

Release Notes: Junos[®] OS Release 15.1X53-D231 for QFX5110 and QFX5200 Switches

Release 15.1X53-D231
October 20, 2017
Revision 3

Contents

Junos OS Release Notes for QFX5110 and QFX5200 Switches	4
New and Changed Features for QFX5110 and QFX5200 Switches	4
New Features in Release 15.1X53-D230	4
DHCP	4
Interfaces and Chassis	5
New Features in Release 15.1X53-D210	6
Hardware	6
Class of Service (CoS)	6
Infrastructure	6
Interfaces and Chassis	7
IPv6	8
Layer 2 Features	8
Layer 3 Features	9
MPLS	9
Multicast	10
Network Management and Monitoring	10
Port Security	10
Security	12
Routing Protocols	12
Software-Defined Networking (SDN)	13
System Management	14
New Features in Release 15.1X53-D30	14
Hardware	14
Infrastructure and Chassis	15
Interfaces and Chassis	16
Layer 2 Features	18
Layer 3 Features	18

MPLS	19
Multicast Protocols	19
Network Management and Monitoring	20
Security	21
Software Installation and Upgrade	21
Storage	21
System Management	22
Traffic Management	22
Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches	23
Authentication and Access Control	23
Known Behavior for QFX5110 and QFX5200 Switches	23
Infrastructure	24
Interfaces and Chassis	24
MPLS	24
Multicast Protocols	24
Known Issues for QFX5110 and QFX5200 Switches	25
Hardware	26
Class of Service (CoS)	26
Infrastructure	26
Interfaces and Chassis	26
Layer 2 Features	27
Layer 3 Features	27
Multicast Protocols	27
Network Management and Monitoring	27
Routing Policy and Firewall Filters	27
Software Installation and Upgrade	27
Spanning Tree Protocols	27
Resolved Issues for QFX5110 and QFX5200 Switches	28
Resolved Issues: Release 15.1X53-D232	28
Infrastructure	28
Interfaces and Chassis	28
Network Management and Monitoring	29
Port Security	29
Software-Defined Networking	29
Resolved Issues: Release 15.1X53-D231	29
Authentication and Access Control	29
Hardware	30
Interfaces and Chassis	30
Layer 2 Features	30
Layer 3 Features	30
Network Management and Monitoring	30
Routing Policy and Firewall Filters	30
Software-Defined Networking (SDN)	30
Security	30
Software Installation and Upgrade	31
Spanning-Tree Protocols	31
Resolved Issues: Release 15.1X53-D230	31
Authentication and Access Control	31
Infrastructure	31

Interfaces and Chassis	31
Layer 2 Features	32
Network Management and Monitoring	32
Resolved Issues: Release 15.1X53-D210	32
Firewall Filters	33
Interfaces and Chassis	33
MPLS	33
Platforms and Chassis	33
Documentation Updates for QFX5110 and QFX5200 Switches	33
Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200	
Switches	34
Downloading Software Files with a Browser	34
Backing Up the Current Configuration Files	35
Installing the Software	36
Product Compatibility for QFX5110 and QFX5200 Switches	36
Hardware Compatibility	36
Documentation Feedback	37
Requesting Technical Support	37
Self-Help Online Tools and Resources	38
Opening a Case with JTAC	38
Revision History	39

Junos OS Release Notes for QFX5110 and QFX5200 Switches

These release notes accompany Junos OS Release 15.1X53-D231 for QFX5110 and QFX5200 switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

New and Changed Features for QFX5110 and QFX5200 Switches

This section describes the new features and enhancements to existing features in Junos OS Release 15.1X53 for QFX5110 and QFX5200 switches. There are no new features or enhancements to existing features for QFX5110 and QFX5200 switches in Release 15.1X53-D232.

- [New Features in Release 15.1X53-D230 on page 4](#)
- [New Features in Release 15.1X53-D210 on page 6](#)
- [New Features in Release 15.1X53-D30 on page 14](#)

New Features in Release 15.1X53-D230

DHCP

- **Support for defining a custom string (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 15.1X53-D230, you can define a custom string for DHCP relay. The new feature of defining a custom string is similar to the existing feature of **use-interface-description** where you send the logical interface or physical interface description on DHCP option-82, but in this case, you have the flexibility of defining a value independent of the interface description and make use of that value as deemed. The configuration has to be done in two places, one where you define the string and one where you enable it.

Definition of the string can be done in three places:

- **#custom string per interface in a group** - Where the value is defined only for that interface. **[edit forwarding-options dhcp-relay group v4 interface irb.100 overrides]**
set user-defined-option-82 string
- **#custom string per group** - Where the value is defined per group
[edit forwarding-options dhcp-relay group v4 overrides]
set user-defined-option-82 string
- **#custom string for global** - Defined globally
[edit forwarding-options dhcp-relay overrides]
set user-defined-option-82 string

Enable the option:

- #Enabling the custom string to go out on circuit-id option82
**[edit forwarding-options dhcp-relay group v4 relay-option-82 circuit-id]
set user-defined**

Interfaces and Chassis

- **Auto-channelization of interfaces (QFX5200 switch)**—Starting in Junos OS Release 15.1X53-D230, you can use the auto-channelization feature to divide and channelize data automatically by detecting the cable type. The mode and number of channels are decided based on the channel link status. On QFX5200, auto-channelization supports three modes of operation with unique port settings:
 - When 4x10G split cables are connected, the 40G port auto-channelizes to four 10G channels.
 - When 2x50G split cables are connected, the 100G port auto-channelizes to two 50G channels.
 - When 4x25G split cables are connected, the 100G port auto-channelizes to four 25G channels.
- **CL74 FEC support for 25-gigabit and 50-gigabit channel speeds (QFX5200 switches)**—Starting with Junos OS Release 15.1X53-D230, you can disable or reenablen clause 74 (CL74)—as well as CL91—forwarding error correction (FEC) support on QFX5200 switches. FEC CL91 is supported for the 100-gigabit port speed and FEC CL74 is supported for both 25-gigabit and 50-gigabit port speeds. FEC CL91 is enabled by default for the 100-gigabit port speed; when the ports are channelized either in 4x25-gigabit or 2x50-gigabit, FEC CL74 is enabled.
 - To disable the FEC mode:
**[edit]
set interfaces *interface-name* gigether-options fec none**
 - To reenablen the FEC mode:
**[edit]
delete interfaces *interface-name* gigether-options fec none**
or
**[edit]
set interfaces *interface-name* gigether-options fec (fec74|fec91)**
 - To check FEC status:
show interfaces *interface-name*

The output for the show command will list FEC statistics for a particular *interface-name*, including the FEC corrected errors count, the FEC uncorrected errors count, and the type of FEC that was disabled or enabled.

