

# Junos<sup>®</sup> OS Software Release Notes for the Juniper Networks QFX Series

Release 13.1X50-D35  
6 April 2015  
Revision 9

## Contents

New Features in Junos OS Release 13.1X50-D30 for the QFX Series . . . . .	3
New Features in Junos OS Release 13.1X50-D20 for the QFX Series . . . . .	3
Layer 3 Features . . . . .	3
Routing Protocols . . . . .	3
New Features in Junos OS Release 13.1X50-D15 for the QFX Series . . . . .	4
Layer 2 Features . . . . .	4
System Management . . . . .	5
New Features in Junos OS Release 13.1X50-D10 for the QFX Series . . . . .	5
Hardware . . . . .	7
High Availability . . . . .	7
Layer 2 Features . . . . .	7
Layer 3 Multicast Features . . . . .	8
Network Management and Monitoring . . . . .	9
QFabric System Deployment . . . . .	10
Routing Protocols . . . . .	12
Security . . . . .	12
Storage and Fibre Channel . . . . .	13
System Management . . . . .	14
Traffic Management . . . . .	14
Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series . . . . .	16
Series . . . . .	16
Class of Service . . . . .	16
QFabric System Administration . . . . .	16
Limitations in Junos OS Release 13.1 for the QFX Series . . . . .	20
Interface and Chassis . . . . .	20
Layer 2 and Layer 3 Features . . . . .	20
Network Management and Monitoring . . . . .	22
QFabric Systems . . . . .	22
Storage and Fibre Channel . . . . .	24
Traffic Management . . . . .	25

Outstanding Issues in Junos OS Release 13.1 for the QFX Series . . . . .	25
Interfaces and Chassis . . . . .	26
Layer 2 and Layer 3 Features . . . . .	26
Network Management and Monitoring . . . . .	27
Platform and Infrastructure . . . . .	27
QFabric System Administration . . . . .	27
User Interface and Configuration . . . . .	30
Resolved Issues in Junos OS Release 13.1 for the QFX Series . . . . .	30
Issues Resolved in Release 13.1X50-D35 . . . . .	31
Interfaces and Chassis . . . . .	31
Issues Resolved in Release 13.1X50-D30 . . . . .	31
Interfaces and Chassis . . . . .	31
Multiprotocol Label Switching (MPLS) . . . . .	32
Platform and Infrastructure . . . . .	32
QFabric Systems . . . . .	33
Routing Protocols . . . . .	33
Issues Resolved in Release 13.1 X50-D25 . . . . .	34
Layer 2 and Layer 3 . . . . .	34
Issues Resolved in Release 13.1 X50-D20 . . . . .	34
Interfaces and Chassis . . . . .	34
Layer 2 and Layer 3 . . . . .	34
QFabric System Administration . . . . .	35
Storage and Fibre Channel . . . . .	35
Issues Resolved in Release 13.1 X50-D15 . . . . .	35
Interfaces . . . . .	35
Layer 2 and Layer 3 . . . . .	36
QFabric Systems Administration . . . . .	36
Storage and Fibre Channel . . . . .	37
Traffic Management . . . . .	37
Errata in Documentation for Junos OS Release 13.1 for the QFX Series . . . . .	38
Routing Protocols . . . . .	38
Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX	
Series . . . . .	38
Procedure for Upgrading CoS from Junos OS Release 11.1 or Release 11.2 to	
Release 11.3 or Later . . . . .	39
CoS Upgrade Requirements from Junos OS Release 12.1 to Junos OS Release	
12.2 . . . . .	40
Basic Procedure for Downloading Junos OS Release 13.1 Software . . . . .	41
Performing a Nonstop Software Upgrade on the QFabric System . . . . .	42
Performing a Standard Software Upgrade on the QFabric System . . . . .	43
Upgrade and Downgrade Support Policy for Junos OS Extended End-of-Life	
Software Releases . . . . .	43
QFX Series Documentation for Junos OS Release 13.1 . . . . .	44
Requesting Support . . . . .	44
Revision History . . . . .	44

## New Features in Junos OS Release 13.1X50-D30 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](#).

There are no new features in 13.1X50-D30 for the QFX Series. Several known issues have been fixed in this release. See “Resolved Issues in Junos OS Release 13.1 for the QFX Series” on page 30.

### Related Documentation

- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

## New Features in Junos OS Release 13.1X50-D20 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](#).

- [Layer 3 Features on page 3](#)
- [Routing Protocols on page 3](#)

### Layer 3 Features

- **IPv6 support (QFabric systems)**—Supported on server Node groups.

### Routing Protocols

- **OSPFv3 (QFabric systems)**—Use for propagating IPv6 routes. The differences in OSPFv3 versus OSPF are primarily caused by the changes in protocol semantics between IPv4 and IPv6 and the increased address size of IPv6. The main differences are:
  - Protocol processing per link, not per subnet.
  - Removal of addressing semantics.
  - Addition of flooding scope.
  - Explicit support for multiple instances per link.

- Use of link-local addresses.
- Packet format changes and LSA format changes.

[See [OSPF Overview](#).]

#### Related Documentation

- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

## New Features in Junos OS Release 13.1X50-D15 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](#).

- [Layer 2 Features on page 4](#)
- [System Management on page 5](#)

### Layer 2 Features

- **Multiple VLAN Registration Protocol (QFabric systems)**—Multiple VLAN Registration Protocol (MVRP) is a Layer 2 messaging protocol that automates the creation and management of virtual LANs, thereby reducing the time you have to spend on these tasks. If your QFabric system connects to servers that host many virtual machines that require their own VLANs, using MVRP can save you the time and effort that would be required to manually configure the VLANs on the appropriate ports. When using MVRP on a QFabric system, you must manually create on the QFabric the VLANs that exist on the attached servers because the QFabric implementation of MVRP does not allow VLANs to be created dynamically. However, you do not need to manually assign VLAN membership to the QFabric ports. MVRP automatically assigns VLAN membership to QFabric ports when it learns about a VLAN from an attached server.

[See [Understanding Multicast VLAN Registration Protocol](#).]

- **Q-in-Q tunneling (QFabric systems)**—Enables data centers and service providers to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Data centers can use Q-in-Q tunneling to isolate customer traffic within a single site or when customer traffic flows between cloud data centers in different geographic

locations. Using Q-in-Q tunneling also allows you to support more than 4089 VLANs in a QFabric system.

[See [Understanding Q-in-Q Tunneling](#).]

## System Management

- **Secure Shell (SSH) authentication keys (QFabric systems)**—Enables you to log in to the QFabric system using SSH authentication key pairs. The QFabric system supports DSA and RSA. SSH keys use public-key cryptography and challenge-response authentication. When you use SSH keys, the server authenticates you without you having to send a password over the network. Because you do not send the password over the network, the risk of an attacker using brute-force password attacks to correctly guess proper credentials is drastically reduced.

[See [Configuring the Root Password](#) ]

### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

## New Features in Junos OS Release 13.1X50-D10 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 7](#)
- [High Availability on page 7](#)
- [Layer 2 Features on page 7](#)
- [Layer 3 Multicast Features on page 8](#)
- [Network Management and Monitoring on page 9](#)
- [QFabric System Deployment on page 10](#)
- [Routing Protocols on page 12](#)
- [Security on page 12](#)
- [Storage and Fibre Channel on page 13](#)

- [System Management on page 14](#)
- [Traffic Management on page 14](#)

## Hardware

- **Support for line-rate performance on access ports of QFX3600 Node devices (QFabric systems)**—You can enable line-rate performance on the access ports (server or network-facing ports) of a QFX3600 Node device by connecting eight uplink ports (labeled Q0 through Q7) of the Node device to Interconnect devices.



**NOTE:** By default, only four ports (labeled Q0 through Q3) operate as 40-Gigabit Ethernet data plane (fte) uplink ports for uplink connections between your QFX3600 Node device and your Interconnect devices. Therefore, you must configure the remaining four ports (labeled Q4 through Q7) as uplink ports before you connect them to Interconnect devices.

[See [Understanding Node Devices](#).]

## High Availability

- **VRRP between two QFabric systems (QFabric systems)**—Configuring hosts (clients or servers) on your network with static default routes minimizes configuration effort and complexity and reduces processing overhead on the hosts because you do not need to configure a dynamic routing protocol or router discovery protocol on each host. However, when hosts are configured with static routes, a failure of the default gateway isolates all hosts that are unable to detect available alternate paths to their gateway. Using the Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for hosts if the primary gateway fails.

VRRP (defined in RFC 3768) provides dynamic failover of IP addresses from one router to another in the event of failure. With Junos OS Release 13.1, you can create a VRRP configuration with two QFabric systems to provide a highly available default gateway without needing to configure dynamic routing or router discovery protocols on your hosts.

QFabric systems configured with VRRP share a virtual IP address, which is the address you configure as the default route on your hosts. At any time, one of the QFabric systems is the VRRP master for the virtual IP address, meaning that it owns the address and is the active default gateway. The other QFabric system is the backup for that virtual IP address. The systems dynamically assign master and backup roles based on priorities that you configure. If the master fails, the backup system becomes the master within a few seconds. The backup system becomes the master without any interaction with the hosts.

You can also create a load-sharing (active-active) configuration in which each QFabric system simultaneously acts as a master for one virtual IP address and a backup for a second virtual IP address. This configuration enables you to maximize your investment in your QFabric systems, because both systems forward traffic simultaneously.

[See [Understanding VRRP](#).]

## Layer 2 Features

- **Port mirroring (QFabric systems)**—Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a system that analyzes the traffic. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Junos OS Release 13.1 increases the number of supported port-mirroring sessions to four per QFabric system Node group, subject to the following limits:

- As many as four of the configurations can be for local port mirroring.
- As many as three of the configurations can be for remote port mirroring.

Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:

- There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
- There can be no more than two configurations that mirror egress traffic.

[See [Understanding Port Mirroring](#).]

- **Private VLANs (QFabric systems)**—Private VLANs (PVLANS) split a broadcast domain into multiple isolated broadcast subdomains. PVLANS restrict traffic flows through their member switch ports (called private ports) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port is usually connected to a router, firewall, server, or provider network. Each PVLAN typically contains many private ports that communicate only with a single uplink, thereby preventing the ports from communicating with one another.

