

# Release Notes: Junos<sup>®</sup> OS Release 14.1X53-D121 for QFabric Systems

Release 14.1X53-D121  
June 1, 2017  
Revision 2

## Contents

Junos OS Release Notes for QFabric Systems . . . . .	3
New and Changed Features . . . . .	3
New Features in Release 14.1X53-D40 . . . . .	4
New Features in Release 14.1X53-D35 . . . . .	11
New Features in Release 14.1X53-D30 . . . . .	11
New Features in Release 14.1X53-D27 . . . . .	15
New Features in Release 14.1X53-D26 . . . . .	16
New Features in Release 14.1X53-D25 . . . . .	17
New Features in Release 14.1X53-D15 . . . . .	18
New Features in Release 14.1X53-D10 . . . . .	23
Changes in Behavior and Syntax . . . . .	34
Authentication and Access Control . . . . .	35
Interfaces and Chassis . . . . .	35
Open vSwitch Database (OVSDb) . . . . .	37
SNMP . . . . .	37
Software Upgrade . . . . .	37
Virtual Chassis and Virtual Chassis Fabric . . . . .	37
Known Behavior . . . . .	38
High Availability . . . . .	38
Infrastructure . . . . .	39
Interfaces and Chassis . . . . .	39
Layer 2 Features . . . . .	41
Layer 3 Protocols . . . . .	41
Layer 3 VPNs . . . . .	41
MPLS . . . . .	41
Multicast Protocols . . . . .	42
Network Management and Monitoring . . . . .	42
OVSDb . . . . .	43
Platform and Infrastructure . . . . .	44
QFabric Systems . . . . .	44

Routing Policy and Firewall Filters . . . . .	45
Routing Protocols . . . . .	45
Security . . . . .	45
Software Installation and Upgrade . . . . .	45
System Management . . . . .	46
Storage and Fibre Channel . . . . .	46
Traffic Management . . . . .	47
Virtual Chassis and Virtual Chassis Fabric . . . . .	51
VXLAN . . . . .	52
Known Issues . . . . .	53
High Availability . . . . .	53
Interfaces and Chassis . . . . .	54
MPLS . . . . .	54
Multicast Protocols . . . . .	54
Routing Protocols . . . . .	54
QFabric Systems . . . . .	54
VXLAN . . . . .	55
Resolved Issues . . . . .	55
Resolved Issues: Release 14.1X53-D43 . . . . .	56
Resolved Issues: Release 14.1X53-D42 . . . . .	58
Resolved Issues: Release 14.1X53-D40 . . . . .	62
Resolved Issues: Release 14.1X53-D35 . . . . .	63
Resolved Issues: Release 14.1X53-D30 . . . . .	65
Resolved Issues: Release 14.1X53-D27 . . . . .	65
Resolved Issues: Release 14.1X53-D26 . . . . .	66
Resolved Issues: Release 14.1X53-D25 . . . . .	66
Resolved Issues: Release 14.1X53-D16 . . . . .	67
Resolved Issues: Resolved Before Release 14.1X53-D16 . . . . .	69
Documentation Updates . . . . .	71
Bridging and Learning . . . . .	71
Network Management and Monitoring . . . . .	71
Virtual Chassis and Virtual Chassis Fabric (VCF) . . . . .	71
Migration, Upgrade, and Downgrade Instructions . . . . .	71
Upgrading to a Controlled Version of Junos OS . . . . .	72
Upgrading Software on QFX5100 Standalone Switches . . . . .	72
Performing an In-Service Software Upgrade (ISSU) . . . . .	75
Preparing the Switch for Software Installation . . . . .	75
Upgrading the Software Using ISSU . . . . .	75
Product Compatibility . . . . .	77
Hardware Compatibility . . . . .	77
Third-Party Components . . . . .	78
Finding More Information . . . . .	78
Documentation Feedback . . . . .	78
Requesting Technical Support . . . . .	78
Self-Help Online Tools and Resources . . . . .	79
Opening a Case with JTAC . . . . .	79
Revision History . . . . .	80

## Junos OS Release Notes for QFabric Systems

---

These release notes accompany Junos OS Release 14.1X53-D121 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

The following QFX Series platforms are supported in Junos OS Release 14.1X53-D121: QFabric systems.



**NOTE:** These release notes include information on all Junos OS Release 14.1X53 releases. Therefore, information about QFX Series devices that are not supported in Junos OS Release 14.1X53-D121 but are supported in other Junos OS Release 14.1X53 releases are included in these release notes.

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 34](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

### New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 14.1X53 for the QFX Series.

- [New Features in Release 14.1X53-D40 on page 4](#)
- [New Features in Release 14.1X53-D35 on page 11](#)
- [New Features in Release 14.1X53-D30 on page 11](#)
- [New Features in Release 14.1X53-D27 on page 15](#)
- [New Features in Release 14.1X53-D26 on page 16](#)
- [New Features in Release 14.1X53-D25 on page 17](#)
- [New Features in Release 14.1X53-D15 on page 18](#)
- [New Features in Release 14.1X53-D10 on page 23](#)

## New Features in Release 14.1X53-D40

---

### *Class of Service*

- **Support for policy drop counters in CLI and SNMP (QFX Series switches)**—Starting in Junos OS Release 14.1X53-D40, the **show interfaces *interface-name* statistics detail** command displays the number of packets dropped on an interface because of policers configured for that interface. The number of packet drops is displayed in the command output in the **Bucket drops** field under **Input errors** and **Output errors**. These statistics are also available through SNMP.

See [show interfaces xe](#).

### *High Availability (HA) and Resiliency*

- **NSSU improvements to optimize total upgrade time and recover from software image copy or reboot failures (QFX5100 Virtual Chassis or Virtual Chassis Fabric [VCF])**—Starting in Junos OS Release 14.1X53-D40, nonstop software upgrade (NSSU) on a Virtual Chassis or VCF supports the following optimizations and error recovery measures:
  - To optimize the time needed to complete NSSU, the master member copies the new software in parallel to multiple members at a time rather than waiting for the copy operation to complete to each member before copying the software image to the next member. By default, the number of parallel copy sessions is based on the Virtual Chassis or VCF size, or you can configure a specific number using the **rcp-count *number*** configuration statement.
  - As before, the master aborts the NSSU process if copying the new software to any member fails. As a new error recovery measure, the master also removes the new software image from all members to which it was already transferred.
  - During NSSU, when each member is rebooted in turn with the new software, if any member fails to reboot, the master aborts the NSSU process. As a new recovery measure, the master automatically brings down and reboots the entire Virtual Chassis or VCF. This recovery action causes downtime for the Virtual Chassis or VCF, but brings it up in a stable state, cleanly running the new software on all members without requiring you to manually recover members individually.

[See [Understanding Nonstop Software Upgrade on a Virtual Chassis and Mixed Virtual Chassis](#) or [Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric](#).]

### *Interfaces and Chassis*

- **LAG local minimum links per Virtual Chassis or VCF member (QFX5100 switches)**—Introduced in Junos OS Release 14.1X53-D40, the local minimum links feature helps avoid traffic loss due to asymmetric bandwidth on link aggregation group (LAG) forwarding paths through a Virtual Chassis or Virtual Chassis Fabric (VCF) member switch when one or more LAG member links local to that chassis have failed. When this feature is enabled, if a user-configured percentage of local LAG member links has failed on a chassis, all remaining local LAG member links on the chassis are forced down, and LAG traffic is redistributed only through LAG member links on *other* chassis. To enable local minimum links for an aggregated Ethernet interface (aex), set the **local-minimum-links-threshold** configuration statement with a threshold value that represents the percentage of local member links that must be up on a chassis for any local LAG member links on that chassis to continue to be active in the aggregated Ethernet bundle. Otherwise all remaining LAG member links on that chassis are also forced down. The feature responds dynamically to bring local LAG member links up or down if you change the configured threshold, or when the status or configuration of LAG member links changes. Note that forced-down links also influence the minimum links count for the LAG as a whole, which can bring down the LAG, so enable this feature only in configurations where LAG traffic is carefully monitored and controlled.

[See [Understanding Local Minimum Links](#).]

### *Layer 2 Features*

- **Support for IRB interfaces on Q-in-Q VLANs (QFX5100 switches and QFX5100 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D40, integrated routing and bridging (IRB) interfaces are supported on Q-in-Q VLANs—you can configure the IRB interface on the same interface as one used by an S-VLAN, and you can use the same VLAN ID for both the VLAN used by the IRB interface and for the VLAN used as an S-VLAN.

Packets arriving on an IRB interface that is using Q-in-Q VLANs will get routed regardless of whether the packet is single tagged or double tagged. The outgoing routed packets contain an S-VLAN tag only when exiting a trunk interface; the packets exit the interface untagged when exiting an access interface.



**NOTE:** You can configure the IRB interface only on S-VLAN (NNI) interfaces, not on C-VLAN (UNI) interfaces.

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

- **Dual VLAN tag translation (QFX5100 switches and QFX5100 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D40, you can use the dual VLAN tag translation (also known as dual VLAN tag rewrite) feature to deploy switches in service-provider domains, allowing dual-tagged, single-tagged, and untagged VLAN

packets to come into or exit from the switch. Operations added for dual VLAN tag translation are swap-push, swap-swap, and pop-push.

Dual VLAN tag translation supports:

- Configuration of S-VLANs (NNI) and C-VLANs (UNI) on the same physical interface
- Control protocols such as VSTP, OSPF, and LACP
- IGMP snooping
- Configuration of a private VLAN (PVLAN) and VLAN on a single-tagged interface
- Use of TPID 0x8100 on both inner and outer VLAN tags

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

- **Support to exclude RVIs from state calculations (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D40, you can exclude a trunk or access interface from the state calculation for a routed VLAN interface (RVI) for member VLANs. An RVI typically has multiple ports in a single VLAN. Excluding trunk and access interfaces from state calculations means that as that soon as the port specifically assigned to the VLAN goes down, the RVI for the VLAN is marked as down. Include the **autostate-exclude** statement at the **[edit interfaces ether-options]** hierarchy level.