New Features in Release 15.1X53-D210

Hardware

- **QFX5110-48S switch**—The QFX5110 line of switches is Juniper Network's versatile fixed-configuration solution for hybrid cloud deployments. The model QFX5110-48S is a 10-Gigabit Ethernet enhanced small form-factor pluggable plus (SFP+) switch with 48 SFP+ ports and four 100-Gbps quad small form-factor pluggable solution (QSFP28) ports. Each SFP+ port (0 through 47) can operate as a native 10-gigabit port or a 1-gigabit port when 1-gigabit optics are inserted. Each QSFP28 port (port numbers 48 through 51) can operate as a native 100-Gigabit Ethernet port, a native 40-Gigabit Ethernet port, or as four independent 10-gigabit ports when using breakout cables. The four QSFP28 ports can be used as either access ports or as uplinks. The QFX5110-48S provides full duplex throughput of 960 Gbps. The QFX5110-48S has a 1U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

Class of Service (CoS)

- **Class-of-service support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, class-of-service (CoS) support on QFX5110 switches is the same as on QFX5100 switches, providing all of the same configuration capabilities and functionality. QFX5110 switches provide a slight increase in buffer memory, which can be seen in the output of **show** commands.

[See [show class-of-service shared-buffer.](#)]

Infrastructure

- **Secure Boot (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, a significant system security enhancement, Secure Boot, has been introduced. 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.
- **Integrated software feature licenses (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, the standard QFX Series premium feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB) software license and the standard QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Multiprotocol Label Switching (MPLS), and Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDB) license are supported.

[See [Software Features That Require Licenses on the QFX Series.](#)]

Interfaces and Chassis

- **Channelizing 40-Gigabit Ethernet QSFP+ ports (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210 on the QFX5110-48S switch, there are four ports labeled 48 through 51, which support QSFP28 ports. The QSFP28 ports support 100-Gigabit Ethernet interfaces and 40-Gigabit Ethernet interfaces. You can channelize the 40-Gigabit Ethernet interfaces into four independent 10-Gigabit Ethernet interfaces by using breakout cables.

[See [Channelizing Interfaces on QFX5110-48S Switches](#).]

- **Multichassis link aggregation group (MC-LAG) (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, MC-LAG enables a client device to form a logical LAG interface using two QFX5110 switches. MC-LAG provides redundancy and load balancing between the two QFX5110 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 QFX5110 switches. Each of these QFX5110 switches has one or more physical links connected to a single client. The QFX5110 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

[See [Multichassis Link Aggregation Features, Terms, and Best Practices](#).]

- **IRB in PVLAN (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, you can configure an integrated routing and bridging (IRB) interface in a private VLAN (PVLAN) so that devices within community VLANs and isolated VLANs can communicate with each other and with devices outside the PVLAN at Layer 3 without requiring you to install a router.

[See [Example: Configuring a Private VLAN Spanning Multiple Switches with an IRB Interface](#).]

- **Link aggregation (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, link aggregation enables you to use multiple network cables and ports in parallel, which increases link speed and redundancy.

[See [Understanding Aggregated Ethernet Interfaces and LACP](#).]

Resilient hashing (QFX5110 switches)—Starting with Junos OS Release 15.1X53-D210, resilient hashing is supported by link aggregation groups (LAGs) and equal-cost multipath (ECMP) sets.

A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient

hashing enhances LAGs by minimizing destination remapping when a new member is added to or deleted from the LAG.

Resilient hashing works in conjunction with the default static hashing algorithm. It distributes traffic across all members of a LAG by tracking the flow's LAG member utilization. When a flow is affected by a LAG member change, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table. Destination paths are remapped when new members are added to or existing members are deleted from a LAG.

[See [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic.](#)]

- **Generic routing encapsulation (GRE) support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, you can use GRE tunneling services on QFX5110 switches 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, the switch that is 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.

IPv6

- **IPv6 feature support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, you can configure Neighbor Discovery Protocol, Virtual Router Redundancy Protocol (VRRP) for IPv6, and Protocol Independent Multicast (PIM) for IPv6. You can also configure BGP and IS-IS for IPv6, as well as OSPFv3. Additionally, unicast IPv6 is supported for virtual-router instances. DHCPv6 is also supported.

[See [Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery and Verifying and Managing DHCPv6 Local Server Configuration.](#)]

Layer 2 Features

- **VLAN support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
- **Link Layer Discovery Protocol (LLDP) support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, LLDP enables a switch to advertise its identity and capabilities on a LAN, as well as receive information about other network devices.

[See [LLDP Overview.](#)]

- **Q-in-Q tunneling support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, QFX5110 switches support Q-in-Q tunneling, which enables 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.

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

- **Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, these protocols enable a switch to advertise its identity and capabilities on a LAN and receive information about other network devices.

[See [Overview of Spanning-Tree Protocols](#).]

Layer 3 Features

- **Support to disable hierarchical ECMP (QFX5200 switches)**—Starting with Junos OS Release 15.1X53-D210, you can disable hierarchical equal-cost multipath (ECMP) groups for LDP forwarding equivalence classes (FECs) at system start time. Hierarchical ECMP is enabled by default. Disabling this feature effectively increases the number of ECMP groups. To disable hierarchical ECMP, include the **no-hierarchical-ecmp** statement at the **[edit forwarding-options]** hierarchical level. Disabling hierarchical ECMP causes the Packet Forwarding Engine to restart. To reenabling hierarchical ECMP, issue the following command: **delete forwarding-options no-hierarchical-ecmp**.

[See [no-hierarchical-ecmp](#).]

