

Release Notes: Junos[®] OS Release 15.1X53-D20 for QFX10002 Switches

Release 15.1X53-D20
6 November 2015
Revision 2

Contents

| | |
|--|----|
| Junos OS Release Notes for QFX10002 Switches | 3 |
| New and Changed Features | 3 |
| New Features in Release 15.1X53-D20 | 3 |
| Hardware | 3 |
| Interfaces and Chassis | 4 |
| Software-Defined Networking (SDN) | 4 |
| New Features in Release 15.1X53-D10 | 4 |
| Hardware | 4 |
| Interfaces and Chassis | 4 |
| Layer 2 Features | 6 |
| Layer 3 Features | 6 |
| Layer 3 Logical Interfaces | 7 |
| Multicast Protocols | 7 |
| Multiprotocol Label Switching (MPLS) | 8 |
| Network Management and Monitoring | 8 |
| Open vSwitch Database (OVSDB) | 9 |
| Security | 9 |
| Storage | 10 |
| System Management | 10 |
| Traffic Management | 10 |
| VXLAN | 11 |
| Changes in Behavior and Syntax | 11 |
| Software-Defined Networking (SDN) | 12 |
| Known Behavior | 12 |
| Interfaces and Chassis | 12 |
| System Management | 12 |
| Known Issues | 13 |
| Hardware | 13 |
| Class of Service | 13 |
| Firewall Filters | 14 |

| | |
|--|----|
| Interfaces and Chassis | 14 |
| IPv4 | 14 |
| Layer 2 Protocols | 14 |
| Layer 3 Protocols | 14 |
| OVSDB | 15 |
| Security | 15 |
| Services | 15 |
| Spanning Tree Protocols | 15 |
| System Management | 15 |
| Traffic Management | 16 |
| Resolved Issues | 16 |
| Resolved Issues: Release 15.1X53-D20 | 16 |
| Class of Service | 17 |
| Interfaces and Chassis | 17 |
| Layer 3 Protocols | 17 |
| Multicast Protocols | 17 |
| OVSDB | 17 |
| Documentation Updates | 18 |
| Migration, Upgrade, and Downgrade Instructions | 18 |
| Downloading Software Files with a Browser | 18 |
| Backing Up the Current Configuration Files | 19 |
| Installing the Software | 19 |
| Product Compatibility | 20 |
| Hardware Compatibility | 20 |
| Documentation Feedback | 20 |
| Requesting Technical Support | 21 |
| Self-Help Online Tools and Resources | 21 |
| Opening a Case with JTAC | 21 |
| Revision History | 22 |

Junos OS Release Notes for QFX10002 Switches

These release notes accompany Junos OS Release 15.1X53-D20 for QFX10002 switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1X53-D20 for QFX10002 switches.

- [New Features in Release 15.1X53-D20 on page 3](#)
- [New Features in Release 15.1X53-D10 on page 4](#)

New Features in Release 15.1X53-D20

Hardware

Qualified 100-gigabit QSFP28 optical transceivers (QFX10002 switch)—Starting with 15.1X53-D20, the 100-gigabit QSFP28 optical transceivers described in [Table 1 on page 3](#) are qualified for use with QFX10002 switches.

Table 1: Qualified 100-Gigabit QSFP28 Optical Transceivers

| Transceiver | Product SKU Number | Range |
|--------------|--------------------|---|
| 100GBASE-SR4 | JNP-QSFP-100G-SR4 | Center wavelengths 840 nm to 860 nm |
| 100GBASE-LR4 | JNP-QSFP-100G-LR4 | Transmit and receive lane wavelengths 1294.53 nm through 1296.59 nm 1299.02 nm through 1301.09 nm 1303.54 nm through 1305.63 nm 1308.09 nm through 1310.19 nm |

[See [Interface Specifications for QSFP28 Optical Transceivers for the QFX10002.](#)]

Interfaces and Chassis

- **Precision Time Protocol (PTP) transparent clock (QFX10002 switch)**—PTP synchronizes clocks throughout a packet-switched network. With a transparent clock, the PTP packets are updated with residence time as the packets pass through the switch. There is no master/slave designation. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that timestamps are sent in one packet. In a two-step process, estimated timestamps are sent in one packet, and additional packets contain updated timestamps. In addition, user UDP over IPv4 and IPv6, and unicast and multicast transparent clocks, are supported. You can configure the transparent clock at the `[edit protocols ptp]` hierarchy level.