As part of a PVLAN configuration, you can configure access ports to be promiscuous, and you can configure trunk ports to carry secondary VLAN (isolated and community VLAN) traffic. These features are useful in VMware deployments in which applications are hosted in virtual machines and customer access to these applications is controlled through private VLANs.

[See [Understanding Private VLANs](#).]

## Layer 3 Multicast Features

- **PIM sparse mode (QFabric systems)**—Enables efficient routing to multicast groups with receivers sparsely spread over multiple networks.
- **MSDP (QFabric systems)**—Enables you to connect multiple IP version 4 PIM sparse mode (PIM-SM) domains to one another. MSDP typically runs on the same routing device as a PIM-SM rendezvous point (RP). Each MSDP routing device establishes adjacencies with internal and external MSDP peers, similar to how BGP peering works. These peers inform each other about active sources within the domain. When they detect active sources, the peers send PIM sparse-mode explicit join messages to the active source.
- **Static RP (QFabric systems)**—Enables the system to be a statically defined rendezvous point (RP) for multicast trees. The RP is the root of a multicast tree and receives



multicast traffic from the source. Configuring a static RP is a simple and convenient way to provide a multicast root for downstream multicast receivers. Using statically defined RPs is particularly suitable if your network requires few RPs or if the RP assignment rarely changes.

- **Anycast RP (QFabric systems)**—Supports multiple rendezvous points (RPs) using anycast addresses (RPs sharing a single routable IP address) in either a PIM or MSDP-enabled network. Using anycast RP provides load balancing and redundancy so that multicast performance is optimized and the failure of a single RP does not result in an interruption in service to multicast receivers.
- **IGMP (QFabric systems)**—Manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members. Junos OS Release 13.1 supports IGMP versions 1 and 2.
- **IGMP filters (QFabric systems)**—Enables you to filter unwanted IGMP reports at the interface level. To filter IGMP group addresses (either IGMPv2 or IGMPv3), include the **route-filter** statement at the **[edit policy-statement policy-name from]** hierarchy level. To filter IGMP source addresses (IGMPv3 only), include the **source-address-filter** statement at the **[edit policy-statement policy-name from]** hierarchy level. Apply the filters by using the **group-policy policy-name** statement at the **[edit protocols igmp interface interface-name]** hierarchy level.

[See [Multicast Protocols on the QFX Series](#).]

## Network Management and Monitoring

- **Support for the Class-of-Service MIB (QFabric systems)**—Extends support for Juniper Networks Class-of-Service MIB (mib-jnx-cos) to QFabric systems. The following MIB table objects are supported: jnxCosIfstatFlagTable, jnxCosQstatTable, jnxCosFcidTable, and jnxCosFcTable. Support for the Class-of-Service MIB does not include SNMP traps.

[See [SNMP MIBs Support](#).]

- **Support for QFabric system Director group information in the SNMP Fabric Chassis MIB (QFabric systems)**—Adds support for the QFabric system Director group information in the Juniper Networks Fabric Chassis MIB (jnx-fabric-chassis). The following MIB table objects contain information about the Director group hardware:
  - jnxFabricContainersTable—Root OID is 1.3.6.1.4.1.2636.3.42.2.2.2
  - jnxFabricContentsTable—Root OID is 1.3.6.1.4.1.2636.3.42.2.2.3
  - jnxFabricDeviceTable—Root OID is 1.3.6.1.4.1.2636.3.42.2.2.1
  - jnxFabricFruTable—Root OID is 1.3.6.1.4.1.2636.3.42.2.2.7
  - jnxFabricOperatingTable—Root OID is 1.3.6.1.4.1.2636.3.42.2.2.5

[See [Fabric Chassis MIB](#).]

- **Enhancements to system log (syslog) support (QFabric systems)**—Provide additional support for syslog messages, including:

- Configuration of syslog output in a structured data format with the **structured-data** statement at the **[edit system syslog file filename]** or **[edit system syslog host hostname]** hierarchy levels.
- Ability to specify the types of syslogs that the QFabric system Director group collects from a component device by configuring the **filter all facility severity** or **filter all match regular-expression** statement at the **[edit system syslog]** hierarchy level.
- Configuration of the output from the **show log** operational mode command by specifying the device type and device ID (or device alias) with the new **device-type** and **device-id | device-alias** optional parameters. Device types include **director-device**, **infrastructure-device**, **interconnect-device**, and **node-device**.

[See [Understanding the Implementation of System Log Messages on the QFabric System.](#)]

## QFabric System Deployment

- **QFabric system backup and restore (QFX3000-G and QFX3000-M QFabric systems)**—Provides a streamlined way to save a copy of the current QFabric system configuration file and a copy of the current software package for recovery purposes. You can use these saved files to restore your QFabric system to full operation after a system failure or shutdown. To save a copy of the configuration and software files locally within the QFabric system, issue the **request system software system-backup** operational mode command.

After you issue this first command, insert a USB drive into your Director device and issue the **request system software system-backup usb-create** operational mode command. The second command saves the configuration and software files from the local copy in the QFabric system to the USB drive. After a system failure, insert the USB drive into the Director device, power on the device, and follow the recovery prompts to restore both the QFabric system software and configuration.

[See [Performing System Backup and Recovery for a QFabric System.](#)]

- **Planned online insertion and removal (OIR) for selected QFabric system components (QFX3000-G and QFX3000-M QFabric systems)**—Enables you to remove Director devices and control plane switches or Virtual Chassis gracefully from your QFabric system.
  - To remove and replace a Director device, issue the **request system halt director-device director-device-name** operational command. When traffic reroutes to the backup Director device, power off the halted Director device, detach the cables, remove the Director device, install a replacement Director device, reattach the cables, and power on the new Director device. The replacement device synchronizes its configuration with the second Director device and resumes normal operations.

[See [Replacing a QFX3100 Director Device in a QFabric System.](#)]

- To remove and replace a control plane Virtual Chassis for a QFX3000-G QFabric system, connect to the master switch in the Virtual Chassis through your management network, save a copy of the Virtual Chassis configuration to a remote server, and issue the **request system halt all-members** operational mode command. When traffic reroutes to the second Virtual Chassis, power off the halted member switches, detach

the cables, remove the members, and install replacement Virtual Chassis members. Reattach the cables, power on the new members, and reload the Virtual Chassis configuration from the remote server. The replacement Virtual Chassis joins the control plane network and resumes normal operations.

[See [Removing and Replacing a QFX3000-G QFabric System Control Plane.](#)]

- To remove and replace a control plane switch for a QFX3000-M QFabric system, connect to the control plane switch through your management network, save a copy of the switch configuration to a remote server, and issue the **request system halt** operational mode command. When traffic reroutes to the second control plane switch, power off the halted switch, detach the cables, remove the switch, and install the replacement switch. Reattach the cables, power on the new switch, and reload the configuration from the remote server. The replacement switch joins the control plane network and resumes normal operations.

[See [Removing and Replacing a QFX3000-M QFabric System Control Plane.](#)]

- **Support for maximum scale in a QFabric system (QFX3000-G QFabric systems)**—Provides full support for 128 Node devices in a QFX3000-G QFabric system. No new configuration statements are required, and the enhanced functionality is enabled by default.

[See [Understanding the QFabric System Hardware Architecture.](#)]

- **QFabric system convergence improvements (QFX3000-G and QFX3000-M QFabric systems)**—Offer enhanced convergence performance for unicast traffic in QFabric systems. Improvements have been made so that Node devices, QSFP+ interfaces, equal-cost multipath members, and link aggregation group (LAG) bundles now converge within 1.5 seconds. No new configuration statements are required, and the enhanced functionality is enabled by default.
- **QFabric system security enhancements (QFX3000-G and QFX3000-M QFabric systems)**—Introduce the following improved security features:
  - Reduce the file size and number of software packages included in the RPM software bundle by removing nonessential files
  - Increase the password encryption strength to permit passwords longer than eight characters in length
  - Tighten file permissions to prevent unauthorized read/writes to files shipped with the product

No new configuration statements are required, and the enhanced functionality is enabled by default.

- **QFabric system control plane link selection enhancements (QFX3000-G and QFX3000-M QFabric systems)**—Provide the QFabric system with a new data collection mechanism to detect minor discrepancies in the control plane links. The system now uses this additional link data to determine and implement appropriate actions to work around operational issues (such as redirecting control plane traffic over an alternate path after failure of the primary path). No new configuration statements are required, and the enhanced functionality is enabled by default.

## Routing Protocols

- **Per-packet load balancing local bias (QFabric systems)**—When you enable per-packet load balancing on a QFabric system, packets might be switched across the fabric even though there is a local port with a same-cost route to the destination. For example, if per-packet load balancing is enabled and a packet arrives at network Node device A, it might be switched to network Node device B and forwarded from there even if there is a same-cost route through a port on Node device A to the destination. In this case, traffic transits the fabric needlessly.

In Junos OS Release 13.1X50-D10, you can configure a QFabric system to choose a locally switched route if one is available. To enable this feature, include the **ecmp-do-local-lookup** statement at the **[edit forwarding-options]** hierarchy level.

[See [Configuring Per-Packet Load Balancing](#).]

## Security

- **Increased numbers of policers (QFabric systems)**—Each Node device on a QFabric system supports the following maximum number of policers:
  - Two-color policers used in ingress firewall filters: 767
  - Three-color policers used in ingress firewall filters: 767
  - Two-color policers used in egress firewall filters: 1022
  - Three-color policers used in egress firewall filters: 512

[See [Overview of Policers](#).]

- **Increased numbers of firewall filters (QFabric systems)**—Each Node device on a QFabric system supports the following maximum number of firewall filter terms per type of attachment point:
  - 768 terms for ingress filters
  - 1024 terms for egress filters

These totals are applied in aggregate. That is, you can commit a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the output direction.



**NOTE:** If you want to create more than 512 egress VLAN filters, your first VLAN ID should be 6 and the subsequent VLAN IDs should increase by 1. For example, to create 1024 egress VLAN filters, the first VLAN ID would be 6, the second ID would be 7, and the sequence would continue through VLAN ID 1029. Similarly, if you want to create fewer than 512 egress VLAN filters but want the total number of terms in those filters to exceed 512, you should number your VLAN IDs in the same manner. If you do not use this approach to create your VLAN IDs, the total number of allowed terms or filters will be less than 1024 and might be 512.

