

Release Notes: Junos OS Release 15.1X53-D69 for QFX10000 Switches

Release 15.1X53-D69
April 26, 2019
Revision 1

Contents

Junos OS Release Notes for QFX10000 Switches	5
New and Changed Features	5
New Features in Release 15.1X53-D69	5
New Features in Release 15.1X53-D61	5
Hardware	5
New Features in Release 15.1X53-D60	5
Interfaces and Chassis	5
Layer 2 VPNs	7
Routing Protocols	7
Software-Defined Networking (SDN)	7
Software Installation and Upgrade	8
New Features in Release 15.1X53-D30	9
Hardware	9
High Availability and Resiliency	10
Infrastructure	10
Interfaces and Chassis	10
Layer 2 Features	13
Layer 3 Features	13
Multicast Protocols	14
Multiprotocol Label Switching (MPLS)	15
Network Management and Monitoring	16
Security	16
Software-Defined Networking (SDN)	17
Software Installation and Upgrade	18
Storage	19
System Management	19
Traffic Management	20

Virtual Private Networks (VPNs)	21
Changes in Behavior and Syntax	21
Infrastructure	21
Interfaces and Chassis	22
Routing Policy and Firewall Filters	22
Software-Defined Networking (SDN)	22
VPNs	23
VXLAN	23
Known Behavior	23
EVPN	24
Interfaces and Chassis	24
Layer 2 Features	24
Layer 3 Features	25
Multicast Protocols	25
Network Management and Monitoring	25
Platform and Infrastructure	25
Routing Protocols	26
Software Installation and Upgrade	26
User Interface and Configuration	27
Known Issues	27
EVPN	27
Interfaces and Chassis	28
Layer 2 Features	28
Multicast	28
Platform and Infrastructure	28
Routing Protocols	29
Resolved Issues	30
Resolved Issues: Release 15.1X53-D69	31
EVPN	31
Platform and Infrastructure	31
Resolved Issues: Release 15.1X53-D68	31
EVPN	31
Layer 2 Features	31
Platform and Infrastructure	32
Software Installation and Upgrade	32
Resolved Issues: Release 15.1X53-D67	32
EVPN	32
Interfaces and Chassis	33
Layer 2 Features	33
Network Management and Monitoring	33
Platform and Infrastructure	33
Routing Policy and Firewall Filters	34
Routing Protocols	34
Resolved Issues: Release 15.1X53-D66	34
EVPN	34
Forwarding and Sampling	35
Interfaces and Chassis	35
Layer 2 Features	35
MPLS	35

Platform and Infrastructure	36
Port Security	36
Routing Policy and Firewall Filters	37
Routing Protocols	37
VPNs	39
Resolved Issues: Release 15.1X53-D65	39
EVPN	39
Infrastructure	39
Interfaces and Chassis	39
Network Management and Monitoring	40
Platforms and Chassis	40
Security	40
Resolved Issues: Release 15.1X53-D64	40
High Availability (HA) and Resiliency	40
Interfaces and Chassis	40
Layer 2 Features	41
Network Management and Monitoring	41
Port Security	41
Routing Protocols	41
Security	41
VXLAN	42
Resolved Issues: Release 15.1X53-D63	42
Authentication and Access Control	42
High Availability (HA) and Resiliency	42
Interfaces and Chassis	43
MPLS	43
Multicast Protocols	44
Platforms and Chassis	44
Routing Protocols	44
Software-Defined Networking (SDN)	45
Software Installation and Upgrade	45
Resolved Issues: Release 15.1X53-D62	45
Interfaces and Chassis	45
Network Management and Monitoring	46
Routing Protocols	46
Resolved Issues: Release 15.1X53-D61	46
Interfaces and Chassis	46
MPLS	47
Routing Protocols	47
Software-Defined Networking (SDN)	47
Resolved Issues: Release 15.1X53-D32	47
Interfaces and Chassis	47
Network Management and Monitoring	47
Documentation Updates	48
Changes to Junos OS for QFX10000 Switches Documentation	48
MPLS	48
Software Defined Networking (SDN)	48

Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches	49
Caveats for Downgrading from Junos OS Release 15.1X53-D60 to Previous	
Software Releases on QFX10000 Switches	49
Upgrading Requires Manual Copy of /var/db/scripts Files	50
Downloading Software Files with a Browser	50
Backing Up the Current Configuration Files	51
Installing the Software on QFX10002 Switches	51
Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS	
Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on	
QFX10008 and QFX10016 Switches	52
Installing the Software on QFX10008 and QFX10016 Switches	53
Product Compatibility	57
Hardware Compatibility	57
Documentation Feedback	58
Requesting Technical Support	58
Self-Help Online Tools and Resources	58
Creating a Service Request with JTAC	59
Revision History	59

Junos OS Release Notes for QFX10000 Switches

These release notes accompany Junos OS Release 15.1X53-D69 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-D69 on page 5](#)
- [New Features in Release 15.1X53-D61 on page 5](#)
- [New Features in Release 15.1X53-D60 on page 5](#)
- [New Features in Release 15.1X53-D30 on page 9](#)

New Features in Release 15.1X53-D69

There are no new features or enhancements to existing features for QFX Series in Junos OS Release 15.1X53-D69 (also true for Releases 15.1X53-D61 through 15.1X53-D68 and 15.1X53-D32 through 15.1X53-D34).

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-level 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 in to 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**, **vlan-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.

In addition, QFX10000 switches that function as Layer 3 VXLAN gateways can route IPv6 data traffic through an EVPN-VXLAN overlay network. With this feature enabled, Layer 2 or 3 data packets from one IPv6 host to another IPv6 host are encapsulated with an IPv4 outer header and transported over the IPv4 underlay network. The Layer 3 VXLAN gateways in the EVPN-VXLAN overlay network learn the IPv6 routes through the exchange of EVPN type-2 and type-5 routes. To enable IPv6 data traffic support, you configure the IRB interfaces on all Layer 3 VXLAN gateways with the same IPv4 and IPv6 anycast virtual gateway addresses (VGAs). To support this feature, no other IPv6 configuration is required in the underlay or overlay networks.

- **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, endpoints (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.

In addition, QFX10000 switches that function as Layer 3 VXLAN gateways can route IPv6 data traffic through an EVPN-VXLAN overlay network. With this feature enabled, Layer 2 or 3 data packets from one IPv6 host to another IPv6 host are encapsulated with an IPv4 outer header and transported over the IPv4 underlay network. The Layer 3 VXLAN gateways in the EVPN-VXLAN overlay network learn the IPv6 routes through the exchange of EVPN type-2 and type-5 routes. To enable IPv6 data traffic support, you configure the IRB interfaces on all Layer 3 VXLAN gateways with the same IPv4 and IPv6 anycast virtual gateway addresses (VGAs). To support this feature, no other IPv6 configuration is required in the underlay or overlay networks.

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 21](#)
- [Known Behavior on page 23](#)
- [Known Issues on page 27](#)
- [Resolved Issues on page 30](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)
- [Product Compatibility on page 57](#)

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.

