

Release Notes: Junos[®] OS Release 15.1X53-D62 for QFX10000 Switches

Release 15.1X53-D62
February 14, 2017
Revision 2

Contents

Junos OS Release Notes for QFX10000 Switches	3
New and Changed Features	3
New Features in Release 15.1X53-D61	3
Hardware	3
New Features in Release 15.1X53-D60	3
Interfaces and Chassis	3
Layer 2 VPNs	5
Routing Protocols	5
Software-Defined Networking (SDN)	5
Software Installation and Upgrade	6
New Features in Release 15.1X53-D30	7
Hardware	7
High Availability and Resiliency	8
Infrastructure	8
Interfaces and Chassis	8
Layer 2 Features	11
Layer 3 Features	11
Multicast Protocols	12
Multiprotocol Label Switching (MPLS)	13
Network Management and Monitoring	14
Security	14
Software-Defined Networking (SDN)	15
Software Installation and Upgrade	16
Storage	17
System Management	17
Traffic Management	18
Virtual Private Networks (VPNs)	18
Changes in Behavior and Syntax	19
Software-Defined Networking (SDN)	19
VXLAN	19

Known Behavior	20
Interfaces and Chassis	20
Layer 3 Features	20
Routing Protocols	20
Routing Proto	20
Known Issues	21
High Availability (HA) and Resiliency	21
Interfaces and Chassis	21
Layer 2 Features	22
MPLS	22
Multicast Protocols	22
Network Management and Monitoring	22
Platform and Infrastructure	23
Routing Protocols	23
Resolved Issues	24
Resolved Issues: Release 15.1X53-D62	24
Interfaces and Chassis	24
Network Management and Monitoring	25
Resolved Issues: Release 15.1X53-D61	25
Interfaces and Chassis	25
MPLS	25
Routing Protocols	25
Software-Defined Networking (SDN)	26
Resolved Issues: Release 15.1X53-D32	26
Interfaces and Chassis	26
Network Management and Monitoring	26
Documentation Updates	26
Changes to Junos OS for QFX10000 Switches Documentation	27
MPLS	27
Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches	27
Caveats for Downgrading from Junos OS Release 15.1X53-D60 to Previous Software Releases on QFX10000 Switches	27
Upgrading Requires Manual Copy of /var/db/scripts Files	29
Downloading Software Files with a Browser	29
Backing Up the Current Configuration Files	29
Installing the Software	30
Product Compatibility	31
Hardware Compatibility	31
Documentation Feedback	31
Requesting Technical Support	32
Self-Help Online Tools and Resources	32
Opening a Case with JTAC	32
Revision History	33

Junos OS Release Notes for QFX10000 Switches

These release notes accompany Junos OS Release 15.1X53-D62 for QFX10000 switches.

New and Changed Features

This section describes the new features in Junos OS Release 15.1X53 for QFX10000 switches.

- [New Features in Release 15.1X53-D61 on page 3](#)
- [New Features in Release 15.1X53-D60 on page 3](#)
- [New Features in Release 15.1X53-D30 on page 7](#)

New Features in Release 15.1X53-D61

Hardware

- The Juniper Networks QFX10016 modular data center spine and core Ethernet switch provides cloud and data center operators with high-levels scale and throughput. The largest of the QFX10000 line of switches, the QFX10016 can provide 96 Tbps of throughput and 32 Bpps of forwarding capacity in a 21 rack unit (21 U) chassis. The QFX10016 has 16 slots for line cards that allow for a smooth transition from 10-Gigabit Ethernet and 40-Gigabit Ethernet networks to 100-Gigabit Ethernet high-performance networks.

New Features in Release 15.1X53-D60

- [Interfaces and Chassis](#)
- [Layer 2 VPNs](#)
- [Routing Protocols](#)
- [Software-Defined Networking \(SDN\)](#)
- [Software Installation and Upgrade](#)

Interfaces and Chassis

- **Configuration support to improve MC-LAG Layer 2 and Layer 3 convergence (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D60, you can configure multichassis link aggregation (MC-LAG) interfaces to improve Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet link goes down or comes up in a bridge domain. To use this feature, ensure that the Inter-Chassis Link (ICL) is configured on an aggregated Ethernet interface. For Layer 2 convergence, configure the **enhanced-convergence** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. For Layer 3 convergence, configure the **enhanced-convergence** statement on an integrated routing and bridging (IRB) interface at the **[edit interfaces irb unit unit-number]** hierarchy level.
- **Configuration synchronization for MC-LAG (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D60, multichassis link aggregation group (MC-LAG) configuration synchronization enables you to easily propagate, synchronize, and commit

configurations from one MC-LAG peer to another. You can log into any one of the MC-LAG peers to manage both MC-LAG peers, thus having a single point of management. You can also use configuration groups to simplify the configuration process.

In addition, you can create conditional groups to specify when a configuration is synchronized with another MC-LAG peer. You can enable the **peers-synchronize** statement at the **[edit system commit]** hierarchy to synchronize the configurations and commits across the MC-LAG peers by default. NETCONF over SSH provides a secure connection between the MC-LAG peers, and Secure Copy Protocol (SCP) copies the configurations securely between them.

[See [Understanding MC-LAG Configuration Synchronization.](#)]

- **Configuration consistency check for MC-LAG (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D60, configuration consistency check uses the Inter-Chassis Control Protocol (ICCP) to exchange MC-LAG configuration parameters (chassis ID, service ID, and so on) and checks for any configuration inconsistencies across MC-LAG peers. An example of an inconsistency is configuring identical chassis IDs on both peers instead of configuring unique chassis IDs on both peers. When there is an inconsistency, you are notified and can take action to resolve it. Only committed MC-LAG parameters are checked for consistency.

[See [Understanding Multichassis Link Aggregation Group Configuration Consistency Check.](#)]

Layer 2 VPNs

- **Ethernet-over-MPLS (L2 circuit) (QFX10000 switches)**—Starting with Junos OS Release 15.X53-D60, you can configure a Layer 2 circuit to create a point-to-point Layer 2 connection using MPLS on the service provider's network. Ethernet-over-MPLS allows sending Layer 2 (L2) Ethernet frames transparently over MPLS. Ethernet-over-MPLS uses a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. It encapsulates Ethernet protocol data units (PDUs) inside MPLS packets and forwards the packets, using label stacking, across the MPLS network. This technology has applications in service provider, enterprise, and data center environments. To enable a Layer 2 circuit, include the `l2circuit` statement at the `[edit protocols mpls labeled-switched-path lsp-name]` hierarchy level.

