

# Release Notes: Junos<sup>®</sup> OS Release 13.2X51 for the QFX Series

Release 13.2X51-D10  
17 December 2013  
Revision 4

## Contents

|  |    |
|--|----|
| Introduction .....                         | 3  |
| New and Changed Features .....             | 3  |
| Hardware .....                             | 3  |
| Interfaces and Chassis .....               | 6  |
| IPv6 Features .....                        | 6  |
| Junos OS Software .....                    | 8  |
| Junos OS XML API and Scripting .....       | 8  |
| Layer 2 Features .....                     | 8  |
| Layer 3 Features .....                     | 9  |
| Multiprotocol Label Switching (MPLS) ..... | 10 |
| Multicast Protocols .....                  | 11 |
| Network Management and Monitoring .....    | 12 |
| Security .....                             | 13 |
| Storage .....                              | 14 |
| System Management .....                    | 14 |
| Traffic Management .....                   | 14 |
| Changes in Behavior and Syntax .....       | 15 |
| Network Management .....                   | 15 |
| Known Behavior .....                       | 15 |
| Network Management and Monitoring .....    | 15 |
| Traffic Management .....                   | 15 |
| Known Issues .....                         | 16 |
| Class of Service (CoS) .....               | 17 |
| Hardware .....                             | 17 |
| Interfaces and Chassis .....               | 17 |
| Multiprotocol Label Switching (MPLS) ..... | 18 |
| Network Management .....                   | 18 |
| Routing Protocols .....                    | 18 |
| Services .....                             | 19 |

|   |    |
|---|----|
| Documentation Updates .....                             | 19 |
| System Management .....                                 | 19 |
| Migration, Upgrade, and Downgrade Instructions .....    | 19 |
| Upgrading Software on QFX5100 Standalone Switches ..... | 20 |
| Product Compatibility .....                             | 21 |
| Hardware Compatibility .....                            | 21 |
| Third-Party Components .....                            | 22 |
| Finding More Information .....                          | 22 |
| Documentation Feedback .....                            | 22 |
| Requesting Technical Support .....                      | 22 |
| Self-Help Online Tools and Resources .....              | 23 |
| Opening a Case with JTAC .....                          | 23 |
| Revision History .....                                  | 24 |

## Introduction

---

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, J Series, M Series, MX Series, PTX Series, QFabric, QFX Series, SRX Series, and T Series.

These release notes accompany Junos OS Release 13.2X51 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

## New and Changed Features

---

This section describes the new features and enhancements to existing features in Junos OS Release 13.2X51-D10 for the QFX Series.

- [Hardware on page 3](#)
- [Interfaces and Chassis on page 6](#)
- [IPv6 Features on page 6](#)
- [Junos OS Software on page 8](#)
- [Junos OS XML API and Scripting on page 8](#)
- [Layer 2 Features on page 8](#)
- [Layer 3 Features on page 9](#)
- [Multiprotocol Label Switching \(MPLS\) on page 10](#)
- [Multicast Protocols on page 11](#)
- [Network Management and Monitoring on page 12](#)
- [Security on page 13](#)
- [Storage on page 14](#)
- [System Management on page 14](#)
- [Traffic Management on page 14](#)

## Hardware

- **QFX5100-48S Switch**—The Juniper Networks QFX5100 line of switches is the next generation of top-of-rack standalone switches. The model QFX5100-48S provides 48 SFP+ access ports and 6 QSFP+ uplink ports that provide a total switching capacity and throughput of 720 Gbps. Each 40-Gigabit QSFP+ port can be configured as either a native 40-Gigabit port or four 10-Gigabit ports using a break-out cable. When using the break-out cable, the switch supports a maximum of 72 10-Gigabit Ethernet ports. The QFX5100-48S has a 1U form factor and comes standard with redundant fans and redundant power supplies. The model can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies. [See [QFX5100 Switch Hardware Overview](#).]
- **QFX5100-24Q Switch**—The Juniper Networks QFX5100 line of switches is the next generation of top-of-rack standalone switches. The model QFX5100-24Q provides 24 QSFP+ access ports and two expansion slots for the QFX-EM-4Q expansion module.

Each 40-Gigabit QSFP+ port can be configured as either a native 40-Gigabit port or four 10-Gigabit ports using a break-out cable. When using the break-out cable, the 24 ports support a 96 10-Gigabit Ethernet ports. Total throughput for the 24 port configuration is 960 Gbps; each additional QFX-EM-4Q expansion module provides an additional 160 Gbps throughput for a total of 1280 Gbps for a fully-populated switch. The QFX5100-24Q has a 1U form factor and comes standard with redundant fans and redundant power supplies. The model can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies. [See [QFX5100 Switch Hardware Overview](#).]