MPLS

- **MPLS support (QFX5110)**—Starting with Junos OS Release 15.X53-D210, the QFX5110 switch supports MPLS. MPLS is a method for engineering traffic patterns by assigning short labels to network packets that describe how to forward the packets through the network. MPLS is independent of routing tables or any routing protocol and can be used for unicast packets. The MPLS framework supports traffic engineering and the creation of VPNs. Traffic is engineered (controlled) primarily by the use of signaling protocols to establish label-switched paths (LSPs). VPN support includes Layer 2 and Layer 3 VPNs and Layer 2 circuits.

[See [MPLS Overview for QFX Series and EX4600 Switches](#).]

- **Equal-cost multipath routing on MPLS label-switching routers (QFX5110)**—Starting with Junos OS Release 15.1X53-D210, the QFX5110 switch supports equal-cost multipath (ECMP) routing on MPLS label-switching routers (LSRs). ECMP is a Layer 3 mechanism for load balancing traffic to a destination over multiple equal-cost next-hops. When a link goes down, ECMP uses fast reroute protection to shift packet forwarding to use operational links, thereby decreasing packet loss.

[See [Understanding ECMP Flow-Based Forwarding](#).]

Multicast

- **Layer 3 multicast support (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, IGMP—including versions 1, 2, and 3—IGMP snooping, PIM SM, and PIM SSM are supported. You can also configure IGMP, IGMP snooping, and PIM in virtual-router instances. MSDP is also supported. Configure IGMP at the **[edit protocols igmp]** hierarchy level. Configure IGMP snooping at the **[edit protocols igmp-snooping]** hierarchy level. Configure PIM at the **[edit protocols pim]** hierarchy level. Configure MSDP at the **[edit protocols msdp]** hierarchy level.

[See [Multicast Overview](#).]

Network Management and Monitoring

- **Port mirroring (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, you can use port mirroring on QFX5110 switches to copy packets entering or exiting a port or entering a VLAN and send the copies to a local interface for local monitoring or to a VLAN for remote monitoring. 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.

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

- **sFlow support (QFX5110)**—Starting with Junos OS Release 15.1X53-D210, the QFX5110 switch supports sFlow. 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 flow** and **clear sflow collector statistics**.

[See [Understanding How to Use sFlow Technology for Network Monitoring on a Switch](#).]

Port Security

- **Access security support (QFX5110)**—Starting with Junos OS Release 15.1X53-D210, the following access security features are supported on QFX5110 switches:
 - **DHCP snooping**—DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information, which it uses to build and maintain a database of valid IP-address-to-MAC-address (IP-MAC) bindings called the DHCP snooping database. Clients on untrusted ports are allowed access to the network only if they are validated against the database.
 - **DHCPv6 snooping**—DHCP snooping for DHCPv6.
 - **DHCP option 82**—You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

- **DHCPv6 option 37**—Option 37 is the DHCPv6 equivalent of the remote ID suboption of DHCP option 82. It is used to insert information about the network location of the remote host into DHCPv6 packets.
- **Dynamic ARP inspection (DAI)**—DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made on the basis of those comparisons.
- **IP source guard**—IP source guard prevents IP address spoofing by examining each packet sent from a host attached to an untrusted access interface on the switch. The IP address, MAC address, VLAN, and interface associated with the host are checked against entries stored in the DHCP snooping database. If the packet header does not match a valid entry in the DHCP snooping database, the packet is discarded.
- **IPv6 source guard**—IP source guard for IPv6.
- **IPv6 router advertisement (RA) guard**—IPv6 RA guard is a mitigation technique based on ICMPv6 Router Advertisement (RA) messages for attack vectors. RA guard is used to validate RA messages on the basis of whether they meet certain criteria, which are configured on the switch using policies. RA guard inspects RA messages and compares the information contained in the message attributes to the configured policy. Depending on the policy, RA guard either drops or forwards the RA messages that match the conditions.
- **IPv6 neighbor discovery (ND) inspection**—IPv6 ND inspection mitigates attacks based on the Neighbor Discovery Protocol by inspecting neighbor discovery messages and verifying them against the DHCPv6 snooping table.
- **MAC limiting**—You can configure MAC limiting on an interface or a VLAN, and specify the action to take on the next packet the interface or the VLAN receives after the limit is reached.
- **MAC move limiting**—You can configure MAC move limiting to track MAC address movements on the switch, so that if a MAC address changes more than the configured number of times within one second, the changes to MAC addresses are dropped, logged, or ignored, or the interface is shut down.
- **Persistent MAC learning**—Persistent MAC addresses (also called sticky MAC addresses) help restrict access to an access port by identifying the MAC addresses of workstations that are allowed access to a given port. Secure access to these workstations is retained even if the switch is restarted.

[See [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity.](#)]

Security

- **Firewall filters (QFX5110)**—Starting with Junos OS Release 15.1X53-D210, the QFX5110 switch supports firewall filters. You can configure firewall filters on the switch to provide rules that define whether to accept or discard packets. You can use firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), LAGs, and loopback interfaces.

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

- **Policers (QFX5110)**—Starting with Junos OS Release 15.1X53-D210, the QFX5110 switch supports policers. A switch polices (or rate-limits) traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service. You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets. You configure policer actions at the `[edit firewall]` hierarchy level.

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

- **Storm control (QFX5110)**—Starting with Junos OS Release 15.1X53-D210, the QFX5110 switch supports storm control. You can enable storm control on 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 the switch to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when a traffic storm occurs.

[See [Understanding Storm Control.](#)]

Routing Protocols

- **Support for advertising multiple paths in BGP (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, you can configure BGP to advertise multiple paths to the same destination, instead of advertising only the active path. The potential benefits of advertising multiple paths for BGP include fault tolerance, load balancing, and maintenance. Include the `add-path` set of statements at the `[edit protocols bgp group group-name family family-type]` hierarchy level.

[See [add-path.](#)]

- **Support for 64 next-hop gateways for ECMP (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, you can configure as many as 64 equal-cost-multipath (ECMP) next hops for RSVP and LDP LSPs. The following Layer 3 protocols are supported as ECMP gateways for both IPv4 and IPv6: OSPF, ISIS, EBGp, and IBGP (resolving over IGP routes). Include the `maximum-ecmp next-hops` statement at the `[edit chassis]` hierarchy level.