Software-Defined Networking (SDN)

OVSDB software in Junos OS software package (QFX10002 switch)—OVSDB software is included in the jsdn package. Prior to Junos OS Release 15.1X53-D20, to install OVSDB software on a QFX10002 switch, you had to install a separate jsdn package (`jsdn-i386-release`) in addition to the Junos OS software (`jinstall`) package. Starting with Junos OS Release 15.1X53-D20, the jsdn package is now included in the jinstall package. The introduction of this new feature means that to install OVSDB software on a QFX10002 switch, you no longer need to download and install the separate jsdn package.

New Features in Release 15.1X53-D10

Hardware

- **QFX10002-72Q switch**—The Juniper Networks QFX10000 line of Ethernet switches provides cloud builders and data center operators with scalable solutions for both core and spine data center deployments. QFX10002-72Q provides 5.76 terabytes of throughput and up to 4 billion packets per second (pps) of forwarding capacity. In native mode, QFX10002-72Q offers 72 ports of 40-gigabit QSFP+. Twenty-four ports are designed to be 100-gigabit capable using QSFP28. Each 40-gigabit QSFP+ port can be configured either as a native 40-gigabit port or as four 10-gigabit ports using a breakout cable. With the breakout cable, the switch supports a maximum of 288 logical 10-Gigabit Ethernet ports.
- **QFX10002-36Q switch**—The QFX10000 line of Ethernet switches provides cloud builders and data center operators with scalable solutions for both core and spine data center deployments. QFX10002-36Q provides 2.88 terabytes of throughput and up to 2 billion pps of forwarding capacity. In native mode, QFX10002-36Q offers 36 ports of 40 gigabit QSFP+. Twelve ports are designed to be 100-gigabit capable using QSFP28. Each 40-gigabit QSFP+ port can be configured either as a native 40-gigabit port or as four 10-gigabit ports using a breakout cable. With the breakout cable, the switch supports a maximum of 144 logical 10-Gigabit Ethernet ports.

Interfaces and Chassis

- **Adaptive load balancing (ALB) for aggregated Ethernet bundles (QFX10002 switch)**—ALB evenly distributes data flows across aggregated Ethernet member links. You use ALB to manage uneven or overloaded data flows on member links. ALB

supports up to 64 member links and up to 50 aggregated Ethernet bundles. The algorithm determines which link to use by considering the scanned packet or bit rate associated with each hash value in conjunction with the mapping of hash values to a given link. ALB can be applied to IPv4, IPv6, and MPLS packet headers. ALB is disabled by default.

Configure ALB by setting the adaptive statement at the **[edit interfaces ae-interface aggregated-ether-options load-balance]** hierarchy level. Under the **load-balance** statement, you can set the following ALB options:

- **scan-interval interval**—Scan interval in multiples of 30 seconds to check the tolerance deviation. The range is 1 to 5. The default is 1.
- **bps**—Scan traffic in bits per second (bps). The default is bits per second.
- **pps**—Scan traffic in packets per second (pps).
- **Channelizing 40-Gigabit Ethernet QSFP+ ports (QFX10002 switch)**—This feature enables you to channelize four 10-Gigabit Ethernet interfaces from the 40-Gigabit Ethernet QSFP+ interfaces. By default, the 40-Gigabit Ethernet QSFP+ interfaces are named **et-fpc/pic/port**. The resulting 10-Gigabit Ethernet interfaces appear in the following format: **xe-fpc/pic/port:channel**, where channel can be a value of 0 through 3. To channelize a 40-Gigabit Ethernet QSFP+ interface into four 10-Gigabit Ethernet interfaces, include the **10g** statement at the **[edit chassis fpc fpc-slot pic pic-slot (port port-number | port-range port-range-low port-range-high) channel-speed]** hierarchy level. To revert the 10-Gigabit Ethernet channels to a full 40-Gigabit Ethernet interface, remove the **10g** statement from the same hierarchy level.
- **Link aggregation (QFX10002 switch)**—Link aggregation enables you to use multiple network cables and ports in parallel to increase link speed and redundancy.
- **Multichassis link aggregation group (MC-LAG) (QFX10002 switch)**—MC-LAG enables a client device to form a logical LAG interface using two QFX10002 switches. MC-LAG provides redundancy and load balancing between the two QFX10002 switches, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG is an MC-LAG client that has one or more physical links in a LAG. This client does not need to be aware of the MC-LAG. On the other side of the MC-LAG are two MC-LAG QFX10002 switches. Each of these QFX10002 switches has one or more physical links connected to a single client. The QFX10002 switches coordinate with each other to ensure that data traffic is forwarded properly.