- **QFX-EM-4Q expansion module (QFX5100 switch)**—Provides four QSFP+ ports that can be hot-inserted into a QFX5100-24Q switch. The 40-Gigabit Ethernet ports in the expansion module provide 160 Gbps throughput to the switch. The expansion module supports both back-to-front and front-to-back airflow. [See [QFX5100 Switch Hardware Overview](#).]
- **Support for 10-Gigabit optical transceivers (QFX5100 switch)**—Provides support for:
  - QFX-SFP-10GE-ER-SFP+—SFP+ 10GBase-ER, 10-Gigabit Ethernet optics; 1550 nm for up to 40 km transmission on single mode fiber-optic (SMF) cables.
  - QFX-SFP-10GE-LR-SFP+—SFP+ 10GBase-LR, 10-Gigabit Ethernet optics; 1310 nm for up to 10 km transmission on multimode fiber-optic (MMF) cables.
  - QFX-SFP-10GE-SR-SFP+—SFP+ 10GBase-SR, 10-Gigabit Ethernet optics; 850 nm for up to 300 m transmission on multimode fiber-optic (MMF) cables.
  - QFX-SFP-10GE-USR—SFP+ 10-Gigabit Ethernet Ultra Short Reach (USR) optics; 850 nm for up to 10 m on OM1 grade fiber, 20 m on OM2 grade fiber, and 100 m on OM3 grade fiber.
  - EX-SFP-10GE-ZR—SFP+ 10GBase-ZR, 10-Gigabit Ethernet optics; 1550 nm for up to 80 km transmission.
  - QFX-SFP-DAC-1M—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), passive, twinax copper cable, 1 m.
  - QFX-SFP-DAC-1MA—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), active, twinax copper cable, 1 m.
  - QFX-SFP-DAC-3M—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), passive, twinax copper cable, 3 m.
  - QFX-SFP-DAC-3MA—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), active, twinax copper cable, 3 m.
  - JNP-QSFP-DAC-5M—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), passive, twinax copper cable, 5 m.
  - JNP-QSFP-DAC-5MA—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), active, twinax copper cable, 5 m.

- JNP-QSFP-DAC-7MA—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), active, twinax copper cable, 7 m.
- QFX-SFP-DAC-10MA—SFP+ 10-Gigabit Ethernet direct attach cable (DAC), active, twinax copper cable, 10 m.

[See [Interface Support for the QFX5100 Switch](#).]

- **Support for 40-Gigabit optical transceivers (QFX5100 switch)**—Provides support for:
  - QFX-QSP-40G-SR4—QSFP+ 40GBase-SR4, 40-Gigabit Ethernet optics; 850 nm for up to 150 m transmission on multimode mode fiber (MMF).
  - QFX-QSP-40G-eSR4—QSFP+ 40GBase-eSR4, 40-Gigabit Ethernet optics; 850 nm for up to 300 m transmission on multimode mode fiber (MMF).
  - QSFP-40GBASE-LR4—QSFP+ 40GBase-LR4, 40-Gigabit Ethernet pluggable optics; 1310 nm for up to 10 km on single mode fiber (SMF).
  - QFX-QSFP-DAC-1M—QSFP+ Direct Attach Cable (DAC) assembly, 1 m, 30 AWG, passive, programmable ID.
  - QFX-QSFP-DAC-3M—QSFP+ Direct Attach Cable (DAC) assembly, 3 m, 30 AWG, passive, programmable ID.
  - JNP-QSFP-DAC-1M—QSFP+ Direct Attach Cable (DAC) assembly, 5 m, 30 AWG, passive, programmable ID.
  - QFX-QSP-DACBO-1M—QSFP+ to SFP+ 10-Gigabit Ethernet, direct attach cable break-out cable (DACBO) assembly, 1 m, 30 AWG twinax copper cable, 1 m.
  - QFX-QSP-DACBO-3M—QSFP+ to SFP+ 10-Gigabit Ethernet, direct attach cable break-out cable (DACBO) assembly, 1 m, 30 AWG twinax copper cable, 3 m.

[See [Interface Support for the QFX5100 Switch](#).]

- **Support for 1-Gigabit optical transceivers (QFX5100 switch)**—Provides support for:
  - QFX-SFP-1GE-LX—SFP 1000Base-LX Gigabit Ethernet optics; 1310 nm for up to 10 Km transmission on single mode fiber (SMF) cable.
  - QFX-SFP-1GE-SX—SFP 1000Base-SX Gigabit Ethernet optics; 850 nm for up to 550 m transmission on multimode fiber (MMF) cable.
  - QFX-SFP-1GE-T—SFP-T Copper transceiver module; for up to 100 m transmission on Category 5 (Cat5) cable.