[See [Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing.](#)]

Software-Defined Networking (SDN)

- **Layer 2 VXLAN gateway (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 15.1X53-D210, you can implement a QFX5110 or a QFX5200 switch as a Virtual Extensible LAN (VXLAN) gateway. VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. You can use VXLAN tunnels to enable migration of virtual machines (VMs) between servers that exist in separate Layer 2 domains by tunneling the traffic through Layer 3 networks. This functionality allows you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See [Understanding VXLANs](#).]

- **EVPN control plane and VXLAN data plane support (QFX5110 and QFX5200 switches)**—By using a Layer 3 IP-based underlay network coupled with an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network, you can deploy larger networks than those possible with traditional Layer 2 Ethernet-based architectures. With overlay networks, endpoints (bare-metal servers [BMSes] or virtual machines [VMs]) can be placed anywhere in the network and can remain connected to the same logical Layer 2 network, enabling the virtual topology to be decoupled from the physical topology.

The physical underlay network over which EVPN-VXLAN is commonly deployed is a two-layer IP fabric, which includes spine and leaf devices. The spine devices provide connectivity between the leaf devices, and the leaf devices function as Layer 2 VXLAN gateways and provide connectivity to the attached endpoints. Starting with Junos OS Release 15.1X53-D210, you can deploy QFX5110 and QFX5200 switches as leaf nodes in the EVPN-VXLAN overlay network.

[See [Understanding EVPN with VXLAN Data Encapsulation](#).]

- **OVSDB support with Contrail (QFX5110 and QFX5200 switches)**—Starting with Junos OS Release 15.1X53-D210, the Open vSwitch Database (OVSDB) management protocol provides a means through which a Contrail controller can communicate with QFX5110 and QFX5200 switches to provision them as Layer 2 VXLAN gateways. In an environment in which Contrail Release 2.22 or later is deployed, a Contrail controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and the reverse.

[See [Understanding the OVSDB Protocol Running on Juniper Networks Devices](#).]

System Management

- **Zero Touch Provisioning (QFX5110 switches)**—Starting with Junos OS Release 15.1X53-D210, Zero Touch Provisioning allows you to provision new Juniper Networks switches in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, the switch attempts to upgrade the Junos OS software image automatically and autoinstall a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version.

[See [Understanding Zero Touch Provisioning](#).]

New Features in Release 15.1X53-D30

Hardware

- **QFX5200-32C switch**—The Juniper Networks QFX5200 line of fixed-configuration access switches is designed for cloud builders and data centers deploying next-generation IP fabric networks. The QFX5200-32C is a highly flexible, 32-port, fixed-configuration switch that can be configured for 10/25/40/50/100-Gigabit Ethernet speeds. The QFX5200-32C provides 100-Gbps spine and leaf connectivity in Layer 3 fabrics for cloud and web services.

The QFX5200-32C is a compact, 1 U standalone switch that provides a throughput of up to 3.2 Tbps, very low latency, and a rich set of Layer 3 features. The Routing Engine and control plane are driven by the 1.8 GHz quad-core Intel CPU with 16 GB of memory and two 32 GB solid-state drives (SSDs) for storage.

- **Support for 100-Gigabit optical transceivers (QFX5200 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 (QFX5200 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; 300m 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, 40Gigabit 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-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 (QFX5200 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
- **Support for QSFP+ direct attach copper (DAC) cables (QFX5200 switch)**—Provides support for:
 - EX-QSFP-40GE-DAC-CM—QSFP+ DAC assembly; 0.5 m, passive
 - QFX-QSFP-DAC-1M—QSFP+ DAC assembly, 1 M, passive
 - QFX-QSFP-DAC-3M—QSFP+ DAC assembly, 3 M, passive
 - QFX-QSFP-DAC-5M—QSFP+ DAC assembly, 5 M, passive
 - QFX-QSFP-DAC-7MA—QSFP+ DAC assembly, 7 M, active
 - QFX-QSFP-DAC-10MA—QSFP+ DAC assembly; 10 M, active

Infrastructure and Chassis

- **Disaggregated Junos OS (QFX5200 switch)**—Starting with the QFX5200 switch, the software has been disaggregated from the hardware. With disaggregated Junos OS, you can now purchase the Junos Base Services (JBS) license to use basic Junos OS functions, the Junos Advanced Services (JAS) license to use Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Virtual Extensible Local Area Network (VXLAN), and the Junos Premium Services (JPS) license to use features supported in the JAS license and the MPLS feature set. The disaggregated Junos OS feature licenses are available on a perpetual basis.



NOTE: You must purchase the JBS license to use basic functions, but you do not need to install the license key in Junos OS Release 15.1X53-D30. JBS basic functions work with this release without installing the license key. However, you will need to install the license key in a future release of Junos OS to be determined, so make sure to retain the authorization code you received from the License Management System to generate a license key for the JBS license.

Interfaces and Chassis

- **Channelizing 100-Gigabit Ethernet QSFP28 interfaces (QFX5200 switch)**—This feature enables you to channelize the 100-Gigabit Ethernet interfaces to two independent 50-Gigabit Ethernet or to four independent 25-Gigabit Ethernet interfaces. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated 40-Gigabit Ethernet interfaces or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables.

There are a total of 32 physical ports on the QFX5200 switch. Any port can be used as either 100-Gigabit Ethernet or 40-Gigabit Ethernet interfaces. You choose the speed by plugging in the appropriate transceiver. They can also be channelized to 50G, 25G or 10G.

By default, the 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces appear in the `et-fpc/pic/port` format. When the 100-Gigabit Ethernet interfaces are channelized as 50-Gigabit Ethernet and 25-Gigabit Ethernet interfaces, the interface names appear in the `et-fpc/pic/port:channel` format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format, where `channel` can be a value of 0 through 3. To channelize the ports, manually configure the port speed using the **set chassis fpc slot-number port port-number channel-speed speed** command, where the speed can be set to 10G, 25G, or 50G. The ports do not support autochannelization.



NOTE: If a 100G transceiver is connected to the switch, channelize the port only to 25G or 50G. If a 40G transceiver is connected, channelize the port only to 10G. Note that there is no commit check for these options.