[See [Excluding a Routed VLAN Interface from State Calculations.](#)]

## MPLS

- **Support for IRB interfaces over an MPLS core network (QFX5100 switches)**—Starting in Junos OS Release 14.1X53-D40, you can configure integrated routing and bridging (IRB) interfaces over an MPLS network on QFX5100 switches. An IRB is a logical Layer 3 VLAN interface used to route traffic between VLANs.

By definition, VLANs divide a LAN's broadcast environment into isolated virtual broadcast domains, thereby limiting the amount of traffic flowing across the entire LAN and reducing the possible number of collisions and packet retransmissions within the LAN. To forward packets between different VLANs, you normally need a router that connects the VLANs. Now you can accomplish this forwarding without using a router by simply configuring an IRB interface on the switch. The IRB interface functions as a logical switch on which you can configure a Layer 3 logical interface for each VLAN. The switch relies on its Layer 3 capabilities to provide this basic routing between VLANs. With IRB, you can configure label-switched paths (LSPs) to enable the switch to recognize which packets are being sent to local addresses, so that they are bridged (switched) whenever possible and are routed only when necessary. Whenever packets can be switched instead of routed, several layers of processing are eliminated.

[See [Example: Configuring IRB Interfaces on QFX5100 Switches over an MPLS Core Network](#) and [Understanding Integrated Routing and Bridging](#) .]

## Multicast Protocols

- **Support for static multicast route leaking for VRF and virtual-router instances (QFX5100 and EX4300 switches)**—Starting with Junos OS Release 14.1X53-D40, you can configure your switch to share IPv4 multicast routes among different virtual routing and forwarding (VRF) instances or different virtual-router instances. Only multicast static routes with a destination-prefix length of /32 are supported for multicast route leaking. Only Internet Group Management Protocol version 3 is supported. To configure multicast route leaking for VRF or virtual-router instances, include the **next-table routing-instance-name.inet.0** statement at the **[edit routing-instances routing-instance-name routing-options static route destination-prefix/32]** hierarchy level. For **routing-instance-name**, include the name of a VRF or virtual-router instance.

On the EX4300 switch, multicast route leaking is supported only when the switch functions as a line card in a Virtual Chassis.

[See [Understanding Multicast Route Leaking for VRF and Virtual-Router Instances](#) .]

## QFabric Systems

- **Support for displaying the Junos OS software version stored in a USB installer key (QFabric systems)**—Starting with Junos OS Release 14.1X53-D40, you can display the version of Junos OS software stored on a standard USB installer key when it is inserted on a Director group device by issuing the **show system software usb-software-version** command.
- **Support for EX4300 switches in a QFabric System control plane**—Starting in Junos OS Release 14.1X53-D40, EX4300 switches can be used as the control plane switches in a QFabric System instead of EX4200 switches.

- The control plane of a QFX3000-G QFabric System can be comprised of two Virtual Chassis with four EX4300-48T switches each for a copper-based control plane, or four EX4300-48P switches for a fiber-based control plane. Four 10-Gigabit Ethernet uplink ports on each Virtual Chassis connect the two Virtual Chassis configurations together.
- The control plane of a QFX3000-M QFabric System can be comprised of two EX4300-48T switches with an SFP+ uplink module installed for a copper-based control plane, or two EX4300-48P switches with an SFP+ uplink module installed for a fiber-based control plane.

You cannot mix EX4300 switches and EX4200 switches in the same QFabric system; the control plane must be comprised of the same type of switch.



**NOTE:** Junos OS Release 15.1R3 is the recommended software version for the EX4300 switches.

[See [Understanding QFX3000-G QFabric System Hardware Configurations](#), [Understanding QFX3000-M QFabric System Hardware Configurations](#), and [Understanding the QFabric System Control Plane](#).]

- **Support for SNMPv3 (QFabric systems)** —Starting in Junos OS Release 14.1X53-D40, QFabric systems support SNMP version 3 (SNMPv3). In contrast to SNMP version 1 (SNMPv1) and SNMP version 2 (SNMPv2), SNMPv3 supports authentication and encryption. With SNMPv3, you can query QFabric systems by using the SNMPv3 request, receive SNMPv3 traps and informs, and query QFabric SNMPv3 MIBs for authentication and encryption. SNMPv3 offers strong authentication to determine whether a message is arriving from a valid source and provides message encryption to prevent the data from being snooped by an unauthorized source.

[See [SNMP v3 Overview](#)]



## Security

- **Distributed denial-of-service (DDoS) protection (QFX5100 switches and Virtual Chassis)**—A denial-of-service (DoS) attack is any attempt to deny valid users access to network or server resources by using up all the resources of the network element or server. Distributed denial-of-service attacks (DDoS) involve an attack from multiple sources, enabling a much greater amount of traffic to attack the network. The attacks typically use network protocol control packets to trigger a large number of exceptions to the switch control plane. This results in an excessive processing load that disrupts normal network operations. Starting in Junos OS 14.1X53-D40, Junos OS DDoS protection enables QFX5100 switches and Virtual Chassis to continue functioning while under attack. It identifies and suppresses malicious control packets while enabling legitimate control traffic to be processed. A single point of DDoS protection management enables network administrators to customize profiles for their network control traffic.

[See [Understanding Distributed Denial-of-Service Protection on QFX Series Switches](#).]

## Software-Defined Networking (SDN)

- **OVSDB-VXLAN support with VMware NSX for vSphere (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D40, the Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which NSX controllers and QFX5100 standalone switches that function as virtual tunnel endpoints (VTEPs) can communicate. In an NSX for vSphere (NSX-v) version 6.2.4 environment, NSX controllers and QFX5100 switches can exchange control and statistical information via the OVSDB schema for physical devices, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to bare-metal servers in a physical network and vice versa. You can set up a connection between the QFX5100 management interface (em0 or em1) and an NSX controller.

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

- **BFD in a VMware NSX for vSphere environment with OVSDB and VXLAN (QFX5100 switches)**—Within a Virtual Extensible LAN (VXLAN) managed by the Open vSwitch Database (OVSDB) protocol, by default, Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic is replicated and forwarded by one or more software virtual tunnel endpoints (VTEPs) or service nodes in the same VXLAN. (The software VTEPs and service nodes are collectively referred to as *replicators*.)

Starting with Junos OS Release 14.1X53-D40, a Juniper Networks switch that functions as a hardware VTEP in a VMware NSX for vSphere (NSX-v) environment uses the Bidirectional Forwarding Detection (BFD) protocol to prevent the forwarding of BUM packets to a nonfunctional replicator.

By exchanging BFD control messages with replicators at regular intervals, the hardware VTEP can monitor the replicators to ensure that they are functioning and are, therefore, reachable. Upon receipt of a BUM packet on an OVSDB-managed interface, the hardware VTEP can choose one of the functioning replicators to handle the packet.

[See [Understanding BFD in a VMware NSX Environment with OVSDB and VXLAN](#).]

- **EVPN-VXLAN support of Virtual Chassis and Virtual Chassis Fabric (QFX5100, QFX5100 Virtual Chassis, Virtual Chassis Fabric)**—Ethernet VPN (EVPN) supports multihoming active-active mode, which enables a host to be connected to two leaf devices through a Layer 2 link aggregation group (LAG) interface. In previous Junos OS releases, the two leaf devices had to be QFX5100 standalone switches. Starting with Release 14.1X53-D40, the two leaf devices can be QFX5100 standalone switches, QFX5100 switches configured as a Virtual Chassis, QFX5100 switches configured as a Virtual Chassis Fabric (VCF), or a mix of these options.

On each leaf device, the LAG interface is configured with the same Ethernet segment identifier (ESI) for the host. The two leaf devices on which the same ESI is configured are peers to each other.

If a host, for example, host 1, is connected to two leaf devices through LAG interfaces, Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic is handled as follows:

- **Sending BUM packets**—Through the control of the LAG interface, only one copy of a BUM packet is forwarded from host 1 to one of the leaf devices to which host 1 is connected.
- **Receiving BUM packets from another host in the Layer 2 overlay**—Per multihoming active-active mode, one of the leaf devices to which host 1 is connected is elected as a designated forwarder (DF). If another host in the Layer 2 overlay—for example, host 2—sends a BUM packet, both leaf devices to which host 1 is connected receive the packet, but only the DF forwards it to host 1. The other leaf device drops the packet.
- **Receiving BUM packets from the host that originated the packets**—If host 1 sends a BUM packet, the packet is received by all other leaf devices in the Layer 2 overlay, including the peer leaf device to which host 1 is also connected. In this case, the peer leaf device drops the packet because the packet must not be forwarded to host 1, which originated the packet.
- **Receiving BUM packets from another host connected to the same leaf device**—If another host—for example, host 3—that is connected to the same leaf device as host 1 sends a BUM packet, the packet is forwarded to both leaf devices to which host 1 is connected. Per a local bias, the same leaf device to which both host 3 and host 1 are connected forwards the packet to host 1. The other remote leaf device to which only host 1 is connected drops the packet.

[See [EVPN-VXLAN Support of Virtual Chassis and Virtual Chassis Fabric](#).]

## New Features in Release 14.1X53-D35

### Interfaces and Chassis

- **PVLAN and Q-in-Q on the same interface (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D35, you can configure a private VLAN and Q-in-Q tunneling on the same Ethernet port. To configure both PVLAN and Q-in-Q on the same physical interface, you must configure flexible Ethernet services to support dual methods of configuring logical interfaces. Q-in-Q requires a service provider configuration method, and PVLAN requires an enterprise configuration method.

To configure a physical interface to support both PVLAN and Q-in-Q:

1. Configure flexible VLAN tagging to enable the interface to transmit packets with two 802.1Q VLAN tags.

```
[edit groups group-name ]
user@switch# set interfaces interface-name flexible-vlan-tagging
```

2. Configure flexible Ethernet services to enable the interface to support PVLAN and Q-in-Q on the same interface.

```
[edit groups group-name ]
user@switch# set interface interface-name flexible-ethernet-services
```

3. Enable VLAN bridge encapsulation on the logical interface.

```
[edit groups group-name]
user@switch# set interfaces interface-name unit unit-number encapsulation vlan-bridge
```

4. Assign the VLAN ID for the logical interface.

```
[edit groups group-name]
user@switch# set interfaces interface-name unit unit-number vlan-id vlan-id
```

### MPLS

- **Support for equal-cost multipath (ECMP) operation on MPLS using firewall filters (QFX5100 switches)**—Starting with Junos OS 14.1X53-D35, QFX5100 switches support ECMP operation on MPLS using firewall filters. Use the following commands to enable the feature:

```
[edit]
user@switch# set policy-options policy-statement load-balancing-policy then
  load-balance per-packet
user@switch# set routing-options forwarding-table export load-balancing-policy
```

## New Features in Release 14.1X53-D30

### Authentication and Access Control

- **Access control and authentication (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D30, QFX5100 switches support controlling access to your network using 802.1X authentication and MAC RADIUS authentication.