[See [Interface Support for the QFX5100 Switch](#).]

- **Online insertion and removal (OIR) (QFX5100 switch)**—Enables you to perform online insertion and removal (“hot swapping”) of fan trays, power modules, and SFPs. [See [QFX5100 Switch Hardware Overview](#).]
- **Support for third-party transceivers (QFX5100 switch)**—Allows you to use fiber-optic transceivers made by other vendors with the QFX5100 switch. This feature is enabled by default, and no configuration is required to use this functionality. [See [QFX5100 Switch Hardware Overview](#).]

## Interfaces and Chassis

- **Channelizing 40-Gigabit Ethernet QSFP+ interfaces (QFX5100 switch)**—Enables you to configure four 10-Gigabit Ethernet interfaces from a 40-Gigabit Ethernet QSFP+ interface. By default, a 40-Gigabit Ethernet QSFP+ interface is 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.



**NOTE:** When you configure 10-Gigabit Ethernet channelization for a 40-Gigabit Ethernet QSFP+ interface, the Packet Forwarding Engine restarts.



**NOTE:** There are restrictions on the ports you can channelize on the QFX5100-24Q switch, depending on the port mode you configure. If you try to channelize ports that are restricted, the configuration is ignored.

[See [Channelizing Interfaces](#).]

- **Layer 3 unicast and multicast support for multichassis link aggregation (QFX5100 switch)**—The following Layer 3 unicast and multicast features are supported:
  - VRRP active-standby support—Enables Layer 3 routing over MC-AE interfaces.
  - Address Resolution Protocol (ARP) synchronization—Enables ARP synchronization to occur independently on the MC-LAG peers.
  - Routed VLAN interface (RVI) MAC address synchronization—Enables an MC-LAG peer to forward Layer 3 packets arriving on MC-AE interfaces with either its own RVI MAC address or its peer's RVI MAC address.
  - Multicast dual designated router (DR)—Enables faster convergence for MC-LAG peer failovers. To configure a dual designated router (DR), issue the **set protocols pim interface interface-name dual-dr** command for VLAN interfaces on each MC-LAG peer.

[See [Understanding Multichassis Link Aggregation](#).]

## IPv6 Features

- **Support for IPv6 on network interfaces (QFX5100 switch)**—Enables IPv6 support for many protocols and features on network interfaces on QFX5100 standalone switches. IPv6 is supported on network interfaces with the following protocols and features:

- Neighbor discovery
- Stateless autoconfiguration
- Router advertisements
- Link aggregation
- Routed VLAN interfaces
- LLDP
- SSH
- Telnet
- Ping
- Traceroute
- Virtual routers
- Static routing
- OSPFv3
- IS-ISv6
- BFD
- BGP
- Graceful restart for routing protocols
- MLD
- CoS
- VRRPv3
- RADIUS
- TACACS+
- AAA
- SNMP
- NTP
- Syslog

- IPv6 statistics and counters
- Path MTU discovery

## Junos OS Software

- **Junos OS Enhanced Layer 2 Software (ELS) configuration statements and operational commands (QFX5100 switch)**—Provides a uniform CLI for configuring and monitoring Layer 2 features on newer model EX Series, MX Series, and QFX Series devices that support ELS, such as the EX4300 and EX9200 switches and MX Series routers in LAN mode (MX-ELM).

The Web-based *ELS Translator tool* is available for registered customers to help them become familiar with the ELS CLI and to quickly translate existing QFX Series switch-based CLI configurations into ELS CLI configurations.

[See [Getting Started with Enhanced Layer 2 Software](#) and [ELS Translator Tool](#).]

## Junos OS XML API and Scripting

- **Packaging Python Scripts**—Python is available on devices running Junos OS when the following packages are installed:
  - jinstall-ex-4300
  - jinstall-qfx
  - jinstall-qfx-5
  - jinstall-dc-re

The Python interpreter only runs scripts that have been installed from signed packages created with the Junos SDK. For more information, see “Using Python on Junos” and “Building a Junos SDK Package Containing Python Scripts” in the *Junos SDK Developer Guide*.