- **Link aggregation (QFX5200 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) (QFX5200 switch)**—MC-LAG enables a client device to form a logical LAG interface using two QFX5200 switches. MC-LAG provides redundancy and load balancing between the two QFX5200 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 QFX5200 switches. Each of these QFX5200 switches has one or more physical links connected to a single client. The QFX5200 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

- **Resilient hashing support for link aggregation groups and equal cost multipath routes (QFX5200 switch)**—Resilient hashing is supported by link aggregation groups (LAGs) and equal cost multipath (ECMP) sets.

A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient hashing enhances LAGs by minimizing destination remapping when a new member is added to or deleted from the LAG.

Resilient hashing works in conjunction with the default static hashing algorithm. It distributes traffic across all members of a LAG by tracking the flow's LAG member utilization. When a flow is affected by a LAG member change, the Packet Forwarding Engine (PFE) rebalances the flow by reprogramming the flow set table. Destination paths are remapped when a new member is added to or existing members are deleted from a LAG.

Resilient hashing applies only to unicast traffic and supports a maximum of 1024 LAGs, with each group having a maximum of 256 members.

An ECMP group for a route contains multiple next-hop equal cost addresses for the same destination in the routing table. (Routes of equal cost have the same preference and metric values.)

Junos OS uses a hash algorithm to choose one of the next-hop addresses in the ECMP group to install in the forwarding table. Flows to the destination are rebalanced using resilient hashing.

Resilient hashing enhances ECMPs by minimizing destination remapping when a new member is added to or deleted from the ECMP group.

- **Ability to create link aggregation groups with interfaces operating at different speeds (QFX5200 switch)**—You can add 10-, 25-, 40-, 50-, and 100-Gigabit Ethernet interfaces into the same link aggregation group (LAG).
- **Support for Layer 3 logical interfaces (QFX5200 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 QFX5200 switch to a Layer 2 switch. Only one physical connection is required between the switches.
- **Generic routing encapsulation (GRE) support (QFX5200 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.

Layer 2 Features

- **VLAN support (QFX5200 switch)**—VLANs enable you to divide one physical broadcast domain into multiple virtual domains.
- **Link Layer Discovery Protocol (LLDP) support (QFX5200 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 (QFX5200 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 (QFX5200 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 (QFX5200 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 (QFX5200 switch)**—The IPv4 OSPF protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). QFX5200 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 (QFX5200 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 QFX5200 switch, you can configure BFD for static routes and for the BGP, IS-IS, OSPF, PIM, and RIP protocols.
- **IS-IS support (QFX5200 switch)**—The IS-IS protocol is an IGP for routing traffic within an AS.
- **Virtual Router Redundancy Protocol (VRRP) support (QFX5200 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.

- **Hierarchical ECMP (QFX5200 switch)**—Hierarchical ECMP resolves route prefixes to two-level ECMP automatically, allowing better load-balancing of traffic. Hierarchical ECMP is enabled by default.

MPLS

- **MPLS support (QFX5200 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 IPv4
 - LDP tunneling over RSVP
 - L2 Circuit (draft Martini) support
 - L3VPN Carrier-Over-Carrier (CoC)
 - ECMP on LSR
 - RSVP auto bandwidth
- **Equal cost multipath (ECMP) groups on label-switching router (LSR) devices for MPLS (QFX5200 switch)**—When a link goes down, ECMP uses fast reroute protection to shift packet forwarding to use operational links, thereby decreasing packet loss.

Multicast Protocols

- **Internet Group Management Protocol (IGMP) support (QFX5200 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 (QFX5200 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 (QFX5200 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 (QFX5200 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 (QFX5200 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 (QFX5200 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 (QFX5200 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.

Network Management and Monitoring

- **Cloud Analytics Engine network device support (QFX5200 switch)**—Cloud Analytics Engine network device support on QFX5200 switches provides flow path data analysis functions to help improve application performance and availability on the network. Cloud Analytics Engine includes components that enable network data collection, analysis, and correlation, helping you better understand the behavior of workloads and applications across the physical and virtual infrastructure.
- **SNMP support (QFX5200 switch)**—SNMP includes versions 1, 2, and 3 for monitoring system activity.

- **System logging (syslog) support (QFX5200 switch)**—Syslog enables you to log system messages into a local directory on the switch or to a syslog server.
- **sFlow technology support (QFX5200 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 (QFX5200 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.

Security

- **Firewall filter support (QFX5200 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 (QFX5200 switch)**—You can use policing to apply limits to traffic flow and to set consequences for packets that exceed those limits.
- **Storm control support (QFX5200 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 Installation and Upgrade

- **Support for FreeBSD 10 kernel for Junos OS (QFX5200 switches)**—On QFX5200 switches, FreeBSD 10 is the underlying OS that enables SMP for Junos OS, rather than the FreeBSD 6.1 that is used in some older Juniper Networks devices. If you compare the QFX5200 to devices that run the older kernel, you will notice that some system commands display different output and a few others are deprecated.

Storage

- **FIP snooping and Data Center Bridging Capability Exchange (DCBX) protocol (QFX5200 switch)**—QFX5200 supports both FIP snooping and DCBX. FIP snooping filters prevent an FCoE device from gaining unauthorized access to a Fibre Channel (FC) storage device or to another FCoE device. DCBX discovers the data center bridging (DCB) capabilities of connected peers. DCBX advertises the capabilities of applications on interfaces by exchanging application protocol information through application time-length-values (TLVs).
- **CEE (QFX5200 switch)**—CEE is an enhanced single interconnect Ethernet technology developed to converge a variety of applications in data centers. CEE's primary focus

is to consolidate the number of cables and adapters connected to servers. You can use data center bridging features on QFX5200 CEE-enabled switches to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) characteristics and other characteristics FC requires for transmitting storage traffic. Only port schedulers are supported; ETS is not supported.

System Management

- **Login authentication using RADIUS and TACACS+ (QFX5200 switch)**—You can use RADIUS and TACACS+ authentication to validate users who attempt to access the switch.
- **System utilization alarms support (QFX5200 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.