The ternary-content addressable memory (TCAM) for firewall filters is divided into slices that accommodate 256 filters (assuming that there is one term per filter), and all the filters in a memory slice must be of the same type and applied in the same direction. A memory slice is reserved as soon as you apply a filter. For example, if you create a port filter and apply it in the input direction, a memory slice is reserved that only stores ingress port filters. If you create and apply only one ingress port filter, the rest of this slice is unused and is unavailable for other filter types.

Continuing with the above example, assume that you create and apply 256 ingress port filters with one term each so that one memory slice is filled. This leaves two more memory slices available for ingress filters. (Remember that the maximum number of ingress filters is 768.) If you then create and apply an ingress Layer 3 filter, another memory slice is reserved for ingress Layer 3 filters. As before, the rest of the slice is unused and is unavailable for different filter types. At this point there is one memory slice available for any ingress filter type.

Now assume that you create and apply a VLAN ingress filter. The final memory slice is reserved for VLAN ingress filters. Memory allocation for ingress filters (once again assuming one term per filter) is as follows:

- Slice 1: Filled with 256 ingress port filters. You cannot apply any more ingress port filters.
- Slice 2: Contains one ingress Layer 3 filter. You can apply 255 more ingress Layer 3 filters.
- Slice 3: Contains one ingress VLAN filter. You can apply 255 more ingress VLAN filters.

Here is another example. Assume that you create 257 ingress port filters with one term per filter—that is, you create one more term than a single memory slice can accommodate. When you apply the filters and commit the configuration, the filter memory allocation is:

- Slice 1: Filled with 256 ingress port filters. You cannot apply any more ingress port filters.
- Slice 2: Contains one ingress port filter. You can apply 255 more ingress port filters.
- Slice 3: This slice is unassigned. You can create and apply 256 ingress filters of any type (port, Layer 3, or VLAN), but all the filters must be of the same type.

[See [Overview of Firewall Filters](#).]

## Storage and Fibre Channel

- **Priority-based flow control (PFC) enhancements (QFabric systems)**—PFC (IEEE 802.1Qbb) prevents frame loss due to congestion. PFC functions similarly to IEEE 802.3x Ethernet PAUSE, but instead of pausing all of the traffic on an entire link, PFC pauses only individually selected IEEE 802.1p priorities. This feature increases the number of supported PFC paused priorities from two priorities to six priorities. In addition, you can configure the maximum receive frame size (MRU) per priority on a per-port basis, and you can configure the length of the attached cable for buffer headroom calculation on a per-port basis. You can also configure which output queue to pause

when a PFC frame is received for a particular priority. On Fibre Channel interfaces, you can configure a fixed classifier so that all of the traffic on the interface is assigned to one forwarding class, and you can remap the priority of the storage traffic.

[See [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#) and [Understanding CoS Fabric Forwarding Class Sets](#).]

- **VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping (FC-BB-6) (QFabric systems)**—Installs filters that help prevent unauthorized access and data transmission on a bridge that connect ENodes in the Ethernet network. A key benefit of VN2VN\_Port FIP snooping is that it enables FCoE initiators and targets to communicate directly through the switch, without going through an FCoE forwarder (FCF) or a Fibre Channel switch. To configure VN2VN\_Port FIP snooping, include the **examine-vn2vn** statement at the **[edit ethernet-switching-options secure-access-port vlan *vlan-name* examine-fip]** hierarchy level.

[See [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch](#).]

- **High availability for FCoE initialization) protocol (FIP) snooping (QFabric systems)**—Restores FIP snooping sessions if the Ethernet switching process is terminated gracefully or ungracefully. The QFabric system keeps FIP snooping session information in a persistent storage module. If the Ethernet switching process terminates, the previously existing ENode and host sessions continue to run and are not lost. Existing logged-in ENodes, discovered FCoE forwarders (FCFs) and Fibre Channel switches, sessions, and filter states are preserved through the restart. This feature is automatic, and there are no new statements or commands.

[See [Understanding FCoE and FIP Session High Availability](#).]

## System Management

- **RADIUS and TACACS+ accounting (QFabric systems)**—Allows you to collect accounting data, such as user login information, length of user sessions, and the amount of network traffic that the user session generated. Additionally, you can configure the QFabric system to send the accounting data to TACACS+ and RADIUS servers.



**NOTE:** TACACS+ and RADIUS authorization is already supported on QFabric systems.

[See [Understanding RADIUS Accounting](#).]

## Traffic Management

- **Layer 3 class of service (QFabric systems)**—Enables you to configure Layer 3 class of service (CoS) by defining classifiers and rewrite rules on a Layer 3 physical interface (an interface configured as family inet or as family inet6) when the Layer 3 physical interface has at least one defined Layer 3 logical interface. Configure Layer 3 CoS by including the **classifiers** and **rewrite-rules** statements at the **[edit class-of-service]** hierarchy level.

[See [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#).]

- **Scheduler configuration on Node device fabric ports (QFabric systems)**—Enables you to configure scheduling on the fabric (fte) ports of QFabric system Node devices. In earlier Junos OS releases, Node device fabric port scheduling was done by default, with no user configuration. In Junos OS 13.1, the default fabric port scheduler configuration is similar to the default scheduler configuration on access interfaces. The default fabric port scheduler supports the five default forwarding classes (best-effort, fcoe, no-loss, network-control, and mcast).

If you configure any new forwarding classes, you must configure scheduling on the fabric ports to allocate bandwidth to those forwarding classes, just as you must configure scheduling on the access ports for user-defined forwarding classes. In addition, if a Node device fabric interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set.

[See [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#) and [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#).]

#### **Related Documentation**

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

## Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series

### Class of Service

- The way the QFX Series handles lossless forwarding classes (the fcoe and no-loss forwarding classes) changes in Junos OS Release 12.2. In Junos OS Release 12.2 and earlier by default, the fcoe and no-loss forwarding classes are mapped to output queue 3 and output queue 4, respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport.

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



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

If you use the default forwarding class settings for the lossless queues (the configuration does not include explicit setting of the fcoe or the no-loss forwarding classes), then the fcoe and no-loss queues behave as lossless queues. When you upgrade to Junos OS Release 12.2, traffic assigned to the fcoe and no-loss queues continues to be treated as lossless traffic.

If your configuration explicitly sets the fcoe or the no-loss forwarding class (**[set class-of-service forwarding-classes class *class-name* queue-num *queue-number*]**), after you upgrade to Junos OS Release 12.2, those queues do *not* receive lossless treatment and behave as lossy (best effort) queues. To retain lossless treatment of the **fcoe** and **no-loss** queues, delete the explicit lossless forwarding class configuration before you upgrade to Junos OS Release 12.2.



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

### QFabric System Administration

- **DHCP traffic rate limited (QFX3000-M and QFX3000-G QFabric systems)**—DHCP packets might get dropped if the transmission rate on the interface exceeds 500 packets per second.
- **Active LAG indicator for the Director group status command (QFX3000-G and QFX3000-M QFabric systems)**—Enables you to view which QFabric system control



plane LAG bundle is active. To view the active control plane LAG links, issue the **show fabric administration inventory director-group status** command and look for the LAG links listed as **Active** in the **Interface Link Status** section.

[See [show fabric administration inventory director-group status](#).]

- **Configure command changes (QFX3000-M and QFX3000-G QFabric systems)**—The default mode for the **configure** command changed from private to shared in Junos OS Release 13.1X50-D15. Support for the **configure exclusive** command was also added.

[See [Forms of the configure Command](#).]

- **Director group status command enhanced (QFX3000-M and QFX3000-G QFabric systems)**—The **show fabric administration inventory director-group status** command now displays errors when the virtual Routing Engines or hard drive cannot respond to a status request from the Director group.

[See [show fabric administration inventory director-group status](#).]

- **Disable multicast make-before-break (QFX3000-M and QFX3000-G QFabric systems)**—Enables you to turn off the default *make-before-break* multicast feature on QFabric systems to prevent the duplication of traffic generated by this default behavior. To disable multicast make-before-break, include the **no-make-before-break** statement at the **[edit fabric routing-options multicast]** hierarchy level.

[See [no make before break](#).]

- **Display the status of a QFabric system software upgrade (QFX3000-M and QFX3000-G QFabric systems)**—Enables you to view the current status of a software upgrade on a QFabric system. This command is available for either a standard *component all* upgrade or a nonstop software upgrade. To display the software upgrade status for your QFabric system, issue the **show system software upgrade status** operational command.

[See [show system software upgrade status](#).]

- **Display messages on LCD screens of QFabric system devices (QFX3000-M and QFX3000-G QFabric systems)**—Enables you to display information on the LCD screens of QFabric system devices. To add a message that appears on the LCD screen of a Node device or an Interconnect device, include the **message** statement at the **[edit chassis display (interconnect-device | node-device) device-id permanent]** hierarchy level. Serial numbers appear on Director devices, Interconnect devices, and Node devices by default, but alias names appear if you configure aliases for your QFabric system devices.

[See [set chassis display message](#) and [show chassis lcd](#).]

- **Enhanced output for the QFabric system component inventory command (QFX3000-M and QFX3000-G QFabric systems)**—Expands the amount of information you can gather about your QFabric system. New output includes details for troubleshooting failed component configurations, status of device mastership for Node groups, and Routing Engine mastership for Interconnect devices. To view the enhanced operational mode command output for your QFabric system, issue the **show fabric administration inventory detail** command.

[See [show chassis lcd.](#)]

- **Enhanced output for the show fabric administration inventory command (QFX3000-M and QFX3000-G QFabric systems)**—The output of the **show fabric administration inventory** command now includes a **Director group** section that shows the serial identifier for each Director device that is part of the Director group in a QFabric system.

[See [show fabric administration inventory.](#)]

- **Expanded permissions for QFabric component access classes (QFX3000-M and QFX3000-G QFabric systems)**—Allows administrators to issue a larger set of operational mode commands than previously when accessing individual components in a QFabric system through the **request component login component-name** command. The *qfabric-operator* class now includes the ability to issue the **monitor** and **show log** commands. The *qfabric-admin* class now includes the ability to issue all operational mode commands except for **configure**.