To configure an MC-LAG, 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

Layer 2 Features

- **VLAN support (QFX10002 switch)**—VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
- **Link Layer Discovery Protocol (LLDP) support (QFX10002 switch)**—LLDP enables a switch to advertise its identity and capabilities on a LAN in addition to receiving information about other network devices.
- **Q-in-Q tunneling support (QFX10002 switch)**—This feature allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag.
- **Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) support (QFX10002 switch)**—These protocols enable a switch to advertise its identity and capabilities on a LAN and to receive information about other network devices.

Layer 3 Features

- **BGP support (QFX10002 switch)**—BGP is an exterior gateway protocol (EGP) for routing traffic between autonomous systems (ASs). You can configure BGP at the `[edit protocols bgp]` hierarchy level.
- **OSPF support (QFX10002 switch)**—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an AS. QFX10002 switches support OSPFv1 and OSPFv2. You can configure OSPF at the `[edit protocols ospf]` hierarchy level.
- **Bidirectional Forwarding Detection (BFD) support for static routes and the BGP, IS-IS, OSPF, PIM, and RIP protocols (QFX10002 switch)**— BFD uses control packets and shorter detection time limits to rapidly detect failures in a network. Hello packets are sent at a specified, regular interval by routing devices. A neighbor failure is detected when a routing device stops receiving a reply after a specified interval.

On a QFX10002 switch, you can configure BFD for static routes and the BGP, IS-IS, OSPF, PIM, and RIP protocols.

- **IS-IS support (QFX10002 switch)**—The IS-IS protocol is an IGP for routing traffic within an AS.
- **Virtual Router Redundancy Protocol (VRRP) support (QFX10002 switch)**—VRRP enables you to provide alternative gateways for end hosts that are configured with static default routes. You can implement VRRP to provide a highly available default path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.

Layer 3 Logical Interfaces

- **Support for Layer 3 logical interfaces (QFX10002 switch)**—A Layer 3 logical interface is a logical division of a physical interface or an aggregated Ethernet interface that operates at the network level and that can receive and forward IEEE 802.1Q VLAN tags. You can use these interfaces to route traffic between multiple VLANs along a single trunk line that connects a QFX10002 switch to a Layer 2 switch. Only one physical connection is required between the switches.
- **Generic routing encapsulation (GRE) support (QFX10002 switch)**—You can use GRE tunneling services to encapsulate any network layer protocol over an IP network. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first adds a GRE header and then adds an outer IP header that is used to route the packet. When it receives the packet, a switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect noncontiguous networks and to provide options for networks that contain protocols with limited hop counts.

Multicast Protocols

- **Internet Group Management Protocol (IGMP) support (QFX10002 switch)**—IGMP manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.
- **IGMP snooping support (QFX10002 switch)**—IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.
- **Protocol Independent Multicast-sparse mode support (QFX10002 switch)**—PIM sparse mode enables efficient routing to multicast groups with receivers that are sparsely spread over multiple networks. To configure PIM sparse mode, include the `pim` statement at the [edit protocols] hierarchy level.
- **PIM source-specific multicast (PIM SSM) support (QFX10002 switch)**—PIM SSM uses a subset of PIM sparse mode and IGMPv3 to enable a client to receive multicast traffic directly from the source. PIM-SSM uses the PIM sparse-mode functionality to create a shortest-path tree (SPT) between the client and the source, but builds the SPT without the help of a rendezvous point.
- **Multicast Source Discovery Protocol (MSDP) support (QFX10002 switch)**—MSDP enables you to connect multiple domains to one another. MSDP typically runs on the same routing device as a PIM sparse mode rendezvous point. Each MSDP routing device establishes adjacencies with internal and external MSDP peers, similar to how BGP peering works. These peers inform each other about active sources within the domain. When they detect active sources, the peers send PIM sparse mode explicit join messages