## Layer 2 Features

- **VLAN support (QFX5100 switch)**—Enables you to divide one physical broadcast domain into multiple virtual domains. [See [Understanding Bridging and VLANs](#).]
- **Link aggregation (LAG) (QFX5100 switch)**—Enables you to use multiple network cables and ports in parallel to increase link speed and redundancy. [See [Understanding Aggregated Ethernet Interfaces and LACP](#).]
- **Link Layer Discovery Protocol (LLDP) (QFX5100 switch)**—Enables a switch to advertise its identity and capabilities on a LAN, as well as receive information about other network devices. [See [Understanding LLDP](#).]
- **Spanning Tree Protocol (STP) (QFX5100 switch)**—Provides you with loop and bridge protocol data unit (BPDU) protection. [See [Overview of Spanning-Tree Protocols](#).]
- **Unified forwarding table (QFX5100 switch)**—Enables you to control the allocation of forwarding table memory available to store the following:
  - MAC addresses



- Layer 3 host entries
- Longest prefix match table entries

You must also use this feature if you want to use more than 16 IPv6 addresses with prefixes in the range /65 through /127. [See [Understanding the Unified Forwarding Table.](#)]

- **Multichassis link aggregation (QFX5100 switch)**—Multichassis link aggregation (MC-LAG) enables a client device to form a logical LAG interface using two QFX5100 devices. MC-LAG provides redundancy and load balancing between the two QFX5100 devices, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG is an MC-LAG client device that has one or more physical links in a LAG. This client device does not need to be aware of the MC-LAG. On the other side of the MC-LAG are two MC-LAG QFX5100 devices. Each of these QFX5100 devices has one or more physical links connected to a single client device. The QFX5100 devices coordinate with each other to ensure that data traffic is forwarded properly.

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

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

[See [Understanding Multichassis Link Aggregation.](#)]

## Layer 3 Features

- **Border Gateway Protocol (QFX5100 switch)**—Border Gateway Protocol (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. [See [BGP.](#)]
- **Open Shortest Path First (QFX5100 switch)**—The IPv4 Open Shortest Path First (OSPF) protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). QFX5100 switches support OSPFv1 and OSPFv2. You can configure OSPF at the **[edit protocols *ospf*]** hierarchy level. [See [OSPF.](#)]
- **Bidirectional Forwarding Detection for static routes and the BGP, IS-IS, OSPF, PIM, and RIP protocols (QFX5100 switch)**—The Bidirectional Forwarding Detection (BFD) protocol 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 QFX5100 switch, you can configure BFD for static routes and the BGP, IS-IS, OSPF, PIM, and RIP protocols. [See [Configuring BFD.](#)]

- **IS-IS (QFX5100 switch)**—The Intermediate System-to-Intermediate System (IS-IS) protocol is an interior gateway protocol (IGP) for routing traffic within an autonomous system (AS). [See [IS-IS Overview.](#)]
- **Virtual Router Redundancy Protocol (QFX5100 switch)**—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. [See [High Availability](#).]

- **Layer 3 logical interfaces (QFX5100 switch)**—A Layer 3 logical interface is a logical division of a physical interface or an aggregated Ethernet interface, which operates at the network level and 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 QFX5100 switch to a Layer 2 switch. Only one physical connection is required between the switches. This topology is often called a “router on a stick” or a “one-armed router” when the Layer 3 device is a router.

To create Layer 3 logical interfaces on a QFX5100 switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface. You can partition one physical interface in up to 4089 different Layer 3 logical interfaces, one for each VLAN. We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. VLAN tagging places the VLAN ID in the frame header, enabling each physical interface to handle multiple VLANs. Junos OS supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces. Double-tagging, which is assigning more than one VLAN ID in the frame header, is not supported. [See [Understanding Layer 3 Logical Interfaces](#).]

- **Generic routing encapsulation (QFX5100 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. [See [Understanding Generic Routing Encapsulation](#) .]

## Multiprotocol Label Switching (MPLS)

- **Multiprotocol Label Switching (QFX5100 switch)**—Enables you to configure Multiprotocol Label Switching (MPLS) to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as voice over IP (VoIP) and other business-critical functions. MPLS has the following advantages over conventional packet forwarding:
  - Packets arriving on different ports can be assigned different labels.
  - A packet arriving at a particular provider edge (PE) switch can be assigned a label that is different from that of the same packet entering the network at a different PE switch. As a result, forwarding decisions that depend on the ingress PE switch can be easily made.

- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than letting it follow the route chosen by the normal dynamic routing algorithm as the packet travels through the network. With MPLS, a label can be used to represent the route so that the packet need not carry the identity of the explicit route.

[See [MPLS on the QFX Series](#).]