[See [show fabric administration inventory.](#)]

- **New QFabric device mode options for QFX3600 switches and QFX3600-I Interconnect (QFX3000-M and QFX3000-G QFabric systems)**—Enables you to convert a QFX3600 switch into an Interconnect device or a QFX3600-I Interconnect device into a Node device or standalone switch. These new options give you more flexibility when expanding or changing your QFX3000-M QFabric system. To convert a QFX3600 switch into an Interconnect device, issue the **request chassis device-mode interconnect-device** operational mode command and reboot the device. To convert a QFX3600-I Interconnect device into a Node device, issue the **request chassis device-mode node-device** operational mode command and reboot the device. To convert a QFX3600-I Interconnect device into a standalone switch, issue the **request chassis device-mode standalone** operational mode command and reboot the device.



**NOTE:** Converting a QFX3600 switch to a Node device requires installation of Jloader 1.1.8. For more information, see [Jloader 1.1.8 Release for QFX-Series Platforms](#)

[See [Converting the Device Mode for a QFabric System Component.](#)]

- **Remove devices from the QFabric system inventory (QFX3000-M and QFX3000-G QFabric systems)**—Enables you to take a disconnected Node device or Interconnect device out of the QFabric system inventory shown in the output of the **show fabric administration inventory** operational mode command. To remove a disconnected device from the QFabric system inventory, issue the **request fabric administration remove (node-device node-device-name | interconnect-device interconnect-device-name)** operational mode command.

[See [request fabric administration remove.](#)]

- **request system software format-qfabric-backup command (QFX3000-M and QFX3000-G QFabric systems)**—Enables you to format the shared directory of your Director group so you can use the QFabric system backup and recovery feature introduced in Junos OS 13.1X50-D10. If you have never used the backup and recovery feature in your QFabric system, you must issue the command before you can use the

feature. You only need to issue the command once on first use. When you issue the command, the software image and directory structure of a QFabric system installation USB flash drive are copied to the Director group, and the QFabric system is ready for all future backup and recovery procedures.

To enable your QFabric system to use the backup and recovery feature, insert a Juniper Networks software QFabric system installation USB flash drive into the master Director device of the Director group, and issue the **request system software format-qfabric-backup** command.

[See [Performing System Backup and Recovery for a QFabric System](#).]

- **Revised output for the show fabric administration inventory node-devices command (QFX3000-M and QFX3000-G QFabric systems)**—The **Configuration** field has been removed from the output of the **show fabric administration inventory node-devices** operational mode command because configurations are assigned to Node groups rather than Node devices. To view configuration status, issue the **show fabric administration inventory node-groups** operational mode command.

[See [show fabric administration inventory node-devices](#).]

- **Summary output for the QFabric system component inventory command (QFX3000-M and QFX3000-G QFabric systems)**—Provides a simplified set of component inventory information for your QFabric system. A new **summary** option for the **show fabric administration inventory** command displays an overview of the number of Node devices and Interconnect devices that are part of the system. To view summary information for your QFabric system, issue the **show fabric administration inventory summary** command.

[See [show fabric administration inventory](#).]

- **Update to a QFabric device conversion command option (QFX3500 and QFX3600 devices)**—Renames the **fabric** option for the **request chassis device-mode** and **show chassis device-mode** commands to **node-device**. This change enables you to identify which device mode your QFX3500 and QFX3600 device assumes in a QFabric system. The **fabric** option is supported for three releases and then deprecated, so we recommend that you use the **node-device** option in Junos OS Release 12.1 and later.

#### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

## Limitations in Junos OS Release 13.1 for the QFX Series

---

This section lists the limitations in Junos OS Release 13.1 for the QFX Series.

### Interface and Chassis

- On a QFabric system, the multicast make-before-break (MBB) feature is enabled by default. Once you disable MBB, the feature becomes permanently disabled.
- When you configure an interface range and make the range a member of a VLAN, a problem occurs if you later remove an interface from the range. In this situation, the interface that you remove from the range is not also removed from the VLAN.
- When you issue the **ping fabric unicast-flow** command, the following message might appear: **error: the fabric-oam subsystem is not responding to management requests**. This message is sent in error. As a workaround, wait a few minutes and reissue the command.
- Multicast frames are not distributed properly among the member links of a link aggregation group. There is no workaround.
- On a QFabric system, some analyzer sessions might not work after a system upgrade from Junos OS Release 12.2X50-D10 to Junos OS Release 12.2X50-D20.

On a QFabric system, if an interface that has not been configured is added to a VLAN, an unrelated error message might be displayed.

- On a QFabric system, mirroring of link aggregation group (LAG) interface traffic on network Node groups might fail if there are changes in port or analyzer configurations. If that happens, error messages containing the string **PFE\_ANALYZER\_TASK\_FAILED** are displayed.
- On a QFabric system, eswd (Ethernet switching daemon) might crash on the master Node device in a redundant server Node group during initial setup if the backup Node device is going offline at the same time.
- On a QFabric system, if you perform a nonstop software upgrade, sometimes a QSFP+ connection (fte- interface) on the Interconnect device will be down while the corresponding link on the redundant server Node group side will be up.

### Layer 2 and Layer 3 Features

- On a QFabric system, momentary loss of Layer 2 multicast traffic might occur when you reboot the backup network Node group Routing Engine.
- On a QFabric system, momentary loss of Layer 2 multicast traffic might occur when you reboot the fabric manager Routing Engine.
- When you disable and enable multiple fabric tunnel 40-Gigabit Ethernet data plane uplink interfaces (fte interfaces) on several Node devices at the same time, the Node devices might receive more than four copies of a traffic stream for a few seconds.

- On a QFabric system, if the network Node group experiences a kernel panic as a result of an update to the routing protocol process (rpd), up to 30 seconds of Layer 3 multicast traffic loss might occur.
- On a QFabric system, if any event related to high availability triggers a routing protocol process (rpd) update, the system experiences Layer 3 multicast traffic loss for a duration of approximately 30 seconds.
- On QFabric systems with the Multiple VLAN Registration Protocol (MVRP) enabled, enabling MVRP traceoptions can disrupt MVRP joins if all the MVRP timers are set to their default values. If you need to enable MVRP traceoptions, set the timers to the following values:
  - join: 3000 ms
  - leave: 10000 ms
  - leave-all: 30000 ms

As a workaround, if MVRP traceoptions need to be enabled, increase join timer to 3s and leave timer to 10s and leave-all timer to 30s.

- On QFabric systems with remote port mirroring enabled, you cannot use Multiple VLAN Registration Protocol (MVRP) to add the output interface to the output VLAN. You must manually configure the output interface to be a member of the output VLAN. As a workaround, configure the mirror port as a static member of the Remote SPAN (RSPAN) VLAN.
- The Q-in-Q tunneling broadcast domain is based on the service VLAN (the S-VLAN outer tag), not on the customer VLAN (the C-VLAN inner tag). For this reason, flooding and learning take place in the context of the S-VLAN. The system does not use the C-VLAN information to determine flooding and learning, which results in unknown unicast traffic being placed on access interfaces that use the same S-VLAN even when the C-VLAN tag mapping is different. As a workaround, map each access interface to only one S-VLAN to prevent unknown unicast traffic from being placed on the wrong interface.
- When you enable Multiple VLAN Registration Protocol (MVRP) and configure it to peer with a device running MVRP in registration mode, a problem occurs if there is a kernel panic on the QFabric device that owns the MVRP interface. In this case, traffic for any VLANs that were created dynamically on the non-QFabric device is dropped for as long as 30 seconds. To prevent this issue, set the MVRP leave timer to more than 30000 milliseconds on the peer (non-QFabric) device.
- When Multiple VLAN Registration Protocol (MVRP) and Spanning Tree Protocol (STP) are enabled on interfaces, and a link failure causes STP convergence, traffic might be dropped for as long as 40 seconds.
- When Multiple VLAN Registration Protocol (MVRP) is enabled on an interface, the system registers VLANs on the interface upon reception of MVRP join messages, even if the port is blocked by xSTP.
- When Multiple VLAN Registration Protocol (MVRP) and Q-in-Q tunneling are enabled on an interface, the tag protocol identifier (EtherType) of packets that egress through

that interface is set to 0x8100 even if you have configured it to be a different value using the **set ethernet-switching-options dot1q-tunneling ether-type** command.

- When you use a range of service VLANs (S-VLANs) for Q-in-Q tunneling on an interface, the S-VLAN tag (the outer tag) is always assigned the EtherType 0x8100.
- On a QFabric system, when an IPV6 route is changed to point to an ECMP next-hop instead of a non-ECMP next-hop for certain prefix lengths, traffic drops might occur. This issue occurs with /92 routes and does not occur with /64 or /128 routes.

## Network Management and Monitoring

- There is a delay in the display of system log messages after the **show log messages** command is issued.
- Disabling all 10-Gigabit Ethernet (xe) interfaces at the same time results in an output displaying a greater number of linkdown traps than the number of interfaces on the system.
- When a TACACS+ accounting event is generated, and the configured TACACS+ server is unreachable, the Junos OS CLI is unresponsive until the configured timeout expires.

The SNMP Fabric Chassis MIB has the following limitations:

- The jnxFabricFruTable table does not include information about network interface cards (NICs) on Director devices.
- Containers for the Director devices in the jnxFabricContainersTable table have a nonhierarchical (or flat) structure. Components in the containers are organized as siblings and not as parents and children.
- Hardware events on the Director group are detected by scanning. As a result, a trap might not be generated until up to 30 seconds after the event has occurred.
- The software does not distinguish between the fan removal and fan failure events on the Director group. In each case, both the jnxFabricFanFailure and jnxFabricFruFailed traps are generated.
- The software does not distinguish between the fan insertion and fan OK events on the Director group. In each case, both the jnxFabricFanOK and jnxFabricFruOK traps are generated.
- During bootup of a Director device, the jnxVccpMemberUp trap might be displayed with unreadable characters.

## QFabric Systems

- On a QFabric system, you cannot check Routing Engine status for the Director group by issuing the **show chassis routing-engine** command. As a workaround, issue the **show fabric administration inventory director-group status** command to check the Routing Engine status of the Director devices.
- On a QFabric system, the Layer 2 and Layer 3 multicast table size is 4 KB. The system derives a hash value based on the traffic pattern and creates entries in the table. When the traffic patterns are similar to each other, there can be hash collisions. When there

are hash collisions, a smaller number of entries are programmed into the table, so the full 4-KB table size might not be used.