to the active source. To configure MSDP, include the **msdp** statement at the **[edit protocols]** hierarchy level and specify groups of local addresses and MSDP peer addresses.

- **Rendezvous point (RP) support (QFX10002 switch)**—This feature supports multiple rendezvous points using anycast addresses (RPs sharing a single routable IP address) in either a PIM or MSDP-enabled network. To configure anycast RP, include the **anycast-pim** statement at the **[edit protocols pim rp local family inet]** hierarchy level.
- **IGMP querier support (QFX10002 switch)**—This feature enables multicast traffic to be forwarded between connected switches in pure Layer 2 networks. If you enable IGMP snooping in a Layer 2 network without a multicast router, the IGMP snooping reports are not forwarded between connected switches. This means that if hosts that are connected to different switches in the network join the same multicast group, and traffic for that group arrives on one of the switches, the traffic is not forwarded to the other switches that have hosts that should receive the traffic. If you enable IGMP querying for a VLAN, multicast traffic is forwarded between switches that participate in the VLAN if they are connected to hosts that are members of the relevant multicast group.

Multiprotocol Label Switching (MPLS)

- **MPLS support (QFX10002 switch)**—MPLS provides both label edge router (LER) and label switch router (LSR) support and provides the following capabilities:
 - Support for both MPLS major protocols, LDP and RSVP
 - IS-IS interior gateway protocol (IGP) traffic engineering
 - Class of service (CoS)
 - Object access method, including ping, traceroute, and Bidirectional Forwarding Detection (BFD)
 - Fast reroute (FRR), a component of MPLS local protection
 - Both one-to-one local protection and many-to-one local protection are supported.
 - Loop free alternate (LFA) FRR
 - 6PE devices
 - Layer 3 VPNs for both IPv4 and IPv6
 - LDP tunneling over RSVP

Network Management and Monitoring

- **SNMP support (QFX10002 switch)**—SNMP includes SNMP versions 1, 2, and 3 for monitoring system activity.
- **System logging (syslog) support (QFX10002 switch)**—Syslog enables you to log system messages into a local directory on the switch or to a syslog server.
- **sFlow technology support (QFX10002 switch)**—This feature provides monitoring technology for high-speed switched or routed networks. You can configure sFlow

technology to monitor traffic continuously at wire speed on all interfaces simultaneously. sFlow technology also collects samples of network packets, providing you with visibility into network traffic information. You configure sFlow monitoring at the **[edit protocols sflow]** hierarchy level. sFlow operational commands include **show sflow** and **clear sflow collector statistics**.

- **Port mirroring (SNMP) support (QFX10002 switch)**—Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Open vSwitch Database (OVSDB)

- **OVSDB support (QFX10002 switch)**—The Junos OS implementation of the OVSDB management protocol provides a means through which VMware NSX controllers and QFX100002 switches can communicate. In an NSX for multi-hypervisor environment, NSX controllers and QFX100002 switches can exchange control and statistical information through the OVSDB schema for physical devices. The ability to exchange this information enables virtual machine (VM) traffic from entities in virtual networks to be forwarded to entities in the physical network and vice versa.

You can set up a connection between the QFX10002 management interface and an NSX controller.

Security

- **Firewall filter support (QFX10002 switch)**—You can provide rules that define whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces.
- **Policing support (QFX10002 switch)**—You can use policing to apply limits to traffic flow and to set consequences for packets that exceed those limits.
- **MAC limiting support (QFX10002 switch)**—You can protect a LAN against flooding by setting a limit on the number of MAC addresses that can be learned from the Layer 2 access interfaces on a switch.
- **MAC move limiting support (QFX10002 switch)**—You can detect MAC movement and MAC spoofing on access ports.
- **Storm control support (QFX10002 switch)**—You can enable the switch to monitor traffic levels and take a specified action when a specified traffic level—called the storm control level—is exceeded, preventing packets from proliferating and degrading service. You can configure a switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs.