- 802.1X authentication provides port-based network access control (PNAC) as defined in the IEEE 802.1X standard. QFX5100 switches support 802.1X features including guest VLAN, private VLAN (PVLAN), server fail fallback, dynamic changes to a user session, RADIUS accounting, and configuration of port-filtering attributes on the RADIUS server using VSAs. You configure 802.1X authentication at the **[edit protocols dot1x]** hierarchy level.
- MAC RADIUS authentication is used to authenticate end devices, whether or not they are enabled for 802.1X authentication. You can permit end devices that are not 802.1X-enabled to access the LAN by configuring MAC RADIUS authentication on the switch interfaces to which the end devices are connected. You configure MAC RADIUS authentication at the **[edit protocols dot1x authenticator interface interface-name mac-radius]** hierarchy level.

[See [Understanding Authentication on Switches](#).]

### **Cloud Analytics Engine**

- **Data Learning Engine (DLE) component APIs to access Network Traffic Analysis (NTA) statistics (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D30 and Network Director 2.5, you can enable devices to generate NTA flow statistics using Network Director, and configure DLE to collect, process, and store the data. DLE NTA APIs are provided to allow access to the NTA data that DLE maintains.

[See [Data Learning Engine API Overview](#).]

- **Data Learning Engine (DLE) streaming flow data subscription service and RESTful APIs (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D30, DLE supports a UDP-based network analytics data subscription service that streams analytics data in bulk to subscribed clients as it is collected. The service supports streaming of application flow path analytics data from active flows on network devices that support Cloud Analytics Engine. DLE clients can subscribe to receive this data using DLE data subscription RESTful APIs, avoiding the overhead of having to periodically request this data from DLE and enabling custom real-time client telemetry.

[See [Data Learning Engine API Overview](#).]

### Ethernet Switching

- **IRB in PVLAN (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D30, you can configure an integrated routing and bridging (IRB) interface in a private VLAN (PVLAN) so that devices in the community 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.](#)]

### Interfaces and Chassis

- **Short-reach mode (QFX5100-48T switch)**—Allows you to use short cable lengths (less than 10 meters) for copper-based 10-Gigabit Ethernet interfaces. Enabling short-reach mode reduces power consumption on these interfaces. You can configure short-reach mode for individual interfaces and for a range of interfaces. Enable short-reach mode for individual interfaces by including the **enable** statement at the `[edit chassis fpc slot-number pic slot-number]` hierarchy level. Enable short-reach mode for a range of interfaces by including the **enable** statement at the `[edit chassis fpc slot-number pic port-range port-range-low port-range-high]` hierarchy level.

### MPLS

- **IPv6 Layer 3 VPNs (QFX5100 switches)**—QFX5100 switch interfaces in a Layer 3 VPN can now be configured to carry IP version 6 (IPv6) traffic. This feature, commonly referred to as 6VPE, allows for the transport of IPv6 traffic across an MPLS-enabled IPv4 backbone to provide VPN service for IPv6 customers.
- **MPLS over Layer 3 subinterfaces (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D30, MPLS over Layer 3 subinterfaces is supported on a QFX5100 switch when the switch is used as a *label switch router (LSR)*. MPLS over Layer 3 subinterfaces has already been supported when a QFX5100 switch is used as a *label edge router (LER)*.

[See [MPLS Limitations on QFX Series and EX4600 Switches.](#)]

- **MPLS features (QFX5100 Virtual Chassis, Virtual Chassis Fabric)**—The following MPLS features are now supported for QFX5100 Virtual Chassis and Virtual Chassis Fabric (VCF):
  - BGP L3 VPN
  - Carrier-over-Carrier and Interprovider
  - Ethernet over MPLS pseudowires based on LDP
  - Static/Dynamic Ethernet pseudowires over LDP/RSVP tunnels
  - Pseudowire over aggregated Ethernet interfaces (core-facing interface)
  - RSVP FRR including link-protection/node-link-protection
  - Junos fast-reroute
  - Ethernet pseudowires over QFX5100 Virtual Chassis and VCF deployments

**Software-Defined Networking (SDN)**

- **Class-of-service support for OVSDB-managed VXLAN interfaces (QFX5100 switches)**—Class-of-service (CoS) features can now be configured on OVSDB-managed VXLAN interfaces on QFX5100 switches. An OVSDB-managed VXLAN interface uses an OVSDB controller to create and manage the VXLAN interfaces and tunnels.
- **Firewall filters on OVSDB-managed interfaces (QFX5100 switches)**—Enables you to configure firewall filters on interfaces managed by a Contrail controller through the Open vSwitch Database (OVSDB) management protocol.

[See [Understanding Firewall Filters on OVSDB-Managed Interfaces.](#)]

- **Policers on OVSDB-managed interfaces (QFX5100 switches)**—Enables you to configure two-rate three-color markers (policers) on interfaces managed by a Contrail controller through the Open vSwitch Database (OVSDB) management protocol.

[See [Understanding Policers on OVSDB-Managed Interfaces.](#)]

- **MAC limiting on OVSDB-managed interfaces (QFX5100 switches)**—Enables you to configure MAC limiting on interfaces managed by a Contrail controller through the Open vSwitch Database (OVSDB) management protocol.
- **NNI and UNI on the same interface (QFX5100 switches)**—Enables you to configure the same interface as a network-to-network interface (NNI) and a user-network interface (UNI) when you use Q-in-Q tunneling.
- **OVSDB in Junos OS software package, ISSU and NSSU support (QFX5100, QFX5100VC)**—Starting with 14.1X53-D30, OVSDB software is included in the Junos OS software package (`jinstall`). The introduction of this new feature results in the following changes:

- To upgrade the OVSDB software on your Juniper Networks switch or Virtual Chassis to a later version, you can now use the in-service software upgrade (ISSU) or nonstop software upgrade (NSSU) process. When upgrading the OVSDB software, be aware that this upgrade requires graceful Routing Engine switchover (GRES) only.
- To install OVSDB on your QFX5100 switch or Virtual Chassis, you no longer need to download and install the `jsdn-i386-release` software package.

[See [Understanding In-Service Software Upgrade \(ISSU\)](#) and [Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric.](#)]

- **OVSDB support with Contrail (QFX5100, QFX5100 Virtual Chassis, Virtual Chassis Fabric)**—Starting with Junos OS Release 14.1X53-D30, the Open vSwitch Database (OVSDB) management protocol provides a means through which a Contrail controller and a QFX5100 switch, QFX5100 Virtual Chassis, or a Virtual Chassis Fabric that includes QFX5100 switches only can communicate. In an environment in which Contrail Release 2.20 or later is deployed, a Contrail controller and a QFX5100 switch, QFX5100 Virtual Chassis, or Virtual Chassis Fabric 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 Open vSwitch Database Management Protocol Running on Juniper Networks Devices](#).]

- **Support for ping and traceroute with VXLANs (QFX5100 switches)**—Enables you to use ping and traceroute to debug the underlay that supports a VXLAN overlay.

[See [ping overlay](#) and [traceroute overlay](#).]

### VPNs

- **EVPN control plane for VXLAN supported interfaces (QFX5100 switches)**—Traditionally, data centers have used Layer 2 technologies such as Spanning Tree Protocol (STP), multichassis link aggregation group (MC-LAG), or TRILL for compute and storage connectivity. As the design of data centers shifts from more traditional to scale-out, service-oriented multitenant networks, a new data center architecture has been provided that allows decoupling of an underlay network from the tenant overlay network with VXLAN. By using a Layer 3 IP-based underlay coupled with a VXLAN-EVPN overlay, you can deploy larger networks than those possible with traditional Layer 2 Ethernet-based architectures. With overlays, endpoints (servers or virtual machines) can be placed anywhere in the network and remain connected to the same logical Layer 2 network. The benefit is that virtual topology, using both MX Series routers and QFX5100 switches, can be decoupled from the physical topology.

## New Features in Release 14.1X53-D27

---

### Hardware

- **QFX5100-24Q-AA switch**—This low-latency, high-performance, top-of-rack switch provides 2.56 Tbps throughput. Each QSFP+ port supports 40-Gigabit Ethernet but can be configured as four independent 10-Gigabit Ethernet ports using breakout cables (channelization mode). The switch can also be configured to support 96 10-Gigabit Ethernet ports using breakout cables (channelization mode) with 1280-Gbps total throughput.

The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

The QFX5100-24Q-AA module bay can accommodate a single double-wide expansion module (QFX-PFA-4Q) and two single-wide optional expansion modules (two or one each of QFX-EM-4Q and EX4600-EM-8F).

- **QFX-PFA-4Q expansion module (QFX5100-24Q-AA switch)**—Starting with Junos OS Release 14.1X53-D27, the QFX5100-24Q-AA switch supports the QFX-PFA-4Q expansion module. This double-wide expansion module provides four additional 40-Gigabit Ethernet QSFP+ ports, a dedicated FPGA, and support for the Precision Time Protocol (PTP).

## New Features in Release 14.1X53-D26

---

### *Network Management and Monitoring*

- **DHCP smart relay (QFX5100)**—Starting with Junos OS Release 14.1X53-D26, you can configure alternative IP addresses for the gateway interface so that if the server fails to reply to the requests sent from the primary gateway address, the switch can resend the requests using alternative gateway addresses. To use this feature, you must configure an IRB interface or Layer 3 subinterface with multiple IP addresses and configure that interface as a relay agent.

### *Open vSwitch Database (OVSDB)*

- **New OVSDB command summaries (QFX5100, QFX5100 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D26, the **show ovssdb commit failures** and **clear ovssdb commit failures** commands are introduced.