- The output of the **show pfe statistics traffic** command does not report statistics for normal traffic discards.
- When you replace an Interconnect device in a QFabric system that has four Interconnect devices, the output of the **show fabric administration inventory** command might not show the new Interconnect device (indicating that the new Interconnect device is not online). As a workaround, follow these steps when replacing Interconnect devices:
  1. Disconnect the Interconnect device that needs to be replaced.
  2. Issue the **request fabric administration remove interconnect-device *interconnect-device-name*** command and wait for 60 minutes before you connect the new Interconnect device to your QFabric system.
  3. Issue the **show fabric administration inventory** command to verify that the new Interconnect device is online.
- When you perform a nonstop software upgrade on a network Node group that contains no members, and a kernel panic occurs on one of the network Node group Routing Engines 8 minutes after the upgrade, all Layer 2 and Layer 3 unicast traffic might be lost even if there is an active Routing Engine present when the panic happens.
- When you configure an alias name for a Node device, configure the serial number of the Node device when adding it into a Node group, and then try to remove the Node device from the Node group by specifying the aliased name, the removal is unsuccessful. As a workaround, when you delete the Node device from the Node group, issue the **show fabric resources node-group *group-name*** operational mode command to view your original configuration, and then use the same alias name or serial number that you used when the Node device was added.
- When you issue the **request fabric administration director-group change-master** command to change mastership within the Director group, and then issue the **show system core-dumps** command on the new master Director device, the resulting output might not display the QFabric component inventory properly. As a workaround, wait a few minutes for the mastership change to settle and then reissue the **show system core-dumps** command.
- If you configure an alias for a Node device and the device is part of an automatically generated server Node group, the hostname for the Node group remains as the serial ID for the device rather than the configured alias name.
- Deleting individual child members from a link aggregation group (LAG) might reset the aggregated Ethernet interface traffic counters.
- When you issue the **show chassis mac-address interconnect-device <interconnect-device-name> | display xml** operational mode command, the output might display incorrect information.
- When you add an NTP server configuration to the QFabric configuration using the **hostname** command, the NTP server hostname is resolved to an IP address. For this reason, the QFabric system configuration shows the IP address instead of the hostname.

To remove an NTP server configuration, you must delete the hostname instead of the IP address. For example, if you configured an NTP server with a hostname of pool.ntp.org and the resolved IP address is 216.160.0.218, you must delete pool.ntp.org to delete the NTP server configuration.

- When you issue the **traceroute unicast-flow** command, the fabric maintenance intermediate points (FMIPs) might send duplicate responses for the same hop on the data flow path. Even though the responses are duplicated, the information they contain is incorrect. There is no workaround.
- When you reboot the Interconnect device, existing FCoE sessions might be dropped.

## Storage and Fibre Channel

- A Fibre Channel fabric supports a maximum of four Fibre Channel over Ethernet (FCoE) VLAN interfaces.
- The maximum number of logins for each FCoE node (ENode) is a range of 32 to 2000 for trusted fabrics and 32 to 376 for untrusted fabrics. (Each ENode can log in to a particular fabric up to the maximum number of configured times. The maximum number of logins is per-fabric, so an ENode can log in to more than one fabric and have its configured maximum number of logins on each fabric.)
- The maximum number of FCoE sessions for the switch, which equals the total number of fabric login (FLOGI) sessions plus the total number of fabric discovery (FDISC) sessions, depends on how you configure the ports in a specified FC fabric. If you configure the ports as FCoE trusted, the maximum number of FCoE sessions (ENode to FCF sessions) the system can support is 2500.
- When you configure FIP snooping filters, if the filters consume more space than is available in the ternary content-addressable memory (TCAM), the configuration commit operation succeeds even though the filters are not actually implemented in the configuration. Because the commit operation checks syntax but does not check available resources, it appears as if the FIP snooping filters are configured, but they are not. The only indication of this issue is that the switch generates a system log message that the TCAM is full. You must check the system log to find out if a TCAM full message has been logged if you suspect that the filters have not been implemented.
- You cannot use a fixed classifier to map FCoE traffic to an interface. The FCoE application type, length, value (TLV) carries the FCoE priority-based flow control (PFC) information when you use an explicit IEEE 802.1p classifier to map FCoE traffic to an interface. You cannot use a fixed classifier to map FCoE traffic to an interface because untagged traffic will be classified in the FCoE forwarding class, but FCoE traffic must have a priority tag (FCoE traffic cannot be untagged).

For example, the following configuration is supported:

```
[edit class-of-service]
user@switch# set congestion notification profile fcoe-cnp input ieee-802.1 code-point
011 pfc
user@switch# set interfaces xe-0/0/24 unit 0 classifiers ieee-802.1 fcoe
```

For example, the following fixed classifier configuration is not supported:



```
[edit class-of-service]
```

```
user@switch# set interfaces xe-0/0/24 unit 0 forwarding-class fcoe
```

## Traffic Management

- You cannot apply classifiers and rewrite rules to routed VLAN interfaces (RVIs) because the members of RVIs 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 RVIs.

### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

## Outstanding Issues in Junos OS Release 13.1 for the QFX Series

---

The following issues are outstanding in Junos OS Release 13.1. 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>.

- [Interfaces and Chassis](#)
- [Layer 2 and Layer 3 Features](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [QFabric System Administration](#)
- [User Interface and Configuration](#)

## Interfaces and Chassis

- On a QFabric system, the output of the **show ospf neighbor instance all neighbor** command displays information regarding all OSPF neighbors instead of displaying the information for only the specified neighbor. As a workaround, use the **show ospf neighbor neighbor** command. [PR/693044]
- On a QFabric system, the **clear ospf neighbor instance instance interface interface** command does not clear the neighbor connections for both the specified routing instance and interface. As a workaround, use the **clear ospf neighbor instance instance** command to clear the neighbor connections for a specific routing instance. Use the **clear ospf neighbor interface interface** command to clear the neighbor connections for a specific interface. [PR/693057]
- On a QFabric system, MAC learning on Node devices is distributed and communicated using the Border Gateway Protocol (BGP). The BGP update interval limits the rate at which a Node device can see MAC addresses learned on other Node devices. This limit is 30 MAC moves during a 5-second MAC move detection interval. If the configured MAC move limit is higher than 30, the higher move count is not detected. [PR/729499]
- On a QFabric system, when you change the time zone, the time zone change is propagated to many QFabric system components. Because of this, there might be a delay in propagating the time zone to all of the components. When you execute the first operation after a time zone change, the timestamp might display the old time zone instead of the new time zone. After the first incorrect display of the time zone, all subsequent commands display the timestamp correctly. [PR/805827]
- On a QFabric system, error messages about VLAN control set failures on network Node groups might be displayed. These messages do not indicate traffic interruptions or resource leaks and can be ignored. [PR/815902]
- On a QFabric system, if you perform a large-scale, multidimensional nonstop software upgrade, there might be some minor packet loss during the network Node group portion of the upgrade. [PR/823980]

## Layer 2 and Layer 3 Features

- On a QFabric system, the **clear ospf neighbor instance instance interface interface** command does not clear the neighbor connections for the specified routing instance and interface. Use the **clear ospf neighbor instance instance** command to clear neighbor connections for a specified routing instance. Use the **clear ospf neighbor interface interface** command to clear the neighbor connections for a specific interface. [PR/693057]
- On a QFabric system, after a reboot of the network Node group, a MAC address might be missing from the fabric routing protocol process (rpdf) and the kernel, but it is available in the Ethernet switching process (eswd). [PR/884070]
- On a QFabric system, the **show pim join extensive** command does not always display the downstream interfaces for multicast routes. The workaround is to issue the **show**

**multicast route** command to display the downstream interfaces for multicast routes. [PR/890214]

- On a QFabric system, if you take the front card of an Interconnect device offline, there might be multicast traffic loss of up to 318 seconds. [PR/894551]

## Network Management and Monitoring

- On a QFabric system, configuration of source addresses in outbound SNMP traps and TACACS+ authentication requests is not supported. The source address in the UDP header of an SNMP trap and the source address in the TCP header of a TACACS+ authentication request are each set to the default IP address that is bound by the socket. [PR/700194]
- On a QFabric system, during a rolling upgrade, the SNMP agent and related components move between Director devices repeatedly, the result is temporary outages in the SNMP service. Once the QFabric system upgrade is completed, normal SNMP service is resumed. [PR/792351]
- On a QFabric system, disabling a 10-Gigabit Ethernet (xe) interface generates the following SNMP traps: jnxVpnIfDown, jnxCmCfgChange, and linkDown. [PR/824799]
- On a QFabric system, the network Node group might generate bgpBackwardTransition traps continuously. [PR/828422]
- On a QFabric system, jnxVpnIfDown traps are generated when a routed VLAN interface (RVI) goes down. [PR/828384]
- On a QFabric system, the network Node group might generate bgpBackwardTransition traps continuously. [PR/828422]

## Platform and Infrastructure

- On a QFabric system, the output displayed for the **show arp expiration-time** operational command might not be properly formatted. The last column (TTE) might be misaligned so that the output might appear to be in the second-to-last column (Flags). [PR/737585]
- On a QFabric system, when you issue the **show fabric administration inventory director-group status** operational command, the LAG interface status might appear as up even when no members of the LAG are present (all members of the LAG are down). [PR/756970]

## QFabric System Administration

- On a QFabric system, autocompletion of interface names does not work in some cases. As a workaround, when pressing Tab to complete the name does not work, type in the entire interface name. [PR/564108]
- On a QFabric system, issuing the **telnet** command to a Director device is not supported. Instead, use the **ssh** command to access the Director device. [PR/588026]
- On a QFabric system, if you issue the **clear log filename** command in the QFabric CLI, the system removes only the contents of the log file stored in the Director group. To

remove the contents of the log file on an individual Interconnect or Node device, log in to the device with the **request component login** command and then issue the **clear log filename** command. [PR/674468]