Storage

- **FCoE transit switch support (QFX10002 switch)**—You can configure a QFX10002 switch as a Fibre Channel over Ethernet (FCoE) transit switch that transports FCoE frames across the Ethernet network and supports the following data center bridging (DCB) standards: priority-based flow control (PFC) and Data Center Bridging Exchange Capability (DCBX) protocol.

System Management

- **Fabric management support (QFX10002 switch)**—Starting with Junos OS Release 15.1X53-D10, you can set up and manage the fabric connections between the Packet Forwarding Engines in the switch. Fabric management collects fabric statistics, monitors hardware health, and responds to CLI queries. It also tracks when you add or remove FRUs from the switch and monitors faults in the data plane. It is enabled by default and can be monitored by using the following operational mode commands:
 - **show chassis fabric summary**—Display summary status information for the fabric.
 - **show chassis fabric errors fpc <fpc-slot>**—Display error information related to an FPC in the fabric.
 - **show chassis fabric fpcs fpc <fpc-slot>**—Display information for FPCs in the fabric.
- **Login authentication using RADIUS and TACACS+ (QFX10002 switch)**—You can use RADIUS and TACACS+ authentication to validate users who attempt to access the switch.
- **System utilization alarms support (QFX10002 switch)**—This feature provides system alarms to alert you of high disk usage in the /var partition on the switch. You can display these alarm messages by issuing the **show system alarms** operational mode command if the /var partition usage is higher than 75 percent. A usage level between 76 and 90 percent indicates high usage and raises a minor alarm condition, whereas a usage level over 90 percent indicates that the partition is full and raises a major alarm condition.
- **Zero Touch Provisioning (ZTP) (QFX10002 switch)**—ZTP allows you to provision new Juniper Networks switches in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network.

Traffic Management

- **CoS rewrite rules support (QFX10002 switch)**—You can use rewrite rules to set the value of the CoS bits within a packet header, so you can alter the CoS settings of incoming packets.
- **Queue shaping support (QFX10002 switch)**—You can manage excess traffic and avoid congestion on a network interface where traffic might exceed the maximum port bandwidth.

- **Priority-based flow control support (QFX10002 switch)**—This feature provides you with PFC (standard IEEE 802.1Qbb) capability, a link-level flow control mechanism that you can use to pause traffic selectively according to its class. You must use PFC for Fibre Channel over Ethernet (FCoE) traffic.
- **Ethernet PAUSE autonegotiation support (QFX10002 switch)**—You can configure asymmetric flow control. To configure PAUSE, include both the **rx-buffers** and **tx-buffers** statements at the **[edit interfaces interface-name ether-options configured-flow-control]** hierarchy level. The **rx-buffers** statement determines whether or not the interface generates and sends PAUSE messages. The **tx-buffers** statement determines whether or not the interface responds to received PAUSE messages.

VXLAN

- **Layer 2 VXLAN gateway (QFX10002 switch)**—You can stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. You can use VXLAN tunnels to enable migration of virtual machines (VMs) between servers that exist in separate Layer 2 domains by tunneling the traffic through Layer 3 networks. This functionality enables you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use STP to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

The QFX10002 switch supports only VXLAN with OVSDB.

Related Documentation

- [Changes in Behavior and Syntax on page 11](#)
- [Known Behavior on page 12](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 20](#)

Changes in Behavior and Syntax

The following are changes in known behavior of Junos OS features or changes in the syntax of Junos OS statements and commands for Junos OS Releases 15.1X53 for QFX10002 switches.