Traffic Management

- **Class of service (CoS) (QFX5200 switch)**—When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default class-of-service (CoS) settings or CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and any rewrite rules to re-mark packets.
- **Class-of-service (CoS) rewrite rules and classifier support (QFX5200 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. Packet classification maps incoming packets to a particular class-of-service (CoS) servicing level. You can use classifiers to map packets to a forwarding class and a loss priority and to assign packets to output queues based on the forwarding class.
- **Port scheduling with queue shaping support (QFX5200 switch)**—You can manage excess traffic and avoid congestion on a network interface where traffic might exceed the maximum port bandwidth. You can manage parameters such as transmit rate, shaping rate, and priority on each queue.
- **Priority-based flow control support (QFX5200 switch)**—This feature provides you with PFC (standard IEEE 802.1Qbb) capability, a link-level flow control mechanism that you can use to pause traffic selectively according to its class. You must use PFC for Fibre Channel over Ethernet (FCoE) traffic.
- **Ethernet PAUSE autonegotiation support (QFX5200 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.

Related Documentation

- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)
- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)

- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)
- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)
- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches

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 QFX5110 and QFX5200 switches.

Authentication and Access Control

- **Change in default value for port ID TLV for QFX5200 switches**—Starting from Junos OS Release 15.1X53-D230, for QFX5200 switches, the default value used for port ID TLV in LLDP messages is interface name, not SNMP index.

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)
- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)
- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)
- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

Known Behavior for QFX5110 and QFX5200 Switches

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Releases 15.1X53 for QFX5110 and QFX5200 switches.

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

- [Infrastructure on page 24](#)
- [Interfaces and Chassis on page 24](#)
- [MPLS on page 24](#)
- [Multicast Protocols on page 24](#)

Infrastructure

- The QFX5200 switch does not support cut-through switching mode. [PR1145359](#)
- On a QFX5200 switch, the **show interface statistics detail** command displays the speed as 1000m for the em0 interface regardless of what the correct speed is. [PR1083176](#)
- On an MC-AE node on a QFX5200 switch, OSPF enabled on a VRRP-based IRB interface might stay in ExStart state if the routing instance has both VRRP-based IRB and **mcae-mac-synchronize**-based IRB. [PR1139558](#)
- On QFX5110 switches, DHCP relay with multiple routing instances does not send the physical interface as the circuit-id value. [PR1239666](#)
- On QFX5110 switches, a custom string that is configured with **user-defined-option-82** does not work with DHCP DISCOVER packets that come from multiple routing instances. [PR1241282](#)

Interfaces and Chassis

- On a QFX5110 switch configured with a LAG created with two interfaces—such as et-0/0/1, speed 40 Gbps and et-0/0/2, speed 40 Gbps—and if one of the interfaces is dynamically hot-swapped with a 100-Gbps interface, then no error message is displayed. However, when you create a mixed-speed configuration from the start with the CLI, the behavior is as expected and appropriate error messages are displayed. As a workaround, if you dynamically hot-swap a LAG member interface for one with a different speed, delete the LAG interface and re-configure it. [PR1204545](#)
- In certain cases the link state might stay as UP between a QFX5200 switch and different NICs (third-party devices) for different combinations of FEC on and off. If the link is taken down from the third-party device end when there is an FEC mismatch, as a workaround, use the CLI command to enable or disable FEC on the QFX5200 to match the FEC mode on the third-party device. [PR1246654](#)

MPLS

- On QFX5200 switches, if you apply either an 802.1p rewrite rule or a DSCP rewrite rule on a network interface that has a Layer 2 circuit that is configured but not yet up, the rewrite rule does not work. If you apply the rewrite rule after the Layer 2 circuit is up, the rewrite rules are applied and work correctly. [PR1105354](#)

Multicast Protocols

- The **set protocols l2-learning disable-vxlan-multicast-transit vxlan-multicast-group multicast-group** command should not be configured on a PIM RP router for QFX5110 or QFX5200 because it will impact Layer 3 multicast protocol behavior. [PR1268167](#)

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)
- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)

- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)
- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

Known Issues for QFX5110 and QFX5200 Switches

This section lists the known issues in hardware and software in Junos OS Release 15.1X53-D232 for QFX5110 and QFX5200 switches.

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

- [Hardware on page 26](#)
- [Class of Service \(CoS\) on page 26](#)
- [Infrastructure on page 26](#)
- [Interfaces and Chassis on page 26](#)
- [Layer 2 Features on page 27](#)
- [Layer 3 Features on page 27](#)
- [Multicast Protocols on page 27](#)
- [Network Management and Monitoring on page 27](#)
- [Routing Policy and Firewall Filters on page 27](#)
- [Software Installation and Upgrade on page 27](#)
- [Spanning Tree Protocols on page 27](#)

Hardware

- On QFX5200 and PTX5000 systems connected using QSFP+4x10G-IR (PSM4 optical transceivers), the interfaces do not link because of a timing issue. When a port is channelized, the link goes down and the optical speed is set before the interface comes up. [PR1307400](#)

Class of Service (CoS)

- On QFX5110 switches, the classifier binding fails if the family configuration and the classifier configuration (manual dscp classifier) are committed in one single commit. To avoid this problem, apply the family configuration and the classifier configuration on an interface in two separate commits. As a workaround, deactivate and then activate the class-of-service classifier configuration—you do not need to deactivate the whole class-of-service configuration; just deactivate the classifier-specific configuration. [PR1241374](#)

Infrastructure

- On QFX5110 switches, on a system reboot, with BFD configured, some of the IPv4 BGP sessions might flap. [PR1180305](#)
- When a QFX5110 switch is acting as a DHCP relay agent, DHCP Reply packets might not be relayed back to the client by the relay agent if the Reply packets enter on a GRE tunneling interface. [PR1198982](#)

Interfaces and Chassis

- On a QFX5110-48S switch, a Gigabit Ethernet interface goes down and comes back up once on a peer as part of a reboot. [PR1237572](#)
- On a QFX5110-48S switch, after a reboot, a 1-Gigabit Ethernet interface on a peer might stay up for approximately four to five seconds even after you issue the **set interfaces ge-interface name disable** command. This issue is not seen on XE or ET interfaces. [PR1237814](#)
- On QFX5200 switches, a link down issue might occur after a port channelization instance in which you try to change the port speed configuration without waiting for ports to come up (100G->25G->100G->25G). The link down issue does not occur if you wait for ports to come up after channelizing or de-channelizing the ports and then do the speed change again. To recover from this link down issue, dechannelize the port to bring it up. [PR1256210](#)