- On a QFabric system, the **disable** statement for aggregated Ethernet interfaces might not be visible as a help option or available for autocompletion at the **[edit interfaces device:ae0 unit 0]** hierarchy level. However, even though it is hidden, you can still disable aggregated Ethernet interfaces by issuing the full **disable** statement at this same hierarchy level. [PR/685964]
- On a QFabric system, when you issue the **request component login** command to access one of the QFabric infrastructure Routing Engines, you might see incorrect names for the Routing Engines as you enter or exit the component. [PR/702622]
- On a QFabric system, you cannot issue an operational mode command to identify the master network Node group Routing Engine and the Director device that is hosting it. [PR/706037]
- On a QFX3000-M QFabric system, the output of the **show lacp interfaces** command displays messages like **lacp subsystem not running - not needed by configuration** for server Node groups and redundant server Node groups that do not have any aggregated Ethernet interfaces. [PR/723700]
- On a QFabric system, when you issue the **show bgp summary** command, the output might show an internal name for the BGP group rather than the external name. (The external name is the displayed name without the ---qfabric tag.) [PR/739290]
- On a QFabric system, if you modify statements at the **[edit interfaces interface-name unit logical-unit-number family ethernet-switching]** hierarchy level (for example, modifying a VLAN range), in some cases there might be a vague QFabric system CLI console error message that says **Check-out failed for Ethernet Switching Process (/usr/sbin/eswd) without details**. [PR/741268]
- On a QFabric system, the **| filter node-group node-group-name** option may not work for some operational commands even though it is available in the CLI. [PR/768821]
- On a QFabric system, the **disable** statement for the OSPF protocol is not visible as a CLI help option or available for autocompletion at the **[edit protocols ospf]** hierarchy level. However, even though it is hidden, you can disable OSPF by issuing the **disable** statement at this hierarchy level. [PR/771990]
- On a QFabric system, the output for the **show bgp neighbor ip-address** and **show policy policy-name** commands displays **qfabric** after the policy name in error. [PR/776390]
- On a QFabric system on which RSTP and MSTP are configured, if you change the port cost of an interface on a nonroot bridge and it results in changing a forwarding port to a blocked port (or vice versa), a temporary loop might be created. [PR/776762]
- On a QFabric system, the output of the **show interfaces routing** command does not display the Node group name of the interface as expected. [PR/777947]
- On a QFabric system, when you issue the **show arp** operational command from the Director device, the output shows only part of the interface name, not the complete interface name. (The interface name is truncated in the output.) [PR/778982]

- On a QFabric system, if you attempt to revise part of a Node device configuration by replacing the device serial number with an alias or vice versa, it might generate a commit error under the following conditions:
  - If you attempt to revise part of a Node device configuration with an alias when all the Node device and related configuration statements were previously configured with the device serial number.
  - If you attempt to remove an alias and instead use the serial number in a few places when all the Node device and related configuration statements were previously configured with the alias. [PR/786481]
- On a QFabric system, during a Director group nonstop software upgrade, if your CLI session ends as the hosting Director device upgrades, you reestablish the CLI session on the other Director device, and then suspend this second CLI session by typing the keyboard shortcut Ctrl+z, the CLI session prompt might not come back up. [PR/789926]
- On a QFabric system, the **show route summary** operational command displays internal names for the routing tables associated with user-defined routing instances. For example, the command displays an internal name such as **virt1---qfabric.inet.0** for the routing table associated with a user-defined routing instance named **virt1**. [PR/805045]
- On a QFabric system, when you issue the **request chassis routing-engine master switch interconnect-device interconnect-device-name** command, you might see the following warning: **Standby Routing Engine is not ready for graceful switchover. Command aborted. Not ready for mastership switch, try after 80 secs..** Even though the command operation should end at this point, you might still be prompted to proceed with the switchover even after the system displays the warning message. As a workaround, type **no** after the prompt, and wait to reissue the command until after the time specified in the warning message. [PR/817524]
- On a QFabric system, the **set date** command is not available. [PR/842466]
- On a QFabric system, during a network Node group switchover, Bridge Data Protocol Units (BPDUs) do not get cleared. [PR/856614]
- On a QFabric system, after a network Node group switchover, sometimes the routing protocol process (rpd) might stop operating on the network Node group master. [PR/878148]
- On a QFabric system, you can use the **track interface** keywords in the CLI to see the operational state of an interface. The priority of a VRRP group is tied to the operational state of an interface. If an interface goes down, the priority of a VRRP group can be reduced and might trigger a VRRP mastership change. On redundant server Node groups and server Node groups, the priority of the VRRP group is not reduced as expected when an aggregated Ethernet interface is operationally down. [PR/882628]

## User Interface and Configuration

- On a QFabric system, if you configure a firewall filter in the **[edit family ethernet-switching]** hierarchy, the QFabric CLI might erroneously allow you to apply the filter to a Layer 3 interface. [PR/677381]
- On a QFabric system, if a route is learned from an eBGP neighbor that is multiple hops from the QFabric system, and the same route is learned from other eBGP routers and multipath is enabled, the routes are not reachable from a server Node device or from a redundant server Node group. The routes remain in the network Node group routing table. Traffic on the network Node group devices can reach the destination route, but traffic on the server Node device or the server Node group cannot reach the destination route. This is the expected system behavior. [PR/682836]
- On a QFabric system, if you establish multiple long-running configuration edit sessions and then issue commands that force a configuration synchronization (such as **show**, **configuration**, **update**, and **commit**) and issue a **commit** command with no changes, you might see an error message that says **warning: no private changes to commit**, and your new active configuration might include some changes made by other users. [PR/736596]
- On a QFabric system, issuing the **show arp** command displays all ARP entries on the system, including those learned on the bme interface of internal components. As a workaround, use the **show arp | except bme** command to filter ARP entries for internal components. [PR/745642]

### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

---

## Resolved Issues in Junos OS Release 13.1 for the QFX Series

The following issues have been resolved in Junos OS Release 13.1 specified maintenance releases listed in this topic.

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>.



**NOTE:** Some issues that apply to EX Series switches may apply to the QFX Series as well. If you are looking for a resolved issue but cannot locate it in this section, see the .

- [Issues Resolved in Release 13.1X50-D35 on page 31](#)
- [Issues Resolved in Release 13.1X50-D30 on page 31](#)
- [Issues Resolved in Release 13.1 X50-D25 on page 34](#)
- [Issues Resolved in Release 13.1 X50-D20 on page 34](#)
- [Issues Resolved in Release 13.1 X50-D15 on page 35](#)

## Issues Resolved in Release 13.1X50-D35

The following issues have been resolved since Junos OS Release 13.1X50-D35. The identifier following the description is the tracking number in our bug database.

### Interfaces and Chassis

- On a QFabric system, if you insert or remove new QSFP optics in a QFX3600 Node device for ports 2 through 7, existing XLE interfaces in ports 8 through 15 might become detached. As a workaround, remove and reinsert the optics for the detached ports. [PR/1060463]

## Issues Resolved in Release 13.1X50-D30

The following issues have been resolved since Junos OS Release 13.1X50-D30. The identifier following the description is the tracking number in our bug database.

### Interfaces and Chassis

- On a QFabric system, sending a **ping** command to an IPv6 address results in the operation terminating after only one packet has been sent, unless the number of packets to send has been specified in the command. [PR/815901]
- On a QFabric system, the traceroute CLI uses the lowest available IP address as its first hop when tracing an IP address behind a server Node group or redundant server Node group. [PR/1018939]
- On a QFabric system, mirroring of LAG interface traffic on Node devices of a network Node group might continue even after the analyzer configuration is disabled. [PR/819942]
- On a QFabric system, if you enable fabric multicast logs, you might experience traffic loss for up to 15 seconds. [PR/813875]

### Multiprotocol Label Switching (MPLS)

---

- Receiving a crafted or malformed Resource Reservation Protocol (RSVP) packet might cause the RPD (routing protocol daemon) to hang or crash. When RPD is unavailable, routing updates cannot be processed, which can lead to an extended network outage. This issue only occurs during the processing of RSVP PATH messages. If RSVP is not enabled on an interface, then this issue is not applicable and cannot be triggered via that interface. Refer to JSA10652 for more information. [PR/954508, PR/945509]

### Platform and Infrastructure

---

- Certain combinations of Junos OS CLI commands and arguments can be exploited to allow root access to the operating system. This allows any user with permissions to run these CLI commands the ability to achieve elevated privileges and gain complete control of the device. Refer to JSA 10608 and JSA10634 for more information. [PR/912707, PR/913328, PR/913449, PR/913831, PR/915313, PR/915957, PR/915961, PR/921219, PR/921499, PR/961397, PR/962834, PR/965758]
- For an established Transmission Control Protocol (TCP) session, TCP input validation only ensures that sequence numbers are within the acceptable window prior to examining whether the SYN flag is set on the segment. If the SYN flag is set, the TCP stack drops the session and sends a TCP reset (RST) segment to the other side. Because the SYN flag only needs to fall within the window, an attacker who can guess an in-window sequence number, source and destination address, and port numbers can exploit this vulnerability to reset any established TCP session. Refer to JSA10638 for more information. [PR/935125]
- Traffic between the routing engine (RE) and transit interfaces is carried over an internal network between the packet forwarding engines (PFEs) and REs. Some REs use management (**em**) interfaces (usually, **em0** and **em1**) to connect to this network. Receiving a carefully crafted set of fragmented packets, destined to the router, can cause the management interface driver to become permanently blocked when trying to formulate a reply. This causes the RE to be unable to communicate over the private network that connects the FPCs and REs, eventually causing all FPCs to go offline and stay offline. Systems with redundant REs will failover, and are subject to the same issue. For systems without modular FPCs (for example, the MX80), the FPC will reboot and clear the management interface output queue. However, additional crafted fragments will cause the issue to reoccur. This issue is applicable to IPv4, IPv6, and Connectionless-mode Network Protocol (CLNP) fragmentation and reassembly scenarios. Transit traffic does not trigger this issue. Additionally, CLNP is only vulnerable if Connectionless-mode Network Service (CLNS)-routing or End System-to-Intermediate System (ES-IS) is explicitly configured. Refer to JSA10655 for more information. [PR/942437]
- When a RADIUS authentication server is configured at the **[edit system radius-server]** hierarchy, an entry is created in the **/var/etc/pam\_radius.conf** directory. When RADIUS accounting servers are configured at the **[edit system accounting destination radius]** hierarchy, the accounting servers are also propagated to **/var/etc/pam\_radius.conf** directory. Authentication requests sent to the RADIUS accounting server might result in unintended successful authentication. If the same