- [Software-Defined Networking \(SDN\)](#)

Software-Defined Networking (SDN)

- **Automatic configuration of trunk interfaces that handle untagged packets in OVSDB-managed VXLANs (QFX10002 switch)**—In previous Junos OS releases, if you specified a VLAN ID of 0 for a logical switch port in VMware NSX Manager or in the NSX API, the QFX10002 switch automatically configured an access interface to handle untagged packets in the associated Open vSwitch Database (OVSDB) –managed Virtual Extensible LAN (VXLAN). Starting with 15.1X53-D20, specifying a VLAN ID of 0 in a logical switch port configuration causes the QFX10002 switch to automatically configure a trunk port. To enable the trunk port to handle untagged packets, the QFX10002 switch also configures a native VLAN with an ID of 4094.

This change supports the division of an OVSDB-managed physical interface into multiple logical interfaces, some of which are associated with VXLANs that have untagged packets and some of which are associated with VXLANs that have tagged packets.

Related Documentation

- [New and Changed Features on page 3](#)
- [Known Behavior on page 12](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 20](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Releases 15.1X53 for QFX10002 switches.

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

- [Interfaces and Chassis on page 12](#)
- [System Management on page 12](#)

Interfaces and Chassis

On a QFX10002 switch, when you apply MC-LAG using **apply-groups** to configure , the commit might fail and display an error message that states **IRB interface(irb.1) and I2-interface(ae0.0) do not belong to the same routing instance**. As a workaround, apply the MC-LAG configuration in the main configuration instead of under **groups**.

System Management

On a QFX10002 switch, the **request system snapshot** command does not work.

- Related Documentation**
- [New and Changed Features on page 3](#)
 - [Changes in Behavior and Syntax on page 11](#)
 - [Known Issues on page 13](#)
 - [Resolved Issues on page 16](#)
 - [Documentation Updates on page 18](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 18](#)
 - [Product Compatibility on page 20](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Releases 15.1X53 for QFX10002 switches.

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

- [Hardware on page 13](#)
- [Class of Service on page 13](#)
- [Firewall Filters on page 14](#)
- [Interfaces and Chassis on page 14](#)
- [IPv4 on page 14](#)
- [Layer 2 Protocols on page 14](#)
- [Layer 3 Protocols on page 14](#)
- [OVSDDB on page 15](#)
- [Security on page 15](#)
- [Services on page 15](#)
- [Spanning Tree Protocols on page 15](#)
- [System Management on page 15](#)
- [Traffic Management on page 16](#)

Hardware

- On a QFX10002 switch, when you insert breakout cables into a port on a QSFP+ interface, the LEDs on that port might not work correctly. [PR1057500](#)

Class of Service

- On a QFX10002 switch, the command **show interfaces queue** does not display packets dropped for loss-priority high. [PR1088514](#)
- Loss-priority settings help determine which packets are dropped from the network during periods of congestion. The four CoS behavior aggregate (BA) classifier packet loss priority (PLP) values for Juniper Networks devices are low, medium-low,

medium-high, and high. On the QFX10002 switch only, CoS rewrite rules do not work for IPV6 traffic when the PLP is set to either medium-high or high. [PR1093543](#)

Firewall Filters

- On a QFX10002 switch, when you configure a filter to discard VLAN flood, the filter also discards known unicast traffic. [PR1054677](#)
- If you configure a discard and log action on an ingress loopback interface of a QFX10002 switch, log messages might display incorrect protocol source and destination addresses. [PR1069494](#)
- On QFX10002 switches, firewall filters will not work on MAC addresses starting with 01:80:C2. The packets are treated as special packets, and the access control lists (ACLs) will not match these addresses. [PR1085374](#)

Interfaces and Chassis

- On a QFX10002 switch, transit statistics are not collected on integrated routing and bridging (IRB) interfaces. Instead, transit statistics show 0. Only local statistics are displayed. [PR1080543](#)
- On a QFX10002 switch, when you reboot a member of an MC-LAG that is in standby mode, after the interchassis link is reestablished, there might be a loss of up to 5 seconds for multicast traffic. [PR1094388](#)
- If a GRE tunnel with an egress interface is configured on a QFX10002 switch and you change the MTU of or add or delete an underlying egress subinterface, the GRE tunnel might drop packets. If this situation occurs, you can deactivate and then reactivate the egress interface to resolve the issue. [PR1114991](#)
- On a QFX10002 switch, if you connect 100G optical transceivers and add the chassis configuration for all 100G ports in a single commit, some of the 100G subinterface's might not be created. To work around this situation, you can configure the 100G optical transceivers first, then connect the transceivers. [PR1130642](#)

