

Junos[®] OS Software Release Notes for the Juniper Networks QFX Series

Release 12.2X50-D70
10 April 2015
Revision 15

Contents

New Features in Junos OS Release 12.2X50-D70 for the QFX Series	4
New Features	4
New Features in Junos OS Release 12.2X50-D65 for the QFX Series	4
New Features	4
New Features in Junos OS Release 12.2X50-D60 for the QFX Series	5
New Features	5
New Features in Junos OS Release 12.2X50-D55 for the QFX Series	5
New Features	5
New Features in Junos OS Release 12.2X50-D50 for the QFX Series	6
New Features	6
New Features in Junos OS Release 12.2X50-D40 for the QFX Series	6
QFabric Systems	7
New Features in Junos OS Release 12.2X50-D30 for the QFX Series	7
QFabric Systems	7
New Features in Junos OS Release 12.2X50-D20 for the QFX Series	8
Hardware	8
Interfaces	8
Layer 2 Protocols	9
Network Management and Monitoring	10
QFabric Systems	11
Security	11
New Features in Junos OS Release 12.2X50-D10 for the QFX Series	14
Hardware	14
High Availability	15
Interfaces	17
IPv6 Features	18
Layer 2 Protocols	18
Layer 3 Protocols	21
Network Management and Monitoring	21
QFabric Systems	23

Security	24
Storage and Fibre Channel	25
System Management	25
Traffic Management	25
Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series	
Series	27
Class of Service	27
Link Layer Discovery Protocol (LLDP)	27
QFabric Systems	28
Limitations in Junos OS Release 12.2 for the QFX Series	29
Interfaces and Chassis	29
Layer 2 Protocols	30
Network Management	30
QFabric Systems	30
Storage	31
Traffic Management	32
Outstanding Issues in Junos OS Release 12.2 for the QFX Series	32
Ethernet Switching	32
Hardware	33
Interfaces and Chassis	33
Layer 2 Features	35
Network Management and Monitoring	36
Platform and Infrastructure	37
QFabric Systems	37
Routing Protocols	42
Storage and Fibre Channel	43
User Interface and Configuration	43
Resolved Issues in Junos OS Release 12.2 for the QFX Series	44
Issues Resolved in Release 12.2X50-D70	44
Interfaces and Chassis	44
Platform and Infrastructure	45
Spanning Tree Protocols	45
Issues Resolved in Release 12.2X50-D65	45
Infrastructure	45
Software Installation and Upgrade	46
Issues Resolved in Release 12.2X50-D60	46
Issues Resolved in Release 12.2X50-D55	46
Interfaces and Chassis	46
QFabric Systems	46
Issues Resolved in Release 12.2X50-D50	46
Interfaces and Chassis	46
QFabric Systems	47
Security	48
Issues Resolved in Release 12.2X50-D40	48
Interfaces and Chassis	48
Network Management and Monitoring	48
QFabric Systems	49
Routing Protocols	50

Issues Resolved in Release 12.2X50-D30	50
Interfaces and Chassis	50
QFabric Systems	51
Issues Resolved in Release 12.2X50-D20	51
Interfaces and Chassis	51
QFabric Systems	52
Errata in Documentation for Junos OS Release 12.2 for the QFX Series	53
QFabric Systems	53
System Log Messages	54
Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series	54
Procedure for Upgrading CoS from Junos OS Release 11.1 or Release 11.2 to Release 11.3 or Later	54
CoS Upgrade Requirements from Junos OS Release 12.1 to Junos OS Release 12.2	55
Upgrading Software on QFX3500 and QFX3600 Standalone Switches	57
Upgrading Software on QFabric Systems	59
Performing a Nonstop Software Upgrade on the QFabric System	60
Performing a Standard Upgrade on the QFabric System	63
Upgrade and Downgrade Support Policy for Junos OS Extended End-of-Life Software Releases	66
QFX Series Documentation for Junos OS Release 12.2	67
Requesting Support	68
Revision History	68

New Features in Junos OS Release 12.2X50-D70 for the QFX Series

- [New Features on page 4](#)

New Features

There are no new features in Junos OS Release 12.2X50-D70.

Related Documentation

- [New Features in Junos OS Release 12.2X50-D65 for the QFX Series on page 4](#)
- [New Features in Junos OS Release 12.2X50-D60 for the QFX Series on page 5](#)
- [New Features in Junos OS Release 12.2X50-D55 for the QFX Series on page 5](#)
- [New Features in Junos OS Release 12.2X50-D50 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D40 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D30 for the QFX Series on page 7](#)
- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D65 for the QFX Series

- [New Features on page 4](#)

New Features

There are no new features in Junos OS Release 12.2X50-D65.

Related Documentation

- [New Features in Junos OS Release 12.2X50-D60 for the QFX Series on page 5](#)
- [New Features in Junos OS Release 12.2X50-D55 for the QFX Series on page 5](#)
- [New Features in Junos OS Release 12.2X50-D50 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D40 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D30 for the QFX Series on page 7](#)
- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)

- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D60 for the QFX Series

- [New Features on page 5](#)

New Features

There are no new features in Junos OS Release 12.2X50-D60.

Related Documentation

- [New Features in Junos OS Release 12.2X50-D55 for the QFX Series on page 5](#)
- [New Features in Junos OS Release 12.2X50-D50 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D40 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D30 for the QFX Series on page 7](#)
- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D55 for the QFX Series

- [New Features on page 5](#)

New Features

There are no new features in Junos OS Release 12.2X50-D55.

**Related
Documentation**

- [New Features in Junos OS Release 12.2X50-D50 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D40 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D30 for the QFX Series on page 7](#)
- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D50 for the QFX Series

- [New Features on page 6](#)

New Features

There are no new features in Junos OS Release 12.2X50-D50.

**Related
Documentation**

- [New Features in Junos OS Release 12.2X50-D40 for the QFX Series on page 6](#)
- [New Features in Junos OS Release 12.2X50-D30 for the QFX Series on page 7](#)
- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D40 for the QFX Series

- [QFabric Systems on page 7](#)

QFabric Systems

- **Support for 40-Gigabit Ethernet extended short-reach QSFP+ transceiver (QFabric systems)**—On QFX3500 and QFX3600 Node devices, you can use the QFX-QSFP-40G-ESR4 transceiver to extend the maximum distance from the Node device to the Interconnect device. The QFX-QSFP-40G-ESR4 transceiver supports distances up to 300 m (984 ft) over OM3-grade fiber, or 400 m (1312 ft) over OM4-grade fiber. However, for lossless traffic, the QFX-QSFP-40G-ESR4 transceiver supports distances up to 150 m (492 ft) over OM3-grade fiber and OM4-grade fiber. For guidelines on lossless transport support, go to the following URL:

http://www.juniper.net/techpubs/en_US/junos12.2/topics/concept/cos-qfx-series-forwarding-classes-understanding.html .

Related Documentation

- [New Features in Junos OS Release 12.2X50-D30 for the QFX Series on page 7](#)
- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D30 for the QFX Series

- [QFabric Systems on page 7](#)

QFabric Systems

- **Support for third-party transceivers (QFX Series)**—Allows you to use fiber-optic transceivers made by other vendors with QFX3500 and QFX3600 standalone switches and QFabric systems. The requirement to use Juniper Networks transceivers exclusively has been removed. This feature is enabled by default, and no configuration is required to use this functionality.
- **reboot command extended to Director devices (QFabric systems)**—You can now reboot a Director device in a QFabric system. To reboot a Director device, issue the `request system reboot director-device device-name` command.

Related Documentation

- [New Features in Junos OS Release 12.2X50-D20 for the QFX Series on page 8](#)
- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)

- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

[New Features in Junos OS Release 12.2X50-D20 for the QFX Series](#)

- [Hardware on page 8](#)
- [Interfaces on page 8](#)
- [Layer 2 Protocols on page 9](#)
- [Network Management and Monitoring on page 10](#)
- [QFabric Systems on page 11](#)
- [Security on page 11](#)

Hardware

- **QFX3600 switch**—The Juniper Networks QFX3600 Switch is a high-speed, multipurpose switch especially designed for next-generation data centers that provides a total switching capacity and throughput of 640 Gbps. The small form-factor and front-facing ports in the switch make it suitable for deployment in high-density server racks and container-based data center deployments.

Sixteen 40-Gbps ports in the switch use QSFP+ transceivers. By default, all 16 ports operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the 40-Gbps ports to operate as four 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.

The QFX3600 switch supports AC or DC power supplies, and is available with FRU-side-to-port-side airflow or port-side-to-FRU-side airflow.

The QFX3600 switch can be converted to a QFabric system Node device using the **request chassis device-mode node-device** operational mode command. Similarly, a QFX3600 Node device running Junos OS Release 12.2X50-D20 or later can be converted to a QFX3600 switch using the **request chassis device-mode standalone** operational mode command.

Interfaces

- **Support for configuring ports on QFX3600 standalone switches to operate as 10-Gigabit Ethernet ports**—On a QFX3600 standalone switch, you can configure ports Q0 through Q15 to operate as 10-Gigabit Ethernet (xe) ports. To configure 10-Gigabit

Ethernet ports, include the **xe port-range *port-range-low port-range-high*** statement at the **[edit chassis fpc 0 pic 0]** hierarchy level.



NOTE: Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports, because one port is reserved. Therefore, you can configure up to 63 (not 64) 10-Gigabit Ethernet ports on ports Q0 through Q15.



NOTE: By default, ports Q0 through Q15 on QFX3600 standalone switches operate as 40-Gigabit Ethernet (xle) ports.

- **Support for configuring ports on QFX3600 Node devices to operate as 40-Gigabit Ethernet ports**—On a QFX3600 Node device, you can configure ports Q2 through Q15 to operate as 40-Gigabit Ethernet (xle) ports. To configure 40-Gigabit Ethernet ports, include the **pic 1 xle port-range *port-range-low port-range-high*** statement at the **[edit chassis node-group name node-device name]** hierarchy level.



NOTE: By default, ports Q0 through Q3 on QFX3600 Node devices operate as 40-Gbps data plane uplink (fte) ports, and ports Q4 through Q15 operate as 10-Gigabit Ethernet (xe) ports.

- **Support for configuring ports on a QFX3500 standalone switch to operate as 40-Gigabit Ethernet ports**—On a QFX3500 standalone switch, you can configure ports Q0 through Q3 to operate as 40-Gigabit Ethernet (xle) ports. To configure a block of 40-Gigabit Ethernet ports, include the **xle port-range *port-range-low port-range-high*** statement at the **[edit chassis fpc 0 pic 2]** hierarchy level. To configure an individual 40-Gigabit Ethernet port, include the **xle port *port-number*** statement at the **[edit chassis fpc 0 pic 2]** hierarchy level.



NOTE: By default, the four 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.

Layer 2 Protocols

- **Change in default configuration for spanning-tree protocols internal to a QFabric system (QFabric systems)**—In Junos OS Release 11.3X30, the default configuration for the spanning tree used by internal QFabric system operations enabled a limited version of RSTP for loop prevention, and BPDU blocking on all ports if there is a physical loop. In Junos OS Release 12.2X50-D20, only BPDU blocking on all ports is enabled by default.

- To prevent loops internal to the QFabric system, enable RSTP by including the **rstp** statement at the **[edit protocols]** hierarchy level. By doing so, you enable full RSTP support on network Node group interfaces and loop prevention on server Node group interfaces.
- To disable BPDU blocking and allow BPDUs to traverse an interface on the network Node group, include the **disable** statement at the **edit ethernet-switching-options bpdu-block interface device-name:xe-fpc/pic/port** hierarchy level.
- To refresh interfaces placed in shutdown mode because of BPDU blocking and bring them back online, issue the **clear ethernet-switching bpdu-error interface interface-name** command.
- To stop sending RSTP BPDUs on a specific network Node group interface, include the **disable** statement at the **[edit protocols rstp interface device-name:xe-fpc/pic/port]** hierarchy level.
- To stop sending loop prevention BPDUs on a specific server Node group or redundant server Node group interface, include the **disable** statement at the **[edit protocols rstp interface device-name:xe-fpc/pic/port]** hierarchy level.

