics polling interval is 2 seconds. The queue statistics polling interval differs from the configured polling interval by approximately 5 milliseconds. [PR911015: This issue has been resolved.](#)
- On a QFX5100 switch, after performing an in-service software upgrade, FIP snooping sessions might stop working. As a workaround, deactivate FIP snooping and then reactivate FIP snooping to restore FIP snooping functionality for all new sessions. [PR965727: This issue has been resolved.](#)

Routing Protocols

- On a QFX5100 switch, if issue the **show interfaces ae** command, the output for aggregated Ethernet logical interfaces might display a value of 0 for bits per second (bps) and packets per second (pps) for the Input and Output fields in the Bundle section. [PR936220: This issue has been resolved.](#)

Issues Resolved in Junos OS Release 13.2X51-D15

This section lists the issues fixed since Junos OS Release 13.2X51-D10 for the QFX Series.

Class of Service (CoS)

- On a QFX5100 switch, issuing the **show interfaces queue *interface-name*** and **show interfaces statistics *interface-name*** commands does not display the correct traffic rates. [PR894390](#)

Interfaces and Chassis

- On a QFX5100 switch, when you enable IGMP snooping on a VLAN, IPv6 multicast traffic is not flooded within the VLAN. [PR925141](#)
- On a QFX5100 switch, if you remove the service ID from a multichassis link aggregation group (MC-LAG) configuration, and then add it back to the configuration, single-homed ARP entries might not synchronize properly with MC-LAG peers. [PR929720](#)
- On a QFX5100 switch, do not use the unified forwarding table **lpm-profile** for IPv6 traffic. This profile does not work for IPv6 traffic. [PR929753](#)
- If you create a virtual routing instance on a QFX5100 switch and configure a routed VLAN interface (RVI) or integrated routing and bridging (IRB) interface under the routing instance, do not configure a multichassis link aggregation group (MC-LAG) interface to participate in the RVI or IRB. This combination is not supported with virtual routing instances. [PR934379](#)
- On a QFX5100 switch, integrated routing and bridging (IRB) MAC address synchronization is not supported, but you can use the Virtual Router Redundancy Protocol (VRRP) instead. As a workaround, configure VRRP on IRB interfaces that host multichassis LAG (MC-LAG) interfaces. [PR936512](#)
- On a QFX5100 switch, the **multichassis-lag-replicate-state** statement is not supported at the **[edit vlans]** CLI hierarchy level. As a workaround, enable the **multichassis-lag-replicate-state** statement globally. [PR937018](#)
- On a QFX5100 switch, you cannot issue interface range commands for channelized interfaces. As a workaround, use interface commands instead of interface range commands. [PR937788](#)
- On a QFX5100 switch, even when traffic is flowing normally, the output of the **show interfaces et-fpc/pic/port** and **show interfaces et-fpc/pic/port:[0-3]** commands does not display accurate bits per second (bps) information for the **Input rate** field. [PR939128](#)
- On a QFX5100 switch, autonegotiation of interfaces is disabled by default for 1-Gigabit Ethernet fiber ports. For these links to be brought online (up), you must disable autonegotiation on the peer interfaces. In addition, if you issue the **show interfaces *interface-name* extensive** command for an SFP access port with a 1-Gigabit optical copper transceiver installed, the output incorrectly shows the media type as **fiber**, if the port parameter is not configured in the **interfaces ge-0/0/port** statement in the **[edit]** hierarchy level. As a workaround, remove and reinsert the transceiver. [PR939439](#)
- On a QFX5100 switch, configuration of the **minimum-interval *milliseconds*** statement for liveness detection on a multichassis link aggregation group (MC-LAG) must be 1000 milliseconds or greater. Subsecond timers are not supported in Junos OS Release 13.2X51-D10. [PR942563](#)

Routing Protocols

- On a QFX5100 switch, BFD timer values of less than 1 second are not supported.
[PR942035](#)

Related Documentation

- [New and Changed Features on page 28](#)
- [Changes in Behavior and Syntax on page 30](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Documentation Updates on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Product Compatibility on page 46](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 13.2X51-D25 documentation.

System Management

- The **request app-engine** and **show app-engine** commands are not documented for the QFX5100 switch in Junos OS Release 13.2X51-D10.

Related Documentation

- [New and Changed Features on page 28](#)
- [Changes in Behavior and Syntax on page 30](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Resolved Issues on page 39](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)
- [Product Compatibility on page 46](#)

Migration, Upgrade, and Downgrade Instructions

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

- [Upgrading Software on QFX5100 Standalone Switches on page 43](#)
- [Performing an In-Service Software Upgrade \(ISSU\) on page 44](#)
- [Preparing the Switch for Software Installation on page 44](#)
- [Upgrading the Software Using ISSU on page 45](#)

Upgrading Software on QFX5100 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.

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 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
source/jinstall-qfx-5-13.2X51-D25-domestic-signed.tgz reboot
```

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

- **/pathname**—For a software package that is installed from a local directory on the switch.

- For software packages that are downloaded and installed from a remote location:
 - `ftp://hostname/pathname`
 - `http://hostname/pathname`
 - `scp://hostname/pathname` (available only for Canada and U.S. version)

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 13.2 jinstall package, you can issue the `request system software rollback` command to return to the previously installed software.

Performing an In-Service Software Upgrade (ISSU)

You can use an in-service software upgrade to upgrade the software running on the switch with minimal traffic disruption during the upgrade.



NOTE: ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 44](#)
- [Upgrading the Software Using ISSU on page 45](#)

Preparing the Switch for Software Installation

Before you begin software installation using ISSU:

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

To verify that nonstop active routing is enabled:



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

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

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

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

Upgrading the Software Using ISSU

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

To upgrade the switch using ISSU:

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

- On the switch, enter:

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

where **package-name.tgz** is, for example, `jinstall-132_x51_vjunos.domestic.tgz`.



NOTE: During the upgrade, you will not be able to access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-qfx-5-13.2X51-D15.4-domestic ...
Install jinstall-qfx-5-13.2X51-D15.4-domestic completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
```

```

Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```



NOTE: An ISSU might stop instead of abort if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).



NOTE: If the ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

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

6. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partitions of all of the switch:

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

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Related Documentation

- [New and Changed Features on page 28](#)
- [Changes in Behavior and Syntax on page 30](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Resolved Issues on page 39](#)
- [Documentation Updates on page 42](#)
- [Product Compatibility on page 46](#)

Product Compatibility

- [Hardware Compatibility on page 47](#)

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 on page 28](#)
- [Changes in Behavior and Syntax on page 30](#)
- [Known Behavior on page 32](#)
- [Known Issues on page 35](#)
- [Documentation Updates on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 42](#)

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

29 July 2014—Revision 5, Junos OS for the EX Series and QFX Series, Release 13.2X51-D25—

5 June 2014—Revision 4, Junos OS for the EX Series and QFX Series, Release 13.2X51-D21—Added support for the D21 image on QFX5100, QFX3600, QFX3500, and EX4300 switches and added Resolved Issues since D20.

30 May 2014—Revision 3, Junos OS for the EX Series and QFX Series, Release 13.2X51-D20—Added support for the unattended boot feature on the EX2200, introduction of an extra SFP management port on the QFX5100, and various other updates.

24 May 2014—Revision 2, Junos OS for the EX Series and QFX Series, Release 13.2X51-D20—Added updates and Known Issues.

30 April 2014—Revision 1, Junos OS for the EX Series and QFX Series, Release 13.2X51-D20

Copyright © 2014, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.