[See [Understanding Ethernet-over-MPLS \(L2 Circuit\)](#).]

Routing Protocols

- **BGP Monitoring Protocol (BMP) version 3 support (QFX10000 switches)**—BMP enables the Junos OS to send BGP route information from the switch to a monitoring application, or station, on a separate device. To deploy BMP in your network, you need to configure BMP on each switch and at least one BMP monitoring station. Only version 3 is supported on QFX10008 and QFX10016 switches starting with Junos OS Release 15.1X53-D60. To configure BMP, configure the `bmp` set of statements at the `[edit routing-options]` hierarchy level. To configure a BMP monitoring station, include the `station-address ip-address` and `station-port number` statements at the `[edit routing-options bmp]` hierarchy level.

[See [Configuring BGP Monitoring Protocol Version 3](#).]

Software-Defined Networking (SDN)

- **EVPN pure type-5 route support (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 15.1X53-D60, you can configure pure type-5 routing in an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which is used to carry the MAC address of the sending switch and to provide next-hop reachability for the prefix. This router MAC extended community provides next-hop reachability without requiring an overlay next-hop or supporting type-2 route. To configure pure type-5 routing, include the `ip-prefix-support advertise direct-nexthop` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level. Pure type-5 routing was previously supported only on QFX10002 switches.

[See [ip-prefix-routes statement](#).]

- **Proxy advertisement of host MAC+IP type 2 routes in EVPN-VXLAN topology with IRB interfaces (QFX10000 switches)**—In an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) topology with integrated routing and bridging (IRB) interfaces, leaf devices typically function as Layer 2 gateways. As such, these devices can advertise only the MAC routes (EVPN type 2 routes) for the attached hosts. Since the Layer 2

gateways are unable to resolve the MAC-to-IP bindings for the hosts, each of the spine devices, which typically function as Layer 3 gateways, must rely on the Address Resolution Protocol (ARP) and the Neighbor Discovery Protocol (NDP) to discover and install the bindings.

Starting with Junos OS Release 15.1X53-D60, QFX10000 switches that function as Layer 3 gateways in this type of topology can advertise the MAC and IP routes (MAC+IP type 2 routes) of hosts. With this feature enabled, after receiving a host MAC route advertisement from a Layer 2 gateway, and ARP and NDP resolve the MAC-to-IP bindings, the QFX10000 switch in turn advertises the host MAC and IP routes along with the next hop, which is set to the Layer 2 gateway to which the host is attached. Upon receipt of this advertisement, Layer 2 and 3 gateways in the topology install the MAC-to-IP bindings along with the associated next hops. When any of these gateways receives a packet with a destination MAC that matches an address in its MAC table, the gateway can check the next hop associated with the MAC address and forward the packet directly to the Layer 2 gateway to which the host is attached. This resulting packet flow eliminates the need for the packet to be forwarded first to a Layer 3 gateway, which then forwards the packet to the Layer 2 gateway.

To enable this feature, specify the **proxy-macip-advertisement** configuration statement at the `[edit interfaces irb unit logical-unit-number]` hierarchy level. The following is a sample command that configures an IRB interface on a QFX10000 switch that functions as a Layer 3 gateway in an EVPN-VXLAN topology that includes both Layer 2 and Layer 3 gateways:

```
user@switch# set interfaces irb unit 0 proxy-macip-advertisement family inet address 192.0.2.100 virtual-gateway-address 192.0.2.125
```

Enabling this feature in an EVPN-VXLAN topology that includes both Layer 2 and Layer 3 gateways is mandatory, while enabling the feature in a topology that includes only Layer 3 gateways is optional.

[See [proxy-macip-advertisement](#).]

Software Installation and Upgrade

- **Support for FreeBSD 10 kernel for Junos OS (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D60, on QFX10000 switches, the base operating system has been upgraded from FreeBSD 6.1 to FreeBSD 10. FreeBSD 10 supports SMP for Junos OS.

Support includes:

- Junos addressable DRAM memory increase from 4G to 12G
- Junos addressable DRAM memory increase from 3.2G to 4G
- Memory increase for rundb from 512MB to 1GB
- SMP support with Junos running on two cores
- 64-bit kernel support

[See [Understanding Junos OS with Upgraded FreeBSD](#).]

New Features in Release 15.1X53-D30

Hardware

- **QFX10008 switch**—The Juniper Networks QFX10000 line of Ethernet switches provides cloud builders and data center operators scalable solutions for both core and spine data center deployments. The QFX10008 switch is an 8-slot, 13 U chassis that supports up to 8 line cards.
- **Support for 100-Gigabit optical transceivers (QFX10008 switch)**—Provides support for:
 - JNP-QSFP 100G-SR4—QSFP28 module 100GBASE-SR4, 100-Gigabit Ethernet pluggable; 850 nm for up to 150 m transmission on multi-mode fiber (MMF) cable.
 - JNP-QSFP-100G-LR4—QSFP28 module 100GBASE-LR4, 100-Gigabit Ethernet pluggable; 1310 nm for up to 10 km single-mode fiber-optic (SMF) cable.
- **Support for 40-Gigabit optical transceivers (QFX10008 switch)**—Provides support for:
 - QFX-QSFP-40G-SR4—QSFP+ module 40GBASE-SR4, 40-Gigabit Ethernet optics; 100 m transmission on OM3, MMF cable and 150 m transmission on OM4, MMF cable.
 - QFX-QSFP-40G-ESR4—Juniper Networks proprietary 4X10G-IR parallel single mode QSFP+ module, 40-Gigabit Ethernet- optics; 300 m transmission on OM3, MMF cable or 400 m transmission on OM4 cable.
 - JNP-QSFP-4X10GE-IR—QSFP+ parallel single mode module 40-Gigabit Ethernet pluggable; 1.4 km transmission on SMF cable.
 - JNP-QSFP-40GE-IR4—Juniper Networks proprietary 40GBASE-IR4, 40-Gigabit Ethernet pluggable; 2 km transmission on SMF cable.
 - JNP-QSFP-40G-LR4—QSFP+ module 40GBASE-LR4, 40-Gigabit Ethernet pluggable; 10 km transmission on SMF cable.
 - JNP-QSFP-4X10GE-LR—Juniper Networks proprietary 4X10G-LR, 40-Gigabit Ethernet; 10 km transmission on SMF cable.
 - JNP-QSFP-40G-LX4—QSFP+ module 40GBASE-LX4, 40-Gigabit Ethernet pluggable; 2 km transmission on SMF cable; 100 m transmission on OM3, MMF cable; or 150 m transmission on OM4, MMF cable
- **Support for 1-Gigabit optical transceivers on the SFP management port (QFX10008 switch)**—Provides support for:
 - QFX-SFP-1GE-SX—SFP module 1000BASE-SX Gigabit Ethernet; 220 m transmission on FDDI, MMF cable; 275 m transmission on OM1, MMF cable; or 550 m transmission on OM2 cable.
 - QFX-SFP-1GE-T—SFP module 1000BASE-T Gigabit Ethernet; 100m transmission on Category 5 cable.