Network Management and Monitoring

- **Support for Juniper Networks enterprise-specific SNMP Fabric Chassis MIB (QFabric systems)**—Provides hardware information about the QFabric system and its component devices in a single MIB. The Fabric Chassis MIB is based on the Juniper Networks enterprise-specific Chassis MIB that provides information for individual devices. Unlike the Chassis MIB, the Fabric Chassis MIB represents the QFabric system component devices as part of the QFabric system. Only the information from the Fabric Chassis MIB (and not from individual Chassis MIBs) is available to SNMP management clients of the QFabric system. SNMP traps are generated from the Fabric Chassis MIB only and not from the Chassis MIB.



NOTE: Information about the Director group is not available in the Fabric Chassis MIB at this time.



NOTE: Information about the Virtual Chassis is not part of the Fabric Chassis MIB.

- **Enhanced configuration support for the default system log messages file (QFabric systems)**—Adds support for the configuration of the default system log file named **messages**, including archive size, and facility and severity levels. You configure the **messages** file archive size at the **[edit system syslog file messages archive]** hierarchy

level, and the *facility* and *severity* options at the **[edit system syslog file messages]** hierarchy level.

QFabric Systems

- **Enhancement to the QFabric USB drive recovery process (QFabric systems)**—Allows you to restore the initial setup settings and system configuration at the same time that you recover the system software. This enhancement saves the additional information on a second partition in the USB drive. To create a recovery USB drive that contains software, initial setup, and configuration settings:
 1. Create a software recovery USB drive by following the instructions in [Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive](#).
 2. In the QFabric default partition CLI, issue the **request system software configuration-backup** command and specify the recovery_configuration path on the shared drive of the Director group:


```
host@qfabric> request system software configuration-backup
/pbdata/recovery_configuration
```
 3. In the Director device CLI, mount a second partition from the USB drive to a mount location:


```
[root@dg0 ~]# mount /dev/sdb2 /mnt/usb/
```
 4. Move the saved backup file from the shared drive on the Director group to the mount location on the USB drive:


```
[root@dg0 ~]# mv /pbdata/recovery_configuration /mnt/usb/
```
 5. Unmount the USB drive in the Director device CLI, but leave the USB drive inserted in the Director device:


```
[root@dg0 ~]# umount /mnt/usb
```
 6. Reboot the Director device and enter **install** at the recovery menu prompt.
 7. After the recovery process, log in to the QFabric default partition CLI and resume normal QFabric system operations.

Security

- **Filter-based forwarding (QFX3500 and QFX3600 switches, QFabric systems)**—You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to which to send packets. You might want to use filter-based forwarding to route specific types of traffic through a firewall or other security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment.

- **Increased numbers of policers (QFX3500 and QFX3600 switches)**—You can configure and commit the following numbers of policers on QFX3500 and QFX3600 devices when they are operating as standalone switches:
 - Two-color policers used in ingress firewall filters: 768
 - Three-color policers used in ingress firewall filters: 768
 - Two-color policers used in egress firewall filters: 510
 - Three-color policers used in egress firewall filters: 255
- **Increased numbers of firewall filters (QFX3500 and QFX3600 switches)**—QFX3500 and QFX3600 switches and Node devices support 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 apply a total of 768 port filters, Layer 3 filters, and VLAN filters in the input direction and 1024 port filters, Layer 3 filters, and VLAN filters in the output direction. These maximum values also assume that each filter has only one term. If you create filters with multiple terms (including implicit terms), the maximum numbers of filters is reduced.



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 memory for 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 will only store 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.

All of the preceding principles also apply to egress filters, but four memory slices are used because IPv4 Layer 3 filters and IPv6 Layer 3 filters are stored in separate slices. The memory slices for egress filters are the same size as those for ingress filters, so the maximum number of egress filter terms is therefore 1024.

If you violate any of these restrictions and commit a configuration that is not in compliance, Junos OS rejects the excessive filters. For example, if you configure 300 ingress port filters and 300 ingress Layer 3 filters and try to commit the configuration, Junos OS does the following (again assuming one term per filter):

- Accepts the 300 ingress port filters (storing them in two memory slices).
- Accepts the first 256 ingress Layer 3 filters it processes (storing them in the third memory slice).
- Rejects the remaining 44 ingress Layer 3 filters.



NOTE: In this situation, be sure to delete excessive filters (for example, the remaining 44 ingress Layer 3 filters) from the configuration before you reboot the device. If you reboot a device that has a noncompliant configuration, you cannot predict which filters are installed after the reboot. Using the example above, the 44 ingress Layer 3 filters that were originally rejected might be installed, and 44 of the port filters that were originally accepted might be rejected.

Related Documentation

- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

New Features in Junos OS Release 12.2X50-D10 for the QFX Series

- [Hardware on page 14](#)
- [High Availability on page 15](#)
- [Interfaces on page 17](#)
- [IPv6 Features on page 18](#)
- [Layer 2 Protocols on page 18](#)
- [Layer 3 Protocols on page 21](#)
- [Network Management and Monitoring on page 21](#)
- [QFabric Systems on page 23](#)
- [Security on page 24](#)
- [Storage and Fibre Channel on page 25](#)
- [System Management on page 25](#)
- [Traffic Management on page 25](#)

Hardware

- **QFX3000-M QFabric system**—The QFX3000-M QFabric system supports up to 16 Node devices and 4 Interconnect devices, which provide a maximum of 768 10-Gigabit Ethernet ports at 3:1 oversubscription. The devices used to build a QFX3000-M QFabric system include the QFX3600-I Interconnect device, the QFX3100 Director device, the QFX3500 or QFX3600 Node device, and two EX4200 switches (for the control plane network). The QFX3000-M QFabric system supports two types of control plane network: copper-based and fiber-based. For a copper-based QFX3000-M QFabric system control plane network, use EX4200-24T switches with an SFP+ uplink module installed. For a fiber-based control plane network, use EX4200-24F switches with an SFP+ uplink module installed.
- **QFX3600-I Interconnect device**—QFabric system device that acts like a backplane for data plane traffic traversing the QFX3000-M QFabric system between Node devices. The QFX3600-I Interconnect device has 16 40-Gbps QSFP+ ports. To support copper-based and fiber-based control plane networks, the QFX3600-I Interconnect device has two RJ-45 management ports and two SFP management ports. The

QFX3600-I Interconnect device supports AC or DC power supplies, and is available with FRU-side-to-port-side airflow or port-side-to-FRU-side airflow.

- **QFX3600 Node device**—QFabric system device that connects either to endpoint systems (such as servers and storage devices) or to external networks. The QFX3600 Node device provides 16 40-Gbps QSFP+ ports. By default, 4 ports (labeled Q0 through Q3) are configured for 40-Gbps uplink connections between your Node device and the Interconnect devices, and 12 ports (labeled Q4 through Q15) use QSFP+ to four SFP+ copper breakout cables to support 48 10-Gigabit Ethernet interfaces for connections to either endpoint systems or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and the Interconnect devices, and ports Q2 through Q15 for 10-Gbps connections to either endpoint systems or external networks. To support copper-based and fiber-based control plane networks, the QFX3600 Node device has two RJ-45 management ports and two SFP management ports. The QFX3600 Node device supports AC or DC power supplies, and is available with FRU-side-to-port-side airflow or port-side-to-FRU-side airflow.
- **QFabric system name changes**—Implements the following name changes:
 - The QFX3008 Interconnect device is renamed QFX3008-I Interconnect device.
 - A QFabric system that uses the QFX3008-I Interconnect device is referred to as the QFX3000-G QFabric system.
 - A QFabric system that uses the QFX3600-I Interconnect device is referred to as the QFX3000-M QFabric system.
- **Support for a fiber-based control plane network (QFX3000-G QFabric systems)**—A QFX3000-G QFabric system can now be deployed with a control plane consisting of eight EX4200-24F switches configured as two Virtual Chassis. This configuration requires SFP management boards in QFX3500 Node devices and SFP network modules in QFX3100 Director devices.
- **Support for QFX-QSFP-DAC-1M and QFX-QSFP-DAC-3M QSFP+ direct-attach copper (DAC) cables**—QFX-QSFP-DAC-1M and QFX-QSFP-DAC-3M DAC cables can be used in QSFP+ ports on QFX3500 and QFX3600 devices to create four 10-Gigabit Ethernet (xe) interfaces over a single cable. The QFX-QSFP-DAC-1M and QFX-QSFP-DAC-3M DAC cables can also be used to create a 40-Gbps data plane connection between QFX3500 or QFX3600 Node devices and QFX3600-I Interconnect devices.

High Availability

- **Nonstop software upgrade (QFabric systems)**—Enables you to upgrade a QFabric system with minimal packet loss and maximum uptime. This feature introduces several high availability improvements to the QFabric system software upgrade process, including:

- Upgrading members of a Director group or Node group one at a time so that one device in the group is always operational.
- Switching mastership of Routing Engine processes to the backup Director device before upgrading the master Director device.
- Rebooting Interconnect devices and fabric control Routing Engines one at a time, so that one Interconnect device and one fabric control Routing Engine are always operational.
- Switching mastership of a Node group to the backup Node device before upgrading the master Node device.
- Specifying an upgrade group if you want sets of Node devices in a network Node group to be upgraded in parallel (which shortens the time needed for the upgrade).
- Rebooting devices automatically as part of the nonstop upgrade process.

Before you begin any nonstop upgrade procedure, download the new software locally to the QFabric system by issuing the **request system software download *package-name*** command.

Nonstop software upgrade for a QFabric system requires that you upgrade the Director group first, followed by the fabric components (Interconnect devices and fabric control Routing Engines) second, and Node groups third. To perform a nonstop software upgrade for Director devices in the Director group, issue the **request system software nonstop-upgrade director-group *package-name*** command. To perform a nonstop software upgrade for Interconnect devices and other fabric-related components, issue the **request system software nonstop-upgrade fabric *package-name*** command. To perform a nonstop software upgrade for a Node group, issue the **request system software nonstop-upgrade node-group *node-group-name package-name*** command.

To create an upgrade group for a network Node group so that sets of Node devices within the Node group are upgraded at the same time, include the **node-devices** statement at the **[edit chassis node-group *node-group-name* nssu upgrade-group *upgrade-group-name*]** hierarchy level and specify which Node devices in the group you wish to be upgraded simultaneously. To view the results of a Node group nonstop software upgrade, issue the **show chassis nonstop-upgrade node-group *node-group-name*** command.



NOTE:

- Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method (such as the **request system software add component all *package-name*** command).
 - When performing a nonstop software upgrade, start with the Director group upgrade, then execute the fabric upgrade, and end with the Node group upgrades.
-
- **BGP and OSPF graceful restart (QFabric systems)**—Enables the graceful restart high availability feature for BGP and OSPF on the network Node group when it connects to external peer routers. To enable graceful restart for BGP and OSPF on the network

Node group, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To enable graceful restart for BGP and OSPF within a routing instance on the network Node group, include the **graceful-restart** statement at the **[edit routing-instances *instance-name* routing-options]** hierarchy level.

Fine-tuning of BGP graceful restart timers is available by including the **disable**, **restart-time**, and **stale-routes-time** statements at the **[edit protocols bgp graceful-restart]** hierarchy level. Fine-tuning of OSPF graceful restart timers is available by including the **disable**, **helper-disable**, **notify-duration**, and **restart-duration** statements at the **[edit protocols ospf graceful-restart]** hierarchy level.

In addition to the external support, graceful restart is enabled by default for the internal components of the QFabric system to provide additional levels of high availability and resiliency. No configuration is required for this internal support.

Interfaces

- **Support for configuring ports on QFX3600 Node devices to operate in 10-Gigabit Ethernet or as 40-Gbps data plane uplink ports (QFabric systems)**—On a QFX3600 Node device, you can configure ports Q0 through Q7 to operate as 40-Gbps data plane (fte) uplink ports, and ports Q2 through Q15 to operate in 10-Gigabit Ethernet (xe) mode.
 - To configure ports to operate in 10-Gigabit Ethernet mode, include the **pic 0 xe port-range *port-range-low port-range-high*** statement at the **[edit chassis node-group *name* node-device *name*]** hierarchy level.
 - To configure ports to operate as 40-Gbps data plane uplink ports, include the **pic 1 fte port-range *port-range-low port-range-high*** statement at the **[edit chassis node-group *name* node-device *name*]** hierarchy level.



NOTE: By default, ports Q0 through Q3 operate as 40-Gbps data plane uplink ports, and ports Q4 through Q15 operate in 10-Gigabit Ethernet mode.