IPv4

- If the **no-gratuitous-arp-reply** statement is specified in a configuration, the QFX10002 switch still receives and processes gratuitous ARP packets. [PR1067593](#)

Layer 2 Protocols

- On a QFX10002 switch, you cannot configure an IRB interface on an interface that is configured for flexible VLAN tagging. This means that you cannot include an IRB interface in a Q-in-Q configuration. [PR1072304](#)

Layer 3 Protocols

- On a QFX10002 switch, the GRE header is not included in the aggregated ethernet and equal-cost multipath (ECMP) hash computation. [PR1039863](#)

- On a QFX10002 switch, enhanced hash key configuration using the enhanced-hash-key statement is not supported. [PR1045559](#)
- If the link between a QFX10002 switch and a leaf switch is down, the IPv6 and IPv4 convergence time can be greater than 1 second. [PR1125748](#)

OVSDB

- On a QFX10002 switch, the collection of statistics for OVSDB-managed interfaces is not supported. As a result, even though an OVSDB-managed interface is up and running, the **show ovbdb statistics** command output for this interface displays 0. [PR1090363](#)
- If a VXLAN tunnel to a service node is over an integrated routing and bridging (IRB) interface with an underlying aggregated Ethernet (AE) interface on a QFX10002 switch, broadcast, unknown unicast, and multicast traffic might not be forwarded to the service node. [PR1133126](#)

Security

- On a QFX10002 switch, if you configure storm control and include the **no-unregistered-multicast** statement to exclude this traffic from storm control, unregistered multicast traffic is still subject to storm control. [PR1079556](#)

Services

- On a QFX10002 switch, when a new interface is added to an existing link aggregation group (LAG) interface that acts as an input analyzer interface, traffic sent to the added interface might not be mirrored. [PR1057527](#)

Spanning Tree Protocols

- On QFX10002 interfaces on which the **flexible-vlan-tagging** statement is specified, STP, RSTP, MSTP, and VSTP are not supported. [PR1075230](#)

System Management

- On a QFX10002 switch, rolling back the software by issuing the **request system software rollback** is not supported. [PR1070892](#)
- On a QFX10002 switch, if you reboot the switch, and then insert a small form-factor pluggable (SFP) transceiver on the management interface, the management interface might not work properly. As a workaround, issue the **request system reboot hypervisor** command after you insert the SFP on the management interface. [PR1075097](#)
- On a QFX10002 switch, loading a scaled configuration might cause disk space issues. [PR1076356](#)
- If you reboot a QFX10002 switch, the switching fabric might not come up, thereby rendering the Flexible PIC Concentrator (FPC) in an offline state. If this issue occurs,

you can power cycle the switch or issue the **request system reboot hypervisor** command to recover from this state. [PR1080877](#)

- On a QFX10002 switch, If you run `gcore` on the `dcpfe` process in the Host OS, the Host OS might become nonfunctional. [PR1126889](#)

Traffic Management

- On a QFX10002 switch, the class-of-service rewrite rule configuration is not supported on the egress interface of a GRE tunnel. If you apply a rewrite rule to the egress interface of a GRE tunnel, it might generate a `dcpfe` core dump. [PR1078849](#)

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 11](#)
- [Known Behavior on page 12](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 20](#)

Resolved Issues

This section lists the issues fixed in the Junos OS 15.1X53 releases.