- **QFX-SFP-1GE-LX**—SFP module 1000BASE-LX Gigabit Ethernet; 10 km transmission on SSF cable; 550 m transmission on OM1, MMF cable; or 550 m transmission on OM2, MMF cable.
- **QFX10000-36Q line card (QFX10008 switches)**—Provides 36 ports of 40-gigabit QSFP+. Twelve ports are designed to be 100-gigabit capable using QSFP28. Each 40-gigabit QSFP+ can be configured as either a native 40-gigabit port or four 10-gigabit ports using a breakout cable. With breakout cables, the line card supports a maximum of 144 logical 10-Gigabit Ethernet ports.
- **QFX10000-30C line card (QFX10008 switches)**—Provides 30 ports of either 100-gigabit or 40-gigabit QSFP28. The ports autodetect the type of transceiver installed and set the configuration to the appropriate speed.

High Availability and Resiliency

- **High availability feature support (QFX10008 switch)**—The QFX10008 switch supports the following high availability features:
 - **Graceful Routing Engine switchover (GRES)**—Enables a switch with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. To configure GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level and the **synchronize** statement at the **[edit system commit]** hierarchy level.
 - **Nonstop active routing (NSR)**—Uses the same infrastructure as GRES to preserve interface and kernel information. NSR also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. To configure NSR, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level.
 - **Nonstop bridging (NSB)**—Uses the same infrastructure as GRES to preserve interface and kernel information. NSB also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine. To configure NSB, include the **nonstop-bridging** statement at the **[edit protocols layer2-control]** hierarchy level.

Infrastructure

- **Secure Boot (QFX10008 switch)**—Junos OS Release 15.1X53-D30 introduces a significant system security enhancement: Secure Boot. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement Secure Boot.

Interfaces and Chassis

- **Adaptive load balancing (ALB) for aggregated Ethernet bundles (QFX10008 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 taking into account 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 (pps). The default is bits per second.
- **pps**—Scan traffic in packets per second (pps).
- **Channelizing 40-Gigabit Ethernet QSFP+ ports (QFX10008 switch)**—This feature enables you to channelize four 10-Gigabit Ethernet interfaces from the 40-Gigabit Ethernet QSFP+ interfaces. Channelization is supported on fiber break-out cable using standard structured cabling techniques.



NOTE: This feature is not supported on the QFX10000-30C line card.

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.

There are 100-Gigabit Ethernet ports that work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet but are recognized as 40-Gigabit Ethernet by default. You cannot channelize the 100-Gigabit Ethernet ports when they are operating as 100-Gigabit Ethernet interfaces. The 40-Gigabit Ethernet ports can operate independently or be channelized into four 10-Gigabit Ethernet ports as part of a port range. Ports cannot be channelized individually. Only the first and fourth port in each 6XQSFP cage is available to channelize as part of a port range. In a port range, the ports are bundled with the next two consecutive ports. For example, if you want to channelize ports 0 through 2, you channelize port 0 only. If you try to channelize a port that is not supported, you receive an error message when you commit the configuration. Auto-channelization is not supported on any ports.

When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports.

- **Link aggregation (QFX10008 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) (QFX10008 switch)**—MC-LAG enables a client device to form a logical LAG interface using two QFX10008 switches. MC-LAG provides redundancy and load balancing between the two QFX10008 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 detect the MC-LAG. On the other side of the MC-LAG are two MC-LAG QFX10008 switches. Each of these QFX10008 switches has one or more physical links connected to a single client. The QFX10008 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
- **Ability to create link aggregation groups with interfaces operating at different speeds (QFX10008 switch)**—You can add 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet interfaces into the same link aggregation group (LAG). Configuring LAGs with interfaces configured at speeds other than 10g, 40g, and 100g is not supported.
- **Support for Layer 3 logical interfaces (QFX10008 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 QFX10008 switch to a Layer 2 switch. Only one physical connection is required between the switches.
- **Generic routing encapsulation (GRE) support (QFX10008 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.
- **Enhanced hash key (QFX10002 switches)**—Starting with Junos OS Release 15.1X53-D30, you can configure the `inet`, `inet6`, `GRE`, `no-mpls`, `vxlan-vnid`, and `hash-seed` values for load-balancing functions. By default, the QFX10002 switches use the system MAC address to generate a `hash-seed` value. You can configure the value for the `hash-seed` statement at the `[edit forwarding-options enhanced-hash-key]` hierarchy level. The `fabric-load-balance` and `user-defined-fields` statements are not supported at the `[edit forwarding-options enhanced-hash-key]` hierarchy level.
- **Support for Micro BFD over child links of AE or LAG bundle (cross-functional Packet Forwarding Engine/kernel/rpd) (QFX10002 switches)**—Provides a Layer 3 BFD liveness detection mechanism for child links of the Ethernet LAG interface. In scenarios

in which you do not have a point-to-point link, and a Layer 1 device fails at one end of the link, Micro BFD detects failures faster than traditional LACP. Micro BFD sessions are independent of each other despite having a single client that manages the LAG interface. Micro BFD is not supported on pure Layer 2 interfaces. To enable failure detection for aggregated Ethernet interfaces, include the `bfd-liveness-detection` statement at the `[edit interfaces aex aggregated-ether-options bfd-liveness-detection]` hierarchy level.