## Multicast Protocols

- **Protocol Independent Multicast sparse mode (QFX5100 switch)**—Enables efficient routing to multicast groups with receivers sparsely spread over multiple networks. To configure PIM sparse mode, include the **pim** statement at the **[edit protocols]** hierarchy level. [See [Multicast Protocols](#).]
- **PIM-SSM (QFX5100 switch)**—Protocol Independent Multicast source-specific multicast (PIM-SSM) uses a subset of PIM sparse mode and IGMP version 3 (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. [See [Multicast Protocols](#).]
- **MSDP (QFX5100 switch)**—Enables you to connect multiple IP version 4 Protocol Independent Multicast sparse mode (PIM-SM) domains to one another. Multicast Source Discovery Protocol (MSDP) typically runs on the same routing device as a PIM-SM rendezvous point (RP). 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. [See [Understanding MSDP](#).]
- **Anycast RP (QFX5100 switch)**—Supports multiple rendezvous points (RPs) using anycast addresses (RPs sharing a single routable IP address) in either a Protocol Independent Multicast (PIM) or Multicast Source Discovery Protocol (MSDP)-enabled network. To configure anycast RP, include the **anycast-pim** statement at the **[edit protocols pim rp local family inet]** hierarchy level. [See [Understanding RP Mapping with Anycast RP](#).]
- **IGMP querier (QFX5100 switch)**—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. [See [IGMP Snooping Overview](#).]

## Network Management and Monitoring

- **Simple Network Management Protocol (SNMP) (QFX5100 switch)**—Provides SNMP versions 1, 2, and 3 for monitoring system activity. [See [Understanding the Implementation of SNMP](#).]
- **System logging (syslog) (QFX5100 switch)**—Enables you to log system messages into a local directory on the switch or to a syslog server. [See [Overview of Junos OS System Log Messages](#).]
- **sFlow technology (QFX5100 switch)**—sFlow technology is a 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, which can provide 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**. [See [Overview of sFlow Technology](#).]
- **Port mirroring (QFX5100 switch)**—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. [See [Understanding Port Mirroring](#).]
- **Uplink failure detection (QFX5100 switch)**—Enables a switch to detect link failures on uplink interfaces and to propagate this information to downlink interfaces, so that servers connected to those downlink interfaces can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all of the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects a disabled downlink interface, it switches over to the secondary link to ensure traffic is not dropped but is sent over the secondary path instead.

When configuring uplink failure detection, you add an uplink interface and a corresponding downlink interface as a pair into a failure detection group. To add uplink interfaces to the group, include the **link-to-monitor *interface-name*** statement at the **[edit protocols uplink-failure-detection group *group-name*]** hierarchy level. To add downlink interfaces to the group, include the **link-to-disable *interface-name*** statement at the **[edit protocols uplink-failure-detection group *group-name*]** hierarchy level.

[See [Overview of Uplink Failure Detection](#).]

- **Network analytics feature (QFX5100 switch)**—Provides visibility into the performance and behavior of the data center infrastructure by enabling microburst and traffic statistics monitoring. You use network analytics to monitor queue and traffic statistics, which include queue depth, latency, and traffic information, at user configurable intervals. The network analytics reports can help identify problem areas in your network

traffic and applications so that you can adjust resources as needed. You configure network analytics at the [\[edit services analytics\]](#) hierarchy level. [See [Network Analytics Overview](#).]

## Security

- **Firewall filters (QFX5100 switch)**—Enable you 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](#).]
- **Policing (QFX5100 switch)**—Enables you to use policing to apply limits to traffic flow and set consequences for packets that exceed those limits. [See [Overview of Policers](#).]
- **MAC limiting (QFX5100 switch)**—Enables you to 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. [See [Understanding MAC Limiting and MAC Move Limiting](#).]
- **MAC move limiting (QFX5100 switch)**—Detects MAC movement and MAC spoofing on access ports. [See [Understanding MAC Limiting and MAC Move Limiting](#).]
- **Storm control (QFX5100 switch)**—Enables 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. [See [Understanding Storm Control](#).]
- **DHCP snooping (QFX5100 switch)**—Filters and blocks ingress DHCP server messages on untrusted ports and builds and maintains an IP address and MAC address binding database (called the DHCP snooping database). [See [Understanding DHCP Snooping](#).]
- **Dynamic ARP inspection (DAI) (QFX5100 switch)**—Prevents ARP spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. [See [Understanding DAI](#).]
- **Sticky MAC address (QFX5100 switch)**—Persistent (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 access port. Secure access to these workstations is retained even if the switch is restarted. [See [persistent-learning](#).]

## Storage

- **FCoE transit switch (QFX5100 switch)**—Enables you to configure a QFX5100 switch as a Fibre Channel over Ethernet (FCoE) transit switch that transports FCoE frames across the Ethernet network, supports FCoE Initialization Protocol (FIP) snooping, and supports the following data center bridging (DCB) standards: enhanced transmission selection (ETS), priority-based flow control (PFC), and Data Center Bridging Exchange Capability (DCBX) protocol. [See [Understanding FCoE Transit Switch Functionality](#).]