- [Infrastructure on page 21](#)
- [Interfaces and Chassis on page 22](#)
- [Routing Policy and Firewall Filters on page 22](#)
- [Software-Defined Networking \(SDN\) on page 22](#)
- [VPNs on page 23](#)
- [VXLAN on page 23](#)

Infrastructure

- On QFX10000 switches, correctable (ccw_lw) and non-correctable (nccw_lw) FEC error counters and BIP-8 counters are reported.
- Starting in Junos OS Release 15.1X53-D63, Junos system memory has been increased from 4GB to 8GB on QFX10002 switches, supporting 128K LSPs. (This change did not occur on QFX10008 or QFX10016 switches.) You can see the difference in the system memory by comparing output from the **show system memory** command in a pre-15.1X53-D63 release to output in 15.1X53-D63:

Pre-15.1X53-D63:

```
user@switch> show system memory
re0:
-----
System memory usage distribution:
  Total memory: 3961716 Kbytes (100%) <<<<
  Reserved memory: 115524 Kbytes ( 2%)
  Wired memory: 686924 Kbytes ( 17%)
  Active memory: 74976 Kbytes ( 1%)
  Inactive memory: 1393384 Kbytes ( 35%)
  Cache memory: 0 Kbytes ( 0%)
  Free memory: 1690908 Kbytes ( 42%)
```

15.1X53-D63:

```
user@switch> show system memory
re0:
-----
System memory usage distribution:
  Total memory: 7961200 Kbytes (100%) <<<<
  Reserved memory: 217328 Kbytes ( 2%)
  Wired memory: 1538968 Kbytes ( 19%)
  Active memory: 2772544 Kbytes ( 34%)
  Inactive memory: 2666176 Kbytes ( 33%)
  Cache memory: 0 Kbytes ( 0%)
  Free memory: 766184 Kbytes ( 9%)
```

Interfaces and Chassis

- On QFX10000 switches, do not configure RPF on IRB interfaces that have a Layer 2 interface configured with a MAC limit packet action of **drop** or **drop-log**.

Routing Policy and Firewall Filters

- The **set firewall filter** configuration for PE-based QFX Series platforms is not supported. The PE-based QFX Series platforms support filters with a specified firewall family—that is, a **set firewall family *family* filter** configuration.

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](#).]

- On QFX10000 switches running Junos OS Release 15.1X53-D65 or later, the local preference setting for an Ethernet VPN (EVPN) pure type-5 route is inherited by IP routes that are derived from the EVPN type-5 route. Further, when selecting an IP route for incoming traffic, the QFX10000 switches consider the local preference of the route. A benefit of the QFX10000 switches including local preference in their route selection criteria is that you can set up a policy to manipulate the local preference, thereby controlling which route the switch selects.

VPNs

- **Support for ping on a virtual gateway address (QFX10000)**—In Junos OS Release 15.1X53-D65, Junos OS supports pinging an IPv4 or IPv6 address on the preferred virtual gateway interface. To set up support for ping, you must include both the **virtual-gateway-accept-data** and the **preferred** statements at the **[edit interfaces irb unit]** hierarchy of the preferred virtual gateway. This enables the interface on the preferred virtual gateway to accept all packets for the virtual IP address, including ping packets.

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 **vxlan ingress-node-replication**.

Related Documentation

- [New and Changed Features on page 5](#)
- [Known Behavior on page 23](#)
- [Known Issues on page 27](#)
- [Resolved Issues on page 30](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)
- [Product Compatibility on page 57](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1X53-D69 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.

- [EVPN on page 24](#)
- [Interfaces and Chassis on page 24](#)
- [Layer 2 Features on page 24](#)
- [Layer 3 Features on page 25](#)
- [Multicast Protocols on page 25](#)
- [Network Management and Monitoring on page 25](#)
- [Platform and Infrastructure on page 25](#)
- [Routing Protocols on page 26](#)
- [Software Installation and Upgrade on page 26](#)
- [User Interface and Configuration on page 27](#)

EVPN

- **L2ALD_IFBD_COUNT_EXCEED: ; current count is** might be logged in syslog in the IFBD scale scenario. IFBD is the logical interface per VLAN or bridge domain. For example, a trunk interface with VLAN members v100, v200, v300 has 3 IFBDs. The current configured maximum IFBD number in TVP schema file for QFX10000 is 262144. [PR1203893](#)
- When EVPN-VXLAN is configured on QFX10002 switches, only 40,000 ARP entries are supported, because additional resources are required for IRB interfaces. [PR1206633](#)

Interfaces and Chassis

- 100G QSFP28 optics not showing in **show chassis hardware** if port is not configured for speed 100g. [PR1169500](#)
- 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](#)
- In an EVPN-VXLAN deployment with QFX10000 switches, when a VXLAN-enabled IRB interface is configured in the same routing instance as that of the underlay VTEP tunnel and if the remote VTEP interface IP address is resolved over the IRB interface using routing protocols or static route, dcpfe core files are generated and all the interfaces would go down. The core files then continue to be generated until the configuration is corrected. [PR1257128](#)
- On QFX10000 switches, FEC counters are not cleared after a power-on or link bringup, because FEC statistics shown in the CLI are the residue of link bringup, and those statistics remain steady throughout. As a workaround, issue the **clear interfaces statistics interface-name** command to clear these FEC statistics. [PR1257751](#)

Layer 2 Features

- On QFX10000 switches, when Layer 2 learning is performed on LAG interfaces, flooding might happen when one of the FPCs that is a part of the LAG restarts. [PR1190822](#)

- When multicast packets are replicated on 1000 IRB interfaces that have multiple receivers over LAG (with multiple child members), it leads to scaled kernel operations, which in turn leads to complete system slowdown for a few minutes for any multicast route changes (add/delete). [PR1195294](#)

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](#)

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

- 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](#)
- If BFD is configured on a LAG interface, disabling a member of the LAG can cause BFD to flap if that member carries BFD packets. [PR1333300](#)
- When the SFlow collector can be reached only through the Routing Engine, large samples due to heavy traffic can cause the Routing Engine CPU to become busy. [PR1332337](#)

Platform and Infrastructure

- On QFX10002 switches, modifying or rolling back a scaled configuration multiple times might cause disk space issues in the configuration partition (/var/run/db). The scale of the configuration depends on the configuration hierarchy. For example, configuring more than 32,000 firewall filters or terms and doing a modification more than four times can result in this issue. As a workaround: 1. Navigate to the Junos OS shell by issuing start shell at the CLI prompt. 2. Run mgd - i. 3. Make a small change in the configuration and commit the configuration. [PR1076356](#)
- A console or SSH session with a QFX10002 switch might hang if the switch is concurrently handling large-scale OVSDB transactions including but not limited to the dynamic creation of a large number of VXLANs or logical interfaces. [PR1087323](#)
- With multihop BFD, traffic loss of around 5 to 10 seconds is observed when an intermediate interface is shut down. After 5 to 10 seconds, traffic recovers and no action is needed. [PR1150695](#)

- On QFX10000 switches, while flapping, LAG members with unknown unicast traffic might cause traffic loss. [PR1156691](#)
- DAC transceivers are not supported. [PR1161345](#)
- 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](#)
- Configuration "set vlan <cli>VLAN name</cli> vlan-id none" is not supported for QFX10000 in the flexible VLAN tagging scenario. [PR1337969](#)

Routing Protocols

- On a QFX10002 switch, when you configure a filter to discard VLAN floods, the filter also discards known unicast traffic. [PR1054677](#)
- On QFX10002 switches, multicast traffic that ingresses from a GRE tunnel is not de-encapsulated and is dropped. [PR1089319](#)
- If there is a failover from one Routing Engine to another on a QFX10008 switch, IPv6 traffic might be dropped. This loss can occur regardless of whether the traffic is for the default routing instance or a virtual routing instance. [PR1134476](#)
- With higher scale of 16,000 logical interfaces, traffic drop might be seen during GRES/master Routing Engine reboot on the QFX10000 line of switches. [PR1148979](#)
- On QFX10008 switches, if VRRP is configured, some VRRP sessions might flap when GRES is performed. [PR1153784](#)
- VRRP is not supported for IPv6 Layer 3 subinterfaces. [PR1176277](#)
- On the QFX10000 line of switches, the Packet Forwarding Engine manager (dc-pfe) might crash continuously and cause traffic drop when adding the IPv6 destination address match condition within the family Ethernet switching filter. [PR1210681](#)
- RSVP refresh reduction configuration must be applied to achieve higher RSVP LSP scale: **set protocols rsvp interface all aggregate set protocols rsvp interface all reliable.** [PR1349626](#)

Software Installation and Upgrade

- On QFX10002 switches, rolling back the software by issuing the **request system software rollback** command is not supported. [PR1070892](#)
- Configuration DB gets formatted when downgrading from 15.1X53-D60 (FreeBSD 10.x-based Junos OS) to 15.1X53-D33 (FreeBSD 6.1-based Junos OS). [PR1186797](#)

User Interface and Configuration

- When entering the "restart r" incomplete command in the CLI, the command "restart routing" is executed. It should throw an error like "error: invalid daemon: r". [PR1075746](#)

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 21](#)
- [Known Issues on page 27](#)
- [Resolved Issues on page 30](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)
- [Product Compatibility on page 57](#)

Known Issues ---

This section lists the known issues in hardware and software in Junos OS Release 15.1X53-D69 for QFX10000 switches.