If you suspect a problem has occurred with the configuration of an OVSDB-managed Virtual Extensible LAN (VXLAN) and associated logical interface(s), you can enter the **show ovssdb commit failures** command. This command describes the OVSDB-managed VXLANs and associated logical interface(s) that the Juniper Networks switch automatically configured but was unable to commit.

After you resolve the problem, you can remove the configuration from the queue and retry committing the configuration by using the **show ovssdb commit failures** command.

- **Storm control on OVSDB-managed interfaces (QFX5100)**—Starting with Junos OS Release 14.1X53-D26, you can configure storm control on VXLAN interfaces that are managed by an OVSDB controller. By default, Layer 2 BUM traffic that originates in an OVSDB-managed VXLAN is replicated and forwarded by a service node in the same VXLAN. Because service nodes can be overloaded if too much BUM traffic is received, you can manually configure storm control on server-facing VXLAN interfaces to control how much of this traffic is allowed into a VXLAN.



## New Features in Release 14.1X53-D25

### MPLS

- **MPLS stitching for virtual machine connections (QFX5100, QFX3500)**—By using MPLS, the stitching feature provides connectivity between virtual machines on opposite sides of data center routers. An external controller, programmed in the data plane, assigns MPLS labels to both virtual machines and servers. Then, the signaled MPLS labels are used between the data center routers, generating static label-switched paths (LSPs), resolved over RSVP or LDP, to provide the routes dictated by the labels. The new CLI command **stitch**, located under the LSP **transit** command, provides this capability.

[See [MPLS Stitching For Virtual Machine Connection](#).]

### Open vSwitch Database (OVSDDB)

- **OVSDDB schema updates (QFX5100 switch, QFX5100 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D25, the Open vSwitch Database (OVSDDB) schema for physical devices version that is implemented on QFX5100 switches is version 1.3.0. In addition, this schema now supports the multicast MACs local table.

[See [Open vSwitch Database Schema for Physical Devices](#).]

### Software Installation and Upgrade

- **Preboot eXecution Environment (PXE) software for Junos Fusion satellite devices (QFX5100 switches)**—Enables you to convert a Junos Fusion satellite device back into a standalone QFX5100 switch. For more information on this feature, please see the Junos OS 14.2R3 Release Notes and the Junos Fusion documentation.

### System Management

- **DHCP relay with DHCP server and DHCP client in separate routing instances**—You can use a stateless DHCP relay agent between a client and server in different virtual routing instances. This feature uses cross-message exchange between the virtual routing instances and supports both DHCPv4 and DHCPv6 packets. This method ensures that:
  - DHCP server network is isolated from the DHCP clients, because there is no direct routing between the client's and server's routing instances.
  - Only DHCP packets, not routine traffic, are relayed across the two routing instances.

[See [DHCP Message Exchange Between DHCP Clients and DHCP Server in Different Virtual Routing Instances](#).]

- **Precision Time Protocol (PTP) transparent clock (QFX5100 switch)**—PTP synchronizes clocks throughout a packet-switched network. With a transparent clock, the PTP packets are updated with residence time as the packets pass through the switch. There is no master/slave designation. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that timestamps are sent in one packet. In a two-step process, estimated timestamps are sent in one packet, and additional packets contain updated

timestamps. In addition, user UDP over IPv4 and IPv6, and unicast and multicast transparent clocks, are supported. You can configure the transparent clock at the **[edit protocols ptp]** hierarchy level.

[See [Understanding Transparent Clocks in Precision Time Protocol.](#)]

## VXLAN

- **Configurable VXLAN UDP port (QFX5100)**—Starting with Junos OS 14.1X53-D25, you can configure the UDP port used as the destination port for VXLAN traffic on a QFX5100 switch. To configure the VXLAN destination port to be something other than the default UDP port of 4789, enter **set protocols l2-learning destination-udp-port port-number**. The port you configure will be used for all VXLANs configured on the switch.



**NOTE:** If you make this change on one switch in a VXLAN, you must make the same change on all the devices that terminate the VXLANs configured on your switch. If you do not do so, traffic will be disrupted for all the VXLANs configured on your switch. When you change the UDP port, the previously learned remote VTEPs and remote MACs are lost and VXLAN traffic is disrupted until the switch relearns the remote VTEPs and remote MACs.

[See [Understanding VXLANs.](#)]

---

## New Features in Release 14.1X53-D15

### Hardware

- **Extended node support (QFX5100-24Q and QFX5100-48T switches)**—Enables you to include a QFX5100-24Q switch and a QFX5100-48T switch as a Node device in a QFabric System. To add the device, first install the QFabric “5” family software package (jinstall-qfabric-5-release.tgz) on the switch, and attach two management ports to the QFabric system control plane. For copper-based control plane systems, use the RJ-45 fixed management port and one SFP management port on the QFX5100 Node device with a copper module. For fiber-based control plane systems, use two SFP management ports on the QFX5100 Node device with fiber modules.

[See [Understanding the QFabric System Hardware Architecture.](#)]

- **Improved online insertion and replacement procedures (QFabric systems)**—Allows for nondisruptive insertion or replacement of server Node groups, network Node groups, redundant server Node groups, Interconnect devices, and front and rear cards of the Interconnect devices.

[See [Powering Off an Existing QFabric Node Device.](#)]

- **QFX5100 Interconnect device (QFabric systems)**—Allows a QFX5100-24Q switch to operate as a QFX3000-M Interconnect device. The interconnect acts like a backplane for data-plane traffic traversing the QFX3000-M QFabric system between Node devices. The QFX5100 Interconnect device has 24 40-Gigabit QSFP+ ports, but only 16 are available as fte ports. The QFX5100 Interconnect device features two RJ-45

management ports and two SFP management ports, which allow connection to either copper-based or fiber-based control-plane networks.

[See [Understanding Interconnect Devices](#).]

### ***Class of Service***

- **Mitigating fate sharing on Interconnect devices by remapping forwarding classes (QFabric systems)**—Enables you to remap traffic assigned to a forwarding class into different, separate forwarding classes to mitigate fate sharing as the traffic crosses the Interconnect device. Separating the traffic into multiple forwarding classes spreads the flows across multiple output queues instead of using one output queue for all of the traffic. (Each forwarding class uses a different output queue, and each output queue has its own dedicated bandwidth resources.) Fate sharing occurs when flows in the same forwarding class (flows that have the same IEEE 802.1p priority code point) use the same output queue on an interface, because the flows share the same path and resources. When one flow becomes congested, the congestion can affect the other flows that use the same output queue even if they are not experiencing congestion, because when the congested flow is paused, the other flows that use the same code point are also paused. Because flows from many Node devices cross the Interconnect device, the flows are aggregated at egress interfaces, which increases the chance of fate sharing. Forwarding class remapping mitigates fate sharing on the Interconnect device by separating the traffic into different forwarding classes that use different output queues, so pausing one congested flow does not affect uncongested flows that have been mapped to different forwarding classes and therefore to different output queues.

[See [Understanding How to Mitigate Fate Sharing on a QFabric System Interconnect Device by Remapping Traffic Flows \(Forwarding Classes\)](#) and [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#).]

- **Scheduler configuration on Interconnect device fabric ports (QFabric systems)**—Enables you to configure scheduling on the fabric (fte and bfte) ports of the QFabric system Interconnect devices. (This complements the Junos OS Release 13.1 feature that provides scheduler configuration on Node device fabric ports. The combination of access port, Node device fabric port, and Interconnect device fabric port scheduling gives you complete control of scheduling across a QFabric system.) In earlier Junos OS releases, Interconnect device fabric port scheduling was done by default, with no user configuration. In Junos OS Release 14.1X53-D15, the default fabric port scheduler on Interconnect devices is the same as it was in earlier releases.

[Understanding CoS Scheduling Across the QFabric System](#) and [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#).]

### **Multicast Features**

- **IGMP querier (QFabric systems)**—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 [Using a Switch as an IGMP Querier.](#)]

- **IGMPv3 (QFabric systems)**—Introduces support for Internet Group Management Protocol version 3 (IGMPv3). IGMPv3 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 routing devices. Multicast routing devices use IGMP to learn which groups have members for each of their attached physical networks.

[See [Understanding IGMP.](#)]

- **IGMPv3 snooping (QFabric systems)**—With IGMP snooping enabled (the default setting), a switch monitors the IGMP traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

[See [IGMP Snooping Overview.](#)]

- **Multicast flow groups (QFabric systems)**—Node devices usually forward multicast traffic on all available Interconnect devices to distribute the load balancing replication load. As a result, redundant multicast streams can flow through one Interconnect device, making that Interconnect device a potential single point of failure for the redundant flows. Some applications require that the redundant multicast streams flow through different Interconnect devices to prevent a single Interconnect device from potentially dropping both streams of multicast traffic during a failure. You can enforce this use of dual Interconnect devices by using the QFabric flow segregation feature.

[See [Understanding QFabric Multicast Flow Groups.](#)]

- **PIM-SSM (QFabric systems)**—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 [PIM SSM.](#)]

### ***Network Management and Monitoring***

- **Cloud Analytics Engine (QFX5100 switches)**—Uses network data analysis to improve application performance and availability. Cloud Analytics Engine includes data collection, analysis, correlation, and visualization, helping you better understand the behavior of workloads and applications across the physical and virtual infrastructure. Cloud Analytics Engine provides an aggregated and detailed level of visibility, tying applications and the network together, and an application-centric view of network status, improving your ability to quickly roll out new applications and troubleshoot problems.

[See [Cloud Analytics Engine](#).]

### ***Open vSwitch Database (OVSDB)***

- **Automatic configuration of OVSDB-managed VXLANs with trunk interfaces (QFX5100 switches)**—In a VMware NSX for Multi-Hypervisor environment for the data center, the QFX5100 switch can automatically configure an OVSDB-managed VXLAN and one or more interfaces associated with the VXLAN, thereby eliminating the need for you to perform these tasks, using the Junos OS CLI. The automatic configuration of the VXLAN and associated interfaces is based on the configuration of a logical switch in NSX Manager or in the NSX API. Starting in Junos OS Release 14.1X53-D15, the switch supports the automatic configuration of trunk interfaces and their association with an OVSDB-managed VXLAN. In this situation, trunk interfaces enable the support of multiple software applications running directly on a physical server that generate traffic that must be isolated by OVSDB-managed VXLANs.