Layer 2 Features

- On QFX5200 switches, if RTG and VSTP are configured on the same VLAN, MAC addresses might or might not be learned on RTG interfaces in that VLAN. If RTG is configured on some interfaces in a VLAN, do not configure VSTP on that same VLAN. [PR1099995](#)

Layer 3 Features

- On a QFX5110 switch with a filter-based forwarding (FBF) configuration, Layer 3 forwarding might stop working if the ARP entry for the destination host address is deleted on the switch. FBF is used to derive the routing instance for this host address. As a workaround, configure static ARP entries for host addresses where FBF is needed. [PR1200707](#)

Multicast Protocols

- On QFX5110 switches, when an interface that functions as a downstream source (S,G) multicast interface is configured with the **targeted-broadcast** statement, multicast routing and forwarding don't work on that or any other downstream interfaces. [PR1237404](#)

Network Management and Monitoring

- On a QFX5110 switch in a scaled configuration, an updated sFlow sample might not be updated in the packet capture at the collector. [PR1233498](#)

Routing Policy and Firewall Filters

- On QFX5110 switches, the dcpfe might create a core file when the applied lo0 firewall filter term is changed in scaled conditions. [PR1241733](#)

Software Installation and Upgrade

- On QFX5110 and QFX5200 switches, ZTP commits the downloaded configuration before installing the downloaded image (obtained thru DHCP). [PR1241412](#)

Spanning Tree Protocols

- On QFX5110 switches, xSTP cannot be configured on a Q-in-Q trunk physical interface. [PR1181789](#)

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)
- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)
- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)

- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

Resolved Issues for QFX5110 and QFX5200 Switches

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

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

- [Resolved Issues: Release 15.1X53-D232 on page 28](#)
- [Resolved Issues: Release 15.1X53-D231 on page 29](#)
- [Resolved Issues: Release 15.1X53-D230 on page 31](#)
- [Resolved Issues: Release 15.1X53-D210 on page 32](#)

Resolved Issues: Release 15.1X53-D232

- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)
- [Port Security](#)
- [Software-Defined Networking](#)

Infrastructure

- On QFX5110 and QFX5200 switches, there is a memory leak in `sysctl net.routetable/sysctl_rtsock()`. [PR1163782](#)
- On QFX5110 and QFX5200 switches, 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](#)

Interfaces and Chassis

- On a QFX5200, the `fxpc` process might crash when an SFP is inserted in a port configured with **flexible-vlan-tagging** or **extended-vlan-bridge**. [PR1159156](#)
- On a QFX5200, a 100-gigabit interface might not come up if connecting to another vendor's switch or an MX Series router after an upgrade to 15.1X53-D210/15.1X53-D230. [PR1292726](#)

Network Management and Monitoring

- After the rebooting of the Virtual Chassis, authentication of SNMPv3 users fails due to the change of the local engine ID. [PR1256166](#)

Port Security

- On a QFX5110 and QFX5200, DHCP Discover/Offer packets might cause memory leaks and create jdhcpd core files. [PR1273452](#)

Software-Defined Networking

- On QFX5110 and QFX5200 switches with EVPN-VXLAN, broadcast and multicast traffic might not be sent to other switches via VTEP interfaces. [PR1293163](#)

Resolved Issues: Release 15.1X53-D231

- [Authentication and Access Control](#)
- [Hardware](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Layer 3 Features](#)
- [Network Management and Monitoring](#)
- [Routing Policy and Firewall Filters](#)
- [Software-Defined Networking \(SDN\)](#) on page 30
- [Security](#)
- [Software Installation and Upgrade](#)
- [Spanning-Tree Protocols](#)

Authentication and Access Control

- On QFX5110 and QFX5200 switches, the auditd daemon might crash after the configuration of tacplus-server is changed. [PR1191527](#)
- On QFX Series switches, SSH key-based authentication is failing. [PR1142992](#)

Hardware

- Fan LEDs on a QFX5200 may work in unexpected ways. [PR1274312](#)

Interfaces and Chassis

- On a QFX5110 and QFX5200, removing force-up causes return-traffic to be dropped by the leaf. [PR1264650](#)

Layer 2 Features

- QFX5110 generated an l2ald core dump for an unknown reason at:
l2ald_mac_process_update_fwd_entry_mask, l2ald_mclag_update_change_for_learn_mask
, logging, vlogging, vlogging_event. [PR1264432](#)

Layer 3 Features

- On QFX Series, EBGp packets with ttl=1 and non-EBGP packets with ttl=1 go to the same queue. [PR1227314](#)

Network Management and Monitoring

- On QFX5200, the error log **ifd ifd-number; does not exist** might appear during an SNMP query and the SNMP query might be delayed. [PR1263794](#)

Routing Policy and Firewall Filters

- A firewall filter using **deny-bgp from port bgp** as a deny term blocks all TCP traffic on a QFX5110. [PR1264373](#)

Software-Defined Networking (SDN)

- Data-plane VXLAN and OVSDb functionality is not supported in Junos OS Release 15.1X53-D230 on QFX5200 and QFX5110 platforms. [PR1267489](#)

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](#)
- 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](#)

Software Installation and Upgrade

- On QFX5110 switches, the **request system software rollback** command is not available. [PR1279767](#)

Spanning-Tree Protocols

- VSTP functionality is not working on QFX5110 switches in a hub scenario. [PR1241456](#)
- QFX5110 and QFX5200 switches do not transfer BPDU packets though xSTP is disabled. [PR1262847](#)

Resolved Issues: Release 15.1X53-D230

- [Authentication and Access Control](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Network Management and Monitoring](#)

Authentication and Access Control

- On QFX5200-32C switches running Junos OS Release 15.1X53-D210, LLDP is not functional when it is configured on the management interface (em0). [PR1227632](#)