- [EVPN on page 27](#)
- [Interfaces and Chassis on page 28](#)
- [Layer 2 Features on page 28](#)
- [Multicast on page 28](#)
- [Platform and Infrastructure on page 28](#)
- [Routing Protocols on page 29](#)

EVPN

- On QFX10000 switches, traffic for a given set of VLANs will not be forwarded to a VTEP if the destination VTEP does not have that set of VLANs configured. As a workaround, in a scale environment on a QFX10000 switch running Junos OS Release 15.1X53-D60, disable ingress-node-replication and reload the switch. [PR1207495](#)
- EVPN will not automatically bring down access side links when the PE is isolated from the BGP core network. In active/active multihoming deployments, this can result in the PE receiving traffic from the multihomed site without any EVPN control plane state available to forward the traffic onward to remote PEs. [PR1334434](#)

Interfaces and Chassis

- On a QFX10000 switch, when a multichassis link aggregation group (MC-LAG) configuration is applied through apply-groups, the commit might fail, and the following error message might be displayed: **IRB interface(irb.1) and l2-interface(ae0.0) do not belong to the same routing instance.** [PR1069782](#)

Layer 2 Features

- 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](#)
- With multihop BFD, traffic loss of around 5 to 10 seconds is observed when an intermediate interface is shut down. After 5 to 10 seconds, traffic recovers and no action is needed. [PR1150695](#)
- On QFX10000 platforms, at a high traffic rate with sFlow enabled, the line card shows high CPU utilization, but there is no functional impact. [PR1160121](#)
- On QFX10000 with QFX10000-36Q or QFX10000-30C inserted, the FPC might crash unexpectedly if the U-Boot release is below 3.0 and Rev us as the following. As a workaround, upgrade U-Boot firmware for those line cards to 3.0. QFX10000-36Q: 750-051354 Rev 32 or below 750-068822 Rev 3 or below QFX10000-30C: 750-051357 Rev 29 or below. [PR1205364](#)
- When **clear ethernet-switching table** was issued on the spine device, the MAC entry got deleted. However, the ARP entry was not deleted for that host. In addition to this, the other spine devices that are part of the same Ethernet segment for the host have both the MAC and ARP entries. [PR1367957](#)

Multicast

- After leave and rejoin are sent in few seconds, L3 multicast traffic will not converge up to 100 percent and a few traffic drops will be seen continuously. This behavior will be seen while you scale beyond 2000 VLANs or 2000 IRB interfaces with VLAN replication in the system. [PR1135045](#)
- In a scenario with more than 2000 VLANs or IRB interfaces, Layer 3 multicast traffic does not converge to 100 percent and continuous drop is observed after the downstream interface is taken down or brought up or while an FPC comes online after FPC restart. [PR1161485](#)

Platform and Infrastructure

- On a QFX10002 switch, insert a small form-factor pluggable (SFP) on the management interface (em1). After a system reboot, replace the SFP with a copper SFP, the management interface might not work properly with speed 10m/100m. [PR1075097](#)
- A two-second traffic loss is seen during rpd restart (BGP protocol reconvergence). [PR1330201](#)

- Traffic is not sent out after you enable **enhanced-hash-key inet no-ipv4-source-address/no-ipv4-destination-address**. DCD informs the kernel to delete ifaddr on the irb.4, then add it later on. As part of processing the delete, the kernel deletes all remote ARP entries and informs l2ald. When the ifaddr add comes, there is no way for l2ald to replay these remote ARP entries. As a workaround, when the configuration is changed, IRB interface must also be deactivated along with the change, and later on reactivated. This recovers the ARP entries. [PR1334529](#)
- With VXLAN configured, when a QFX10000 functions as a spine device, it drops IPv6 traffic from the leaf switches. [PR1367548](#)
- All child interfaces of the aggregated Ethernet interface on a device share a single L2D ht entry unlike in the case of a physical port. So if PE ports 0, 12, 24, 36 are part of an aggregated Ethernet interface, then all will point to single bit 0 in the VLAN membership. [PR1373955](#)
- On all QFX Series platforms where processes related to Packet Forwarding Engine and processes related to platforms run on the Linux host OS and Routing Engine runs as Junos VM (QFX10000), the RPC output does not show failure when **request system software add** command is run with software already staged. [PR1353466](#)
- On a Junos device in the multichassis link aggregation group (MC-LAG) scenario with integrated routing and bridging (IRB) interface and enhanced-convergence enabled, if the MC-LAG has only one member link, after taking offline the FPC hosting that member link and then clearing ARP, the traffic which is expected to egress the interchassis link (ICL) might get dropped, due to the nexthop being incorrectly set as Discard by code in Junos kernel. [PR1381446](#)
- On QFX10000/PTX1000/PTX10000 devices, when 6PE (the **protocols mpls ipv6-tunneling** statement) is configured, the IPv6 Next-Hop installation might be incorrect and it might cause IPv6 traffic drop. [PR1393937](#)
- On QFX10008/QFX10016 platforms, when the FPC come online after a restart, the intra-VLAN traffic ingressing on the aggregated Ethernet interface might be permanently lost if MC-LAG enhanced-convergence is configured and there is only one member link in MC-LAG on other FPC. [PR1409631](#)

Routing Protocols

- On QFX Series switches, if equal-cost routes are flapping, some unilist next hops may not be deleted, even if they are not referenced. This might result in overrunning the ECMP group limit and failing to install new next-hops. [PR1096600](#)
- On a QFX10008 switch deployed in a spine and leaf topology, temporary loops can occur when Layer 3 links are disabled between a spine and a leaf with large numbers of routes. The duration of the loop depends on the number of route prefixes in the leaf and the spine. For example, with 20K total routes, the loop is on the order of 1-2 seconds; with 70–100K total routes, the loop is on the order of 15 seconds. [PR1140086](#)
- On a QFX10008 switch, if maximum ECMP (16) and BGP multipath are configured, the switch might install 32 paths instead of 16 paths. [PR1141454](#)

- On QFX10008 switches, if VRRP is configured, some VRRP sessions might flap when GRES is performed. [PR1153784](#)
- Modifying the route-distinguisher (RD) identifier flushes all the routes, including interface routes, from the virtual routing and forwarding (VRF) routing table, also known as a routing information base (RIB). As a result, the VRF RIB remains empty. As a workaround, deactivate the VRF before modifying the route-distinguisher identifier. [PR1155647](#)
- Traffic drop may be seen while bringing up the 16 ports out of 64-way ECMP which was made down , before BGP routes getting converged. [PR1178158](#)
- While restarting an FPC with 32 members with 64-way ECMP , you might see high convergence time. This behavior will be seen with 64-way ECMP with 32 members in two different fpcs. [PR1186979](#)
- On QFX10000 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. [PR1187169](#)
- On QFX10000 line switches, traffic drop is seen with IS-IS version 6 traffic during convergence in either of the following two scenarios: 1. While doing port unshutdown (that is, bringing up the ports after bringing them down). 2. While an FPC comes online after an FPC restart. This behavior is seen while one of the IS-IS version 6 session is flapped. [PR1190180](#)
- Reboot from shell on master Routing Engine (for GRES) is not supported on FreeBSD 10.x-based Junos OS platforms. [PR1190815](#)
- Stale IPv6 entries are seen where the entries are deleted from evpn databases These will either be eventually deleted, or if the hosts re-sent NS/ND, they will be updated again across all PEs. [PR1367947](#)

**Related
Documentation**

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 21](#)
- [Known Behavior on page 23](#)
- [Resolved Issues on page 30](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)
- [Product Compatibility on page 57](#)