## System Management

- **Login authentication using RADIUS and TACACS+ (QFX5100 switch)**—Enables you to use RADIUS and TACACS+ authentication to validate users who attempt to access the switch. [See [Junos OS User Authentication Methods](#).]
- **System utilization alarms (QFX5100 switch)**—Provide 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. [See [System Utilization Alarms](#).]

## Traffic Management

- **CoS rewrite rules (QFX5100 switch)**—Enable you to 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. [See [Understanding CoS Rewrite Rules](#).]
- **Queue shaping (QFX5100 switch)**—Enable you to manage excess traffic and avoid congestion on a network interface where traffic may exceed the maximum port bandwidth. [See [Understanding CoS Output Queue Schedulers](#).]
- **Priority-based flow control (QFX5100 switch)**—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. [See [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\)](#).]
- **Ethernet PAUSE autonegotiation (QFX5100 switch)**—Enable you to configure asymmetric flow control. To configure PAUSE, include both the **rx-buffers** and **tx-buffers** statements at the `[edit interfaces interface-name ether-options configured-flow-control]` hierarchy level. The **rx-buffers** statement determines whether or not the interface generates and sends PAUSE messages. The **tx-buffers** statement determines whether or not the interfaces responds to received PAUSE messages. [See [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\)](#).]

- Related Documentation**
- [Known Behavior on page 15](#)
  - [Known Issues on page 16](#)

- [Documentation Updates on page 19](#)
- [Migration, Upgrade, and Downgrade Instructions on page 19](#)
- [Product Compatibility on page 21](#)

## 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 in Junos OS Release 13.2X51-D10 for the QFX Series.

- [Network Management](#)

### Network Management

- On QFX5100 switches, egress firewall filters are not applied to sFlow sampling packets. On this platform, the software architecture is different from that on other QFX Series devices, and sFlow packets are sent by the Routing Engine (not the line card on the host) and are not transiting the switch. Egress firewall filters affect data packets that are transiting a switch but do not affect packets sent by the Routing Engine. As a result, sFlow sampling packets are always sent to the sFlow collector.

#### Related Documentation

- [New and Changed Features on page 3](#)
- [Known Behavior on page 15](#)
- [Known Issues on page 16](#)
- [Documentation Updates on page 19](#)
- [Migration, Upgrade, and Downgrade Instructions on page 19](#)
- [Product Compatibility on page 21](#)

## Known Behavior

---

This section lists the limitations in Junos OS Release 13.2X51-D10 for the QFX Series.

### Network Management and Monitoring

- If a QFX5100 switch drops traffic because of an ingress firewall filter, the switch does not generate an sFlow technology monitoring flow sample packet that contains this dropped traffic.
- On the QFX5100 switch, the J-Web interface is not supported. As a result, the **web-management** configuration statement in the **[edit system services]** hierarchy level is not available in the CLI.

### Traffic Management

- On a QFX5100 switch, CPU-generated host outbound traffic is forwarded on the network-control forwarding class, which is mapped to queue 7. If you use the default

scheduler, the network-control queue receives a guaranteed minimum bandwidth (transmit rate) of 5 percent of port bandwidth. The guaranteed minimum bandwidth is more than sufficient to ensure lossless transport of host outbound traffic.

However, if you configure a scheduler, you must ensure that the network-control forwarding class (or whatever forwarding class you configure for host outbound traffic) receives sufficient guaranteed bandwidth to prevent packet loss.

If you configure a scheduler, we recommend that you configure the network-control queue (or the queue you configure for host outbound traffic if it is not the network-control queue) as a strict-high priority queue. Strict-high priority queues receive the bandwidth required to transmit their entire queues before other queues are served.



**NOTE:** As with all strict-high priority traffic, if you configure the network-control queue (or any other queue) as a strict-high priority queue, you must also create a separate forwarding class set (priority group) that contains only strict-high priority traffic, and apply the strict-high priority forwarding class set and its traffic control profile (hierarchical scheduler) to the relevant interfaces.

- You cannot apply classifiers and rewrite rules to IRB interfaces because the members of an IRB are VLANs, not interfaces. You can apply classifiers and rewrite rules to Layer 2 logical interfaces and Layer 3 physical interfaces that are members of VLANs that belong to IRB interfaces.

#### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 15](#)
- [Known Issues on page 16](#)
- [Documentation Updates on page 19](#)
- [Migration, Upgrade, and Downgrade Instructions on page 19](#)
- [Product Compatibility on page 21](#)

---

## Known Issues

The following issues are outstanding in Junos OS Release 13.2X51-D10. The identifier following the description is the tracking number in our bug database.