[See [Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment](#).]

- **OVSDB support with NSX (QFX5100 Virtual Chassis, Virtual Chassis Fabric)**—Starting with Junos OS Release 14.1X53-D15, the Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which VMware NSX controllers and a QFX5100 Virtual Chassis or a Virtual Chassis Fabric that includes QFX5100 switches only can communicate. In an NSX multi-hypervisor environment, NSX version 4.0.3 and later controllers and a QFX5100 Virtual Chassis or Virtual Chassis Fabric can exchange control and statistical information via the OVSDB schema for physical devices, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to entities in a physical network and vice versa.

You can set up a connection between the QFX5100 management interface (**em0** or **em1**) and an NSX controller.

[See [Setting Up Open vSwitch Database Connections Between Junos OS Devices and Controllers](#).]

### ***QFabric Systems***

- **QFabric system software downgrade support (QFabric systems)**—Starting with Junos OS 14.1X53-D15, downgrading software provides a quick recovery mechanism to a previous software version and configuration file in cases where a software upgrade or configuration changes have made the QFabric system unstable or inoperable. The recovery mechanism consists of a “restore-point,” which is a snapshot of the software

on the QFabric system as well as the configuration that can be rolled back to. Downgrade support does not replace the existing backup and restore functionality.

- To enable software downgrade:
- Create a restore-point.



**NOTE:** You can only create one restore-point at a time. Creating a new restore-point deletes the existing restore-point if there is one. Also, all CLI commands are blocked while creating a restore-point.

To create a restore-point, issue the **request system software restore-point** command.

- To roll back to the restore-point, issue the **request system software recover-from-restore-point** command.
- To display the status of the Director group after creating a restore-point for the QFabric system, issue the **show system software restore-point status** command.

### **Security**

- **Error message displayed when TCAM is full (QFX5100 switches)**—Firewall filters are stored in ternary content addressable memory (TCAM). With previous versions of Junos OS, if you configure a firewall filter that cannot fit into the available TCAM space, the filter defaults to "permit any," and no error message is displayed in the CLI. With Junos OS Release 14.1X53-D15, an error message is displayed in the CLI if this occurs.

[See [Planning the Number of Firewall Filters to Create.](#)]

- **Media Access Control Security (MACsec) support (QFX5100-24Q switches)**—Starting with Junos OS Release 14.1X53-D15, MACsec is supported on all eight SFP+ interfaces on the EX4600-EM-8F expansion module when it is installed in a QFX5100-24Q switch. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is capable of identifying and preventing most security threats, and can be used in combination with other security protocols to provide end-to-end network security. MACsec is standardized in IEEE 802.1AE.

[See [Understanding Media Access Control Security \(MACsec\).](#)]

### *Virtual Chassis and Virtual Chassis Fabric*

- **Increase vmember limit to 512k support (Virtual Chassis Fabric)**—Increases the number of vmembers to 512k. For example, to calculate how many interfaces are required to support 4000 VLANs, divide the maximum number of vmembers (512,000) by the number of configured VLANs (4000). In this case, 128 interfaces are required.

[See [Understanding Bridging and VLANs](#).]

### *VLAN Infrastructure*

- **Support for private VLANs (QFX5100 switches)**—VLANs limit broadcasts to specified users. Private VLANs (PVLANS) take this concept a step further by splitting the broadcast domain into multiple isolated broadcast subdomains and essentially putting secondary VLANs inside a primary VLAN. PVLANS restrict traffic flows through their member switch ports (called “private ports”) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port is usually connected to a router, firewall, server, or provider network. Each PVLAN typically contains many private ports that communicate only with a single uplink, thereby preventing the ports from communicating with each other.

Just like regular VLANs, PVLANS are isolated on Layer 2 and require that a Layer 3 device be used to route traffic among them. PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts. Service providers use PVLANS to keep their customers isolated from one another.

[See [Understanding Private VLANs](#).]

---

## **New Features in Release 14.1X53-D10**

### *Authentication and Access Control*

- **IPv6 for RADIUS AAA (QFX5100 switch and Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D10, QFX5100 switches and QFX5100 Virtual Chassis support IPv6, along with the existing IPv4 support, for user authentication, authorization, and accounting (AAA) using RADIUS servers.

RADIUS authentication is a method of authenticating users who attempt to access the router or switch. To use RADIUS authentication on the switch, configure information about one or more RADIUS servers on the network by including one **radius-server** statement at the **[edit system]** hierarchy level for each RADIUS server.

When you configure a source address for each configured RADIUS server, each RADIUS request sent to a RADIUS server uses the specified source address.

- **Authentication**—Specify which source address Junos OS uses when accessing your network to contact an external RADIUS server for authentication. You configure the IPv6 source address for RADIUS authentication at the **[edit system radius-server server-address source-address]** hierarchy level.
- **Accounting**—Specify which source address Junos OS uses when contacting a RADIUS server for sending accounting information. You configure the IPv6 source address

for RADIUS authentication at the **[edit system accounting destination radius server server-address source-address]** hierarchy level.

[See [source-address](#).]

### ***Bridging and Learning***

- **MAC notification (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, MAC notification is supported on QFX5100 switches. The switches track clients on a network by storing MAC addresses in the Ethernet switching table on the switch. When switches learn or unlearn a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC Notification MIB controls MAC notification for the network management system.

The MAC notification interval defines how often these SNMP notifications are sent to the network management system. The MAC notification interval works by tracking all MAC address additions or removals on the switch over a period of time and then sending all tracked MAC address additions or removals to the network management server at the end of the interval.

Enabling MAC notification allows you to monitor the addition and removal of MAC addresses from the Ethernet switching table remotely using a network management system. The advantage of setting a high MAC notification interval is that the amount of network traffic is reduced because updates are sent less frequently. The advantage of setting a low MAC notification interval is that the network management system is better synchronized with the switch.

Two new MIBs related to MAC notification are provided at Junos OS Release 14.1X53-D10. See [“Documentation Updates” on page 71](#).

[See [Configuring MAC Notification \(CLI Procedure\)](#).]

- **Default VLAN and multiple VLAN range support (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, the default VLAN and multiple VLAN range are supported on QFX5100 switches. They provide the ability for the switch to operate as a *plug and play* device and connect to various Ethernet-enabled devices in a small, scaled enterprise network. When the switch boots, a VLAN named **default** is created. The default VLAN is automatically created for every routing instance that belongs to a type of **virtual-switch** and for the default routing instance named **default-switch**. All interfaces on the switch are automatically configured as access interfaces and are part of the default VLAN.

The default VLAN accepts and forwards untagged packets only and is preconfigured with a VLAN ID (**vlan-id**) of 1. The default VLAN does not support a VLAN ID list (**vlan-id-list**), **vlan-id** set to **all**, or **vlan-id** set to **none**. You can configure the VLAN ID to be another value, but the value must be between 1 and 4093.

Access interfaces that are VoIP-enabled or 802.1X-enabled are internally converted to trunk interfaces, so that the interfaces can belong to multiple VLANs. If the interfaces do not belong to a valid VLAN, the interfaces automatically become part of the default VLAN.



You can configure more than one VLAN range, and each range can contain unique VLAN properties.



**NOTE:** Virtual Chassis interfaces cannot be preconfigured to belong to the default VLAN or any other VLAN.



**NOTE:** For interfaces to be part of the default VLAN, you must configure the interfaces to be part of the Ethernet switching family. You can configure Ethernet switching at the [edit interfaces *interface-name* unit family] CLI hierarchy level.

- **Ethernet ring protection switching (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, Ethernet ring protection switching (ERPS) is supported on QFX5100 switches. ERPS helps achieve high reliability and network stability. Links in the ring never form loops that fatally affect the network operation and services availability.

[See [Understanding Ethernet Ring Protection Switching Functionality](#).]

### *High Availability*

- **Resilient hashing support for link aggregation groups and equal cost multipath routes (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, resilient hashing is now 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.

[See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk Groups](#).]

### *Infrastructure*

- **Licensing enhancements (QFX Series)**—Starting with Junos OS Release 14.1X53-D10, licensing enhancements on QFX Series switches enable you to configure and delete license keys in a Junos OS CLI configuration file. The license keys are validated and installed after a successful commit of the configuration file. If a license key is invalid, the commit fails and issues an error message. You can configure individual license keys or multiple license keys by issuing Junos OS CLI commands or by loading the license key configuration contained in a file. All installed license keys are stored in the `/config/license/` directory.

To install an individual license key in the Junos OS CLI, issue the **set system license keys key *name*** command, and then issue the **commit** command.

For example:

```
[edit]
root@switch# set system license keys key "JUNOS_TEST_LIC_FEAT testabc123"
root@switch# commit
```

commit complete

To verify that the license key was installed, issue the **show system license** command.

For example:

```
root@switch> show system license
License usage:

```

	Licenses	Licenses	Licenses	Expiry
Feature name	used	installed	needed	
sdk-test-feat1	0	1	0	
permanent				

```

Licenses installed:
License identifier: JUNOS_TEST_LIC_FEAT
License version: 2
Features:
  sdk-test-feat1  - JUNOS SDK Test Feature 1
  permanent

```

To install multiple license keys in the Junos OS CLI, issue the **set system license keys key *name*** command, and then issue the **commit** command.

For example:

```
[edit]
root@switch# set system license keys key "key_1"
set system license keys key "key_2"
set system license keys key "key_2"
set system license keys key "key_4"
root@switch# commit
commit complete
```

To verify that the license key was installed, issue the **show system license** command.

To install an individual license key configuration in a file, issue the **cat** command:

For example:

```
[edit]
root@switch%cat license.conf
system {
  license {
    keys {
      key "JUNOS_TEST_LIC_FEAT testabc123";
    }
  }
}
```

Load and merge the license configuration file.

For example:

```
[edit]
root@switch# load merge license.conf
load complete
```

Issue the **show | compare** command to see the configuration, and then issue the **commit** command.