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-D69 on page 31](#)
- [Resolved Issues: Release 15.1X53-D68 on page 31](#)
- [Resolved Issues: Release 15.1X53-D67 on page 32](#)
- [Resolved Issues: Release 15.1X53-D66 on page 34](#)
- [Resolved Issues: Release 15.1X53-D65 on page 39](#)
- [Resolved Issues: Release 15.1X53-D64 on page 40](#)
- [Resolved Issues: Release 15.1X53-D63 on page 42](#)
- [Resolved Issues: Release 15.1X53-D62 on page 45](#)
- [Resolved Issues: Release 15.1X53-D61 on page 46](#)
- [Resolved Issues: Release 15.1X53-D32 on page 47](#)

Resolved Issues: Release 15.1X53-D69

EVPN

- On EVPN setups, incorrect destination MAC addresses starting with 45 might show up when using the **show arp hostname** command. This is a cosmetic issue with no impact. [PR1392575](#)

Platform and Infrastructure

- Prior to Junos OS Release 15.1X53-D68, Layer2 traffic with destination mac matches with irb mac-address will get punted to CPU and dropped in kernel. [PR1389191](#)

Resolved Issues: Release 15.1X53-D68

- [EVPN on page 31](#)
- [Layer 2 Features on page 31](#)
- [Platform and Infrastructure on page 32](#)
- [Software Installation and Upgrade on page 32](#)

EVPN

- If the EVPN type-5 route advertises an IPv6 default route, an infinite loop of entries is created when you issue the **show evpn ip-prefix-database** command. [PR1369166](#)
- On QFX10000 Series Switches, rpd might crash and generate a core file because of a memory allocation problem. [PR1369705](#)

Layer 2 Features

- If a port is the last one of an ASIC chipset, after it is deleted from an aggregated Ethernet interface having multiple interfaces and VLANs configured, it might still forward traffic with the corresponding VLANs. [PR1372577](#)

- On QFX10000 and QFX10008 with a GRE tunnel deployed, the maximum transmission unit (MTU) change of physical interface for a GRE tunnel might cause the dcpfe process to crash and impact traffic. [PR1384517](#)

Platform and Infrastructure

- On QFX10000 line of switches, sFlow monitoring output might display a negative number of samples after a long run. [PR1244080](#)
- On QFX10000 platforms, Netconf SSH TCP port 830 traffic hitting the host path or an unclassified queue causes DDoS violations in the unclassified queue. The following log appears: **DDOS_PROTOCOL_VIOLATION_SET: Warning: Host-bound traffic for protocol/exception.** [PR1345744](#)
- On the QFX10000 line of switches with DDoS, incorrect DDoS counter values and syslog messages might be seen after DDoS statistics for a specific protocol statistics are manually cleared. [PR1351212](#)
- QFX10000 platform drops the vendor wireless access point (AP) heartbeat packets; as a result, the vendor wireless access point does not work. [PR1352805](#)
- FIP packets will be forwarded across a trunk interface that does not have the vlan associated with the FIP packet. This can cause looping depending on the topology. [PR1360830](#)
- On the QFX10000 line of switches, the Packet Forwarding Engine (PFE) might crash if the GRE tunnel destination IP address is resolved over another GRE tunnel. [PR1382727](#)

Software Installation and Upgrade

- When QFX10000 switches are upgraded to Junos OS Release 15.1X53-D60 from previous Junos OS 15.1X53 releases, the contents of the `/var/db/scripts` directory are not preserved. So the scripts related features cannot take effect for scripts missing. [PR1209576](#)

Resolved Issues: Release 15.1X53-D67

EVPN

- Error message - JPRDS_DLT_ALPHA KHT- shows as failed, but the entries in hardware are programmed correctly. This may cause confusion between working and non-working condition. [PR1258933](#)
- In EVPN/VXLAN scenario with VGA configured, if the packets' destination MAC is virtual gateway MAC and the packets don't have IP header, then these packets are looped between virtual gateways. This behavior causes high link utilization between virtual gateways. [PR1318382](#)
- In Ethernet VPN (EVPN)/Virtual Extensible Local Area Network (VXLAN) environment with multiple layers of Bidirectional Forwarding Detection (BFD) setting, if there are BFD flaps on the overlay Border Gateway Protocol (BGP) peering and the virtual tunnel endpoint (VTEP) flaps, the Packet Forwarding Engine might crash. It is a timing issue due to communication delays between the Kernel and the Packet Forwarding Engine.

The issue recovers itself in the sense that the Packet Forwarding Engine reboots and recovers. [PR1339084](#)

- RPD has unreproducible core file generated with scaling EVPN-VXLAN configuration on QFX10000 platform due to memory depletion on EVPN MAC route entries queue for L2ALD. L2ALD closed the IPC connection which caused RPD cumulated EVPN MAC route entries in the queue and ends up running out of memory. [PR1339979](#)
- After upgrading to Junos OS Release 15.1X53-D66 on a QFX10002 switch, debug messages might start flooding the system log message output if the following CLI command is executed: **set system syslog file *filename* any**. [PR1343876](#)

Interfaces and Chassis

- In a multichassis link aggregation group (MC-LAG) scenario on a QFX Series switch, if an aggregated Ethernet (AE) interface is configured as interchassis link (ICL) and Internet Group Management Protocol (IGMP) snooping is enabled, then a traffic loop might appear. [PR1338278](#)

Layer 2 Features

- On QFX10000 platforms, if the QFX10000 works as MC-LAG and DHCP-Relay role, the Dynamic Host Configuration Protocol (DHCP) Discover packets might be looped back to the DHCP client. As a result, the clients are not able to obtain an IP address. [PR1325425](#)
- A warning would be issued if "flexible-vlan-tagging" and "family ethernet-switching" are both configured on the same interface on QFX10000, as this configuration is not supported on this device. [PR1337311](#)
- During IP/MAC move, old mac+ip entries are deleted from the global db. But under certain conditions, one of these entries missed the corresponding deletion from the Bridge Domain's hbt tree. [PR1339543](#)

Network Management and Monitoring

- Error messages are seen after configuring multiples interfaces under the protocol sflow, and these errors are related to sFlow Bindpoint set error and related to lookup failed IFD_EGRESS_IMPL_FILTER on doing commit. Whenever any commit is performed on QFX10002, these errors messages are not logged in syslog. [PR1346493](#)

Platform and Infrastructure

- In a very rare scenario, during a TAC accounting configuration change, the auditd process crashes due to a race condition between auditd and its sigalarm handler. [PR1191527](#)
- In case that license keys are activated in the system via the configuration, which would mean under "system license keys" configuration stanza, certain events/changes can make them non-effective. Those events/changes include Routing Engine mastership switchover or "group" related configuration changes. [PR1259460](#)
- Configuring an IRB interface as an underlay on a Layer 2 gateway/leaf node may impact forwarding. [1267201](#)

Routing Policy and Firewall Filters

- On QFX Series switches, the command show policy which has a parameter of "load-balance consistent-hash" might cause rpd to crash. [PR1200997](#)

Routing Protocols

- Remotely received traffic was not flooded to AC on FPC 1 when FPC 0 was offline. [PR1290500](#)
- With protocol-independent load balancing for Layer 3 VPNs enabled (that is, configure 'routing-instances <routing instance name> routing-options multipath') in a virtual routing and forwarding (VRF) routing instance, when toggling a TTL action statement (that is, vrf-propagate-ttl/no-vrf-propagate-ttl) for this VRF routing instance, if BGP receives a VPN route update for the VRF during the processing of the reconfiguration, the rpd might crash. This is a timing issue due to the race condition. [PR1302504](#)
- On QFX10000 platform, TNP(Trivial Network Protocol) Ethertype 0x8850 is used for Control link communication. It might not get processed correctly as it is an unknown ethertype. Additional 2 bytes of 00 are inserted between the SMAC and the ethertype 0x8850. The packet received has DMAC, SMAC, 0000 (2bytes), 8850 (ether_type). It might cause communication issues. [PR1343575](#)