For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

- [Class of Service \(CoS\)](#)
- [Hardware](#)
- [Interfaces and Chassis](#)



- [Multiprotocol Label Switching \(MPLS\)](#)
- [Network Management](#)
- [Routing Protocols](#)
- [Services](#)

## Class of Service (CoS)

- On a QFX5100 switch, issuing the **show interfaces queue *interface-name*** and **show interfaces statistics *interface-name*** commands does not display the correct traffic rates. [PR894390](#)

## Hardware

- On a QFX5100 switch, a system reboot with Routing Engine FPGA version 7.1 might not successfully boot the Junos OS software. In case of a system reboot failure, you need to power cycle the switch. To check the current FPGA version, issue the **show chassis firmware** command. [PR926940](#)

## Interfaces and Chassis

- On a QFX5100 switch, the multicast PIM dual-DR mode is not supported in Junos OS 13.2X51-D10, even though this feature is supported on other QFX Series devices. [PR897618](#)
- On a QFX5100 switch, when you enable IGMP snooping on a VLAN, IPv6 multicast traffic is not flooded within the VLAN. [PR925141](#)
- On a QFX5100 switch, if you remove the service ID from a multichassis link aggregation group (MC-LAG) configuration, and then add it back to the configuration, single-homed ARP entries might not synchronize properly with MC-LAG peers. [PR929720](#)
- On a QFX5100 switch, do not use the unified forwarding table **lpm-profile** for IPv6 traffic. This profile does not work for IPv6 traffic. [PR929753](#)
- On a QFX5100 switch, the output of the **show interfaces *gr-fpc/pic/port*** command for a GRE tunnel interface does not display the proper statistics for encapsulated traffic. [PR932561](#)
- Running **tcpdump** on a QFX5100 switch might cause system instability or cause protocols such as STP or LACP to fail. [PR932592](#)
- On a QFX5100 switch, you cannot enable the **minimum-links** statement at the **[edit groups]** CLI hierarchy level. As a workaround, enable the **minimum-links** statement at the **[edit interfaces]** CLI hierarchy level. [PR932622](#)
- If you create a virtual routing instance on a QFX5100 switch and configure a routed VLAN interface (RVI) or integrated routing and bridging (IRB) interface under the routing instance, do not configure a multichassis link aggregation group (MC-LAG) interface to participate in the RVI or IRB. This combination is not supported with virtual routing instances. [PR934379](#)
- On a QFX5100 switch, integrated routing and bridging (IRB) MAC address synchronization is not supported, but you can use the virtual router redundancy protocol

(VRRP) instead. As a workaround, configure VRRP on IRB interfaces that host multichassis LAG (MC-LAG) interfaces. [PR936512](#)

- On a QFX5100 switch, the **multichassis-lag-replicate-state** statement is not supported at the **[edit vlans]** CLI hierarchy level. As a workaround, enable the **multichassis-lag-replicate-state** statement globally. [PR937018](#)
- On a QFX5100 switch, if you apply a firewall filter to the loopback interface, the filter might not discard BFD packets headed to the CPU as expected. [PR937408](#)
- On a QFX5100 switch, you cannot issue interface range commands for channelized interfaces. As a workaround, use interface commands instead of interface range commands. [PR937788](#)
- On a QFX5100 switch, you cannot configure multichassis link aggregation when Multiple Spanning Tree Protocol (MSTP) or VLAN Spanning Tree Protocol (VSTP) is enabled. [PR939049](#)
- On a QFX5100 switch, even when traffic is flowing normally, the output of the **show interfaces et-fpc/pic/port** and **show interfaces et-fpc/pic/port:[0-3]** commands does not display accurate bits per second (bps) information for the **Input rate** field. [PR939128](#)
- On a QFX5100 switch, autonegotiation of interfaces is disabled by default for 1-Gigabit Ethernet fiber ports. For these links to be brought online (up), you must disable autonegotiation on the peer interfaces. In addition, if you issue the **show interfaces interface-name extensive** command for an SFP access port with a 1-Gigabit optical copper transceiver installed, the output incorrectly shows the media type as **fiber**, if the port parameter is not configured in the **interfaces ge-0/0/port** statement in the **[edit]** hierarchy level. As a workaround, remove and reinsert the transceiver. [PR939439](#)
- On a QFX5100 switch, configuration of the **minimum-interval milliseconds** statement for liveness detection on a multichassis link aggregation group (MC-LAG) must be 1000 milliseconds or greater. Subsecond timers are not supported in Junos OS Release 13.2X51-D10. [PR942563](#)