Layer 2 Features

- **VLAN support (QFX10008 switch)**—VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
- **Link Layer Discovery Protocol (LLDP) support (QFX10008 switch)**—LLDP enables a switch to advertise its identity and capabilities on a LAN, as well as receive information about other network devices.
- **Q-in-Q tunneling support (QFX10008 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 (QFX10008 switch)**—These protocols enable a switch to advertise its identity and capabilities on a LAN and receive information about other network devices.

Layer 3 Features

- **BGP support (QFX10008 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 (QFX10008 switch)**—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). QFX10008 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 (QFX10008 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 QFX10008 switch, you can configure BFD for static routes and for the BGP, IS-IS, OSPF, PIM, and RIP protocols.

- **IS-IS support (QFX10008 switch)**—The IS-IS protocol is an IGP for routing traffic within an AS.

- **Virtual Router Redundancy Protocol (VRRP) support (QFX10008 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.
- **IPv4 address conservation method for hosting providers (QFX10008 switch)**—If your company hosts servers for customers, you might be using many routable IP addresses when you assign addresses for servers. For example, you need to assign network and broadcast IP addresses, the address for the gateway that the server is connected to, and the address of the individual server, all of which are publicly routable addresses. When this approach is multiplied across thousands of customers, you end up using a large number of publicly routable addresses.

Starting with Junos OS Release 15.1X53-D30, this issue can be resolved by configuring an interface on the gateway switch with an address from the reserved IPv4 prefix for shared address space (RFC 6598) and by creating static routes that use that interface as the next hop. (The shared address space address range is 100.64.0.0/10.) You also configure the network and broadcast addresses from this range. You then configure the server with a static route that points to the RFC 6598 address used on the switch interface. With this approach, you can significantly reduce the number of routable IPv4 addresses that you use for your hosting customers.

- **IPv6 VPN Provider Edge (6VPE) routing (QFX10000 switches)**—IPv6 VPN Provider Edge (6VPE) routing functionality provides IPv6 forwarding over IPv4-based MPLS networks. Starting with Junos OS Release 15.1X53-D30, QFX10000 switches support 6VPE.

Multicast Protocols

- **Internet Group Management Protocol (IGMP) support (QFX10008 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 (QFX10008 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 (PIM) sparse mode support (QFX10008 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 (QFX10008 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 (QFX10008 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 (QFX10008 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 (QFX10008 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 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 (QFX10008 switch)**—MPLS provides both label edge router (LER) and label switch router (LSR) 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

- **Auto-bandwidth and dynamic LSP count sizing (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D30, auto-bandwidth and dynamic label-switched path (LSP) count sizing are supported on QFX10000 switches. Auto-bandwidth allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. Dynamic LSP count sizing provides an ingress router with the capability of acquiring as much network bandwidth as possible by creating parallel LSPs dynamically.

Network Management and Monitoring

- **SNMP support (QFX10008 switch)**—SNMP includes versions 1, 2, and 3 for monitoring system activity.
- **System logging (syslog) support (QFX10008 switch)**—Syslog enables you to log system messages into a local directory on the switch or to a syslog server.
- **sFlow technology support (QFX10008 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 support (QFX10008 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.
- **Virtual-router aware DHCP server/DHCP relay agent (QFX10008 switch)**—The QFX10008 switch can be configured to act as a DHCP server or DHCP relay agent for IPv4 and IPv6. If you have virtual router instances on the switch, the DHCP implementation can work with them.

Security

- **Firewall filter support (QFX10008 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 (QFX10008 switch)**—You can use policing to apply limits to traffic flow and to set consequences for packets that exceed those limits.
- **MAC limiting support (QFX10008 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 (QFX10008 switch)**—You can detect MAC movement and MAC spoofing on access ports.
- **Storm control support (QFX10008 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.

Software-Defined Networking (SDN)

- **Layer 2 VXLAN gateway and OVSDB support (QFX10008 switch)**—In a physical network, a Juniper Networks device that supports a Virtual Extensible LAN (VXLAN) can function as a hardware virtual tunnel endpoint (VTEP). In this role, the Juniper Networks device encapsulates in VXLAN packets Layer 2 Ethernet frames received from software applications that run directly on a physical server. The VXLAN packets are tunneled over a Layer 3 fabric. Upon receipt of the VXLAN packets, software VTEPs in the virtual network de-encapsulate the packets and forward the packets to virtual machines (VMs).

In this VXLAN environment, you can also include SDN (VMware NSX or Contrail) controllers and implement the Open vSwitch Database (OVSDB) management protocol on the Juniper Networks device that functions as a hardware VTEP. The Junos OS implementation of OVSDB provides a means through which SDN controllers and Juniper Networks devices can exchange MAC addresses of entities in both physical and virtual networks. This exchange of MAC addresses enables the Juniper Networks device that functions as a hardware VTEP to forward traffic to software VTEPs in the virtual network and software VTEPs in the virtual network to forward traffic to the Juniper Networks device in the physical network.

- **Integrated routing and bridging support for EVPN-VXLAN (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D30, QFX10000 switches support integrated routing and bridging (IRB) interfaces that route packets between Virtual Extensible LANs (VXLAN)s in an Ethernet VPN (EVPN)-VXLAN topology. This functionality is typically needed to provide Layer 3 connectivity between physical servers and virtual machines (VMs) on servers in the virtual network. Use the **set interfaces irb** command to configure an IRB interface for each VXLAN that needs to exchange packets with a host in another VXLAN, and specify a default gateway address for the hosts in the VXLAN to use by including the **virtual-gateway-address** configuration statement. Configuring this default gateway sets up a redundant default gateway for the hosts in the VXLAN.
- **EVPN control plane for VXLAN supported interfaces (QFX10000 switches)**—Traditionally, data centers have used Layer 2 technologies such as Spanning Tree Protocol (STP), multichassis link aggregation groups (MC-LAGs), or TRILL for compute and storage connectivity. As the design of data centers shifts from more traditional to scale-out, service-oriented multitenant networks, a new data center architecture allows decoupling of an underlay network from the tenant overlay network with VXLAN. By using a Layer 3 IP-based underlay coupled with a VXLAN-EVPN overlay, you can deploy larger networks than those possible with traditional Layer 2 Ethernet-based architectures. With overlays, end points (servers or virtual machines) can be placed anywhere in the network and remain connected to the same logical

Layer 2 network. The benefit is that virtual topology, using both MX Series routers and QFX10000 switches, can be decoupled from the physical topology.