Resolved Issues: Release 15.1X53-D66

EVPN

- In an EVPN IRB deployment, when a given IP address initially bound to MAC M1 is later moved to be bound to MAC M2 instead, there might be a period of time where multiple IP routes exist for the IP address (one route associated with MAC M1, and one route associated with MAC M2) if M1 and M2 are hosted behind different EVPN PE devices on different Ethernet segments. Additionally, after such IP movement, multiple EVPN PEs might have ARP bindings for the IP address in question, though only one PE will have the latest binding. Other PEs might transiently have earlier, stale bindings until they age out via normal ARP procedures. [PR1141336](#)
- On QFX10000 platforms, when ARP learning happens, NH (Next-HOP) installation error messages are seen. [PR1258930](#)
- Ethernet A-D per Ethernet segment route (Type-1 PER ES) is NOT generated with a new route target after changing the route target. [PR1279529](#)
- In Ethernet Virtual Private Network/Virtual Extensible Local Area Networks (EVPN/VxLAN) scenario with dual-homing Layer3 gateway configuration, if one Layer3 gateway receives the changed Address Resolution Protocol (ARP) route update entry (IP+MAC route packet of EVPN) from another Layer3 gateway, it might delete the ARP/host route related (it should update this ARP route). In such error state, it might cause traffic drop. [PR1306024](#)
- In Ethernet VPN (EVPN) and Virtual Extensible VLAN (VXLAN) scenario, after clearing Address Resolution Protocol (ARP) and Media Access Control (MAC) on spine node, the mac address is stuck with "DR" flag on spine node even though packets are received

on interface from source MAC. Then MAC+IP Type-2 route is not generated due to "DR" flag of the MAC in ethernet-switching table. [PR1320724](#)

Forwarding and Sampling

- Error messages "SNMP_EVLIB_FAILURE: PFED ran out of transfer credits with Packet Forwarding Engine. Failed to get stats. ifl index:" seen in syslog. [PR1270686](#)

Interfaces and Chassis

- On QFX10002 or QFX10008 platforms, Packet Forwarding Engine (PFE) might crash after changing analyzer configuration if the "output" stanza includes a VLAN which belongs to a Link Aggregation Group (LAG) interface, or the "output" stanza includes a LAG interface. [PR1316245](#)
- If you use 100G-PSM4 optics in ULC30C line card and the firmware on the line card is below, you need to upgrade to the new firmware. FPC 4 U-Boot Bank A: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52) CTRL FPGA 3.2 <<<<<< PORT FPGA 2.1 <<<<<< New versions below - FPC 4 U-Boot Bank A: U-Boot 2011.12-gfbea47a (Feb 26 2016 - 22:56:52) CTRL FPGA 3.3 PORT FPGA 2.2 [PR1323321](#)

Layer 2 Features

- On the QFX10000 line of switches with Rapid Spanning Tree Protocol (RSTP) configured, when "bpdu-block-on-edge" is configured on the edge interface, and if the "fast-tune" is also enabled, the edge interface might not work correctly. [PR1307440](#)
- NLB heartbeat packets are not one of the known Ethernet types (0x886f) on QFX10000, which causes NLB packet to become corrupted and may be dropped. Two bytes header is added to the frame erroneously. [PR1322183](#)

MPLS

- For BGP-pipe mode OAM the MPLS echo reply is sent via inet.3 route. [PR1164406](#)
- The rpd might crash upon receiving a TLE (Tag Label Element) delete notification arriving during a cleanup sequence. When adaptive teardown is configured and TLE delete notification comes during a cleanup sequence, this triggers a recursive clean up and since the same cleanup routines are called and them being non-reentrant causes the code to assert. [PR1172567](#)
- This PR fixes an FD (file descriptor) leak problem in MGD process when NETCONF traceoptions are set. If <commit> rpc is executed via a NETCONF session, there is an FD leak in the corresponding MGD pid. [PR1174696](#)
- In LDP-signaled LSPs scenario, if LDP statistic is configured or the command **show ldp traffic-statistics** is executed, the device processes statistics for every LDP-signaled LSP. If there is an LSP with scaled next hops, it might take too much time to look up all the next hops and overloading the rpd. [PR1191406](#)
- When using RSVP-TE protocol to establish LSPs, make before break (MBB) might not quit and will start again when there is a failure on PSB2 (RSVP Path State Block for new LSP) in some cases where PathErr is not seen. (For example, for a PSB2 that is

already up and there is PathErr processing for it in place already; in this case, no PathErr is seen owing to local-reversion and a quick flap.) As a result, no rerouting happens even if the TE metric cost is raised. This issue has more chances of occurring when there is non-default optimize switchover delay. [PR1205996](#)

- In a scaled environment, when there are many unicast next hops related to the same LSP (for example, the same RSVP or LDP label), MPLS traffic statistics collection might take too much CPU time in kernel mode. This can in turn lead to various system impacting events, like scheduler slips of various processes and losing connection toward the backup Routing Engine and FPCs. [PR1214961](#)
- If the link or node failure that triggered a bypass persists for a long time, and there are LSPs that do not get globally repaired, as a result multiple stale LSP entries are shown and get listed multiple times in the MPLS LSP. [PR1222179](#)
- The routing protocol process crash might be seen if egress-policy is configured in LDP and the same route prefixes are in both inet.0 and inet.3. [PR1266358](#)
- When MPLS builds the next hop for an mpls.0 route for the scenario with IDP over RSVP LSP over bypass tunnel and the IDP label is implicit-NUL, the label stack constructed for the next hop might be incorrect, with an invalid bottom label value of 1048575. [PR1270877](#)
- When performing traceroute to a remote host for an MPLS (Multiprotocol Label Switching) label-switched path signaled by the LDP (Label Distribution Protocol), the rpd process might crash. [PR1299026](#)
- In some cases, it is seen that the label states are getting deleted twice, which results in Routing Protocol Process (rpd) crash. This is applicable only when ultimate-hop popping (UHP) based label-switched paths (LSPs) are configured. [PR1309397](#)

Platform and Infrastructure

- In rare cases, the Packet Forwarding Engine might drop the TCP RST (reset) packet from the Routing Engine side while doing GRES or flapping an interface, and traffic might be dropped. [PR1269202](#)
- QFX10000 drops packets on Packet Forwarding Engine running on Junos OS Release 15.1X53-D63. [PR1306435](#)
- The Flexible PIC Concentrator (FPC) memory might be exhausted on QFX10002 or QFX10008 or QFX10016 platforms, after reaching 100% utilization. The FPC stops processing the control traffic for that FPC. The SHEAF leak messages would be seen in the syslog if hitting this issue. [PR1311949](#)

Port Security

- When Junos OS devices use the Link Layer Discovery Protocol (LLDP), the command **show lldp neighbor** displays the contents of PortID type, length, and value (TLV) received from the peer in the field 'Port Info', and it could be the neighbor's port identifier or port description. A Junos OS CLI configuration statement can select which "interface-name" or "SNMP ifIndex" to generate for the PortID TLV, so you do not have any problem as long as two Junos OS devices are connected for LLDP, but you might have an

interoperability issue if another vendor device that can map the configured 'port description' in the PortID TLV is used. In this case, Junos OS displays the neighbor's PortDescription TLV in the Port info field, and if the peer sets the port description whose TLV length is longer than 33 bytes (included), Junos OS is not able to accept the LLDP packets and discards the packets as errors. The PortID TLV is given as : "the port id tlv length = port description field length + port id subtype(1B)". [PR1126680](#)

- On a VM based QFX Series switch running Junos OS, when traffic flood on ingress of management port exceeds ~200-250Mbps, it generates an interrupt (irq11). If the traffic flood continues, the kernel can go high and slow responsive since irq11 takes higher priority. If this continues, the CMLC (pfe-chassisd) connection to Routing Engine drops and FPC reboots. [PR1149867](#)