- **Support for 32 members in a LAG (QFabric systems)**—You can configure up to 32 members in a LAG for server Node groups and network Node groups.
- **Support for configuring ports on QFX3600 Node devices to operate in 10-Gigabit Ethernet or as 40-Gbps data plane uplink ports (QFabric systems)**—On a QFX3600 Node device, you can configure ports Q0 through Q7 to operate as 40-Gbps data plane (fte) uplink ports, and ports Q2 through Q15 to operate in 10-Gigabit Ethernet (xe) mode.
 - To configure ports to operate in 10-Gigabit Ethernet mode, include the **pic 0 xe port-range *port-range-low port-range-high*** statement at the **[edit chassis node-group *name* node-device *name*]** hierarchy level.
 - To configure ports to operate as 40-Gbps data plane uplink ports, include the **pic 1 fte port-range *port-range-low port-range-high*** statement at the **[edit chassis node-group *name* node-device *name*]** hierarchy level.



NOTE: By default, ports Q0 through Q3 operate as 40-Gbps data plane uplink ports, and ports Q4 through Q15 operate in 10-Gigabit Ethernet mode.

IPv6 Features

- **Support for IPv6 on management interfaces (QFX3500 switches)**—Enables IPv6 support for the following protocols and features on management interfaces on QFX3500 switches:
 - Neighbor discovery
 - SSH
 - Telnet
 - Ping
 - Traceroute
 - Static routing
 - Firewall filters
 - SNMP

Layer 2 Protocols

- **Virtual Ethernet port aggregator (VEPA) and reflective relay (QFabric Systems)**—Virtual Ethernet port aggregator (VEPA) technology aggregates packets generated by virtual machines located on the same server and relays them to a physical switch. The physical switch then provides connectivity between the virtual machines located on the server, so the virtual machines do not communicate with one another. Offloading switching activities from a virtual switch to a physical switch reduces the computing overhead on the virtual servers and takes advantage of the security, filtering, and management features of the physical switch. Reflective relay, also known as “hairpin turn,” enables the physical switch to receive aggregated packets from the virtual machines hosted on the server through the VEPA on the downstream port and send those packets out the same downstream port from which the physical switch received them.
 - **VEPA**—Even though virtual machines are capable of sending packets directly to one another, it is more efficient to pass these aggregated packets from the VEPA to a physical switch. The switch can then send any packets destined for a virtual machine located on the same server to the VEPA.
 - **Reflective relay**—Reflective relay, also known as a “hairpin turn” or “hairpin mode,” returns aggregated packets to the VEPA by using the same downstream port that initially delivered the aggregated packets from the VEPA to the switch. Reflective relay must be configured on the interface located on the physical switch that receives aggregated packets, such as VEPA packets, because some of these packets might need to be sent back to the server if they are destined for another virtual machine on the same server.

Reflective relay only occurs in two situations:

- When the destination address of the packet was learned on that downstream port
- When the destination has not yet been learned

Reflective relay does not otherwise change the operation of the switch. If the interface to which the virtual machine is connected and the MAC address of the virtual machine packet are not yet included in the Ethernet switching table for the virtual machine's associated VLAN, an entry is added. If the source MAC address of an incoming packet under the respective VLAN is not yet present in the Ethernet switching table, the switch floods the packet on all the other ports that are members of the same VLAN, including the port on which the packet arrived.

- **Multichassis link aggregation (MC-LAG) (QFX3500 switches)**—MC-LAG enables a client device to form a logical LAG interface using two QFX3500 devices. MC-LAG provides redundancy and load balancing between the two QFX3500 devices, multihoming support, and a loop-free Layer 2 network without running STP. At this time, MC-LAGs support only Layer 2 features.

On one end of an MC-LAG is an MC-LAG client device that has one or more physical links in a LAG. This client device does not need to be aware of the MC-LAG. On the other side of the MC-LAG are two MC-LAG QFX3500 devices. Each of these QFX3500 devices has one or more physical links connected to a single client device. The QFX3500 devices coordinate with each other to ensure that data traffic is forwarded properly.

To create a multichassis aggregated Ethernet (MC-AE) interface, you need to perform the following high-level tasks on each QFX3500 device hosting the MC-LAG:

- Configure the number of devices participating in an aggregated Ethernet interface for the interchassis link (ICL).

The ICL provides redundancy when a link failure (for example, MC-LAG trunk failure) occurs on one of the active links. The ICL can be either a 10-Gigabit Ethernet or an aggregated Ethernet interface. You can configure only one ICL between the two QFX3500 devices.

- Configure the MC-LAG interface properties.
- Configure ICCP for the peers participating in the MC-LAG.

ICCP replicates control traffic and forwarding states across the QFX3500 devices. It is also used to validate configuration across QFX3500 devices and communicate the operational state of the MC-LAG members.

Configure BFD on the peers hosting the MC-LAG. BFD enables rapid detection of communication failures between peers. When an ICCP connection is operationally down, the status of the peers hosting an MC-LAG is detected by sending liveness detection requests between the peers.

- Configure multichassis protection for the interchassis link protection link (ICL-PL)
This link is used to carry data between the peers hosting the MC-LAG if the MC-LAG trunk fails.

To configure an MC-LAG, you need to include the following statements:

- **mc-ae** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level
- **iccp** statement at the **[edit protocols]** hierarchy level
- **multi-chassis** statement at the **[edit]** hierarchy level

Use the following **show** commands to verify the MC-LAG configuration:

- **show iccp**
- **show interfaces mc-ae**
- **Spanning-tree protocols (STPs) (QFabric systems)**—Provide Layer 2 loop prevention through STP, RSTP, MSTP, and VSTP. On a network Node group, STP, RSTP, MSTP, and VSTP are supported.
- **Port mirroring (QFabric systems)**—Copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Junos OS Release 12.2 adds support for remote analysis, which means that you can send the mirrored traffic to a VLAN that the analyzer is connected to. You configure port mirroring at the **[edit ethernet-switching-options analyzer]** hierarchy level.
- **Private VLANs (QFX3500 switches)**—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.

In Junos OS Release 12.1, promiscuous ports must be trunk ports, and trunk ports cannot carry secondary VLANs. With Junos OS Release 12.2 for the QFX Series, you can configure access ports to be promiscuous, and trunk ports can 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. You configure private VLANs at the **[edit vlans]** hierarchy level.

Layer 3 Protocols

- **Proxy ARP (QFX3500 switches)**—Enables the switch to respond to ARP queries for network addresses by offering its own MAC address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP allows hosts on different subnets to reach one another without requiring that you configure routing or a default gateway. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts do not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to communicate transparently with one another through the switch. You configure proxy ARP at the **[edit interfaces]** hierarchy level.

Network Management and Monitoring

- **sFlow monitoring technology (QFabric systems)**—Enables you to configure sFlow technology on a QFabric system to monitor traffic continuously at wire speed on all interfaces simultaneously. An sFlow monitoring system consists of an sFlow agent embedded in the QFabric system and up to four external collectors. The sFlow agent performs packet sampling and gathers interface statistics, and then combines the information into UDP datagrams that are sent to the sFlow collectors. sFlow monitoring technology uses packet-based and time-based sampling mechanisms, as well as an adaptive sFlow sampling algorithm that dynamically adapts the sampling rate of interfaces to traffic conditions. You configure sFlow monitoring at the **[edit protocols sflow]** hierarchy level. sFlow operational commands include **show sflow** and **clear sflow collector statistics**.
- **Junos OS automation scripts support (QFabric systems)**—Enables you to configure and deploy Junos OS commit and operational (op) scripts from the QFabric system Director group. Commit scripts automate the commit process and enforce custom configuration rules. Op scripts automate operational and troubleshooting tasks by executing Junos OS operational commands and inspecting the resulting output. Scripts are stored in the shared media in the **/pbdata/mgd_shared/** directory. You configure commit scripts at the **[edit system scripts commit]** hierarchy level, and configure op scripts at the **[edit system scripts op]** hierarchy level.
- **Enhanced SNMP support (QFabric systems)**—Extends SNMP support on the QFabric system to include the following MIBs:
 - RFC 2665, *Definitions of Managed Objects for the Ethernet-like Interface Types*
 - Juniper Networks enterprise-specific Interface MIB Extension (mib-jnx-if-extensions)
 - Juniper Networks enterprise-specific Configuration Management MIB Extension (mib-jnx-cfgmgmt)
 - Juniper Networks enterprise-specific Utility MIB (mib-jnx-util)

You configure SNMP at the **[edit snmp]** hierarchy level.

- **Support for Juniper Networks enterprise-specific SNMP Utility MIB (QFabric systems)**—Provides you with SNMP MIB container objects of the following types: 32-bit

counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations. You use the **request snmp utility-mib set instance *name* object-type *type* object-value *value*** operational mode command to store objects. You use the **request snmp utility-mib clear instance *name* object-type *type*** operational mode command to clear the stored value.

- **Internal fabric OAM monitoring (QFabric systems)**—Enables fabric OAM unicast ping, multicast ping, and traceroute operations for a given VLAN in the QFabric system when you configure the flow specification, fabric maintenance association (FMA), and fabric maintenance endpoints (FMEPs) parameters. You configure the internal fabric monitoring feature at the **[edit protocols oam fabric flow-specs]** and **[edit protocols oam fabric fabric-maintenance-associations]** hierarchy level. Use the **show flow-specification** and **show oam fabric interfaces** operational mode commands to view the configuration and interfaces of the flow specifications, FMAs, and FMEPs.
- **Debugging and diagnostics support (QFabric systems)**—Uses the following commands to provide diagnostic and debugging support for the Interconnect device and the Control Boards and Flexible PIC Concentrators installed on the Interconnect device:
 - **show chassis ethernet-switch interconnect-device *name* cb**
 - **show chassis ethernet-switch interconnect-device *name* cb detail port *port-number***
 - **show chassis ethernet-switch interconnect-device *name* cb detail**
 - **show chassis ethernet-switch interconnect-device *name* cb slot-number**
 - **show chassis ethernet-switch interconnect-device *name* cb detail port *port-number* slot-number**
 - **show chassis ethernet-switch interconnect-device *name* cb detail slot-number**
 - **show chassis ethernet-switch interconnect-device *name* fpc**
 - **show chassis ethernet-switch interconnect-device *name* fpc detail slot-number**
 - **show chassis ethernet-switch interconnect-device *name* fpc detail**

QFabric Systems

- **Aliases for Director devices and Interconnect devices (QFabric systems)**—Enable you to map serial identifiers of Director devices and Interconnect devices to user-defined alias names.

To provide a user-defined name for a Director device, include the **director-device** statement at the **[edit fabric aliases]** hierarchy level and specify the serial identifier (**aliasable-item-name**) and desired alias (**item-name**). To provide a user-defined name for an Interconnect device, include the **interconnect-device** statement at the **[edit fabric aliases]** hierarchy level and specify the serial identifier (**aliasable-item-name**) and desired alias (**item-name**).

To find the serial identifier of a QFabric device, you can use the context-sensitive help question mark (?) with the aliasing syntax. You can also issue operational mode commands to identify the identifier before the aliasing configuration task. To display the serial identifier for a Director device, issue the **show fabric administration inventory director-group status** command. To display the serial identifier for an Interconnect device, issue the **show fabric administration inventory** command.

In conjunction with extended alias support, several operational mode commands have been enhanced to display alias names instead of serial identifiers.