- **Layer 3 connectivity between data centers (QFX10002 switch)**—Starting with Junos OS Release 15.1X53-D30, you can create pure Layer 3 connections between data centers with VXLAN encapsulation by using the EVPN type-5 IP prefix routes. If you do not have VLANs that stretch between data centers, you do not need to advertise MAC and IP routes between your data centers, so a pure Layer 3 approach is feasible. EVPN pure type-5 routes decouple MAC addresses from IP addresses and advertise only IP prefixes. Include the `ip-prefix-support forwarding-mode symmetric` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level to configure EVPN pure type-5 routes between QFX10002 switches.

Software Installation and Upgrade

- **Firmware upgrade (QFX10008 switch)**—Starting with Junos OS Release 15.1X53-D30, you can upgrade the system firmware. There are several firmware components that you can upgrade.

On a line card, you upgrade the following firmware components:

- Uboot—Responsible for loading the operating system on the line card
- FPGA—Controls all functions of the line card

You can also upgrade the following firmware components:

- RE- FPGA—The RE-FPGA is located on the control board and manages board initialization, reboot, and other functions.
- FTC FPGA—The FTC FPGA is located on the fan controllers and controls the fan controllers.
- FPD FPGA—The FPD FPGA is located on the LED board and is responsible for the LED board.
- SIB FPGA—The SIB FPGA is located on the SIB and handles the SIBs.

Before you can upgrade the firmware components, you need to install a software package that contains the firmware images that you want to upgrade. The `jloader-qfx-10` package contains the `uboot` binary (bootloader), and the `qfx-10-m-firmware` package contains the FPGA images.

To install these packages, issue the `request system software add` command. The package that contains the `uboot` binary is a `jloader-qfx-10` package. To upgrade the `uboot` binary (bootloader), issue the `request system firmware upgrade fpc slot slot-number` command. To upgrade the FPGA components, issue the `request system firmware upgrade fpga (cb | ftc | fpd)` command or the `request system firmware upgrade fpga (fpc | sib) slot slot-number` command. Upgrading the firmware takes between 2 and 3 minutes, depending on which firmware components you are upgrading.



NOTE: The `request system firmware upgrade` command is not visible in the CLI. To use the command to upgrade the bootloader or the FPGA components, type the command after the operational-mode prompt (`>`)—for example:

```
user@switch> request system firmware upgrade fpga sib
```



CAUTION: Do not reboot the system during a firmware upgrade because the FPGA might get corrupted. You cannot recover the FPGA if it is corrupted.

Storage

- **FCoE transit switch support (QFX10008 switch)**—You can configure a QFX10008 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 (QFX10008 switch)**—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 fpcs fpc fpc-slot**—Display information for Flexible PIC Concentrators (FPCs) in the fabric.
 - **show chassis fabric plane-location**—Display the fabric plane location of each Switch Interface Board (SIB).
 - **show chassis fabric sibs**—Display the state of the switch fabric link between the SIBs and the FPCs.
 - **show chassis fabric topology**—Display the input-output link topology.
- **Login authentication using RADIUS and TACACS+ (QFX10008 switch)**—You can use RADIUS and TACACS+ authentication to validate users who attempt to access the switch.
- **System utilization alarms support (QFX10008 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.

- **FATAL and MAJOR FAULT information support (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D30, QFX10000 switches support the ability to report FATAL and MAJOR errors in the output of the **show chassis fpc errors** command.

Traffic Management

- **Class-of-service (CoS) rewrite rules support (QFX10008 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 (QFX10008 switch)**—You can manage excess traffic and avoid congestion on a network interface where traffic might exceed the maximum port bandwidth.
- **Ethernet PAUSE autonegotiation support (QFX10008 switch)**—You can configure symmetric flow control. To configure PAUSE, include the **flow-control** statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level
- **CoS command to detect the source of RED-dropped packets (QFX10008 switch)**—If traffic on the switch is congested, you can use the **show interfaces voq *interface-name*** CLI command to identify which ingress Packet Forwarding Engine is the source of random early detection (RED)-dropped packets that are contributing to congestion. The command output displays RED drop statistics from all ingress Packet Forwarding Engines associated with the specified physical egress interface. In the VOQ architecture on the switch, egress output queues (shallow buffers) buffer data in virtual queues on ingress Packet Forwarding Engines.
- **DCB standards support (QFX10008 switch)**—The switch supports these data center bridging standards:
 - Priority-based flow control (PFC) allows you to select traffic flows within a link and pause them, so that the output queues associated with the flows do not overflow and drop packets.
 - Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP-based networks.

Virtual Private Networks (VPNs)

- **Layer 2 Ethernet virtual private network control plane support (QFX10000 switches)**—Ethernet VPNs (EVPNs) enable you to connect groups of dispersed customer sites to one another using Layer 2 virtual bridges. Layer 2 EVPN control planes support is supported on QFX10000 switches starting in Junos OS Release 15.1X53-D30. You configure the feature on QFX10000 switches under the global **[edit switching-options]** and **[edit protocols evpn]** hierarchy levels.

Related Documentation

- [Changes in Behavior and Syntax on page 19](#)
- [Known Behavior on page 20](#)

- [Known Issues on page 21](#)
- [Resolved Issues on page 24](#)
- [Documentation Updates on page 26](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)
- [Product Compatibility on page 31](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands for Junos OS Release 15.1X53 for QFX10000 switches.