RADIUS server is used for both authentication and accounting, which is a common configuration, the issue is less severe because RADIUS authentication is sent to the intended server despite the duplicate entries. However, if the RADIUS authentication server is later removed from the configuration, the duplicate entry created by configuring the RADIUS accounting server will remain in `/var/etc/pam_radius.conf`, which can result in possible unintended authentication success. Refer to JSA10654 for more information. [PR/947307]

- On a QFabric system, when a Virtual Router Redundancy Protocol (VRRP) backup member is configured, VRRP might flap while receiving VRRP packets from a redundant server Node group or server Node group. Also, the kernel might crash when a network Node group prefers the Virtual IP (VIP) route of the redundant server Node group or server Node group if VRRP is configured with the **accept-data** option. [PR/1019046]

### QFabric Systems

- On a QFabric system, a Node device that is being rebooted is displayed as **disconnected** for 5 to 6 seconds in the output of the **show fabric administration inventory node-devices** command before its status changes to **connected**.
- On a QFabric system with firewall filters configured, the firewall filter might not be pushed to Node devices when all of the first terms in the firewall filter are configured with **interfaces** match conditions. [PR/962915]
- On a QFabric system, when you perform a non-stop software upgrade (NSSU), the upgrade might fail on the backup Director group. [PR/1015286]
- On a QFabric system, when Multiple Spanning Tree Protocol (MSTP) is configured, server Node groups and redundant server Node groups might have VLANs and RVIs configured even though there are no local VLAN members. [PR/1030244]

### Routing Protocols

- Multiple vulnerabilities have been reported against earlier versions of OpenSSL included within Junos OS. Refer to JSA10575 for more information. [PR/853724]
- In a router link-state advertisement (LSA), the **lsa-id** and the **advertising-router-id** must be the same. If they are different, Open Shortest Path First (OSPF) routes might be purged in a neighbor. Please refer to JSA10582 (PSN-2013-08-987) for more information. [PR/878639]
- Under specific time-sensitive circumstances, if Border Gateway Protocol (BGP) determines that an update is too large to be sent to a peer, and immediately attempts to send a withdraw message, the routing protocol daemon (RPD) might crash. An example of an oversized BGP UPDATE is one in which a very long AS\_PATH causes the packet to exceed the maximum BGP message size (4096 bytes). Using a very large number of BGP communities can also be used to exceed the maximum BGP message size. Please refer to JSA10609 for additional information. [PR/918734]
- When a Border Gateway Protocol (BGP) update contains a specifically crafted set of transitive attributes, memory might be corrupted and cause the routing protocol daemon (RPD) to crash and restart. Crashes were only observed through in-house routing protocol fuzz testing. This issue only affects routers supporting 4-byte autonomous

system (AS) numbers, which were introduced starting with Junos OS Release 9.1. Additionally, the router is only vulnerable if the BGP peer does not support 4-byte AS numbers. Refer to JSA10653 for more information. [PR/953037]

- On a QFX Series switch, the routing protocol daemon (RPD) scheduler might slip because of a memory allocation issue. This might occur when routing entries are updated. [PR/998911]
- The OpenSSL project released a security advisory on 2014-08-06 that contained nine security issues. The following four issues affect Junos OS: See JSA10649 for more information.
  - CVE-2014-5139: Crash with SRP ciphersuite in Server Hello message
  - CVE-2014-3509: Race condition in ssl\_parse\_serverhello\_tlsext
  - CVE-2014-3511: OpenSSL TLS protocol downgrade attack
  - CVE-2014-3512: SRP buffer overrun [PR/1016458]

## Issues Resolved in Release 13.1 X50-D25

The following issues have been resolved since Junos OS Release 13.1X50-D20. The identifier following the description is the tracking number in our bug database.

### Layer 2 and Layer 3

---

- On a QFabric system, if an analyzer output port is present on a Node device, and you clear the MAC table (but the ARP entry still exists), then any Layer 3 traffic on that Node device is also flooded to the analyzer output port. [PR/695171]

## Issues Resolved in Release 13.1 X50-D20

The following issues have been resolved since Junos OS Release 13.1X50-D15. The identifier following the description is the tracking number in our bug database.

### Interfaces and Chassis

---

- On a QFabric system, if you perform a *fabric* nonstop software upgrade on a system has a network Node group that contains no Node devices, the upgrade might cause Layer 2 traffic loss lasting for 6 to 7 seconds. [PR/808605]

### Layer 2 and Layer 3

---

- On a QFabric system, during a power cycle of the network Node group, rpdf (fabric routing protocol process) might crash. The system recovers after the crash. [PR/894521]
- On a QFabric system, performing a nonstop software upgrade on a redundant server Node group with the multicast make-before-break feature disabled, Layer 3 multicast traffic is duplicated for up to 35 seconds, and traffic is dropped for up to 10 seconds. [PR/899250]
- On a QFabric system, performing a nonstop software upgrade on a redundant server Node group with the multicast make-before-break feature enabled, Layer 3 multicast

traffic is duplicated for up to 35 seconds, and traffic is dropped for up to 3 minutes. [PR/899274]

### QFabric System Administration

- On a QFabric system, if LACP is enabled on one Node device but not on another Node device, this message appears: **warning: lACP subsystem not running - not needed by configuration then this message is seen on the Fabric**. The message is displayed in error, and you can ignore it. [PR/743032]
- On a QFabric system, you might be able to commit a configuration for nonstop software upgrade groups that contain fabric-based components (such as Interconnect devices, and fabric control or fabric manager Routing Engines) in error. As a workaround, remove any configuration that lists fabric-based components in an upgrade group and only configure upgrade groups that contain Node devices associated with the network Node group NW-NG-0. [PR/756511]
- On a QFabric system, the **show configuration** command might display incorrect version information. To see the correct version number, use the **show version** command instead. [PR/793885]
- On a QFabric system, if the system has a routed VLAN interface (RVI) configured for the network Node group, a server Node group has a VLAN that connects to the RVI, and the network Node group reboots, a kernel panic on the server Node group might occur. [PR/859968]
- On a QFabric system, assigning IPv6 addresses to a QFabric management network is not currently supported, but will be available in a future software release. [PR/911193]

### Storage and Fibre Channel

- On a QFabric system, if you reboot front cards on an Interconnect device, FCoE traffic might be dropped. [PR/883705]

## Issues Resolved in Release 13.1 X50-D15

The following issues have been resolved since Junos OS Release 13.1X50-D10. The identifier following the description is the tracking number in our bug database.

### Interfaces

- On a QFabric system, before you perform a recovery installation of the QFabric system using the **recovery\_configuration** file on the USB install media, you must validate the **recovery\_configuration** file manually. The QFabric system does not validate the **recovery\_configuration** file during installation. [PR/810242; PR/810244: This issue has been resolved.]
- On a QFabric system, the initial configuration script allows multicast and broadcast MAC addresses to be configured. Configuration of multicast and broadcast MAC addresses should not be allowed. [PR/811032: This issue has been resolved.]

### Layer 2 and Layer 3

---

- On a QFabric system, activating the OSPF protocol and removing a default static route that points to an RVI which does not have a network Node group membership, creates vmcore and ksyncd core files. This issue occurs when:
  - Traffic is running in the QFabric system that has a static default route pointing to an RVI interface which has no network Node group membership, with OSPF deactivated.
  - You clear the system ARP table.
  - You delete the static default route that points to the RVI interface, activate OSPF protocol, and commit the new configuration.
  - Clearing the ARP table and deleting the static default route creates the vmcore and ksyncd core files. [PR/821778: This issue has been resolved.]
- On a QFabric system, if you issue the **show fabric logging** command on a redundant server Node group, the fabric routing protocol process (rpdf) might stop operating and generate a core file. [PR/848664: This issue has been resolved.]
- On a QFabric system, with scaled configurations, deleting protocol, interface, and VLAN configurations, and then issuing a rollback, causes traffic loss. In this scenario, the system does not recover. Traffic loss occurs because 1) Layer 2 multicast entries are programmed immediately, and the hardware scale limit is met, and 2) Layer 2 multicast entries are not deleted when the PIM and routed VLAN instance configurations are rolled back. As a result, there are no hardware entries available to program Layer 3 multicast entries. [PR/899659: This issue has been resolved.]
- On a QFabric system, during a network Node group switchover, there might be permanent traffic loss on multicast groups with scaled multicast routes. The workaround is to restart the fabric daemon on the network Node group. [PR/900485: This issue has been resolved.]

### QFabric Systems Administration

---

- On a QFabric system, in very rare cases after you install QFabric software from a USB drive, the system might fail to boot and display the following message: **switchroot: mount failed: No such file or directory**. As a workaround, restart the affected Director device. [PR/664511: This issue has been resolved.]
- On a QFabric system, if a Director group loses power unexpectedly, the shared storage might become corrupted. If this happens, the Director group devices perform a file-system check when they reboot, and system services do not start until the check finishes (which might take several hours). As a workaround, if your Director group loses power unexpectedly, restart one Director device at a time and allow the device to finish booting before restarting the other device. [PR/674625: This issue has been resolved.]
- On a QFabric system, if you issue the **show | compare | display inheritance** operational mode command, the output might display an error message. [PR/706924: This issue has been resolved.]

- On a QFabric system, if you perform a large scale, multidimensional nonstop software upgrade, there might be some minor packet loss during the Director group portion of the upgrade. [PR/823978: This issue has been resolved.]
- On a QFabric system, if you enable and disable routed VLAN interface (RVI) interfaces multiple times, the sFlow technology process might generate a core file. [PR/869066: This issue has been resolved.]
- On a QFabric system, if there is a core dump of the kernel synchronization process (ksyncd) and a replication error for the backup network Node group, you might see the following system log message: **ksyncd[1729]: KSYNCD: resync error, issu\_state[0], type ROUTE subtype 18: Protocol not available** [PR/876576: This issue has been resolved.]
- On a QFabric system, if you issue a command to reboot all components of the network Node group (master, backup and all its line cards), when the network Node group comes back online, a few MAC addresses might be missing. As a workaround, issue the **clear bgp neighbor fabric** command to restore the proper MAC address tab. [PR/892533: This issue has been resolved.]
- On a QFabric system, the upgrade to Junos OS Release 12.2X50-D40 from an earlier Junos OS release might fail while the system is migrating syslog data from one release to the other. As a workaround, clear the archived syslog data in the database of the current release before you upgrade to Junos OS Release 12.2X50-D40. [PR/860982: This issue has been resolved.]