- **Node group filter for standard Junos OS operational mode commands (QFabric systems)**—Commands such as **show vlans** that do not include QFabric system components as optional qualifiers (such as **node-device**, **interconnect-device**, **director-group**, and so on) have been extended with a pipe character (|) to filter output based on Node groups and sets of Node groups. To filter a command by a Node group, issue the **operational-mode-command | filter node-group node-group-name** command. To filter on a set of Node groups, include multiple Node group names (for example, **show vlans | filter node-group [node0 node1 node2]**).
- **Support for four Interconnect devices (QFabric systems)**—A QFabric system can now support up to four Interconnect devices. Increasing the number of Interconnect devices lowers the oversubscription ratio for your QFabric system.
- **Support for IPv6 (QFabric systems)**—Enables IPv6 support for the following protocols and features on QFabric network Node groups:
 - Static routing
 - Firewall filters
 - RADIUS
 - SNMP
 - NTP
 - SSH
 - Neighbor discovery
 - Telnet

- Ping
- Traceroute
- **QFX Series standalone device to QFabric system Node device conversion kit (QFabric systems)**—When you order QFX3500 or QFX3600 devices from Juniper Networks, you can add the free QFX Series standalone device to QFabric system Node device conversion kit (Juniper Networks model number QFX-NODE-KIT) to your order so that the devices are shipped to you as fabric-ready Node devices (indicated by a QFX-NODE label on the shipping carton and the top of the chassis for the Node devices). This saves you the time required to convert the devices into fabric-ready Node devices.
- **Updates to QFabric system control plane switch configurations and software support (QFabric systems)**—The following additions and changes have been made in Junos OS Release 12.2 for the control plane network in QFabric systems:
 - **Software version**—All EX4200 switches that support QFabric systems now require Junos OS Release 12.1R1.10 to support additional high availability features such as nonstop bridging and LLDP. To upgrade your control plane switches to this software package, download and install the following:
<https://download.juniper.net/software/junos/12.1R1.10/jinstall-ex-4200-12.1R1.10-domestic-signed.tgz>
 - **Configurations**—There are now three configurations available for QFabric systems. Download and load the configuration most suited for your system as follows:
 - **Copper-based QFX3000-G QFabric system (EX4200-48T):**
<https://download.juniper.net/software/junos/12.2X50-D10.3/qfx3000g-ex4200-cpe.cfg>
 - **Fiber-based QFX3000-G QFabric system (EX4200-24F):**
<https://download.juniper.net/software/junos/12.2X50-D10.3/qfx3000g-ex4200-24f-cpe.cfg>
 - **Copper-based or fiber-based QFX3000-M QFabric system (EX4200-24F or EX4200-24T):**
<https://download.juniper.net/software/junos/12.2X50-D10.3/qfx3000m-ex4200-cpe.cfg>

Security

- **MAC move limiting (QFX3000-G QFabric systems)**—Detects and prevents MAC movement and MAC spoofing on access ports. You configure MAC move limiting at the `[edit ethernet-switching-options secure-access-port]` hierarchy level.
- **Firewall filters (QFX3500 switches)**—Enable you to provide rules that define whether to accept, police, or discard packets. With previous releases of Junos OS for the QFX Series, you can apply firewall filters to network interfaces, VLANs, routed VLAN interfaces (RVIs), LAGs, and loopback interfaces. With Junos OS Release 12.2, you can also apply firewall filters to management interfaces. You configure firewall filters at the `[edit firewall]` hierarchy level.

Storage and Fibre Channel

- **VN_Port-to-VN_Port FIP snooping (QFX3500 switches)**—Enables you to configure virtual links between VN_Ports without sending traffic through an FCoE forwarder (FCF). Without VN_Port-to-VN_Port FIP snooping, communication from one VN_Port must traverse the FCF before it goes to the targeted VN_Port (VN_Port→ QFX3500 switch→ FCF→ QFX3500 switch→ VN_Port). When you enable VN_Port-to-VN_Port FIP snooping, traffic between VN_Ports does not need to traverse the FCF, so the path is VN_Port→ QFX3500 switch→ VN_Port. To configure VN_Port-to-VN_Port FIP snooping, include the **examine-vn2vn** statement at the **[edit ethernet-switching-options secure-access-port vlan *vlan-name* examine-fip]** hierarchy level.
- **Graceful restart for FIP snooping (QFX3500 switches)**—Enables the system to recover gracefully from Ethernet protocol restarts and crashes. If the Ethernet protocol restarts, 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.

System Management

- **Licensing (QFabric systems)**—Provides ability to add, delete, and view licenses using the following operational mode commands:
 - **request system license add (*filename* | *url*)**
 - **request system license add terminal**
 - **request system license delete *license-id***
 - **show system license**

Traffic Management

- **Layer 3 class of service (QFX3500 switches)**—Enables you to configure Layer 3 class of service (CoS) on QFX3500 switches 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.
- **DCBX application protocol TLV exchange (QFabric systems)**—Extends support to QFabric systems; this feature is already supported on QFX3500 switches. Enables you to configure Layer 2 and Layer 4 application protocol TLV exchange on DCBX-enabled interfaces, including iSCSI TLV exchange. In addition to FCoE and iSCSI TLVs, you can configure up to 14 DCBX application TLVs. Define applications by mapping them to the IEEE 802.1p code points of incoming traffic, and applying the maps to the interfaces. You define Layer 2 applications at the **[edit applications application]** hierarchy level by including the **ether-type** statement, and you define Layer 4 applications by including the **protocol** and **destination-port** statements. You map applications to code points by including the **application** and **code-points** statements at the **[edit policy-options]** hierarchy level.

application-maps] hierarchy level. You apply the application map to an interface by including the **application-map** statement at the **[edit dcbx interface]** hierarchy level.

- **DCBX standards support (QFX Series)**—Enables you to configure the DCBX version that each port supports on a per-port basis. You can configure a port to use IEEE DCBX or DCBX version 1.01, or you can configure a port to autonegotiate with the connected peer to set the DCBX version. The default mode is autonegotiation. Configure the DCBX version by including the **mode** statement at the **[edit protocols dcbx interface *interface-name*]** hierarchy level. You can also disable the ETS recommendation TLV advertisement by including the **recommendation-tlv no-auto-negotiate** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.

**Related
Documentation**

- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

Changes in Default Behavior and Syntax in Junos OS Release 12.2 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 in earlier releases, 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*** command), after you upgrade to Junos OS Release 12.2, those queues do *not* receive lossless treatment and instead 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 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.

Link Layer Discovery Protocol (LLDP)

- **Change in output for the show lldp local-information command on QFX3500 switches**—With QFX3500 switches running Junos OS Release 12.2X50 version 4 and higher (excepting Junos OS Release 12.3X50-D35), the **Interface description** field displays the advertised LLDP description as the interface name (for example, **xe-3/0/1.0**).

For QFX3500 switches running Junos OS Release 12.3X50-D35, the **Interface description** field displays the advertised LLDP description as the user-configured port description (for example, **rr7-14-ws-dis-f1-t1-r24**).

QFabric Systems

- **Changes to the Director group mastership command (QFabric systems)**—The **request fabric administration director-group change-master** command now displays the prompt **Do you intend to switchover mastership?** You can answer **yes** or **no**, with **no** being the default choice to prevent accidental switchovers. Previously, the switchover was executed the moment you issued the command.
- **Updates to the Director group status command (QFabric systems)**—The **show fabric administration inventory director-group status** command now displays serial numbers or aliases for Director devices. These values are accepted when the **target** option of this Director group status command is used. This command also displays standard QFabric system terminology for items such as the Director group processes, the Director group manager, and the control plane interfaces (displayed as **CP Link [slot/port]**).
- **New configuration statement to modify the access password for individual QFabric components (QFabric systems)**—Allows you to modify the password used to access individual QFabric components when you issue the **request component login component-name** command. To change the component-level password, issue the **device-authentication** statement at the **[edit system]** hierarchy level and specify a plain-text or an encrypted password.



NOTE: Configuring the device-authentication statement overrides the password for individual QFabric system components that was set during the initial QFabric system setup procedure.

- **Expanded permissions for QFabric component access classes (QFabric systems)**—Allows administrators to issue a larger set of operational mode commands than available previously when you access individual components in a QFabric system using the **request component login component-name** command. Permissions for the **qfabric-operator** class now include the ability to issue the **monitor** and **show log** commands. Permissions for the **qfabric-admin** class now include the ability to issue all operational mode commands except **configure**.
- **Software upgrade validation options (QFabric systems)**—The **request system software add** command now supports the **no-validate** and **validate** options on QFabric systems. When you do not specify one of these two options, the default value for a software upgrade is **validate**.

**NOTE:**

- The validation step adds up to 10 minutes to the overall QFabric system upgrade time.
- If the validation fails, the upgrade does not proceed, and the QFabric system automatically issues the **request system software rollback** command to restore the current software image. If you upgrade more than one component (as with the **component all** option), validation failure on one device stops the upgrade process for other devices.
- Because the nonstop software upgrade option includes validation by default, there is no change for this type of upgrade.

- **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.
- **Specifying the QFabric system platform (QFabric systems)**—During the initial setup script run on the QFX3100 Director device to initialize your QFabric system, you are now required to specify which QFabric system you are installing. Indicating the correct QFabric system ensures that the proper software and components are installed and ready. Current options include the QFX3000-G QFabric system (which uses the QFX3008-I Interconnect device) and the QFX3000-M QFabric system (which uses the QFX3600-I Interconnect device).

Related Documentation

- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

Limitations in Junos OS Release 12.2 for the QFX Series

This section lists the limitations in Junos OS Release 12.2 for the QFX Series.

Interfaces and Chassis

- On a QFX3500 switch, configuring the **traceoptions** statement at the **[edit forwarding-options helpers]** hierarchy level has no effect. As a workaround, configure the **traceoptions** statement with the **level all** and **flag all** options at the **[edit system services dhcp]** hierarchy level.

- On a QFabric system, 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.

Layer 2 Protocols

- On a QFabric system, you can configure spanning-tree protocols (STP, RSTP, MSTP, and VSTP) on network Node groups only. Spanning-tree protocols are not supported on server Node groups.

Network Management

- If the QFX3500 switch drops traffic because of an ingress firewall filter, the switch does not generate an sFlow monitoring technology flow sample packet that contains this dropped traffic.
- sFlow monitoring ingress samples are not captured for a port on the QFabric system network Node group if the ingress traffic on that port is being copied to the CPU at the same time. This can happen during an ICMP redirect operation.
- On a QFabric system, the following SNMP scalar variables are not supported: ifStackStatus, ifTableLastChange, and ifStackLastChange.
- On a QFabric system, the SNMP Fabric Chassis MIB JnxFabricRedundancyTable object does not contain Interconnect device information for the QFX3008-I and QFX3600-I Interconnect devices.

QFabric Systems

- When you use the SecureCRT client to connect to a Director device for the initial setup of a QFabric system, the Backspace key does not work. As a workaround, use the Shift + Delete key combination in SecureCRT as a Backspace key equivalent, or use a different UNIX client to support the Backspace key natively.
- On a QFabric system, if the static ARP configuration includes a multicast MAC address on a routed VLAN interface (RVI), and the RVI spans multiple Node groups (server Node group, network Node group, and redundant server Node group), IP traffic sent to the configured static IP address floods traffic only on the Node groups that the VLAN spans.
- On a QFX3500 switch, issuing the **restart routing immediate** operational command causes temporary traffic loss, as follows:
 1. Traffic loss occurs while the routing process restarts.
 2. Traffic is recovered.
 3. Traffic is lost again for a brief period.
- On a QFX3500 switch with a large number of VLANs configured (thousands of VLANs), if you enable or disable IGMP snooping on the VLANs, the system may take several minutes to implement the change.

Storage

- 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. If the ports are not FCoE trusted, the maximum number of FCoE sessions is 376.
- When you configure FIP snooping filters, if the filters consume more space than is available in the 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 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 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 12.2X50-D10 for the QFX Series on page 14](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

Outstanding Issues in Junos OS Release 12.2 for the QFX Series

The following issues are outstanding in Junos OS Release 12.2. 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>.