Routing Policy and Firewall Filters

- On all Junos OS platforms with "vrf-target auto" configured under routing-instance, the rpd might crash after an unrelated configuration change. [PR1301721](#)

Routing Protocols

- The rpd crash might be seen during deletion of address family on an interface while rpf check is configured. The fix removes the possible inconsistency chance (that can trigger this type of rpd crash) between rpf-check flags in KRT table and interface family data structure for the same interface which has rpf-check enabled/disabled. [PR1127856](#)
- The rpd process might crash if restarting the interface control with LDP configured scenario. [PR1130494](#)
- On Junos OS based products, changes in routing-instance, like changing route-distinguisher or routing-option changes in some corner cases might lead to rpd crash. [PR1134511](#)
- When we have a route received from different eBGP neighbors, for this specific route, if all BGP selection criteria is matching, we might end up using router ID. As this is eBGP route, so BGP will use active route as the preferred one. Now if this specific route flapped with sequence from the non-preferred to the preferred path, RPD runs the path selection. During RPD path selection we might generate a core file. This issue has no operational impact. [PR1180307](#)
- In large-scale BGP route environments with multipath configured, if BGP sessions go down simultaneously, the rpd might crash because it cannot finish multipath cleanup within a 10-minute limit. [PR1209695](#)
- When multiple labels become stale in stale-label-holddown-duration (default 60 seconds), it restarts the timer and accumulates all the stale-labels without getting deleted. This might cause memory for allocating labels to be exhausted and then MPLS traffic might be affected due to abnormal/failing label allocation. [PR1211010](#)
- The routing protocol process (rpd) on a backup Routing Engine might restart unexpectedly in a large BGP NLRI environment. [PR1220651](#)
- On all platforms, if MPLS goes down due to link flap, FPC reboot, or restart, rpd core files could be seen. [PR1228388](#)

- When adding or deleting a dynamic-tunnel destination network for IPv6 over IPv4 dynamic UDP tunnels, an rpd core file might be seen. [PR1230152](#)
- An incorrect PE router is attached to an ESI when the router receives two copies of the same AD/ESI route (for example, one through eBGP and another one received from an iBGP neighbor). This causes a partial traffic black hole and stale MAC entries. You can confirm the issue by checking the members of the ESI: `user@router> show evpn instance extensive ...` Number of ethernet segments: 5 ESI: 00:13:78:00:00:00:00:00:01 Status: Resolved Number of remote PEs connected: 3 Remote PE MAC label Aliasing label Mode 87.233.39.102 0 0 all-active 87.233.39.1200 0 all-active <<<< this PE is not part of the ESI 87.233.39.101 200 0 all-active. [PR1231402](#)
- Routes learned over EBGp multipath peering might not get installed in the forwarding table resulting in traffic black-holing for the affected destinations. This happens only if in addition to EBGp multipath there is as well multihop configuration statement enabled for that peering and uRPF ((Unicast Reverse Path Forwarding)) check is enabled over the involved interfaces. Corresponding routes would end up stuck in the KRT queue and related KRT log messages containing error code 'EINVAL -- Bad parameter in request' would be seen in the logs. [PR1241501](#)
- If the add-path statement is configured for BGP, when deactivating the BGP peer, which has the add-path statement configured as well, and then active this BGP peer, after committing previous two steps separately, the routes learned from this BGP peer might not be advertised to other peers. This is a timing issue and the traffic sending to the destination associated with the peer might be broken if this issue happens. [PR1246349](#)
- On rare occasions during the route add or delete or change operation, the kernel might encounter a crash with the error "rn_clone_unwire no ifclone parent". [PR1253362](#)
- Rpd memory leak is seen when NG-MVPN type 6 and type 7 route adds/deletes/changes. The leak is 36 byte block size on Junos OS releases prior to Junos OS Release 15.1, and 44 byte block size on Junos OS releases 15.1 and later. [PR1259579](#)
- If IS-IS segment routing but certain interface is not enabled RSVP, then it might cause corrupted TLV 22 of IS-IS (the size of the value part of the TLV exceeds 255), and it might cause rpd to crash for parsing the LSP (labeled switchover path). [PR1262612](#)
- When the policy with damping is applied on BGP, the rpd might crash after deactivating or activating protocol bgp, which can result in protocol flap or traffic drop. [PR1272202](#)
- During LDP shutdown, route added and deleted by LDP in the inet.0 table might be in the process of being deleted but still in the inet.0 table. The **show route extensive** CLI command might cause RPD to crash when trying to display the task name for such LDP route. [PR1272993](#)
- In an IS-IS segment routing (SR) scenario, during interoperability with another vendor's device, which includes new subTLV for supporting SID (segment identifier) within LSPs, due to a issue in handling them, LSPs received from the other vendor's device might be parsed unsuccessfully and then be dropped on the Juniper Networks side. [PR1280522](#)

- If the statement "egress-te" is configured for BGP and the BGP flaps, the rpd might crash. If the statement "switchover-on-routing-crash" and NSR are not enabled on the device, while rpd recovers, there is an unexpected routing protocol disruption. [PR1295062](#)
- With BGP prefix independent convergence (PIC) enabled, the routing protocol process (rpd) might crash, generating a core file while deleting a multipath route. [PR1302395](#)
- On the QFX10000 line of switches in large scale routes scenario, when programming Generic Routing Encapsulation (GRE) information, packets drop might be seen due to egress next-hop error. [PR1308438](#)
- IS-IS SPF gets triggered by LSP updates containing changes in Reservable Bandwidth in TE extensions. [PR1313147](#)
- On QFX10000, Bidirectional Forwarding Detection (BFD) might be stuck in init state when BFD is deployed in load-balance over VXLAN scenario. [PR1326100](#)

VPNs

- An rpd crash might be observed with a segmentation fault after applying an L2VPN configuration followed by the "ping mpls l2vpn" command. [PR1272612](#)

Resolved Issues: Release 15.1X53-D65

EVPN

- In an IP-CLOS topology with a QFX10000 switch as a leaf node using EVPN-VxLAN and acting as a Layer 2 or Layer 3 gateway (in a collapsed scenario), rebooting the leaf node might cause a PFE process core. [PR1285705](#), [PR1294055](#)

Infrastructure

- An FPC major alarm might be seen with error messages **DLU: ilp memory cache error** and **DLU: ilp prot1 detected_imem_even error**. [PR1251154](#)
- When **console log-out-on-disconnect** is enabled, system reboot or switchover can result in processes remaining in the wait state and failure of the syslog feature. [PR1253544](#)
- The DCPFE might create a core file when an OSPF interface goes down. [PR1289716](#)
- The DCPFE might crash after a period of idle time. [PR1294055](#)

Interfaces and Chassis

- On QFX10000 switches with a large number of LAG interfaces (for example, more than 400), if a few of those LAG interfaces are flapped (for example, 10 LAG interfaces are disabled and then the disabling is committed), other LAG interfaces might also flap. [PR1250741](#)
- Interfaces do not come up randomly after a line card is rebooted. [PR1262839](#)
- LAG interface input bytes counter continuously decreases when no packets come in. [PR1266062](#)

- Multicast data packets are looping in MC-LAG. [PR1281646](#)
- GRE tunnel traffic does not switch over to the alternate path if the primary path to the tunnel destination changes. [PR1287249](#)
- The switch is learning its own MAC address on the network interface. [PR1291184](#)
- Traffic might be transmitted from a member interface of an AE interface with BFD and LACP down states in an mBFD scenario. [PR1300309](#)
- Disabled 10-gigabit interfaces might stay up. [PR1300775](#)

Network Management and Monitoring

- bgpPeerState/bgpPeerTable returns an invalid value for an IPv6 peer. [PR1233790](#)

Platforms and Chassis

- Some processes might not work after the switch is power-cycled. [PR1222504](#)

Security

- A DDoS protocol group of host-path cannot be configured, which might impact traffic of http/https. [PR1226937](#)

Resolved Issues: Release 15.1X53-D64

- [High Availability \(HA\) and Resiliency](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Network Management and Monitoring](#)
- [Port Security](#)
- [Routing Protocols](#)
- [Security](#)
- [VXLAN](#)