### Storage and Fibre Channel

- On a QFabric system, if you configure a classifier at the Node device ingress port that marks traffic for queue 3 (default fcoe forwarding class queue), or queue 4 (default no-loss forwarding class queue) with a loss priority of medium-high (yellow), then on fabric (fte) ports, queues 3 and 4 might experience packet loss. By default, queues 3 and 4 transmit lossless traffic.

Because the loss priority is not significant for lossless traffic, this issue does not arise when you use queues 3 and 4 for lossless traffic. If you use queues 3 and 4 for lossy traffic, you might experience this issue. You can control the drop pattern on access (xe and xle) interfaces by applying a drop profile (weighted random early detection (WRED) drop profile) for medium-high priority traffic. [PR/898840: This issue has been resolved.]

### Traffic Management

- On a QFabric System, if and only if there is no class-of-service configuration, the **show class-of-service forwarding-table scheduler-map** operational command shows some queue numbers incorrectly. The incorrect output has no functional impact and does not affect the default queue assignments. The issue is only that the **show** command displays incorrect queue numbers. [PR/697075: This issue has been resolved.]

#### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)

- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

---

## Errata in Documentation for Junos OS Release 13.1 for the QFX Series

---

This section lists outstanding issues with the documentation.

### Routing Protocols

- On the QFabric system, documentation for the **ecmp-do-local-lookup** statement at the **[edit forwarding options]** hierarchy is not documented in Junos OS Release 13.1X50-D10, Junos OS Release 13.1X50-D15, Junos OS Release 13.1X50-D20, Junos OS Release 13.1X50-D25.

#### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series on page 38](#)

---

## Upgrade and Downgrade Instructions for Junos OS Release 13.1 for the QFX Series

---

This section discusses the following topics:

- [Procedure for Upgrading CoS from Junos OS Release 11.1 or Release 11.2 to Release 11.3 or Later on page 39](#)
- [CoS Upgrade Requirements from Junos OS Release 12.1 to Junos OS Release 12.2 on page 40](#)
- [Basic Procedure for Downloading Junos OS Release 13.1 Software on page 41](#)
- [Performing a Nonstop Software Upgrade on the QFabric System on page 42](#)

- [Performing a Standard Software Upgrade on the QFabric System on page 43](#)
- [Upgrade and Downgrade Support Policy for Junos OS Extended End-of-Life Software Releases on page 43](#)

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

Before you upgrade to Junos OS Release 11.3 or later, you must deactivate the CoS configuration on the QFX3500 switch and QFX3600 switch and the QFabric system if the CoS configuration uses the **excess-rate** option, **strict-high** or **high** priority queues, or any of the default multidestination forwarding classes. For full information about this topic, see [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)](#). A summary of the upgrade steps is included here.

After you upgrade to Junos OS Release 11.3 or later, modify the CoS configuration on the QFX3500 switch and QFX3600 switch and the QFabric system to conform to the Junos OS Release 11.3 or later CoS requirements. Then activate the CoS configuration and commit the changes:

1. Deactivate the CoS configuration.

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

2. Upgrade to Junos OS Release 11.3 or later.

3. Make the following changes to the CoS configuration:

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

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

5. Activate the CoS configuration.

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

6. Commit the CoS configuration.



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

## CoS Upgrade Requirements from Junos OS Release 12.1 to Junos OS Release 12.2

Before you upgrade to Junos OS Release 12.2, you might need to edit the class-of-service (CoS) configuration because the way the QFX3500 and QFX3600 switches and the QFabric system handle lossless forwarding classes has changed in Junos OS Release 12.2.

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

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



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

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

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

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

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

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



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



forwarding classes is not treated as lossless traffic until you remove the explicit forwarding class configuration.

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

- To delete the explicit **fcoe** forwarding class configuration:

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

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

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



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

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

## Basic Procedure for Downloading Junos OS Release 13.1 Software

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

- Using a Web browser, navigate to the <http://www.juniper.net/support>.
- Click the **By Task** tab.
- Click **Download Software**.
- In the Switching box, click **Junos Platforms**.
- In the QFX Series section, click the name of the platform for which you want to download software.
- Click the **Software** tab and select **13.1** from the Release drop-down list.
- In the Install Package section of the Software tab, select the **Install Package** for the 13.1X50-D20.1 release.

A login screen appears.

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

9. Download the software to a local host.
10. Copy the software to the device or to your internal software distribution site.

## Performing a Nonstop Software Upgrade on the QFabric System

---



**NOTE:** Before you can perform a nonstop software upgrade to Junos OS Release 13.1X50-D10, you must have Junos OS Release 12.2X50-D42 or later installed. You cannot perform a nonstop software upgrade with Junos OS Release 12.2X50-D41 or earlier. Contact the Juniper Technical Assistance Center for information on how to download Junos OS Release 12.2X50-D42. Performing a standard software upgrade (that is, issuing the `request system software add component all` command) does not require that you upgrade to an intermediate Junos OS software release.

To perform a nonstop software upgrade to Junos OS Release 13.1X50-D10:

1. First perform a standard software upgrade to Junos OS Release 12.2X50-D42 or later.
2. Then perform a nonstop software upgrade to Junos OS Release 13.1X50-D10.

To perform a nonstop software upgrade from Junos OS Release 13.1X50-D10 to Junos OS Release 13.1X50-D15, issue the `request system software nonstop-upgrade` command.

---

Nonstop software upgrade enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. Nonstop software upgrade is the recommended method for upgrading the QFabric system.

The QFabric system software package contains software for all of the different devices in the QFabric system infrastructure, Director group, Interconnect devices, and Node devices.

To download the software package to various locations—for example, USB memory stick, remote server, or FTP site—issue the `request system software download /path/package-name` command.

Additionally, you can back up your current QFabric system configuration file and installation-specific parameters using the `request system software configuration-backup path` command. Although you can save this file locally, we recommend that you save it to an external location, like an FTP site or USB device.

The following CLI commands enable you to install the software for the Director group, fabric control Routing Engines, fabric manager Routing Engine, Interconnect devices, and the network and server Node groups.

- `request system software nonstop-upgrade director-group package-name`—This command installs software for the Director devices in a Director group.

- **request system software nonstop-upgrade fabric *package-name***—This command installs the software for the Interconnect device and other components in the fabric.
- **request system software nonstop-upgrade node-group *node-group-name package-name***—This command installs software for a redundant server Node group or a network Node group.

QFabric system components are rebooted automatically as part of the nonstop upgrade process.

## Performing a Standard Software Upgrade on the QFabric System

The QFabric system software package contains software for all of the different devices in the QFabric system infrastructure, Director group, Interconnect devices, and Node devices.

To download the software package to various locations—for example, USB memory stick, remote server, or FTP site—issue the **request system software download */path/package-name*** command.

Additionally, you can back up your current QFabric system configuration file and installation-specific parameters using the **request system software configuration-backup** command. Although you can save this file locally, we recommend that you save it to an external location, like an FTP site or USB device.

To install the software package for the Director group, fabric control Routing Engines, fabric manager Routing Engine, Interconnect devices, and the network and server Node groups, and reboot the QFabric system, issue the **request system software add *package-name* component all reboot** command.

## Upgrade and Downgrade Support Policy for Junos OS Extended End-of-Life Software Releases

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

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

### Related Documentation

- [New Features in Junos OS Release 13.1X50-D20 for the QFX Series on page 3](#)
- [New Features in Junos OS Release 13.1X50-D15 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 13.1X50-D10 for the QFX Series on page 5](#)
- [Limitations in Junos OS Release 13.1 for the QFX Series on page 20](#)

- [Changes in Default Behavior and Syntax in Junos OS Release 13.1 for the QFX Series on page 16](#)
- [Outstanding Issues in Junos OS Release 13.1 for the QFX Series on page 25](#)
- [Resolved Issues in Junos OS Release 13.1 for the QFX Series on page 30](#)
- [Errata in Documentation for Junos OS Release 13.1 for the QFX Series on page 38](#)

## QFX Series Documentation for Junos OS Release 13.1

Title	Description
<i>QFX3000-G QFabric System Hardware Documentation</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for QFX3000-G QFabric systems
<i>QFX3000-M QFabric System Hardware Documentation</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for QFX3000-M QFabric systems
<i>QFX3500 Device Hardware Documentation</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for QFX3500 devices
<i>QFX3600 Device Hardware Documentation</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for QFX3600 devices
<i>QFX3000-G QFabric System Deployment Guide</i>	Hardware and software information for deploying QFX3000-G QFabric systems
<i>QFX3000-M QFabric System Deployment Guide</i>	Hardware and software information for deploying QFX3000-M QFabric systems
<i>Complete Software Guide for Junos OS for the QFX Series, Release 13.1</i>	Software feature descriptions, configuration examples, and tasks for Junos OS for the QFX Series
<i>Junos OS Software Release Notes for the Juniper Networks QFX Series, Release 13.1X50-D25 (this document)</i>	Summary of hardware and software features, and known problems with the software and hardware

## Requesting Support

For technical support, open a support case with the Case Manager link at <http://www.juniper.net/customers/support/>, email the technical assistance center (TAC) at [support@juniper.net](mailto:support@juniper.net), or call 1-888-314-JTAC (from the United States, Canada, or Mexico) or 1-408-745-9500 (from elsewhere).

## Revision History

6 April 2015—Revision 9, Junos OS for the QFX Series, Release 13.1X50-D35

16 March 2015—Revision 8, Junos OS for the QFX Series, Release 13.1X50-D35  
31 October 2014—Revision 7, Junos OS for the QFX Series, Release 13.1X50-D30  
1 October 2014—Revision 6, Junos OS for the QFX Series, Release 13.1X50-D20  
4 August 2014—Revision 5, Junos OS for the QFX Series, Release 13.1X50-D25  
27 May 2014—Revision 4, Junos OS for the QFX Series, Release 13.1X50-D20  
7 January 2014—Revision 3, Junos OS for the QFX Series, Release 13.1X50-D15  
19 November 2013—Revision 2, Junos OS for the QFX Series, Release 13.1X50-D15  
25 July 2013—Revision 1, Junos OS for the QFX Series, Release 13.1X50-D10

Copyright © 2015, 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.