- [Software-Defined Networking \(SDN\) on page 19](#)
- [VXLAN on page 19](#)

Software-Defined Networking (SDN)

- **Change in configuration statement for pure type-5 routing for EVPN (QFX10000 switches)**—Starting in Junos OS Release 15.1X53-D60, to configure pure type-5 routes in an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment, include the **ip-prefix-routes advertise direct-nexthop** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. When this feature was introduced in Junos OS Release 15.1X53-D30 on QFX10002 switches only, you included the **ip-prefix-support forwarding-mode symmetric** statement. This statement has been deprecated and is no longer supported. Any configuration with the original statement is automatically upgraded to the new **ip-prefix-routes** statement when you upgrade to Junos OS Release 15.1X53-D60. Pure type-5 routing is now also supported on QFX10008 and QFX10016 switches.

[See [ip-prefix-routes statement](#).]

VXLAN

- **Best practice for EVPN-VXLAN configuration (QFX10000 switches)**—Starting with Junos OS Release 15.1X53-D60, in an EVPN-VXLAN configuration on QFX10000 switches, you no longer need to configure **vlan ingress-node-replication**.

Related Documentation

- [New and Changed Features on page 3](#)
- [Known Behavior on page 20](#)
- [Known Issues on page 21](#)
- [Resolved Issues on page 24](#)
- [Documentation Updates on page 26](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)
- [Product Compatibility on page 31](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1X53-D62 for QFX10000 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 20](#)
- [Layer 3 Features on page 20](#)
- [Routing Protocols on page 20](#)
- [Routing Proto on page 20](#)

Interfaces and Chassis

- On a QFX10000 switch configured as an MC-LAG peer, approximately 7 through 9 seconds of traffic might be impacted when the ICCP/ICL interface goes down. [PR1191337](#)
- On QFX10008 and QFX10016 switches, when ICCP cannot reach the mclag-cfgchkd daemon on a remote MC-LAG peer, configuration consistency check cannot compare the configurations, and all of the MC-LAG parameters are marked as PASS. When the remote MC-LAG peer and the mclag-cfgchk daemon become reachable, configuration consistency check exchanges the configurations and compares them. As a workaround, make sure that the mclag-cfgchkd daemon is running on both MC-LAG peers and that ICCP can reach the MC-LAG peers. [PR1220630](#)
- When you use a QFX10000 switch as a PE device in a 6VPE topology, you must enable family inet6 on the core-facing interfaces that will be used for the IPv6 traffic. If you do not do so, these interfaces drop the IPv6 traffic. [PR1226014](#)

Layer 3 Features

- When you upgrade the software on a QFX10000 switch to Junos OS Release 15.1X53-D60 or 15.1X53-D61, BFD sessions might flap after the upgrade. As a workaround, restart the periodic packet management process (ppmd) after the upgrade. [PR1210316](#)

Routing Protocols

- When EVPN-VXLAN is configured on QFX10002 switches, the number of ARP entries supported depends on the number of VTEPs, number of VNIs and number of IFBDs (port * vlan). The above mentioned objects along with ARP entries share a global resource pool. [PR1206633](#)

Routing Proto

- Related Documentation**
- [New and Changed Features on page 3](#)
 - [Changes in Behavior and Syntax on page 19](#)

- [Known Issues on page 21](#)
- [Resolved Issues on page 24](#)
- [Documentation Updates on page 26](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)
- [Product Compatibility on page 31](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1X53-D62 for QFX10000 switches.

- [High Availability \(HA\) and Resiliency on page 21](#)
- [Interfaces and Chassis on page 21](#)
- [Layer 2 Features on page 22](#)
- [MPLS on page 22](#)
- [Multicast Protocols on page 22](#)
- [Network Management and Monitoring on page 22](#)
- [Platform and Infrastructure on page 23](#)
- [Routing Protocols on page 23](#)

High Availability (HA) and Resiliency

- On QFX10008 and QFX10016 switches, IPv4 and IPv6 distributed BFD sessions might flap in high availability due to a master reboot or an unplanned Routing Engine switchover. [PR1176745](#)

Interfaces and Chassis

- On a QFX10000 switch, when a new interface is added to an existing link aggregation group (LAG) interface which acts as an input analyzer interface, traffic sent to the added interface might not be mirrored. As a workaround, delete the configuration and then reconfigure it. [PR1057527](#)
- On QFX10008 and QFX10016 switches, during a Routing Engine switchover, BGP on the IRB interface might flap when the IRB interface and the underlying Layer 2 logical interface are configured with different MTU values. As a workaround, configure the same MTU value on both the Layer 2 logical interface and the IRB interface. [PR1187169](#)
- On QFX10000 switches with enhanced MC-LAG IRB next hops, member links of the aggregate underlying Layer 2 interfaces might not be present on all Packet Forwarding Engine instances in a given FPC. Under this condition, during IRB next-hop installation for the Packet Forwarding Engine instance where the underlying Layer 2 interface link is not present, failure logs are generated for the PFE uKernel. Those failure logs do not impact traffic or performance on the switch, and they are harmless. [PR1221831](#)

- On QFX10008 and QFX10016 switches, an error message such as **expr_cos_rw_nh_qix_get @ 150: Unable to get chip num for ill:994 on mc-ae status-control active node** might be displayed after an ARP request is sent. These messages are only for information and have no functional impact on the operation of the switches. [PR1228080](#)
- On QFX10000 switches, removal or insertion of a transceiver for a LAG member when the LAG bundle is configured as a member of thousands of VLANs (for example, 4093 VLANs for the ICL in this PR) might cause high CPU utilization in the Packet Forwarding Engine for a few seconds, preventing critical protocols from running in a timely manner and causing timeouts for BFD sessions, LACP, and so on. Such timeouts might lead to ICL or ICCP flaps and ARP flushes in the MC-LAG topology. As a workaround, avoid unplanned removals or insertions of transceivers for LAG members. If the transceiver removal or insertion is necessary, remove the corresponding interface from the LAG bundle by using CLI configuration commands before you remove or insert transceivers. [PR1229547](#)

Layer 2 Features

- On QFX10000 switches, MAC limiting does not take effect on the Packet Forwarding Engine. [PR1131502](#)

MPLS

- On QFX10008 and QFX10016 switches, after a high availability (HA) graceful Routing Engine switchover (GRES), the **show mpls lsp autobandwidth** command might display an incorrect bandwidth on the new Routing Engine. As a workaround, use the **show rsvp session extensive** command. If the LSP is resigaled, the problem is corrected. [PR1149863](#)