High Availability (HA) and Resiliency

- The VRRP virtual address might be lost on the VRRP backup when a new logical-interface VRRP group is added or a VRRP group for a logical interface is reconfigured. [PR1255978](#)

Interfaces and Chassis

- `expr_nh_fwd_create_arp_ndp_egress_descr()`, 1237:nh 131650 type Compst, failed to create L2 descr failure log message; no impact on traffic or performance. [PR1221831](#)
- During a firewall script run, a switchover is performed. The new master takes ownership and stays up, but the old master goes to `db>`. [PR1222582](#)

- On QFX10000, IPv4 traffic drops when a member interface of a LAG is changed. [PR1270011](#)

Layer 2 Features

- On QFX10000, an IPv6 double-tagged frame does not pass through the switch if the service-provider configuration style has been used. [PR1254492](#)

Network Management and Monitoring

- The eventd process stops sending syslog messages to a configured syslog server. [PR1246712](#)
- QFX10002 is not sending syslog messages. [PR1259603](#)

Port Security

- Storm control might not be programmed correctly in the Packet Forwarding Engine if it is applied with a port-speed configuration in a single commit. [PR1255562](#)

Routing Protocols

- A BGP export policy with **from protocol** might cause issue on a 64-bit rpd. [PR1206511](#)
- BGP routes might be seen in the advertising-protocol table on the local end but not be seen in the receive-protocol table on the remote end. [PR1231707](#)

Security

- NTP.org and FreeBSD have published security advisories for vulnerabilities resolved in ntpd (NTP daemon). Server-side vulnerabilities are only exploitable on systems where NTP server is enabled within the [edit system ntp] hierarchy level. A summary of the vulnerabilities that may impact Junos OS is in JSA10776. Refer to JSA10776 for more information. [PR1159544](#), [PR1234119](#)
- Incorrect signedness comparison in the ioctl(2) handler allows a malicious local user to overwrite a portion of the kernel memory. Refer to JSA10784 for more information, <https://kb.juniper.net/JSA10784>. [PR1184592](#)
- When an IPv6 node receives an ICMPv6 PTB (Packet Too Big) message with MTU < 1280, the node will emit atomic fragments. This behavior might result in denial of service attack. And please refer to JSA10780 for more information. [PR1250832](#)

VXLAN

- QFX100002 generated an L2ALD core file for an unknown reason at `l2ald_mac_process_update_fwd_entry_mask`, `l2ald_mclag_update_change_for_learn_mask`, `logging`, `vlogging`, `vlogging_event`. [PR1264432](#)

Resolved Issues: Release 15.1X53-D63

- [Authentication and Access Control](#)
- [High Availability \(HA\) and Resiliency](#)
- [Interfaces and Chassis](#)
- [MPLS](#)
- [Multicast Protocols](#)
- [Platforms and Chassis](#)
- [Routing Protocols](#)
- [Software-Defined Networking \(SDN\)](#)
- [Software Installation and Upgrade](#)

Authentication and Access Control

- On QFX Series switches, SSH authentication may fail due to improper file ownership. [PR1142992](#)
- In instances when an SSH-key is longer than 256 characters, the router can go into amnesiac mode and any login is denied. For example, if the authentication `ssh-dsa` value (the DSA public key) is configured at the `[edit system login user username]` hierarchy is longer than 256 characters, the router login may be denied. [PR1169516](#)

High Availability (HA) and Resiliency

- When nonstop routing (NSR) is configured in a group, and that group applied to routing-options, NSR sometimes fails. [PR1168818](#)
- On QFX10000 switches, when nonstop active routing (NSR) is configured with a Label Distribution Protocol (LDP) export policy or an L2 smart policy, the routing protocol process (rpd) on the backup RE may crash when LDP tries to delete a filtered label binding. [PR1211194](#)
- On QFX10000 switches operating with Layer 3 VPN and configured to allow chained composite next hops for devices handling ingress or transit traffic in the network, packets may not be forwarded after they pass through the generic routing encapsulation (GRE) tunnel. This issue is observed on routers operating with Layer 3 VPN that also include the statement `chained-composite-next-hop ingress` at the `[edit routing-options forwarding-table]` hierarchy level. When configured in this manner, the Packet Forwarding Engine cannot push VPN labels for packets. As a result, packets arriving at the next-hop destination cannot be forwarded. [PR1215382](#)

Interfaces and Chassis

- In MC-LAG environments on QFX10000 switches, partial packet loss may occur due to a delay in the Address Resolution Protocol (ARP)/Neighbor Discovery (ND) state being synchronized between MC-LAG peers. This issue has been observed between two routers (Router A and Router B). During Graceful Routing Engine switchover (GRES)/In-Service Software Upgrade (ISSU) on Router A, if an ARP entry ages out for hosts/servers on Router B, Router B re-arms to Router A. A possible reply may be received, but during the GRES/ISSU window, this state is not synchronized. Inter-Chassis Control Protocol (ICCP) is used to exchange control information between the MC-LAG peers, and ICCP is not operating during the GRES switchover. The ARP request is not resolved until Router A is fully rebooted. During this switchover window, packet loss may occur. [PR1079736](#)
- On QFX10000 switches, the **show interfaces interface-name extensive** output does not display the Physical Coding Sublayer (PCS) statistics. [PR121160](#)
- 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. [PR1229547](#)
- On QFX10000 switches configured with MC-LAG, Cisco Discovery Protocol (CDP) packets with destination address 01:00:0c:cc:cc:cc loop. [PR1237227](#)

MPLS

- RSVP local revertive mode is supported by default on all Juniper Networks routers running Junos OS. In instances when global revertive mode is configured to override the default RSVP local revertive mode by including the **no-local-reversion** statement at the **[edit protocols rsvp]** hierarchy, it is observed that sometimes during link failure, a link-protected route is associated indefinitely with the bypass label-switched path (LSP). This occurs when an interface is brought down on which the packet state block (PSB or new path) is established before the RSVP PSB switchover. This is a timing issue. [PR1091774](#)
- On QFX10000 switches, when changing the **routing-options forwarding-table chained-composite-next-hop** configuration while there are active MPLS LSPs, an LSP traffic loss may be observed afterwards. [PR1243088](#)
- On QFX10000 switches running in a virtual routing and forwarding (VRF) environment and configured for Dynamic Host Configuration Protocol (DHCP) Relay, DHCP offer packets (with an MPLS header) are dropped on the ingress MPLS interface. [PR1243936](#)

Multicast Protocols

- For devices populated with a master and backup routing engines (RE) and configured for nonstop active routing (NSR) and Protocol Independent Multicast (PIM) configuration, the routing protocol process (RPD) may crash on the backup RE due to a memory leak. This leak occurs when the backup RE handling mirror updates about PIM received from the master RE deletes information about a PIM session from its database. But due to a software defect, a leak of 2 memory blocks (8 or 16 bytes) may occur for every PIM leave. If the memory is exhausted, the rpd may crash on the backup RE. There is no impact seen on the master RE when the rpd crashes on the backup RE. [PR1155778](#)

Platforms and Chassis

- On QFX10000 switches, the routing protocol process (rpd) may eventually become exhausted and crash when Layer 2 Circuit, Layer 2 VPN, or virtual private LAN service (VPLS) configurations are committed. These commit activities may create a small memory leak of 84 bytes in the rpd. If the rpd memory is exhausted, recovery can be accomplished by retarting rpd. If nonstop routing (NSR) is configured, the master Routing Engine can be switched over to the standby Routing Engine, causing the master rpd to exit and restart and free the leaked memory. [PR1220363](#)
- When ICMP traffic is directed towards a local interface on a QFX10000 switch, high latency and jitter may be observed. While this issue is not service impacting, it can indicate an incorrect performance metric when troubleshooting traffic concerns. [PR1221053](#)
- On QFX10000 switches, the routing protocol process (rpd) sometimes is interrupted and halted when it tries to free a session reference block. This can occur when the memory redzone check fails when attempting to free reference memory block. The fail is caused when the redzone check receives an address that is not the beginning of a memory block. [PR1232742](#)
- 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](#)
- On QFX10000 switches, a power entry module (PEM) may be wrongly detected as offline, repeatedly triggering an SNMP trap. Shortly after the SNMP traps are generated, the PEM is detected as being online again. [PR1233537](#)
- On QFX10000 switches, the routing protocol process (rpd) sometimes crashes and produces a core-dump. This issue is observed when there is a full internet feed and a BGP peer goes down. [PR1250978](#)