For example:

```
[edit]
root@switch# show | compare
[edit system]
+   license {
+       keys {
+           key "JUNOS_TEST_LIC_FEAT testabc123";
+       }
+   }
[edit]
root@switch# commit
```

To verify that the license key was installed, issue the **show system license** command.

For example:

```
root@switch> show system license
License usage:

Feature name          Licenses used  Licenses installed  Licenses needed  Expiry
-----
sdk-test-feat1       0             1                   0
permanent

Licenses installed:
License identifier: JUNOS_TEST_LIC_FEAT
License version: 2
Features:
  sdk-test-feat1  - JUNOS SDK Test Feature 1
  permanent
```

To install multiple license keys in a file, issue the **cat** command:

For example:

```
[edit]
root@switch%cat license.conf
system
{
  license
  {
    keys
    {
      key "key_1"
      key "key_2"
      key "key_3"
      ...
      key "key_n"
    }
  }
}
```

Load and merge the license configuration file, and then issue the **commit** command.

For example:

```
[edit]
root@switch# load merge license.conf
load complete
[edit]
root@switch# commit
```

To verify that the license key was installed, issue the **show system license** command.

You can also delete or deactivate individual and multiple license keys in the Junos OS CLI by issuing the **delete system license keys** or **deactivate system license keys** commands. Do not use the **request system license delete** command to delete the license keys.

For example, to issue the **delete system license keys** command:

```
[edit]
root@switch# delete system license keys
root@switch# commit
```

### *Interfaces and Chassis*

- **Fast reboot option (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, you can enhance the reboot time on a QFX5100 by issuing the new **fast-boot** option with the **request system reboot** command (**request system reboot fast-boot**). The switch reboots in such a way as to minimize downtime of network ports by not bringing the network ports down immediately as in the normal reboot option. There is minimal traffic loss while the forwarding device is reprogrammed.

[See [request system reboot](#).]

- **Keep a link up on a multichassis link aggregation group (MC-LAG) when LACP is not configured on one of the MC-LAG peers (QFX5100 switch)**—Junos OS Release 14.1X53-D10 provides connectivity from provider edge devices to customer edge devices when LACP is not configured on a customer edge device. The customer edge device must have one link connected to the provider edge device, though, and multichassis link aggregation must be configured between the provider edge devices in the MC-LAG. You can configure the force-up feature in Link Aggregation Control Protocol (LACP) on the provider edge device for which you need connectivity. Additionally, only one member interface in the aggregated Ethernet interface can be active, otherwise the provider edge device will receive duplicate packets.

[See [Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#).]

### *Layer 3 Features*

- **Loop-free alternates (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, QFX5100 switches support loop-free alternates (LFA) to compute backup next hops for IS-IS routes, providing IP fast-reroute capability for IS-IS routes. These routes, with precomputed backup next hops, are preinstalled in the Packet Forwarding Engine, which performs a local repair and switches to the backup next hop when the link for the primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives recomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. You can configure loop-free alternates (LFA) for IS-IS at the **[edit protocols isis]** hierarchy level.
- **IS-IS support (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, on QFX5100 switches, the IS-IS protocol has extensions to differentiate between different sets of routing information sent between routers and switches for unicast and multicast. IS-IS routes can be added to the RPF table when special features such as traffic engineering

and shortcuts are turned on. You configure the feature under the **[edit protocols isis]** hierarchy level.

### **MPLS**

- **MPLS-based Layer 3 VPNs (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, MPLS-based Layer 3 VPNs are supported on QFX5100 switches.

Customer networks are private and can use either public addresses or private addresses. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with private addresses being used by other network users. MPLS BGP VPNs solve this problem by adding the route distinguisher prefix to the route.

You can configure the switch as a CE or PE using Layer 3 MPLS/BGP VPN for interprovider and carrier-of-carrier VPNs. The key difference between interprovider and carrier-of-carriers VPNs is whether the customer sites belong to the same autonomous system (AS) or to a separate AS:

- Interprovider VPNs—The customer sites belong to different ASs. You need to configure EBGP to exchange the customer's external routes.
  - Carrier-of-carriers VPNs—The customer sites belong to the same AS. You need to configure IBGP to exchange the customer's external routes.
- **Ethernet-over-MPLS (L2 circuit) (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, Ethernet-over-MPLS is supported on QFX5100 switches. Ethernet-over-MPLS enables you to send Layer 2 Ethernet frames transparently over an MPLS cloud. Ethernet-over-MPLS uses a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. It encapsulates Ethernet protocol data units (PDUs) inside MPLS packets and forwards the packets, using label stacking, across the MPLS network.

This technology has applications in service provider, enterprise, and data center environments. For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network. These data centers require Layer 2 connectivity between them for the following reasons:

    - To replicate the storage over Fibre Channel over IP (FCIP). FCIP works only on the same broadcast domain.
    - To run a dynamic routing protocol between the sites.
    - To support high availability clusters that interconnect the nodes hosted in the various data centers.
  - **MPLS LSP protection (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, the following types of MPLS LSP protection are supported on QFX5100 switches:
    - Fast reroute (FRR)
    - Link protection
    - Node link protection

[ See [MPLS Overview](#).]

***Network Management and Monitoring***

- **Chef for Junos OS (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, Chef for Junos OS is supported on all QFX5100 switches, not just QFX5100 switches that are running Junos OS with automated enhancements for QFX5100 switches.
- **Puppet for Junos OS (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, Puppet for Junos OS is supported on QFX5100 switches that are not running Junos OS with automated enhancements for QFX5100 switches.
- **IEEE 802.3ah (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, QFX5100 switches support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. You configure the feature under the **[edit protocols oam ethernet]** hierarchy level.

### *OpenFlow*

- **Support for OpenFlow v1.0 and v1.3.1 (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, QFX5100 switches support OpenFlow v1.0 and v1.3.1. OpenFlow v1.0 enables you to control traffic in an existing network by adding, deleting, and modifying flows in the switch. You can configure one OpenFlow virtual switch and one active OpenFlow controller under the **[edit protocols openflow]** hierarchy on each QFX5100 switch in the network.

In addition to the OpenFlow v1.0 functionality, OpenFlow v1.3.1 allows the action specified in one or more flow entries to direct packets to a base action called a group. The purpose of the group action is to further process these packets and assign a more specific forwarding action to them. You can view groups that were added, modified, or deleted from the group table by way of the OpenFlow controller using the **show openflow groups** command. You can view group statistics using the **show openflow statistics groups** command.

OpenFlow v1.0 and v1.3.1 are not supported on MX Series routers or EX9200 switches in Junos OS Release 14.1X53-D10. OpenFlow v1.0 is supported in Junos OS Release 14.1 on these platforms.

[See [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS.](#)]

### *Open vSwitch Database (OVSDB)*

- **OVSDB support with NSX (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, the Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which VMware NSX controllers and QFX5100 switches that support OVSDB can communicate. In an NSX multi-hypervisor environment, NSX version 4.0.3 controllers and QFX5100 switches can exchange control and statistical information via the OVSDB schema for physical devices, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to entities in a physical network and vice versa.

You can set up a connection between the QFX5100 management interface (**em0** or **em1**) and an NSX controller.

[See [Setting Up Open vSwitch Database Connections Between Junos OS Devices and Controllers.](#)]

### *Security*

- **Port mirroring to IP address (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, you can send mirrored packets to an IP address over a Layer 3 network (for example, if there is no Layer 2 connectivity to the analyzer device). This feature also enables you to apply an IEEE-1588 timestamp to the mirrored packets.

### *Software Installation*

- **Open Source Python modules supported in automation enhancement (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, these Open Source Python modules are pre-installed in the jinstall-qfx-5-flex-x.tgz software bundle:



- **ncclient**—Facilitates client scripting and application development through the NETCONF protocol.
- **lxml**—Combines the speed and XML feature completeness of the C libraries libxml2 and libxslt with the simplicity of a native Python API.
- **jinja2**—Serves as a fast, secure, designer-friendly templating language.

[See [Overview of Python with QFX5100 Switch Automation Enhancements.](#)]

#### ***Virtual Chassis and Virtual Chassis Fabric***

- **Alias support for Virtual Chassis and Virtual Chassis Fabric (VCF) nodes**—Starting with Junos OS Release 14.1X53-D10, an alias can be used to label nodes in a Virtual Chassis and VCF. An alias allows you to more clearly identify a member switch in your Virtual Chassis or VCF by assigning a text label to it. The text label appears alongside the switch's serial number whenever operational commands, such as **show virtual-chassis**, are used to monitor Virtual Chassis status.

[See [aliases.](#)]

- **Local link bias support for Virtual Chassis with QFX Series member switches**—Starting with Junos OS Release 14.1X53-D10, Virtual Chassis Local Link Bias is available on Link Aggregation Group (LAG) bundles on QFX3500 Virtual Chassis, QFX3600 Virtual Chassis, and mixed QFX3500 and QFX3600 Virtual Chassis. Virtual Chassis local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis that has a LAG bundle composed of member links on different member switches in the same Virtual Chassis. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis using a different member link in the LAG bundle.

[See [Understanding Local Link Bias.](#)]

- **Adaptive load balancing support (Virtual Chassis Fabric)**—Starting with Junos OS Release 14.1X53-D10, adaptive load balancing (ALB) is supported in Virtual Chassis Fabric (VCF). ALB improves traffic management within a VCF by using dynamic load information to make traffic forwarding decisions. ALB introduces a method to better manage extremely large traffic flows—*elephant flows*—by splicing them into smaller flows—*flowlets*—and individually forwarding the flowlets across the VCF to the same destination device over different paths.

[See [Understanding Traffic Flow Through a Virtual Chassis Fabric.](#)]

## VXLAN

- **Layer 2 VXLAN gateway (QFX5100)**—Starting with Junos OS Release 14.1X53-D10, VXLAN is an overlay technology that enables 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 between servers that exist in separate Layer 2 domains by tunneling the traffic through Layer 3 networks. This functionality enables you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use STP to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See [Understanding VXLANs](#).]