- [Ethernet Switching](#)
- [Hardware](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [QFabric Systems](#)
- [Routing Protocols](#)
- [Storage and Fibre Channel](#)
- [User Interface and Configuration](#)

Ethernet Switching

- On a QFX3500 switch, the maximum amount of bandwidth available on a 10-Gigabit Ethernet interface (xe) port is 10 Gbps. The maximum bandwidth configuration for a queue (the shaping rate configured in the scheduler mapped to the queue forwarding class) or for a priority group (the shaping rate configured in the traffic control profile mapped to the forwarding class set) should not be greater than the total port

bandwidth. If you set a shaping rate that is greater than 10 Gbps for a queue or for a priority group, the shaping rate is invalid. However, if you configure an invalid shaping rate, the switch does not perform the proper commit check, and the invalid configuration commit operation succeeds. [PR/684041]

- On a QFX3500 switch, when you configure Q-in-Q tunneling on an interface, persistent MAC learning does not happen. The MAC entries are learned as dynamic entries. [PR/720380]

Hardware

- On a QFX3500 switch, the QSFP+ module might not be recognized by the switch. This happens only when you insert the QSFP+ module and does not happen during operation. As a workaround, remove and reinsert the QSFP+ module. [PR/793124]

Interfaces and Chassis

- On a QFabric system, momentary loss of Layer 2 multicast traffic might occur when you reboot the backup network Node group Routing Engine. [PR/691729]
- On a QFabric system, momentary loss of Layer 2 multicast traffic might occur when you reboot the fabric manager Routing Engine. [PR/692006]
- 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 the 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, 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. [PR/699091]
- On a QFX3000-M QFabric system, the **show lacp interfaces** operational command outputs extraneous messages such as **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. Ignore these types of messages. [PR/723700]
- On a QFabric system, support for the system log (syslog) configuration is limited. You cannot configure a filename using the **file filename** statement in the **[edit system syslog]** hierarchy. In addition, you cannot configure the **archive size** or **time-format milliseconds** statements in the **[edit system syslog file]** hierarchy. [PR/726727]
- On a QFabric system, MAC learning on Node devices is distributed and communicated using 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 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. [PR/734824]
- On a QFX3500 switch, if you enable the **flag** option for PIM trace options, the trace options log might display the wrong source IP address for outgoing register messages. [PR/735035]
- On a QFabric system, if you reboot the Interconnect device, existing FCoE sessions might be dropped. [PR/735423]
- On a QFX3500 switch, the **show iccp** operational command output shows the registered client daemons even if no ICCP peer is configured. The **show iccp** operational command always shows registered modules regardless of whether or not ICCP peers are configured. [PR/741964]
- On a QFX3500 switch, if you enable backup liveness detection for a multichassis LAG by including the **iccp peer peer-ip-address backup-liveness-detection backup-peer-ip ip-address** at the **[edit protocols]** hierarchy level, and if the backup liveness packets are lost because of a temporary failure on the link, then both of the Node devices in the multichassis LAG remain active. If this happens, both of the Node devices send packets to the connected server. [PR/748755]
- If you connect Juniper Networks MX Series routers to a QFabric system and enable VRRP on the MX interfaces that connect to the QFabric system, a Routing Engine failover on the QFX3008-I interconnect device might cause VRRP to fail over on the MX Series routers. [PR/750044]
- On a QFabric system, even if the BDPU timeout action feature is configured on a port, the port does not move to a loop-inconsistent state and does not generate an alarm as expected. [PR/756942]
- 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.[PR/770964]
- 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 multicast make-before-break (MBB) feature is enabled by default. Once you disable MBB, the feature becomes permanently disabled. [PR/779220]
- On a QFX3500 switch, when you update members of a LAG interface, there might be traffic loss for other receivers. [PR/779772]
- On a QFabric system, the output of the **show route** command does not indicate the selected next hops using an angle bracket (>). [PR/787385]
- On a QFX3600 switch, when FCS error frames are sent continuously through a 40-Gigabit Ethernet (xle) port, the egress queue of the intended destination port for those frames becomes stuck, resulting in a soft reset of the egress port. [PR/796199]
- On a QFX3500 switch, if the following events happen in this order—ICCP goes down, and the MC-AE interface on the active node goes down—a double failover occurs. In this scenario, the standby node does not detect what happens on the active node, proceeds as if the active node were up, and blocks the interchassis link (ICL) traffic. The ICL traffic is not forwarded. [PR/797950]
- On a QFabric system configured for at-scale Layer 3 routing with RVIs, ARP routes for the QFabric system are sometimes not detected by an EX Series switch connected through the network Node group. As a workaround, issue a **ping** command; the first ping packet resolves this situation. [PR/803085]
- On a QFabric system, the format of system log (syslog) messages for some Interconnect devices might make the messages difficult to read and interpret. [PR/804985]
- 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, there are some conditions and topologies that might lead to traffic loss during a nonstop system software upgrade. Contact Juniper Networks customer support for assistance with a nonstop software upgrade on a QFabric system. [PR/815358]
- 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. [PR/815390]
- On a QFabric system, after an active Director device is powered off, certain operations (such as fabric OAM ping operations) might take up to a minute to work. During this interval, critical services are being moved to the peer Director device. [PR/819377]

Layer 2 Features

- On a QFX3500 switch, an analyzer configured to match on ingress packets in a VLAN does not match protocol data units (PDUs) for Layer 2 protocols (for example, STP or LACP). These Layer 2 control PDUs are not mirrored to the output interface or to the VLAN configured in the analyzer. [PR/725710]

- If you enable Layer 2 protocol tunneling on a QFX3500 switch and there are a large number of VLANs configured, the **show ethernet-switching layer2-protocol-tunneling statistics** command might not work. In this case, Layer 2 tunneling works properly, but the CLI command does not display any output. [PR/739027]
- On a QFX3500 switch, the SFID daemon might be unable to send and receive control packets. The workaround is to restart the SFID daemon. [PR/792662]

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, the **request chassis fabric fpc interconnect-device interconnect-device-name slot slot-number online** command does not set the specified Interconnect device FPC online. As a workaround, issue the **request chassis fpc interconnect-device interconnect-device-name slot slot-number online** command. [PR/719345]
- On a QFabric system, during a rolling upgrade, the SNMP agent and related components move between Director devices repeatedly, causing temporary outages in the SNMP service. Once the QFabric system upgrade is completed, normal SNMP service is resumed. [PR/792351]
- On a QFabric system, SNMP service might be disrupted if you delete and then add a fabric resource configuration. The SNMP service should recover within 60 seconds. [PR/802047]
- On a QFabric system, if a loopback IP address is configured on a network Node group with no line cards installed, SNMP service might be disrupted temporarily for the IP MIB. The SNMP service resumes within 60 seconds. [PR/803550]
- On a QFabric system, even if multiple line cards are configured in the network Node group or redundant server Node group, the SNMP Fabric Chassis MIB `jnxFabricContentsTable` object provides the `jnxFabricContentsType` row for one line card only. [PR/809752]
- 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, 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. [PR/824838]
- On a QFabric system, there is a delay in the display of system log messages after the **show log messages** command is issued. [PR/826830]
- On a QFabric system, the Director software might stop fetching system logs or traps from some Node devices. If this happens, reboot the Node device or the Director devices. [PR/827213]

- On a QFabric system, during bootup of a Director device, the `jnxVccpMemberUp` trap might be displayed with unreadable characters. [PR/828253]
- On a QFabric system, `jnxVpnIfDown` traps are generated when an RVIs 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 LED display on the Director device does not show the name shown in the CLI. [PR/709261]
- 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 show as up even when no members of the LAG are present (all members of the LAG are down). [PR/756970]
- On a QFabric system, if you configure new filters for a system log file that contains a large number of messages, processing of the **`show log filename`** command is slow. [PR/815565]

QFabric Systems

- On a QFabric system, autocompletion of interface names does not work in some cases. As a workaround, when pressing the Tab key 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 QFX3500 switch, if you move the aggregated Ethernet interfaces from one routing instance to another, it is possible that one or more of the LAGs might be declared down even though all its member interfaces are up. As a workaround, disable and enable each member link of the affected LAG to bring the LAG to the up state. [PR/612277]
- On a QFabric system, in very rare cases after you install QFabric software from a USB drive, the system might fail to boot and might display the following message: **`switchroot: mount failed: No such file or directory`**. As a workaround, restart the affected Director device. [PR/664511]
- On a QFabric system, you can configure interfaces on Node devices that belong to a network Node group as Layer 3 interfaces using the **`family inet address ip-address`** statement at the **`[edit interfaces interface-name unit unit-number]`** hierarchy level. You cannot configure interfaces on Node devices that belong to a server Node group as Layer 3 interfaces. However, if you configure a Node device in a server Node group as a Layer 3 interface, the **`commit`** command succeeds even though it should fail. [PR/674010]

- 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, 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]
- On a QFabric system, a configured QFabric log file might contain a large amount of data. If you issue a matching command such **show log | grep**, the system might take a long time to process the request. As a workaround, you can reduce the time required by these commands by configuring a log file to contain only the necessary facilities and severities. [PR/681413]
- 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**. [PR/683442]
- 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 they are 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, if you issue the **help reference** or **help topic** operational mode commands, the output might display the following incorrect items: (1) a bracketed number in the Syntax and Related Topics fields, and (2) an unsupported References field. [PR/687527]
- On QFX3500 switches, ingress and egress sFlow technology sampling can be enabled only on interfaces that are configured with the logical unit 0. sFlow monitoring technology sampling does not work on 802.1q subinterfaces with a nonzero logical unit number. [PR/693879]
- On a QFabric system, if a Director device is not listed in the output of the **show fabric administration inventory director-group status** command, the device might not be operational. [PR/697626]
- On a QFabric system, deleting individual child members from a LAG might reset the aggregated Ethernet interface traffic counters. [PR/700039]
- On a QFabric system, if you configure a system log message archive size with the **archive size** statement at the **[edit system syslog]** hierarchy level, the configuration has no effect on the output, which can generate up to 100 MB of data. [PR/700511]
- 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, if you attempt to commit an incorrect or incomplete configuration, the system might include error messages in the commit failure notification such as **warning: from xxx: [<xnm:error xmlns:xnm="http://xml.juniper.net/xnm/1.1/ ...** These messages are extraneous, they do not impact the operation of the QFabric system, and they can be ignored. [PR/703134]
- On a QFabric system, if 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. [PR/705888]
- 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 QFabric system, if you issue the **show | compare | display inheritance** operational mode command, the output might display an error message. [PR/706924]
- On a QFabric system, if the backup Control Board in an Interconnect device is disconnected and you issue the **request chassis routing-engine master switch interconnect-device device-name** command, the mastership switchover might be executed in error. To prevent this situation, issue the **show fabric administration inventory** command to verify that the backup Control Board is in the Connected state before you issue the switchover command. [PR/708832]
- On a QFX3500 switch, MAC table entries might not age out properly if the MAC address is assigned to a private VLAN. To work around this issue, you can restart the Ethernet switching process (eswd) to flush out all of the stale MAC entries. [PR/707487]
- On a QFabric system, if you issue the **show fabric inventory node-devices** command, the resulting output might not show the status of the Configuration field. [PR/717082]
- On a QFabric system, when you remove the primary Ethernet module from the master Director device in the Director group, the default partition of the QFabric system CLI might not be accessible immediately. As a workaround, wait a few minutes and log back in to the QFabric system CLI when it becomes available. [PR/725468]
- On a QFabric system, when a Director device fails to communicate with its peer in the Director group, you might not be able to access that Director device by using SSH or by logging in with the console port immediately after a reboot. Incorrect wiring, port failure, and cable failure are some of the issues that can cause communication failure between members of a Director group. [PR/731883]
- On a QFabric system, if 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. [PR/735549]

- On a QFabric system, the output of the **show pfe statistics traffic** command does not report statistics for normal traffic discards. [PR/736966]
- 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, 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, if you configure remote port mirroring with the input and output interfaces on different Node devices, the system might add an extra VLAN tag to the mirrored packets. This occurs if the Node device that has the output interface also has an interface or VLAN that is the input for another remote port mirroring configuration. In this case, the VLAN tag for the analyzer VLAN is added to the mirrored packets twice. For example, assume the following configuration:
 - Remote port mirroring configuration 1: input interface on Node device A, output interface on Node device B, analyzer VLAN ID is 100
 - Remote port mirroring configuration 2: input interface on Node device BWhen packets mirrored by remote port mirroring configuration 1 leave the output interface on Node device B, they are tagged with VLAN ID 100 twice. [PR/751776]
- 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 QFX3500 switch, if you enable multicast VLAN registration (MVR) and also apply a firewall filter to a VLAN (in either the input or output direction), the filter does not take effect. [PR/776071]
- 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, the multicast VLAN registration route traffic does not flow through the ICCP interface if the ICCP interface is configured as an MVLAN source port. [PR/777007]
- 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 QFX3000-M QFabric system, if you connect two QSFP+ data plane connections from a single Node device to the same Interconnect device, performing a nonstop software upgrade might not work. As a workaround, disconnect one of the QSFP+ interfaces, perform the nonstop software upgrade, and then reconnect the disconnected interface after the upgrade is completed. [PR/786391]
- 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 fabric unicast ping and fabric unicast traceroute operations cannot perform data path validation from a redundant server Node group or network Node group if the destination MAC address defined in the flow specification is learned on any of its member Node devices. [PR/792786]
- 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, when you show the configuration of a firewall filter that includes a policer or counter, the system might display an internal name for the policer or counter instead of the configured name. [PR/799243]
- On a QFabric system, if you create a routing instance named *virt1* and then issue the **show route table virt1.inet.0** operational command to view the route entries in the routing table associated with the routing instance, the **No routing tables matching specification** error message is displayed. As a workaround, specify the name of the routing instance as the routing table name in the **show route table** command. For example, issue the **show route table virt1** command to display the route entries for the routing instance named *virt1*. [PR/805043]
- On a QFabric system, 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. [PR/805406]