Infrastructure

- Upon restarting rpd on a QFX5200-32C switch, you might observe a multicast traffic drop of about 30 seconds. [PR1224639](#)
- For Junos OS Release 14.1R1 and later releases, when a broadcast packet is sent in a scenario of integrated routing and bridging (IRB) over Virtual Tunnel End Point (VTEP) over IRB, the packet is dropped in the kernel as it is looping due to a software issue. The error log message **if_pfe_vtep_ttp_output: if_pfe_ttp_output failed with error 50** is observed when the issue occurs. [PR1145358](#)

Interfaces and Chassis

- On QFX5200-32C switches, when you insert a JNP-QSFP-100G-SR4 optical transceiver into a 100 Gbps port, then channelize the 100 Gbps port and then delete the configuration, the port might go down. [PR1159546](#)
- On QFX5200 switches, after multiple link flaps have occurred, randomly some aggregated Ethernet (ae) member links might remain in the detached state. [PR1243421](#)
- On a QFX5200 switch with FEC support on a 100-gigabit port, if you channelize the 100-gigabit port into 25-gigabit or 50-gigabit ports, the FEC statistics are carried over to the first channelized port. [PR1256221](#)

Layer 2 Features

- On QFX5110 switches, when the same VLAN tag ID is configured on the NNI and UNI interfaces belonging to the same bridge domain, the traffic on the NNI exits with a single tag instead of dual tags. As a workaround, use different VLAN tag IDs on the NNI and UNI interfaces. [PR1192760](#)

Network Management and Monitoring

- In a sampling feature, certain scenarios force handling of the sampled packet at the interrupt context, which might corrupt the BMEB packet context and lead to BMEB FDB corruption. [PR1156464](#)

Resolved Issues: Release 15.1X53-D210

- [Firewall Filters](#)
- [Interfaces and Chassis](#)
- [MPLS](#)
- [Platforms and Chassis](#)

Firewall Filters

- On a QFX5200 switch, if a firewall filter applied on a loopback interface is also applied to a management interface (em0), all traffic on the management interface is dropped by default. You must explicitly configure an accept term to allow traffic to the management interface. [PR1225137](#)

Interfaces and Chassis

- On a QFX5200 switch, the **show chassis led** command displays incorrect status for the Link/Activity LED. For example, when an interface is administratively disabled, **show chassis led** shows the LED status as green even though the Link/Activity LED indicates that the port is disabled. [PR1081459](#)

MPLS

- QFX5200 switches do not support having the same interface as part of both an MPLS configuration and a routing-instance configuration. When the same interface is configured for MPLS and for a routing instance, a commit does not work and an error occurs. [PR1097427](#)

Platforms and Chassis

- On QFX5200 switches, periodic polling of fans occurs in intervals of less than a second. For some frequencies of polling, the presence of the fan module is not detected, and an alarm is logged. This alarm is corrected and cleared immediately in the next poll cycle. This behavior does not affect the working of the fans. [PR1217426](#)

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)
- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)
- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)
- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

Documentation Updates for QFX5110 and QFX5200 Switches

There are no errata or changes in Junos OS Releases 15.1X53 for QFX5110 and QFX5200 switch documentation.

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)

- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)
- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)
- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

[Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches](#)

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

- [Downloading Software Files with a Browser on page 34](#)
- [Backing Up the Current Configuration Files on page 35](#)
- [Installing the Software on page 36](#)

Downloading Software Files with a Browser

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



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

This procedure shows you how to upgrade software on QFX5200 and QFX5110 switches. The upgrade process is the same for both switches.

1. Using a Web browser, navigate to <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the By Technology box, click **Switching | QFX Series | QFX5200**.
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-5e<version>-domestic-signed.tgz** file on your computer.
9. Open or save the installation package either to the local system in the **var/tmp** directory or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.

Backing Up the Current Configuration Files

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

To back up your current configuration files:

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

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

Installing the Software



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

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

For example:

```
user@switch> request system software add
/var/tmp/jinstall-qfx-5e-15.1X53-D220.n-domestic.tgz reboot
```

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

For example:

```
user@switch> request system software add
ftp://ftpsrvr/directory/jinstall-qfx-5e-15.1X53-D220.n-domestic.tgz reboot
```

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

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

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)
- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)
- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)
- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)
- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Product Compatibility for QFX5110 and QFX5200 Switches on page 36](#)

Product Compatibility for QFX5110 and QFX5200 Switches

- [Hardware Compatibility on page 36](#)

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 QFX5110 or QFX5200 switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

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

Related Documentation

- [New and Changed Features for QFX5110 and QFX5200 Switches on page 4](#)
- [Changes in Behavior and Syntax in QFX5110 and QFX5200 Switches on page 23](#)
- [Known Behavior for QFX5110 and QFX5200 Switches on page 23](#)
- [Known Issues for QFX5110 and QFX5200 Switches on page 25](#)
- [Resolved Issues for QFX5110 and QFX5200 Switches on page 28](#)
- [Documentation Updates for QFX5110 and QFX5200 Switches on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions for QFX5110 and QFX5200 Switches on page 34](#)

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 J-Care or JNASC 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.

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service 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 <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.

- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

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

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

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

Revision History

20 October 2017—Revision 3, Junos OS for QFX5110 and QFX5200 switches, Release 15.1X53-D231—update to Known Issues

11 August 2017—Revision 2, Junos OS for QFX5110 and QFX5200 switches, Release 15.1X53-D231—update to Changes in Behavior and Syntax

12 July 2017—Revision 1, Junos OS for QFX5110 and QFX5200 switches, Release 15.1X53-D231

4 April 2017—Revision 1, Junos OS for QFX5110 and QFX5200 switches, Release 15.1X53-D230

13 January 2017—Revision 2, Junos OS for QFX5110 and QFX5200 switches, Release 15.1X53-D210—updates to New Features

6 January 2017—Revision 1, Junos OS for QFX5110 and QFX5200 switches, Release 15.1X53-D210

8 November 2016—Revision 4, Junos OS for QFX5200 switches, Release 15.1X53-D30—update to New Features

5 August 2016—Revision 3, Junos OS for QFX5200 switches, Release 15.1X53-D30

21 December 2015—Revision 2, Junos OS for QFX5200 switches, Release 15.1X53-D30—Added item to Known Issues.

11 December 2015—Revision 1, Junos OS for QFX5200 switches, Release 15.1X53-D30

Copyright © 2017 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.