## Multiprotocol Label Switching (MPLS)

- On a QFX5100 switch acting as an MPLS penultimate hop popping (PHP) router, an MPLS label route with a BOS (bottom of the stack) next hop will consume one filter entry and one extra next-hop entry. The total number of filter entries used will be the number of next hops in the system with PHP and BOS operation. [PR922969](#)

## Network Management

- On a QFX5100 switch, the minimum configurable queue statistics polling interval is 10 milliseconds. However, the actual queue statistics polling interval differs from the configured polling interval by approximately 5 milliseconds. [PR911015](#)

## Routing Protocols

- On a QFX5100 switch, BFD timer values of less than 1 second are not supported. [PR942035](#)

- On a QFX5100 switch, when you delete the OSPF configuration from an interface, the OSPF and ISIS routing protocols transition down and up on all other configured interfaces. [PR933536](#)
- On a QFX5100 switch, if you issue the **show interfaces ae** command, the output for aggregated Ethernet logical interfaces might display incorrect information for input and output rates. [PR936220](#)

## Services

- If you enable STP on a QFX5100 switch on which you also configure port mirroring, port mirroring might not work. [PR931633](#)

### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 15](#)
- [Known Behavior on page 15](#)
- [Documentation Updates on page 19](#)
- [Migration, Upgrade, and Downgrade Instructions on page 19](#)
- [Product Compatibility on page 21](#)

## Documentation Updates

---

This section lists the errata and changes in Junos OS Release 13.2X51-D10 documentation.

### System Management

- The **request app-engine** and **show app-engine** commands are not documented for the QFX5100 switch in Junos OS Release 13.2X51-D10.

### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 15](#)
- [Known Behavior on page 15](#)
- [Known Issues on page 16](#)
- [Migration, Upgrade, and Downgrade Instructions on page 19](#)
- [Product Compatibility on page 21](#)

## Migration, Upgrade, and Downgrade Instructions

---

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

- [Upgrading Software on QFX5100 Standalone Switches on page 20](#)

## Upgrading Software on QFX5100 Standalone Switches

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

The download and installation process for Junos OS Release 13.2 is the same as for previous Junos OS releases.

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

1. In a browser, go to <http://www.juniper.net/support/downloads/junos.html> .  
The Junos Platforms Download Software page appears.
2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **13.2** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 13.2 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.



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

---

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

```
user@host> request system software add  
source/jinstall-qfx-5-13.2X51-D10.6-domestic-signed.tgz reboot
```

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

- **/pathname**—For a software package that is installed from a local directory on the switch.

- For software packages that are downloaded and installed from a remote location:
  - `ftp://hostname/pathname`
  - `http://hostname/pathname`
  - `scp://hostname/pathname` (available only for Canada and U.S. version)

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

Rebooting occurs only if the upgrade is successful.



**NOTE:** After you install a Junos OS Release 13.2 jinstall package, you can issue the `request system software rollback` command to return to the previously installed software.

---

**Related  
Documentation**

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 15](#)
- [Known Behavior on page 15](#)
- [Known Issues on page 16](#)
- [Documentation Updates on page 19](#)
- [Product Compatibility on page 21](#)

---

## Product Compatibility

- [Hardware Compatibility on page 21](#)

### 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 QFX Series switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:  
<http://pathfinder.juniper.net/feature-explorer/>

**Related  
Documentation**

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 15](#)
- [Known Behavior on page 15](#)
- [Known Issues on page 16](#)
- [Documentation Updates on page 19](#)

- [Migration, Upgrade, and Downgrade Instructions on page 19](#)

## Third-Party Components

---

This product includes third-party components. To obtain a complete list of third-party components, see [Copyright and Trademark Information](#).

For a list of open source attributes for this Junos OS release, see [Open Source: Source Files and Attributions](#).

## Finding More Information

---

For the latest, most complete information about known and resolved issues with Junos OS, see the Juniper Networks Problem Report Search application at:  
<http://prsearch.juniper.net>.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at:  
<http://pathfinder.juniper.net/feature-explorer/>.

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF format. Find Content Explorer at:  
<http://www.juniper.net/techpubs/content-applications/content-explorer/>.

## 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](mailto: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 postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/customers/support/downloads/710059.pdf>.

- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

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

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

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

## Opening a Case with JTAC

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

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

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

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

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

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

## Revision History

---

17 December 2013—Revision 4, Junos OS for the QFX Series, Release 13.2X51-D10

16 December 2013—Revision 3, Junos OS for the QFX Series, Release 13.2X51-D10

6 December 2013—Revision 2, Junos OS for the QFX Series, Release 13.2X51-D10

22 November 2013—Revision 1, Junos OS for the QFX Series, Release 13.2X51-D10

Copyright © 2013, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.