Multicast Protocols

- During a graceful Routing Engine switchover (GRES) on QFX10000 switches, some IPv6 groups might experience momentary traffic loss. This issue occurs when IPv6 traffic is running with multiple paths to the source, and the **join-load-balance** statement for PIM is also configured. [PR1208583](#)

Network Management and Monitoring

- If a QFX10008 or a QFX10016 switch is fully populated with line cards and all the interfaces are configured, SNMP timeouts can occur if the switch is simultaneously polled by multiple network management stations. [PR1147934](#)
- On QFX10008 and QFX10016 switches, sFlow external router data is not updated in sFlow datagrams if sampling is in the egress direction for flows going across FPCs. The updating issue does not occur if samples are taken at the ingress FPC for such flows or if flows are within the same FPC. [PR1241362](#)
- On QFX10000 switches, sFlow monitoring technology output might display a negative number of samples after a long run. As a workaround, issue the **clear sflow collector** command to show or reset the count. [PR1244080](#)

Platform and Infrastructure

- On QFX10008 and QFX10016 switches, micro Bidirectional Forwarding Detection (BFD) sessions might flap during a graceful Routing Engine switchover (GRES). [PR1165509](#)
- On QFX10008 switches, if you reboot a QFX10000-36Q line card or a QFX10000-30C line card with traffic running, sometimes framing errors are displayed in the CLI output. This is only a display issue. No actual framing errors have occurred, and traffic is unaffected. [PR1223330](#)
- On QFX10002 switches, when you plug in a USB, FRU insertion messages such as **REO & ?CAMGETPASSTHRU ioctl failed cam_lookup_pass: Inappropriate ioctl for device?** might be displayed. These are harmless messages and will not be displayed after you have removed the USB. [PR1233037](#)

Routing Protocols

- When VRF route leak is configured along with a static route next hop over an IRB interface on QFX10000 switches, ARP requests for these routes might not be resolved. [PR1147176](#)
- Modifying the route-distinguisher identifier flushes all the routes, including interface routes, from the virtual routing and forwarding (VRF) routing information base (RIB), also known as the routing table. As a result, the VRF RIB remains empty. As a workaround, deactivate the VRF before you modify the route-distinguisher identifier and then reactivate the VRF. [PR1155647](#)
- The **no-indirect-next-hop** configuration statement that was mandatory in Junos OS Release 15.1X53-D30 and earlier releases is no longer required on QFX10000 switches in Release 15.1X53-D60. Delete it from the configuration before you upgrade to Release 15.1X53-D60. [PR1187482](#)
- On QFX10000 switches configured as route type-5 peers, when only peer 1 advertises routes, that peer might not install the decap next-hop (NH) route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any route type-5. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. [PR1191092](#)
- On QFX10000 switches, when a type-5 route is also learned through another protocol with a lower preference or is statically configured, the switch installs the type-5 route even though it is not the preferred route. [PR1200519](#)
- On QFX10000 switches, the Packet Forwarding Engine sorts next-hop constituent element next hops by their weights in ascending order before installing them in the forwarding plane. The inefficiency of the sorting algorithm means that the unicast next-hop programming acknowledgments are delayed. As a result, BGP route updates can fail because no buffer is available. [PR1225309](#)
- On QFX10002 switches, if the MAC age timer is set to a value greater than that of the ARP age timer, after the ARP ages out, MAC and MAC+IP is advertised by all ESI peers regardless of which device learns ARP locally. As a workaround, set the MAC age timer to a value less than that of the ARP age timer. [PR1238718](#)

**Related
Documentation**

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 19](#)
- [Known Behavior on page 20](#)
- [Resolved Issues on page 24](#)
- [Documentation Updates on page 26](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)
- [Product Compatibility on page 31](#)

Resolved Issues

This section lists the issues fixed in Junos OS Release 15.1X53 for QFX10000 switches.

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-D62 on page 24](#)
- [Resolved Issues: Release 15.1X53-D61 on page 25](#)
- [Resolved Issues: Release 15.1X53-D32 on page 26](#)

Resolved Issues: Release 15.1X53-D62

- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)

Interfaces and Chassis

- On QFX10000 switches, the kernel might fail to allocate IFBD tokens, with the error message **IFBD hw token couldn't be allocated for <interface>**, even though there are enough IFBD tokens, and thus you might be unable to assign some VLANs to the related interfaces. [PR1216464](#)
- On a QFX10002 switch, 40 Gigabit Ethernet ports can take up to 4 seconds to link when using JNP-QSFP-40G-LR4 optical transceivers. [PR1219336](#)
- On QFX10000 switches, on aggregated Ethernet interfaces with adaptive load balancing enabled, frequent link flaps might result in zero active members in the LAG bundle, causing memory leaks and eventually causing an FPC crash. The FPC restarts automatically after the crash. [PR1236046](#)

Network Management and Monitoring

- On QFX10000 switches, IPv6 MIB statistics for `jnxIpv6IfInOctets` and `jnxIpv6IfOutOctets` for an aggregated Ethernet (AE) bundle show double the count that is shown in CLI output. [PR1230923](#)

Resolved Issues: Release 15.1X53-D61

- [Interfaces and Chassis](#)
- [MPLS](#)
- [Routing Protocols](#)
- [Software-Defined Networking \(SDN\)](#)

Interfaces and Chassis

- On QFX10000 switches, in a multichassis link-aggregation group (MC-LAG) configuration, the `all` option at the `[edit protocols igmp-snooping vlan]` hierarchy level does not work. As a workaround, enable IGMP snooping on a per-VLAN basis on each of the MC-LAG peers. [PR1180494](#)
- On QFX10000 switches, traffic might drop on an aggregated Ethernet interface in the following scenario:

Topology:

- The AE has two child members connected to the same PFE.
- The child port numbers should be < 32.

Trigger and symptoms: When an AE member is removed by a physical OIR of a transceiver or by deactivating the member port configuration, traffic is lost in the AE interface. The problem does not happen if the AE members are spread across multiple PFEs or across FPCs. As a workaround, disable the AE interface and then reenabling it.