### Related Documentation

- [Changes in Behavior and Syntax on page 34](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

## 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 from Junos OS Release 14.1X53 for QFX Series.

- [Authentication and Access Control](#)
- [Interfaces and Chassis](#)
- [Open vSwitch Database \(OVSDB\)](#)
- [SNMP](#)
- [Software Upgrade](#)
- [Virtual Chassis and Virtual Chassis Fabric](#)

## Authentication and Access Control

---

- **Increase in TACACS message length (QFX Series)**—Starting with Junos OS Release 14.1X53-D40, the length of TACACS messages allowed on Junos devices has been increased from 8150 to 65535 bytes.

## Interfaces and Chassis

---

- **ARP and MAC table synchronization during MC-LAG troubleshooting (QFX Series switches and EX4300 switches)**—Starting in Junos OS Release 14.1X53-D40, the **arp-l2-validate** CLI statement is supported at the **[edit interfaces irb]** hierarchy level for QFX Series switches and EX4300 switches. This command can be used to help maintain ARP and MAC table synchronization in an MC-LAG to prevent traffic loss while troubleshooting network problems that cause inconsistencies between the two tables.

[See [Troubleshooting Multichassis Link Aggregation](#) and [arp-l2-validate](#).]

- **Configuring unified forwarding table profiles (EX4600 Virtual Chassis, QFX5100 Virtual Chassis, and QFX Series Virtual Chassis Fabric)**—Starting in Junos OS Release 14.1X53-D40, Packet Forwarding Engines on switches in a Virtual Chassis or Virtual Chassis Fabric (VCF) do not automatically restart upon configuring and committing a unified forwarding table profile change using the **set chassis forwarding-options** statement. Instead, a message is displayed at the CLI prompt and logged to the switch's system log, prompting you to reboot the Virtual Chassis or VCF for the change to take effect. This change avoids Virtual Chassis or VCF instability that might occur with these switches if the profile update propagates to member switches and otherwise causes multiple Packet Forwarding Engines to automatically restart at the same time. This behavior change does not apply to other switch types or to EX4600 and QFX5100 switches not in a Virtual Chassis or VCF; in those cases, the switch continues to restart automatically when a unified forwarding table profile change is committed.

We recommend that you plan to make profile changes in a Virtual Chassis or VCF comprised of these switches only when you can perform a Virtual Chassis or VCF system reboot shortly after committing the configuration update, to avoid instability if one or more member switches restart unexpectedly with the new configuration (while the remaining members are still running the old configuration).

[See [Configuring the Unified Routing Table](#) and [forwarding-options \(chassis\)](#).]

- **New vc-path command display for Virtual Chassis Fabric (VCF)**—Starting in Junos OS Release 14.1X53-D40, the output from the **show virtual-chassis vc-path** command displays additional fields when showing the forwarding path from a source interface to a destination interface in a Virtual Chassis Fabric (VCF), including details of multiple possible next hops. The **vc-path** command display for a forwarding path in a Virtual Chassis remains unchanged.

[See [show virtual-chassis vc-path](#).]

- **Gigabit interface speeds (QFX5100 switches)**—Starting with Junos OS Release 14.1X53-D43, QFX5100 switches correctly interpret and display the interface speed as 1000mbps (1 Gbps) for **ge-** interfaces on 1-Gigabit Ethernet SFP ports. In prior releases

from Junos OS Release 13.2X52-D20 up until 14.1X53-D43, the system incorrectly interprets and displays the speed of these interfaces as 10 Gbps. [See [show interfaces ge.](#)]

---

## Open vSwitch Database (OVSDB)

---

- **Automatic configuration of trunk interfaces that handle untagged packets in OVSDB-managed VXLANs (QFX5100, QFX5100VC)**—In previous Junos OS releases, if you specified a VLAN ID of 0 for a logical switch port in VMware NSX Manager or in the NSX API, the QFX5100 switch automatically configured an access interface to handle untagged packets in the associated Open vSwitch Database (OVSDB)-managed Virtual Extensible LAN (VXLAN). Starting with 14.1X53-D26, specifying a VLAN ID of 0 in a logical switch port configuration causes the QFX5100 switch to automatically configure a trunk port. To enable the trunk port to handle untagged packets, the QFX5100 switch also configures a native VLAN with an ID of 4094. Upon receipt of an untagged packet, the trunk interface adds a VLAN tag of 4094 to the packet and removes the tag as the packet exits the interface, thereby rendering the packet as untagged again.

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

---

## SNMP

---

- **Change in value for a QFabric SNMP object**—The `jnxFabricDeviceEntryName` object now displays the alias of the device and the `jnxFabricDeviceEntryDescription` object contains the serial number only.

---

## Software Upgrade

---

- A controlled version of Junos OS is introduced for the QFX Series in Junos OS Release 14.1X53-D15. The controlled version of Junos OS is required to enable Media Access Control security (MACsec) on a switch. The controlled version of a Junos OS release contains all features and functionality available in the standard version of the Junos OS release while also supporting MACsec. The controlled version of Junos OS is not, by default, shipped on any QFX Series switch. You can download the controlled version of Junos OS from the Software Download Center, provided that you are located in a geography where you are allowed to download the controlled version of Junos OS. If you are unsure of which version of Junos OS is running on your switch, enter the **show version** command. If the “JUNOS Crypto Software Suite” description appears in the output, you are running the controlled version of Junos OS.

The controlled version of Junos OS contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS is also subject to controls imposed under the laws of other countries.

If you have questions about acquiring the controlled version of Junos OS in your country, contact the Juniper Networks Trade Compliance group at [compliance\\_helpdesk@juniper.net](mailto:compliance_helpdesk@juniper.net).

## Virtual Chassis and Virtual Chassis Fabric

---

- **New VCF multicast distribution tree configuration option**—Starting with Junos OS Release 14.1X53-D35, a new Virtual Chassis Fabric (VCF) configuration option, **fabric-tree-root**, is available on EX Series and QFX Series devices in an autoprovisioned or preprovisioned VCF. This option changes how the VCF builds the multicast distribution trees (MDTs) used for forwarding and load-balancing broadcast, unknown unicast, and multicast (BUM) traffic within the VCF. By default, a VCF builds MDTs with each VCF member as the root of a tree, creating as many MDTs as members in the VCF. Setting the **fabric-tree-root** option for one or more members preempts this behavior. Instead, for each member configured with this option, the VCF only builds MDTs with those members as root nodes (referred to as the fabric tree roots). The recommended usage of this option is to set all spine devices in the VCF, and only spine devices, as fabric tree roots.

Using this option avoids traffic interruption in a VCF when a leaf device becomes unavailable and the VCF needs to redistribute traffic within the VCF over the available MDTs. Using only spine-rooted MDTs provides a redistribution path to any destination leaf member directly through a spine member, and prevents traffic from flowing redundantly over paths to and from leaf members (which happens with leaf-rooted MDTs, creating excess traffic load in large VCFs).

[See [fabric-tree-root](#).]

### Related Documentation

- [New and Changed Features on page 3](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

## Known Behavior

This section lists the limitations in Junos OS Release 14.1X53 for the QFX Series.

### High Availability

---

- On a Virtual Chassis Fabric that has more than two Routing Engines configured, a nonstop software upgrade (NSSU) might not succeed. [PR1081786](#)
- On a QFX5100 switch, when you perform an in-service software upgrade (ISSU), interfaces configured with the Link Aggregation Control Protocol (LACP) and with the speed set to fast will come up and go down, causing all protocols on the interface to come up and go down. As a workaround, before you perform an ISSU, configure the LACP speed setting to slow on the switch and its peers. [PR1106510](#)

- On an EX4300 or a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On a QFX Series Virtual Chassis or Virtual Chassis Fabric, an NSSU to Release 14.1X53-D35 might cause a traffic loss of a few seconds for BUM traffic. [PR1128208](#)
- On a QFX5100 Virtual Chassis, when you perform a nonstop software upgrade (NSSU) from Junos OS Release 14.1X53-D30.6 to Junos OS Release 14.1X53-D32, there might be traffic loss for up to one second. [PR1154635](#)
- During a QFabric NSSU from Junos OS Release 12.2X50 to Release 14.1X53, multicast traffic might be impacted (loss or duplication) for up to 60 seconds while upgrading ICs. The variation in loss duration depends on the number of front/back cards on the IC, number of distribution trees passing through those ICs, and so on, because all forwarding paths need to be set up afresh after an upgrade. [PR1225870](#)
- On QFabric systems, when the RSNG backup has a lower MAC/SYSID at reboot, you might observe an abrupt mastership switchover.

As a workaround for the RSNG backup reboot:

1. ssh login to the RSNG master.
2. Execute this command to get access to the RSNG backup: **request session member fpc-number backup**.
3. Go to CLI mode on the RSNG backup, go to edit mode, and execute the following commands: **set interfaces me5 disable**; **set interfaces me6 disable**.
4. On commit it will ask for **confirmation ?confirm ?yes?** and press **enter**. This will automatically reboot the RSNG backup.

[PR1240951](#)

## Infrastructure

- On QFX5100 switches, Ethernet VPN (EVPN) routes from compute nodes can be withdrawn when no change has taken place on either the compute node or the QFX5100 switch. [PR1106510](#)
- On a QFX5100 switch that has performed a pseudowire switchover, traffic might drop for 10 seconds immediately after the switchover. [PR1049606](#)

## Interfaces and Chassis

- On a QFX5100 switch, high ICMP delays are experienced when pinging directly connected integrated routing and bridging (IRB) interfaces. [PR966905](#)
- On a QFX5100 switch, if you configure MC-LAG, IRB mac sync, and LACP force up, the number of packets received (rx) might be twice the amount sent (tx) from the customer edge to the core. [PR1015655](#)

- If a transceiver is removed from a port on a QFX5100, QFX3600, or EX4300 switch within 30 seconds of converting the port into a Virtual Chassis port (VCP), the port might not get initialized as a VCP. [PR1029829](#)
- QFX5100-48T 10gbase-T copper ports support 100m speed in both autonegotiation and force speed mode. To configure 100m speed in autonegotiation mode, you must configure 100m speed on the peer. To configure 100m speed in force speed mode, you must delete the interface by using the **delete interface ether-options auto-negotiation configuration** command. Deleting the interface sets the port into force speed mode, in which 10gbase-T copper ports support only 100m speed. Therefore, the port speed mode is automatically set to 100m. That is, you do not need to explicitly configure the speed as 100m in force speed mode. [PR1044860](#)



**NOTE:** In Junos OS Release 14.1X53-D15, the **ether-options** statement is not available in the **[edit interfaces *interface-name*]** hierarchy.