Routing Protocols

- When a BGP speaker (router) has multiple peers configured in a BGP group, there is sometimes an inaccurate count of prefixes. This occurs when the BGP speaker receives a route from a peer and re-advertises the route to another peer within the same group. In such instances, the MIB object **jnxBgpM2PrefixOutPrefixes** for peers in the same

group reports the total number of advertised prefixes in the group. MIB value `jnxBgpM2PrefixOutPrefixes` is defined as being used on a per-peer basis. However, it is instead being used to report prefixes on a per-group basis. To display an accurate number of advertised prefixes, use the **show bgp neighbor** command. [PR1116382](#)

- 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](#)

Software-Defined Networking (SDN)

- On QFX Series switches with Virtual Ethernet VPN (EVPN) deployed, the routing protocol process (rpd) may crash if the following commands are executed:
 - **show evpn database neighbor *neighbor-name* vlan-id *vlan-id* mac-address *address***
 - **show evpn database vlan-id *vlan-id* mac-address *address***
 - **show evpn database vlan-id *vlan-id* mac-address *address* instance *instance-name***

[PR1119301](#)

Software Installation and Upgrade

- In some rare instances on QFX10002 switches, no network ports are detected following a software upgrade and the subsequent reboot sequence. The switch can experience this state due to a hardware failure or CPU memory issue that triggers an Inter-integrated Circuit (I2C) transaction failure. If it is not a hardware failure, rebooting the switch clears the issue. If it is a hardware failure, rebooting the switch will not provide recovery and a Return Material Authorization (RMA) for the affected part must be made. [PR1247753](#)

Resolved Issues: Release 15.1X53-D62

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

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](#)

Routing Protocols

- 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 unilist next-hop programming acknowledgments are delayed. As a result, BGP route updates can fail because no buffer is available. [PR1225309](#)

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. [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.

[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. [PR1210195](#)

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 ovbdb 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 5](#)
- [Changes in Behavior and Syntax on page 21](#)
- [Known Behavior on page 23](#)
- [Known Issues on page 27](#)
- [Documentation Updates on page 48](#)

- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)
- [Product Compatibility on page 57](#)

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 48](#)

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.

Software Defined Networking (SDN)

- For up-to-date information about configuring EVPN-VXLAN on QFX10000 switches, see [Example: Configuring EVPN-VXLAN In a Collapsed IP Fabric Topology Within a Data Center](#).
- The following updates will be added to the *EVPN Control Plane and VXLAN Data Plane Feature Guide for QFX Series Switches*:
 - QFX10000 switches that are deployed in an EVPN-VXLAN environment do not support an IPv6 physical underlay network.
 - In an EVPN-VXLAN environment, if you have manually configured a specific route target for each VNI by using the **vrf-target** statement in the **[edit protocols evpn vni-options vni]** hierarchy, you do not need to additionally configure the **vrf-target** statement with the **auto** option. The **auto** option automatically derives route targets for each VNI from the autonomous system (AS) that you specified by using the **autonomous-system** statement in the **[edit routing-options]** hierarchy.
 - In a topology in which EVPN-VXLAN is deployed over a two-layer IP fabric, a leaf device is multihomed to two QFX10000 switches, which function as spine devices. If the spine devices also function as redundant gateways, these devices must be configured with the same AS number that is specified in the **autonomous-system** statement in the **[edit routing-options]** hierarchy.

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 21](#)
- [Known Behavior on page 23](#)
- [Known Issues on page 27](#)
- [Resolved Issues on page 30](#)

- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)
- [Product Compatibility on page 57](#)

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 49](#)
- [Upgrading Requires Manual Copy of /var/db/scripts Files on page 50](#)
- [Downloading Software Files with a Browser on page 50](#)
- [Backing Up the Current Configuration Files on page 51](#)
- [Installing the Software on QFX10002 Switches on page 51](#)
- [Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches on page 52](#)
- [Installing the Software on QFX10008 and QFX10016 Switches on page 53](#)

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.
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.

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-D30	No	<p>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.</p> <p>Yes, but you must downgrade to Junos OS Release 15.1X53-D33 before you downgrade to Release 15.1X53-D30.</p>
Releases prior to 15.1X53-D30	No	<p>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.</p> <p>Yes</p>

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. [PR1209576](#)

Downloading Software Files with a Browser

To download the software package from the Juniper Networks Support website, go to <https://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 <https://www.juniper.net/support>.
- Click **Download Software**.
- In the By Technology box, click **Switching | QFX Series | QFX10008**.

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-10-m-flex-<version>-secure-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 on QFX10002 Switches



.....

NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

.....

Install the software in one of two ways:

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

For example:

```
user@switch> request system software add  
/var/tmp/jinstall-host-qfx-10-17.2R1.n-secure-signed.tgz reboot
```

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

For example:

```
user@switch> request system software add  
ftp://ftpserver/directory/jinstall-host-qfx-10-17.2R1.n-secure-signed.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
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches



NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

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>** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.4-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-host-qfx-10-m-15.1X53-D60.4-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.4-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-host-qfx-10-m-15.1X53-D60.4-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
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.



NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI `delete chassis redundancy` command when prompted. If GRES is enabled, it will be removed with the `redundancy` command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the `[edit routing-options]` hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-17.2R1.n-secure-signed.tgz
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the **request system reboot** command:

```
user@switch> request system reboot
```



NOTE: You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete <package-name>** command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

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

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software:

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

Current state	Backup
Election priority	Master (default)

```
Routing Engine status:
```

```
Slot 1:
```

Current state	Master
Election priority	Backup (default)

14. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-17.2R1.n-secure-signed.tgz
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

15. Reboot the Routing Engine using the **request system reboot** command:

```
user@switch> request system reboot
```



NOTE: You must reboot to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your system. Instead, finish the installation and then issue the **request system software delete jinstall <package-name>** command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the **show version** command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

```
Current state           Master
Election priority       Master (default)
```

```
Routing Engine status:
```

```
Slot 1:
```

```
Current state           Backup
Election priority       Backup (default)
```

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 21](#)
- [Known Behavior on page 23](#)
- [Known Issues on page 27](#)
- [Resolved Issues on page 30](#)
- [Documentation Updates on page 48](#)
- [Product Compatibility on page 57](#)

Product Compatibility

- [Hardware Compatibility on page 57](#)

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:

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

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 21](#)
- [Known Behavior on page 23](#)
- [Known Issues on page 27](#)
- [Resolved Issues on page 30](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX10000 Switches on page 49](#)

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>

- Join and participate in the Juniper Networks Community Forum:
<https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- 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, see <https://support.juniper.net/support/requesting-support/>.

Revision History

April 26, 2019—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D69

November 14, 2018—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D68

June 25, 2018—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D67

March 5, 2018—Revision 2, Junos OS for QFX10000 Switches, Release 15.1X53-D66

February 26, 2018—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D66

September 25, 2017—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D65

August 28, 2017—Revision 3, Junos OS for QFX10000 Switches, Release 15.1X53-D64—Moved PR1225309 to Resolved Issues.

July 21, 2017—Revision 2, Junos OS for QFX10000 Switches, Release 15.1X53-D64

June 26, 2017—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D64

April 24, 2017—Revision 2, Junos OS for QFX10000 Switches, Release 15.1X53-D63

April 4, 2017—Revision 1, Junos OS for QFX10000 Switches, Release 15.1X53-D63

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 © 2019 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be 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.