- On a QFabric system, when you execute the **show igmp-snooping membership detail** command, you might see the following error message: **error: from NW-NG-0: Current state is not among the states START_ELEMENT , A ...** This message is harmless and can be safely ignored. [PR/808686]
- 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]
- 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]
- On a QFabric system, sending a **ping** command to an IPv6 address results in the operation terminating after sending only one packet unless the number of packets to send has been specified in the command. [PR/815901]
- On a QFabric system, if an interface that has not been configured is added to a VLAN, an unrelated error message might be displayed. [PR/816600]
- On a QFabric system, mirroring of 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. [PR/819569]
- 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, after you commit a configuration change, the system inventory might not be updated because of a database failure. There is no workaround. [PR/828144]
- On a QFabric system, the Director software might stop fetching system logs (syslogs) or SNMP traps from Node devices. If syslogs are affected, the Juniper Message Bundles (JMBs) used by the Advanced Insight Solutions feature for troubleshooting are also not fetched. In that case, reboot the Node device or Director devices. [PR/832159]
- 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]

Routing Protocols

- On a QFX3500 switch, the output of the **show msdp brief** command might display incorrect values for the SA count field. [PR/732115]
- On a QFX3500 switch, if you configure a firewall filter with more than 128 policers and attempt to apply the filter to a Layer 3 interface in the output direction, the commit operation fails, and the filter is not created. [PR/745327]

Storage and Fibre Channel

- On a QFX3500 switch, if the switch is in autonegotiation mode but a peer switch is disabled, the **show dcbx neighbors interface** command displays the protocol mode incorrectly with the IEEE version instead of the DCBX version. As a workaround, enable the peer switch to display the output correctly. [PR/802001]

User Interface and Configuration

- 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]
- You can use the **vlan vlan-name interface interface-name mapping-range** statement to map a range of customer VLANs to a range of service VLANs instead of using multiple **vlan vlan-name interface interface-name mapping (push | swap)** statements to configure Q-in-Q tunneling or VLAN translation on a per-VLAN basis. If you enter the first statement and configure a range in which the first value is larger than the second value (which is unsupported), you see an error message that does not clearly describe the problem. [PR/728938]
- On a QFX3500 switch, the log messages of IP addresses assigned by the DHCP server are not displayed unless the level flag is set in the **traceoptions** file. [PR/729571]
- On QFX3500 switches, although the **igmp-snooping** statement at the **[edit routing-instances instance-name protocols]** hierarchy level is visible, it is not supported. When you configure IGMP snooping in a routing instance, the configuration does not work. [PR/729629]
- 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 the QFX Series, if 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. [PR/780290]
- 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]

Related Documentation

- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)

- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

Resolved Issues in Junos OS Release 12.2 for the QFX Series

The following issues have been resolved in Junos OS Release 12.2 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 “Resolved Issues in Junos OS Release 12.2 for the EX Series” section in the [Junos OS 12.2 Release Notes](#).

-
- [Issues Resolved in Release 12.2X50-D70 on page 44](#)
 - [Issues Resolved in Release 12.2X50-D65 on page 45](#)
 - [Issues Resolved in Release 12.2X50-D60 on page 46](#)
 - [Issues Resolved in Release 12.2X50-D55 on page 46](#)
 - [Issues Resolved in Release 12.2X50-D50 on page 46](#)
 - [Issues Resolved in Release 12.2X50-D40 on page 48](#)
 - [Issues Resolved in Release 12.2X50-D30 on page 50](#)
 - [Issues Resolved in Release 12.2X50-D20 on page 51](#)

Issues Resolved in Release 12.2X50-D70

The following issues have been resolved since Junos OS Release 12.2X50-D65. The identifier following the description is the tracking number in our bug database.

Interfaces and Chassis

- On a QFabric system, if the system is operating under a heavy load and the system polling counters are not retrieved within the specified time limit, the **PFED_NOTIFICATION_STATS_FAILED** message might appear. [PR/807372: This issue has been resolved.]
- On a QFabric system, the fabric control protocol (rpdf), responsible for distributing routing information between all system devices, might stop operating after the rpdf

process restarts, the Director group restarts, or you perform a nonstop software upgrade. [PR/1013756: This issue has been resolved.]

Platform and Infrastructure

- On devices running Junos OS software, traffic between the Routing Engine and transit interfaces is carried over an internal network between the Packet Forwarding Engine and the Routing Engine. Some Routing Engines use em interfaces (usually, em0 and em1) to connect to this network. If the device receives a carefully crafted set of fragmented packets that are destined to the device itself, the em interface driver might become permanently blocked when trying to formulate a reply, the Routing Engine might be unable to communicate over the private network that connects the FPCs and the Routing Engines, and eventually all FPCs might go offline and stay offline. Systems with redundant Routing Engines will switchover, but are still subject to this same issue. For systems without modular FPCs (for example, an MX80 router), the FPC will reboot and clear the em0 interface output queue. However, additional crafted fragments can cause the issue to reoccur.

This issue is applicable to IPv4, IPv6, and CLNP fragmentation and reassembly scenarios. Transit traffic does not trigger this issue. Additionally, CLNP is only vulnerable if clns-routing or ES-IS is explicitly configured. Refer to [JSA10655](#) for more information. [PR/942437: This issue has been resolved.]

- On a QFabric system, if you issue the **show system users** command on the two different Director devices in the Director group, you might receive different output from each of the Director devices. [PR/1021831: This issue has been resolved.]

Spanning Tree Protocols

- On a QFX Series switch running MSTP, a non-edge interface might be marked incorrectly as a boundary interface. [PR/996108: This issue has been resolved.]

Issues Resolved in Release 12.2X50-D65

The following issues have been resolved since Junos OS Release 12.2X50-D60. The identifier following the description is the tracking number in our bug database.

Infrastructure

- On a QFabric system configured with multicast, if some events (such as an interface transition) cause reference counters to increase but not decrease, the fabric control protocol (rpdf) on the master member of a redundant server Node group might stop operating and generate a core file. [PR/961327: This issue has been resolved.]
- If QFSP+ DAC cables are used on a QFX3500 switch, then the `xcvr_qsfp_get_info` diagnostic for port messages is not implemented. [PR/976373: This issue has been resolved.]

Software Installation and Upgrade

- During a redundant server Node group (RSNG) nonstop upgrade, if a fan is not operational in one of the FPCs, the upgrade may not complete successfully. [PR/969104: This issue has been resolved.]

Issues Resolved in Release 12.2X50-D60

No issues have been resolved since Junos OS Release 12.2X50-D55.

Issues Resolved in Release 12.2X50-D55

The following issues have been resolved since Junos OS Release 12.2X50-D50. The identifier following the description is the tracking number in our bug database.

Interfaces and Chassis

- On a QFabric system, the Fabric Control Protocol rpdf process might be using a high percentage of CPU resources on the redundant server Node group because obsolete core key entries are not being deleted. [PR/919027: This issue has been resolved.]

QFabric Systems

- On a QFabric system, the SNMP process (sfcsnmp) might stop responding to requests because it is operating at 100 percent of capacity. [PR/913855: This issue has been resolved.]

Issues Resolved in Release 12.2X50-D50

The following issues have been resolved since Junos OS Release 12.2X50-D40. The identifier following the description is the tracking number in our bug database.

Interfaces and Chassis

- On a QFabric system, if any interface is in a disabled state, this information is stored in the analyzer entry and is shown in the output when the **show analyzer** command is executed. [PR/805364: This issue has been resolved.]
- On a QFabric system, when you configure the analyzer with an input ingress interface with the **all** option, LAG member links are added to the analyzer. [PR/807302: This issue has been resolved.]
- On a QFabric system, issuing the **ping fabric unicast-flow** command between Node devices in a network Node group does not work. The request times out, and no responses are received from the destination Node devices. [PR/818952: This issue has been resolved.]
- On a QFX3500 switch, after the initial software installation, the message **/kernel: kern.maxfiles limit exceeded by uid 0, please see tuning(7)** might appear on the console. This problem occurs if the kdm_savecore process exceeds the maximum open file allowance in the kernel. As a workaround, issue the **savecore -C** command to stop

further file processing and clear the kernel crash flag, and then reboot the switch.
[PR/871506: This issue has been resolved.]

- On a QFX3500 switch, disabling an RVI might result in the RVIs of any VLAN sending ARP entries for all configured IP addresses. As a workaround, deactivate the RVI instead of disabling it. [PR/871874: This issue has been resolved.]

QFabric Systems

- On a QFabric system, if you connect to the console port of a Node device that is a nonmaster member of the network Node group, you might be able to log in without being authenticated. [PR/710660: This issue has been resolved.]
- On a QFabric system, when you issue the **show fabric administration inventory infrastructure detail** command, sometimes you see no output displayed. As a workaround, reissue the command (because this might be an intermittent issue). [PR/790341: This issue has been resolved.]
- On a QFabric system, the configuration at the **[edit fabric]** hierarchy level is visible to users with read-only permission, who should not be able to see the configuration information. [PR/812935: This issue has been resolved.]
- On a QFabric system, after you reboot a Director device, system log messages (syslogs) that were generated by system component devices before the reboot are sent to the syslog server. [PR/859547: This issue has been resolved.]
- On a QFabric system, the BPDU block feature does not disable a PVST interface on a redundant network Node group or server Node group as expected. [PR/835548: This issue has been resolved.]
- On a QFabric system, if you remove the connection between the Director devices (dg0 and dg1) and power off the master Director device (dg0), mastership switches to the backup Director device (dg1) as expected. However, if you restore the connections and power on dg0, then remove the Director device interconnection a second time and power off the new master Director device (dg1), the second mastership change to dg0 incorrectly shows in the output of the **show fabric administration inventory infrastructure** command that both Director devices are hosting virtual Routing Engines, even though one Director device (dg1) is powered off. [PR/855371: This issue has been resolved.]
- On a QFabric system, the lldpd process on a redundant server Node group might crash after a commit operation if there are multiple unknown TLV elements included in the LLDP PDUs. [PR/882778: This issue has been resolved.]
- On a QFabric system, RADIUS authentication for users is not working properly. User permissions are not granted as configured. [PR/887239: This issue has been resolved.]

Security

- On a QFX3500 switch, filter-based forwarding is not supported on an egress interface. If you configure filter-based forwarding on an egress interface, the configuration is committed successfully, but the filter is not applied. [PR/781606: This issue has been resolved.]

Issues Resolved in Release 12.2X50-D40

The following issues have been resolved since Junos OS Release 12.2X50-D30. The identifier following the description is the tracking number in our bug database.

Interfaces and Chassis

- On a QFX3500 switch, when the switch sends ICMP redirect messages to the sending host to redirect the traffic to a different IP address in the same LAN segment, the switch also forwards the data from both software and hardware processes and could result in duplicate packets. Because the sending host reroutes the data to the newly redirected IP address after it receives the ICMP redirect messages, this condition resolves itself momentarily. [PR/730549: This issue has been resolved.]
- On a QFabric system, issuing the **traceroute fabric unicast-flow** command might result in duplicate responses being displayed in the output. The number of responses sent by the destination Node device is correct, but a buffering issue at the source Node device can cause duplications in the output. [PR/817439: This issue has been resolved.]
- On a QFabric system, even if you configure the fabric OAM ping traceoptions feature with the **all** flag, no output appears in the **traceoptions** log file. [PR/821051: This issue has been resolved.]
- On a QFabric system, if you include the **trace-route** statement at the **[edit protocols oam fabric faboam-trace-options flag]** hierarchy level, the commit operation might not succeed. [PR/821058: This issue has been resolved.]

Network Management and Monitoring

- On a QFabric system, if you create a system log (syslog) message file that matches all syslog traffic and should be equivalent to the **messages** log file, some of the log messages might be missing in the log file you configured. As a workaround, send the message file to an external log server, which should contain all the log messages. [PR/786388: This issue has been resolved.]
- On a QFabric system, if the rear cards on an Interconnect device are brought online after being offline, traffic loss lasting up to 0.3 seconds occurs. [PR/789813: This issue has been resolved.]
- On a QFabric system, when you change the time zone, you might receive older timestamps from the previous time zone. [PR/805841: This issue has been resolved.]