[PR1210220](#)

MPLS

- On QFX10000 switches, when MPLS automatic bandwidth allocation is configured for an LSP, disabling the configuration might generate an RPD core file. [PR1152449](#)

Routing Protocols

- On QFX10000 switches, VRRPv2 for IPv4 is not working correctly. A router with a physical interface with the highest IPv4 address preempts mastership even in case of a priority tie. The feature works correctly for IPv6 address families. [PR1204969](#)
- On QFX10000 switches, whenever a host moves from one leaf switch to another leaf switch, the ARP entry for that host is not updated in the remote leaf switch or switches. As a workaround, restart the `l2ald` process. [PR1210195](#)

- With an EVPN-VXLAN configuration, QFX10000 switches do not forward DHCP relay packets from remote VTEPs. [PR1209499](#)

Software-Defined Networking (SDN)

- On QFX10000 switches, during an upgrade to Junos OS Release 15.1X53-D60, OVSDB-based MAC learning might fail and traffic loss might occur. The output from the **show ovsdb commit failures** CLI command might show commit failures. [PR1207165](#)

Resolved Issues: Release 15.1X53-D32

- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)

Interfaces and Chassis

- If you commit a huge configuration on a QFX10000 switch, in rare cases some ports are not activated. [PR1160220](#)
- On a QFX10008 switch, a 100-Gigabit optical interface might not activate if the interface is disabled and enabled several times. [PR1160236](#)
- On a QFX10002 switch, the major alarm LED may light even though there are no alarms. [PR1160248](#)

Network Management and Monitoring

- On QFX10000 switches, when sFlow is configured and traffic is routed out of a link aggregation interface, the SNMP index of the output port might not be displayed, which means that the traffic flows cannot be monitored. [PR1161197](#)

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 19](#)
- [Known Behavior on page 20](#)
- [Known Issues on page 21](#)
- [Documentation Updates on page 26](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)
- [Product Compatibility on page 31](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1X53 for QFX10000 switch documentation.

- [Changes to Junos OS for QFX10000 Switches Documentation on page 27](#)

Changes to Junos OS for QFX10000 Switches Documentation

MPLS

- The new-feature item “Auto-bandwidth (QFX10000 switches)” has been removed from *New and Changed Features in Junos OS Release 15.1X53-D60*. The auto-bandwidth feature is supported starting with Release 15.1X53-D30 on QFX10000 switches.

Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 19](#)
- [Known Behavior on page 20](#)
- [Known Issues on page 21](#)
- [Resolved Issues on page 24](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)
- [Product Compatibility on page 31](#)

Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches

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

- [Caveats for Downgrading from Junos OS Release 15.1X53-D60 to Previous Software Releases on QFX10000 Switches on page 27](#)
- [Upgrading Requires Manual Copy of /var/db/scripts Files on page 29](#)
- [Downloading Software Files with a Browser on page 29](#)
- [Backing Up the Current Configuration Files on page 29](#)
- [Installing the Software on page 30](#)

Caveats for Downgrading from Junos OS Release 15.1X53-D60 to Previous Software Releases on QFX10000 Switches

Table 1: Caveats for Downgrading from Junos OS Release 15.1X53-D60 to Previous Software Releases

Junos OS Releases	Using the CLI	Using a USB Stick
15.1X53-D34	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded.	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded.

Table 1: Caveats for Downgrading from Junos OS Release 15.1X53-D60 to Previous Software Releases (*continued*)

Junos OS Releases	Using the CLI	Using a USB Stick
15.1X53-D33	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded.	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded.
15.1X53-D32	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded. NOTE: You must downgrade to Junos OS Release 15.1X53-D33 before you downgrade to Release 15.1X53-D32.	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded. Yes, but you must downgrade to Junos OS Release 15.1X53-D33 before you downgrade to Release 15.1X53-D32.
15.1X53-D30	No	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded. Yes, but you must downgrade to Junos OS Release 15.1X53-D33 before you downgrade to Release 15.1X53-D30.
Releases prior to 15.1X53-D30	No	Yes, but the configuration is lost, and there is no workaround. We recommend that you save the configuration to an external location, such as a remote server, and then load the configuration after you have successfully downgraded. Yes

Upgrading Requires Manual Copy of /var/db/scripts Files

- When QFX10000 switches are upgraded to Junos OS Release 15.1X53-D60 from previous 15.1X53 releases, the contents of the `/var/db/scripts` directory are not preserved. As a workaround, after you upgrade the software to Release 15.1X53-D60, manually copy files from the `/config/db/scripts` directory to the `/var/db/scripts` directory. [PRI209576](#)

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 QFX10008 switch.

- Using a Web browser, navigate to <http://www.juniper.net/support>.
- Click **Download Software**.
- In the By Technology box, click **Switching | QFX Series | QFX10008**.
- In the QFX Series section, click the name of the platform for which you want to download software.
- Click the **Software** tab and select the install package from the Install Package box. A login screen appears.
- Enter your name and password and press **Enter**.
- Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
- Save the `jinstall-qfx-10-m-flex-<version>-secure-domestic-signed.tgz` file on your computer.
- 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.

The switch contains two routing engines, so you will need to install the software on each routing engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> (re0 | re1)** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-qfx-10-m-flex-15.1X53-D60-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-qfx-10-m-flex-15.1X53-D60-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-qfx-10-m-flex-15.1X53-D60-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-qfx-10-m-flex-15.1X53-D60-secure-domestic-signed.tgz re1
```

Reboot both routing engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

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 19](#)

- [Known Behavior on page 20](#)
- [Known Issues on page 21](#)
- [Resolved Issues on page 24](#)
- [Documentation Updates on page 26](#)
- [Product Compatibility on page 31](#)

Product Compatibility

- [Hardware Compatibility on page 31](#)

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 QFX10000 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 19](#)
- [Known Behavior on page 20](#)
- [Known Issues on page 21](#)
- [Resolved Issues on page 24](#)
- [Documentation Updates on page 26](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 27](#)

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

February 14, 2017—Revision 2, Junos OS for QFX10000 Switches, Release 15.1X53-D62

January 18, 2017—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D62

November 18, 2016—Revision 2, Junos OS for QFX10000 Switches, Release 15.1X53-D61

November 14, 2016—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D61

September 14, 2016—Revision 2, Junos OS for QFX10000 Switches, Release 15.1X53-D60—Updates.

August 30, 2016—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D60

Copyright © 2017, 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.