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

- [Resolved Issues: Release 15.1X53-D20 on page 16](#)

Resolved Issues: Release 15.1X53-D20

- [Class of Service](#)
- [Interfaces and Chassis](#)
- [Layer 3 Protocols](#)
- [Multicast Protocols](#)
- [OVSDB](#)

Class of Service

- On a QFX10002 switch, when you configure a custom EXP (experimental bits) classifier on a link aggregation group (LAG), the default DiffServ code point (DSCP) classifier stops working. [PR1097142](#)

Interfaces and Chassis

- On a QFX10002 switch, if the system time changes to the year 2050 and later, because of the Network Time Protocol (NTP) or for any other reason, Link Aggregation Control Protocol (LACP) might go down and stay down. As a workaround, set the system time to the current time and day. [PR1095658](#)
- On a QFX10002 switch, when you issue the **monitor interface statistics** command for Layer 3 tagged subinterfaces, the packet per second field does not display the correct value. Instead, it displays a value of 0. As a workaround, issue the **monitor interface statistics** command for physical interfaces associated with the aggregated Ethernet interface. [PR1096024](#)

Layer 3 Protocols

- On a QFX10002 switch, when you add and delete a Bidirectional Forwarding Detection (BFD) configuration multiple times, BFD sessions might go down. [PR1066780](#)
- On a QFX10002 switch, in a scaled environment, when the Packet Forwarding Engine adds, modifies, or deletes routes, these actions can be slow. If there are large numbers of routes and there is a sudden change that results in many routes per next hop, convergence issues might occur. [PR1088832](#)

Multicast Protocols

- On a QFX10002 switch, multicast traffic that ingresses from a GRE tunnel is not de-encapsulated and is dropped. [PR1089319](#)

OVSDB

- On a QFX10002 switch, OVSDB-managed interfaces do not support mixing access and tagged logical interfaces. [PR1093061](#)

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 11](#)
- [Known Behavior on page 12](#)
- [Known Issues on page 13](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 20](#)

Documentation Updates

There are no errata or changes in Junos OS Releases 15.1X53 for QFX10002 switch documentation.

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 11](#)
- [Known Behavior on page 12](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 16](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 20](#)

Migration, Upgrade, and Downgrade Instructions

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

- [Downloading Software Files with a Browser on page 18](#)
- [Backing Up the Current Configuration Files on page 19](#)
- [Installing the Software on page 19](#)

Downloading Software Files with a Browser

To download the software package from the Juniper Networks Support website, go to <http://www.juniper.net/support/>.



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>.

This procedure shows you how to upgrade software on a QFX10002 switch.

1. Using a Web browser, navigate to <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the By Technology box, click **Switching | QFX Series | QFX10002**.
4. In the QFX Series section, click the name of the platform for which you want to download software.
5. Click the **Software** tab and select the install package from the Install Package box.
A login screen appears.

6. Enter your name and password, and press **Enter**.
7. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
8. Save the `jinstall-qfx-<version>-domestic-signed.tgz` file on your computer.
9. Open or save the installation package either to the local system in the `var/tmp` directory or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.

Backing Up the Current Configuration Files

Before you install the new installation package, we strongly recommend that you back up your current configuration files, because the upgrade process removes all of the stored files on the switch.

To back up your current configuration files:

```
user@switch# save filename
```

Executing this command saves a copy of your configuration files to a remote location such as an external USB device.

Installing the Software



NOTE: On the switch, use the `force-host` option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the `force-host` option.

If the installation package resides locally on the switch, execute the `request system software add validate <pathname> <source> reboot` command.

For example:

```
user@switch> request system software add validate
/var/tmp/jinstall-qfx-10-f-15.1X53-D10-domestic.tgz reboot
```

If the Install Package resides remotely from the switch, execute the `request system software add validate <pathname> <source> reboot` command.

For example:

```
user@switch> request system software add validate
ftp://ftpserverserver/directory/jinstall-qfx-10-f-15.1X53-D10-domestic.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the `show version` command.

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

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 11](#)
- [Known Behavior on page 12](#)

- [Known Issues on page 13](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 18](#)
- [Product Compatibility on page 20](#)

Product Compatibility

- [Hardware Compatibility on page 20](#)

Hardware Compatibility

To obtain information about the components that are supported on the devices, and special compatibility guidelines with the release, see the Hardware Guide for the product.

To determine the features supported on QFX10002 switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

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

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 11](#)
- [Known Behavior on page 12](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 16](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

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

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

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

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

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

Revision History

6 November 2015—Revision 2, Junos OS for QFX10002 Switches, Release 15.1X53-D20

29 October 2015—Revision 1, Junos OS for QFX10002 Switches, Release 15.1X53-D20

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.