- On a QFabric system, SNMP responses sent from a Node device to the Director group might get lost if there is an overflow of the socket receive buffer. [PR/844604: This issue has been resolved.]
- On a QFabric system, unnecessary log files can fill the space in the `/tmp` directory if an external application issues the `show system syslog stream` command. [PR/847681: This issue has been resolved.]

QFabric Systems

- On a QFabric system, storm control does not work for multicast traffic. [PR/774169: This issue has been resolved.]
- On a QFabric system, after you replace a hard drive on a Director device, the `director-group status` command displays erroneous CLI output:
 - SCSI IDs are not shown.
 - Volume status is shown as **degraded** until the RAID synchronizes with the hard drive.
 - Commas are shown after the online entries in the Physical ID fields.
 [PR/799286: This issue has been resolved.]
- On a QFabric system, the SNMP service might be disrupted after a change in the fabric configuration (such as changing from using Layer 3 to Layer 2 protocols). To recover the SNMP service, restart the SNMP agent. [PR/805793: This issue has been resolved.]
- On a QFabric system, the `show log` command displays error messages for operational failures even when the related operational command has been executed successfully. [PR/807331]
- On a QFabric system, a CLI session in operational mode can occur on any Director device in the Director group because of system load balancing. A configuration that is saved in one Director device is not automatically loaded on another Director device, and the operational mode commands for the new configuration might not be available on the second Director device. As a workaround, make a configuration change on the affected Director device so that its configuration gets refreshed. This issue does not affect configuration mode. [PR/807775: This issue has been resolved.]
- On a QFabric system, if remote port mirroring is enabled and you change the maximum transmission unit (MTU) size of the output interface, port mirroring might stop working. [PR/808071: This issue has been resolved.]
- On a QFabric system, if you enter the command `show igmp-snooping membership vlan vlan detail`, the command might fail and display the message **local part cannot be "null" when creating a QName**. If this occurs, executing the command again might resolve the issue. If reentering the command does not work, you can possibly work around the problem by appending a Node group filter to the command as follows: `show igmp-snooping membership vlan vlan detail | filter node-group node-group`. [PR/809006: This issue has been resolved.]
- On a QFabric system, there might be traffic loss with a multichassis configuration if the host has not been sending traffic and the host MAC address has aged out. In this

scenario, any Layer 3 traffic directed to the host might not be flooded in the VLAN. [PR/821911: This issue has been resolved.]

- On a QFX3500 switch, sFlow monitoring data incorrectly shows traffic not generated locally (traffic in-transit through the switch) with a source or destination IfIndex of zero. [PR/824423: This issue has been resolved.]
- On a QFabric system, after an upgrade of the network Node group (part of a nonstop system software upgrade), there might be complete traffic loss on some static Layer 2 multicast routes if IGMP snooping is enabled. [PR/824443: This issue has been resolved.]
- On a QFabric system, after you perform a nonstop system software upgrade, you might not be able to log in to a Node device by specifying the device alias name as part of the **request component login** command. As a workaround, issue the **request component login** command and specify the serial ID number of the Node device. [PR/828304: This issue has been resolved.]
- On a QFabric system, the wrong loss priority is assigned to low-priority drop profiles configured under the **[edit class-of-service schedulers]** hierarchy. [PR/836116: This issue has been resolved.]
- On a QFabric system, class IDs for ERACL and ERACL for IPv6 are not cleared when a firewall is deleted, resulting in TCP and UDP traffic being dropped from egress ports. [PR/846374: This issue has been resolved.]

Routing Protocols

- On QFX3500 and QFX3600 switches, policers applied to a loopback interface (to protect the Routing Engine) do not work properly. [PR/839453: This issue has been resolved.]

Issues Resolved in Release 12.2X50-D30

The following issues have been resolved since Junos OS Release 12.2X50-D20. The identifier following the description is the tracking number in our bug database.

Interfaces and Chassis

- On a QFabric system, the output of the **show vlans summary** command indicates that the default VLAN is a configured VLAN instead of an internal VLAN. [PR/739204: This issue has been resolved.]
- If you connect Juniper Networks MX Series routers to a server Node group on a QFabric system and enable VRRP on the MX interfaces that connect to the Node group, a failover on the Node group might cause VRRP to fail over on the MX Series routers. [PR/751375: This issue has been resolved.]
- On a QFabric system, if you restart the master device in a network Node group, you might experience traffic loss for as long as 10 minutes. [PR/804815: This issue has been resolved.]

- On a QFabric system, when you add a new device, the status of the device might be shown as disconnected after you perform an auto-upgrade. [PR/813380: This issue has been resolved.]
- On a QFabric system, when you reboot a Node device in a network Node group, there might be multicast traffic loss on the LAG interface on the Node device for approximately 60 seconds. [PR/813486: This issue has been resolved.]

QFabric Systems

- If you create 4090 VLANs on a QFX3500 switch (the maximum number of VLANs) and any of them are private secondary VLANs (community or isolated VLANs), a problem occurs if you later reduce the number of VLANs. In this case, traffic for the secondary VLANs is not forwarded between switches that participate in those VLANs. [PR/671831: This issue has been resolved.]
- On a QFabric system, after a USB installation, the backup Control Board of the Interconnect device might stay disconnected, and the Interconnect device does not get upgraded. [PR/811171: This issue has been resolved.]
- On a QFabric system, there might be traffic loss with a multichassis configuration if the host has not been sending traffic and the host MAC address has aged out. In this scenario, any Layer 3 traffic directed to the host may not be flooded in the VLAN. [PR/821911: This issue has been resolved.]

Issues Resolved in Release 12.2X50-D20

The following issues have been resolved since Junos OS Release 12.2X50-D10. The identifier following the description is the tracking number in our bug database.

Interfaces and Chassis

- After you install the Junos SDK on a QFX3500 switch, you must reboot the switch to complete the installation. [PR/700167: This issue has been resolved.]
- On QFX3500 switches, if a multicast route entry in the forwarding table is listed as *discard*, MSDP traffic might not be forwarded. [PR/728619: This issue has been resolved.]
- When installing a Junos OS image on a QFX3500 switch, you might see the error message **LA : fips-mode:: not found**. This does not indicate a problem with the installation and may be safely ignored. [PR/739453: This issue has been resolved.]
- On a QFabric system, if you issue the **request chassis fpc interconnect-device device-name slot slot-number online** command, the resulting message prompt might list the incomplete command **show chassis fpc** as a follow-on action. As a workaround, issue the **show chassis fpc device-type device-name** command as the correct follow-on action. [PR/741212: This issue has been resolved.]
- On a QFabric system, if any devices are disconnected during an snmpwalk, the remaining MIB data for the disconnected devices is skipped. [PR/744885: This issue has been resolved.]

- On a QFabric system, when a Node group nonstop software upgrade is in progress and you issue the **request chassis routing-engine master switch node-group** command, you might be prompted to proceed with or cancel the mastership change. Instead, you should receive the warning message **warning: ISSU in progress**, and the request should be denied. As a workaround, issue the **request chassis routing-engine master switch node-group no-confirm** command to see the correct warning message and system behavior. [PR/749832: This issue has been resolved.]
- On a QFabric system, if you configure an alias name for the Interconnect device, the LCD panel displays the default alias instead of the configured alias. The QFabric system CLI operational commands show the correct alias names for the Interconnect devices. [PR/773380: This issue has been resolved.]
- On a QFX3500 switch, after a software upgrade and during the boot process, the switch displays the following system log error messages for all physical interfaces: *** fpc0 ETH: ifd (xe-0/0/10) unknown boolean option 129 * /kernel: if_pfe: Error 5 (Invalid) on IF command 55 (Ether set boolean)**. These errors are harmless and do not adversely impact the operation of the switch. [PR/778094: This issue has been resolved.]
- On a QFabric system, a QFX3500 Node device might not be able to participate as a Node device in the QFabric system if the Node device is running a different version of software from that of the Director group. This mismatch of software versions between the Node device and the Director group can occur when the Node device is introduced into the setup and both Director devices go offline before the Node device completes its auto-upgrade process to upgrade its software version to the same software version running on the Director group. [PR/791133: This issue has been resolved.]

QFabric Systems

- On a QFabric system, if you issue the **show igmp-snooping membership** command, the system might not execute the command properly, the CLI session hangs, and no output is returned. [PR/734299: This issue has been resolved.]
- On a QFabric system, if you issue the **show fabric administration system mac-pool** command, the output might not display the correct values for the available MAC address blocks. However, there is no adverse operational impact on the QFabric system. [PR/738645: This issue has been resolved.]
- If you connect Juniper Networks MX Series routers to a server Node group on a QFabric system and enable VRRP on the MX interfaces that connect to the Node group, a failover on the Node group might cause VRRP to fail over on the MX Series routers. [PR/751375: This issue has been resolved.]
- On a QFabric system, interactive commands are not logged. [PR/753564: This issue has been resolved.]
- On a QFabric system, when the diagnostic Routing Engine upgrades during a Director group nonstop software upgrade, you might not be able to issue the **show system core-dumps** command for 5 to 10 minutes. [PR/775223: This issue has been resolved.]
- On a QFabric system, when you use the **vstp vlan-group group vlan-group** configuration statement at the **[edit protocols]** hierarchy level, the configuration might fail to commit. [PR/778800: This issue has been resolved.]

- On a QFabric system, when you perform a Director group nonstop software upgrade, after the system upgrades the first Director device, you might not be able to issue operational mode commands. As a workaround, log out of your CLI session and reestablish your connection. [PR/783408: This issue has been resolved.]
- On a QFabric system, MSTP is not supported on server Node groups. If you configure MSTP on a server Node group and commit the configuration, you do not receive an error message. [PR/783737: This issue has been resolved.]
- On a QFabric system, if you commit a configuration exceeding 1000 fabric maintenance associations, the system generates a **commit complete** message (even though the commit operation should fail and generate a error message). [PR/786326: This issue has been resolved.]
- On a QFabric system, if you enable sFlow monitoring packet sampling on network Node group interfaces and issue the **show sflow collector** command, the collector statistics in the resulting output display **0** for some of the collectors even when all the collectors receive the sFlow monitoring samples as expected. [PR/787443: This issue has been resolved.]
- On a QFabric system, when a Director group nonstop software upgrade is in progress and you open a new SSH session on the second Director device, you might not be able to issue operational mode commands for a few minutes. As a workaround, wait for the Director group upgrade to finish and log in again. [PR/789956: This issue has been resolved.]

Related Documentation

- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

Errata in Documentation for Junos OS Release 12.2 for the QFX Series

This section lists outstanding issues with the published documentation for Junos OS Release 12.2X50 for the QFX Series.

QFabric Systems

- The enhanced USB drive recovery process for the QFabric system is not documented.
- The Rapid Spanning Tree Protocol (RSTP) default behavior changes on the QFabric system are not documented.

System Log Messages

- System log messages originating from the QFabric system Director group are currently undocumented. Messages originating from other QFabric system components are documented in the [Junos OS System Log Messages Reference](#).

Related Documentation

- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for the QFX Series on page 54](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.2 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 54](#)
- [CoS Upgrade Requirements from Junos OS Release 12.1 to Junos OS Release 12.2 on page 55](#)
- [Upgrading Software on QFX3500 and QFX3600 Standalone Switches on page 57](#)
- [Upgrading Software on QFabric Systems on page 59](#)
- [Upgrade and Downgrade Support Policy for Junos OS Extended End-of-Life Software Releases on page 66](#)

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

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

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

1. Deactivate the CoS configuration.
`user@switch# deactivate class-of-service`
2. Upgrade to Junos OS Release 11.3 or later.

3. Make the following changes to the CoS configuration:

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

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

5. Activate the CoS configuration.

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

6. Commit the CoS configuration.



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

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

Before you upgrade to Junos OS Release 12.2, you might need to edit the class-of-service (CoS) configuration because the way the QFX Series handles 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 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.

Upgrading Software on QFX3500 and QFX3600 Standalone Switches

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



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



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

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

The download and installation process for Junos OS Release 12.2 is the same as for previous Junos OS releases.

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

1. Using a Web browser, navigate to <http://www.juniper.net/support/downloads/junos.html>.
2. In the QFX Series section, click the name of the platform for which you want to download software.
3. Click the Software tab and select **12.2** from the Release drop-down list.
4. In the Install Package section, select the install package for the 12.2X50-D70.1 release.
An Alert box appears.
5. In the Alert box, read the information about the software alert, and click the link to download the software.
A login screen appears.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.



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

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

```
user@host> request system software add validate
source/jinstall-qfx-12.2X50D70.1-domestic-signed.tgz reboot
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 12.2 jinstall package, you cannot issue the `request system software rollback` command to return to the previously installed software. Instead you must issue the `request system software add validate` command and specify the jinstall package that corresponds to the previously installed software.

Upgrading Software on QFabric Systems

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



NOTE: Before you upgrade from Junos OS Release 11.3 to Junos OS Release 12.2, you must replace your existing upgrade script on both Director devices with one that you download from the PSN when you download the software. For more information, see [“Downloading Software Files for the QFabric System Using a Browser”](#) on page 64.



NOTE: If you upgrade the QFabric system Director group software directly from Junos OS Release 12.2X50-D10 to 12.2X50-D30 without first upgrading to 12.2X50-D20, the system might experience massive traffic disruptions. This is because not all the network Node groups are upgraded at the same time.

As a workaround, we recommend that you:

1. Use the `request system software add package-name` component all reboot command to upgrade all QFabric system components from Junos OS Release 12.2X50-D10 to 12.2X50-D20.
2. Use the nonstop system software upgrade (NSSU) method to upgrade the QFabric system from Junos OS Release 12.2X50-20 to 12.2X50-D30.

You can upgrade the software on the QFabric system using one of the following methods:

- [Performing a Nonstop Software Upgrade on the QFabric System](#) on page 60
- [Performing a Standard Upgrade on the QFabric System](#) on page 63

Performing a Nonstop Software Upgrade on the QFabric System

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.



NOTE: Before you can perform a nonstop software upgrade in your QFabric system, you must first upgrade your system to Junos OS Release 12.2 by using a conventional upgrade method such as issuing the **request system software add component all** command.

Perform the following tasks:

- [Backing Up the Current Configuration Files on page 61](#)
- [Downloading Software Files Using a Browser on page 61](#)
- [Retrieving Software Files for Download on page 62](#)
- [Performing a Nonstop Software Upgrade for Director Devices in a Director Group on page 62](#)

- [Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components on page 62](#)
- [\(Optional\) Creating Upgrade Groups for Node Groups on page 62](#)
- [Performing a Nonstop Software Upgrade on a Node Group on page 63](#)

Backing Up the Current Configuration Files

To back up your current configuration files:

```
user@qfabric> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@qfabric> request system software configuration-backup /media/USB/
```

Downloading Software Files Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to <http://www.juniper.net/support/downloads/junos.html>.
2. In the QFX Series section, click the name of the platform for which you want to download software.
3. Click the Software tab and select **12.2** from the Release drop-down list.
4. In the Install Package section, select the install package for the 12.2X50-D70.1 release.
An Alert box appears.
5. In the Alert box, read the information about the software alert, and click the link to download the software.
A login screen appears.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.

Retrieving Software Files for Download

Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download** command. The software package is copied from where you downloaded it and is placed locally on the QFabric system.

- To retrieve the software, issue the **request system software download /path/package-name** command.

For example:

```
user@qfabric> request system software download
ftp://server/files/jinstall-qfabric-12.2X50-D70.1.rpm
```

Performing a Nonstop Software Upgrade for Director Devices in a Director Group



NOTE: If you reboot any server Node groups or Interconnect devices after you perform a nonstop software upgrade on the Director group, these devices are upgraded to the same version of software that is running on the Director group.

To upgrade the software on the Director devices in a Director group:

- Issue the **request system software nonstop-upgrade director-group package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade director-group
jinstall-qfabric-12.2X50-D70.1.rpm
```

Performing a Nonstop Software Upgrade for Interconnect Devices and Other Fabric-Related Components

Before you perform a nonstop upgrade on the Interconnect devices and other fabric-related components, verify that both Director devices in the Director group are online. Both Director devices must be online before you attempt to perform a nonstop upgrade. To do this, issue the **show fabric administration inventory director-group status** command.

To install the software on the Interconnect device and other components in the fabric:

- Issue the **request system software nonstop-upgrade fabric package-name** command.

For example:

```
user@qfabric> request system software nonstop-upgrade fabric
jinstall-qfabric-12.2X50-D70.1.rpm
```

(Optional) Creating Upgrade Groups for Node Groups

Upgrade groups enable two or more Node devices in a Node group, or an entire Node group, to be rebooted at the same time. If you do not create an upgrade group, the Node devices are upgraded one at a time. Before performing a nonstop upgrade on a Node

group, create an upgrade group and include the devices you want to reboot at the same time.



NOTE: If you add Node devices that have links to the same LAG, there might be traffic loss.

- Create the upgrade group by issuing the **set chassis node-group node-group-name nssu upgrade-group upgrade-group-name node-devices name** command at the [edit chassis] hierarchy.

For example:

```
user@qfabric# set chassis node-group nodegroup1 nssu upgrade-group upgrade1 node-devices
[node1 node2]
```

Performing a Nonstop Software Upgrade on a Node Group

When you perform a nonstop software upgrade on a network Node group, the Node devices in the network Node group are upgraded in a serial fashion except when upgrade groups are configured. If you perform a nonstop software upgrade on a redundant server Node group, both Node devices must be online for a successful upgrade. If one of the Node devices is no longer available, remove it from the configuration before you perform the nonstop software upgrade. If you perform a nonstop software upgrade on a Node group with only one Node device, traffic loss occurs while the Node device is rebooting.



NOTE: You can upgrade multiple Node groups with this command: **request system software nonstop-upgrade node-group node-group-name package-name**. However, if more than one Node group is specified, there may be traffic loss depending on the topology of the network.

To upgrade software on a Node group:

- Issue the **request system software nonstop-upgrade node-group node-group-name package-name** command.

- To perform a nonstop software upgrade on one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group nodegroup1
jinstall-qfabric-12.2X50-D70.1.rpm
```

- To perform a nonstop software upgrade on more than one Node group:

```
user@qfabric> request system software nonstop-upgrade node-group [nodegroup1
nodegroup2 nodegroup3] jinstall-qfabric-12.2X50-D70.1.rpm
```

Performing a Standard Upgrade on the QFabric System

Use the **request system software add component all reboot** CLI command 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.

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.

Perform the following tasks:

- [Backing Up the Current Configuration Files on the QFabric System on page 64](#)
- [Downloading Software Files for the QFabric System Using a Browser on page 64](#)
- [Installing the Software Package on the Entire QFabric System on page 65](#)
- [Reinstalling and Booting the QFabric System from a USB Device on page 65](#)

Backing Up the Current Configuration Files on the QFabric System

To back up your current configuration files:

```
user@switch> request system software configuration-backup path
```

Back up the configuration files to a local directory, remote server, or removable drive (for example, an external USB flash drive).

For example:

```
user@switch> request system software configuration-backup /media/USB/
```

Downloading Software Files for the QFabric System Using a Browser



NOTE: To access the download site, you must have a service contract with Juniper Networks and an access account. If you need help obtaining an account, complete the registration form at the Juniper Networks website <https://www.juniper.net/registration/Register.jsp>.

1. Using a Web browser, navigate to <http://www.juniper.net/support/downloads/junos.html>.
2. In the QFX Series section, click the name of the platform for which you want to download software.
3. Click the Software tab and select 12.2 from the Release drop-down list.
4. In the Install Package section, select the install package for the 12.2X50-D70.1 release.
An Alert box appears.
5. In the Alert box, read the information about the software alert, and click the link to download the software.
A login screen appears.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.

Installing the Software Package on the Entire QFabric System

To install the software on all of the QFabric system components:

1. Issue the **request system software add package-name component all reboot** command.

For example:

```
user@switch> request system software add jinstall-qfabric-12.2X50-D70.1.rpm component
all reboot
```



NOTE: If you receive an error message after issuing the **request system software add package-name component all reboot** command that says that the configuration file cannot be loaded as is, you need to enter configuration mode, make any necessary changes to the configuration file, and then commit the changes.

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

Reinstalling and Booting the QFabric System from a USB Device

You can boot and reinstall software on your Director group. This is especially useful in the event of disaster recovery.

The first section of this procedure describes how to boot and reinstall the software when you do *not* have access to the default partition. The second section of this procedure describes how to boot and reinstall the software when you do have access to the default partition. Before you begin this procedure, make sure you have the software installed on a USB device.



NOTE: After you perform a software installation from a USB drive, some system devices might become disconnected, and the upgrade might not be completed for these devices. As a workaround, turn off the disconnected devices and power them back on to activate the new software.

- To boot and reinstall software on a Director group when you do *not* have access to the default partition:
 1. Connect to one of the Director devices in the Director group by using the management console connection.
 2. Insert the USB device in the USB port in the Director device.

The following menu appears once you are connected to the Director group:

Juniper Networks QFabric Director Install/Recovery Media

- To boot from the USB device, wait 10 seconds or press the <ENTER> key.
- To reinstall the QFabric software on this Director device, type: **install <ENTER>**.
- To perform a network installation on this Director device, type: **network <ENTER>**.

3. Reinstall the software on the Director group by typing **install**, and then pressing **Enter**.

The Director group copies the software from the USB device, occasionally displaying status messages. Copying the software can take up to 10 minutes.

4. Remove the USB device when prompted and then press **Enter**.

The Director group then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the Director group displays the login prompt.

5. Create a new configuration file as you did when the Director group was shipped from the factory, or restore the previously saved configuration file to the Director group.

- To boot and reinstall software on a Director group when you have access to the default partition:

1. Log in to the default partition.

2. Copy the software from the external USB device to the internal storage of the Director group by issuing the **request system software download** command, and specify the path and the package name on the external USB device. For example:

```
user@switch> request system software download  
/media/usbdisk/jinstall-qfabric-12.2X50-D70.1.rpm
```

3. Install the software by issuing the **request system software add package-name component director-group reboot** command, and specify the name of the software package:

```
user@switch> request system software add jinstall-qfabric-12.2X50-D70.1.rpm  
component director-group reboot
```

The Director group installs the software, occasionally displaying status messages. Copying the software can take up to 10 minutes.

4. Remove the external USB device when prompted, and then press **Enter**.

The Director group then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the Director group displays the login prompt.

5. Create a new configuration as you did when the Director group was shipped from the factory, or restore the previously saved configuration file to the Director group.

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

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

downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. However, you cannot upgrade directly from a non-EEOL release that is more than three releases before or after.

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

Related Documentation

- [New Features in Junos OS Release 12.2X50-D10 for the QFX Series on page 14](#)
- [Limitations in Junos OS Release 12.2 for the QFX Series on page 29](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for the QFX Series on page 27](#)
- [Outstanding Issues in Junos OS Release 12.2 for the QFX Series on page 32](#)
- [Resolved Issues in Junos OS Release 12.2 for the QFX Series on page 44](#)
- [Errata in Documentation for Junos OS Release 12.2 for the QFX Series on page 53](#)

QFX Series Documentation for Junos OS Release 12.2

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 12.2</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 12.2X50-D70 (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, or call 1-888-314-JTAC (from the United States, Canada, or Mexico) or 1-408-745-9500 (from elsewhere).

Revision History

10 April 2015—Revision 15, Junos OS for the QFX Series, Release 12.2X50-D70—update to Changes

22 January 2015—Revision 14, Junos OS for the QFX Series, Release 12.2X50-D70

30 July 2014—Revision 13, Junos OS for the QFX Series, Release 12.2X50-D65

13 March 2013—Revision 12, Junos OS for the QFX Series, Release 12.2X50-D60

22 November 2013—Revision 11, Junos OS for the QFX Series, Release 12.2X50-D55

8 August 2013—Revision 10, Junos OS for the QFX Series, Release 12.2X50-D50

11 March 2013—Revision 9, Junos OS for the QFX Series, Release 12.2X50-D40

7 January 2013—Revision 8, Junos OS for the QFX Series, Release 12.2X50-D30

30 November 2012—Revision 7, Junos OS for the QFX Series, Release 12.2X50-D30

27 September 2012—Revision 6, Junos OS for the QFX Series, Release 12.2X50-D20

26 September 2012—Revision 5, Junos OS for the QFX Series, Release 12.2X50-D20

19 September 2012—Revision 4, Junos OS for the QFX Series, Release 12.2X50-D20

17 September 2012—Revision 3, Junos OS for the QFX Series, Release 12.2X50-D20

5 September 2012—Revision 2, Junos OS for the QFX Series, Release 12.2X50-D10

6 July 2012—Revision 1, Junos OS for the QFX Series, Release 12.2X50-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.