- When an EX4600 or a QFX5100 switch is downgraded from Junos OS Release 14.1X53-D15 or later to Junos OS Release 14.1X53-D10 or earlier, the 40-Gbps Ethernet interfaces on QSFP+ transceivers might not return to the UP state. As a workaround, power cycle the switch after the Junos OS upgrade. [PR1061213](#)
- On 40-gigabit links between EX4300 and QFX5100 switches, you must disable **auto-negotiation** on both ends of the link for the interfaces to remain up. On each switch, issue the **set interface et-x/y/z ether-options no-auto-negotiation** command. Also, because **auto-negotiation** is disabled, you must also explicitly configure the **link-mode** and **speed** options on those interfaces. [PR1118318](#)
- In a Q-in-Q tunneling configuration on a QFX5100 switch, if you configure a VLAN ID list on an ingress customer-VLAN (C-VLAN) interface but only configure a VLAN ID without **vlan-id-list** on the egress C-VLAN interface, packets that are sent to the egress C-VLAN interface might be dropped. As a workaround, configure the same VLAN ID list on both the ingress and the egress C-VLAN interfaces. [PR1216732](#)
- In a Q-in-Q tunneling configuration on a QFX5100 switch that is running under Junos OS Release 14.1X53-D40, if you configure a VLAN ID on the egress UNI interface that is the same as the SVLAN ID, and if the **vlan-id-list** statement is not configured on the logical interface on that UNI interface, Q-in-Q packets might be forwarded out with dual tags after they exit from the UNI interface. We recommend that you always include **vlan-id-list** in the Q-in-Q configuration.
- On QFX5100 switches, if the loopback interface does not have an accept term for the management IP address out of the box, all traffic on the management port will be dropped. You must configure an explicit accept term on the loopback filter to "allow all traffic" to pass through the management interface. [PR1215384](#)



---

## Layer 2 Features

---

- On a QFabric system, system log messages might be flooded during the mapping of interfaces to VLANs. You can ignore these system log messages. [PR1200853](#)

## Layer 3 Protocols

---

- On QFX5100 switches, if you use the **next-table** statement in the configuration of a static route that is part of a virtual routing instance, the switch does not forward ICMP packets destined to a route that is present in the inet.0 routing table. [PR970895](#)

## Layer 3 VPNs

---

- In a Layer 3 VPN, if IRB is used between the penultimate hop and the PE node, you cannot check VRF connectivity using PE to PE ping. Pinging to the PE loopback address or interface IP address from the remote PE does not work. As a workaround, use CE to CE ping to verify VRF connectivity. [PR1211462](#)

## MPLS

---

- If a link failure occurs when multiple LSPs are using a link-protected, fast-rerouted link, the convergence time is proportional to the number of LSPs sharing the protected link. [PR1015806](#)
- In a scaled configuration for MPLS FRR and L2 circuit, the convergence time for FRR might increase. For L2 circuit, there might be packet drops. [PR1016146](#)
- A pseudowire is a port-based Layer 2 circuit that emulates a service over a packet switched network (PSN). You can emulate any circuit end to end using a pseudowire. In the event of a link failure on a transit router that hosts a Layer 2 circuit over an RSVP tunnel, the traffic convergence time is approximately 350 milliseconds for a single pseudowire. [PR1016992](#)
- On a QFX5100 switch, if an MPLS link is in hot standby mode and a pseudowire switchover is triggered by the event **remote site local interface signaled down**, traffic flowing through the pseudowire is dropped. [PR1027755](#)
- On a QFX5100 switch that is using the Ethernet tagged mode of operation on a pseudowire, L2 control protocols can fail to come up between customer edge devices (CEs) across the pseudowire. This issue is not seen when the pseudowire mode of operation is Ethernet raw mode. [PR1028537](#)
- For an L2 circuit on QFX5100 switches, when IS-IS is used as an IGP between CEs connected to an L2 circuit, the CEs fail to form an IS-IS adjacency over the pseudowire. As a workaround, consider using an alternative IGP protocol, such as OSPF. [PR1032007](#)
- On a QFX5100 switch, the enhanced hash key does not work for MPLS-IP packets. [PR1095136](#)
- On QFX5100 switches, MPLS ECMP with penultimate hop popping (PHP) does not work with single labels. [PR1212113](#)

### Multicast Protocols

---

- When an IGMP leave is sent from a host to a QFX5100 switch, one packet per multicast group is dropped during route programming. [PR995331](#)
- On a QFabric system, when you use the **set interfaces vlan disable** command to disable VLAN interfaces that are running multicast at Layer 3, wait up to 4 minutes to enable the VLAN interfaces. Waiting to reenabling the interfaces provides the time necessary to clear Layer 3 multicast groups, preventing the loss of multicast traffic. [PR1194388](#)
- On a QFabric system, rate-limiting is enabled by default for multicast traffic. [PR1198892](#)

### Network Management and Monitoring

---

- This issue applies to the Cloud Analytics Engine feature. The Compute Agent Web API does not provide the option to configure the VXLAN destination port. [PR1036372](#)
- This issue applies to the Cloud Analytics Engine feature. Compute Agent CPU utilization goes up when 50 flows are sent. As a workaround, use staggered probe initiation or lower the number of probes per second. [PR1041190](#)
- This issue applies to the Cloud Analytics Engine feature. The Cloud Analytics Engine Compute Agent process (cagent) does not start after the server is rebooted. As a workaround, restart the cagent process manually or with customer automation framework scripts. [PR1041931](#)
- On a QFX5100 switch, the DHCP relay bindings of clients bound via secondary addresses may get cleared when the primary address on the gateway interface configured as DHCP Smart Relay is modified or deactivated. [PR1084911](#)
- On a QFabric system, an SNMP query might not complete because it is querying a nonexistent node-group. This node-group is present in the QFabric configuration but the corresponding node-device has been removed or powered off.
  1. From the shell, issue **mysql -uroot -ppassword sfcdb -hdb -e "Call GetSNMPFabricChassisINES();" to list devices in the database.**
  2. Determine if there are devices listed that are not supposed to be there.
  3. Delete these devices' configurations from the CLI.
  4. From the shell, confirm removal of the devices using the command in Step 1.

### [PR1235326](#)

- After upgrading from Junos OS Release 12.2 to Junos OS Release 14.1X53-D40 in a QFabric system, queries for jnxFabricOperatingTable entries in the enterprise-specific SNMP Fabric Chassis MIB might not return results for the two Director groups DG0 and DG1. You can follow these steps to work around this issue:
  1. Remove the file **/opt/dcf/config/add\_director** from Director groups DG0 and DG1.

2. Run the script `/opt/DCF/scripts/move_cores.sh` on Director group DG0 and on Director group DG1. This script is available by default on switches running Junos OS and used at QFabric system setup time to add devices in the system to the MIB database; running it after a Junos OS upgrade adds DG0 and DG1 if they were not already present.
3. Confirm the Director group entries are present using the **show snmp mib walk ... mib-object** CLI command or other SNMP MIB query tools to retrieve `jnxFabricOperatingTable` entries.

[PR1214737](#)

## OVSDB

- On QFX5100 switches, the amount of time that it takes for other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs) to learn a new MAC address after the first packet is sent from this MAC address is a maximum of 4.5 seconds. (The amount of time depends upon the server configuration on which VMware NSX is running.) During this time, traffic destined for this MAC address is flooded into the VXLAN. [PR962945](#)
- After the connections with NSX controllers are disabled on a Juniper Networks device, interfaces that were configured to be managed by OVSDB continue passing traffic. [PR980577](#)
- QFX5100 switches do not support multiple service nodes for the handling of Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic within an OVSDB-managed VXLAN. [PR985872](#)
- If an entity with a particular MAC address is moved so that its traffic is handled by a different Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP), this MAC address is not learned by entities served by the new hardware VTEP until the hardware VTEP that previously handled its traffic ages out the MAC address. During this transitional period, traffic destined for this MAC address is dropped. [PR988270](#)
- On QFX5100 switches, an NSX controller occasionally overrides an existing local MAC with a remote MAC of the same address. If the Junos OS hardware VTEP detects such a condition (that is, it receives a remote MAC from the NSX controller that conflicts (matches) with an existing local MAC), the hardware VTEP in a Junos OS network accepts the remote MAC and stops publishing the local MAC to the NSX controller. [PR991553](#)
- On QFX5100 switches, an active path in the OVSDB overlay, which you can view by using the **show ovssdb mac** operational command, does not always match the active path in the Layer 3 network underlay, which you can view by using the **show route** operational command. [PR1015998](#)
- On QFX5100 switches, in NSX Manager, when a logical switch is deleted, the corresponding VXLAN on a QFX5100 switch might not be automatically deleted and might still appear in the output of the **show vlans** command. [PR1024169](#)

- On a QFX5100 switch on which OVSDB-managed interfaces are automatically configured, if you delete the configuration of one or more of the interfaces from the switch using the **delete vlans interfaces** command, the interfaces will be automatically reconfigured per the logical switch, gateway service, and logical switch port configurations that still reside in NSX. Despite the automatic reconfiguration of the OVSDB-managed interfaces, an 8- to 12-second loss of traffic might occur. This loss is because local MAC addresses learned by the interfaces and port-to-logical switch bindings were cleared when the interfaces were deleted and must be re-learned after the interfaces are up and running again. [PR1069889](#)
- If an NSX or Contrail controller pushes a large logical-system configuration to a QFX5100 switch, the existing Bidirectional Forwarding Detection (BFD) sessions with aggressive timers might flap. As a workaround, configure the BFD timer to be at least 1 second. [PR1084780](#)
- On QFX5100, during GRES, FPC reboot, or OVSDB server restarts, an OVSDB session flaps between the TOR and the TOR agent, causing approximately 30 seconds of BUM traffic loss. [PR1254188](#)

---

### Platform and Infrastructure

- On a QFX5100 switch, Bidirectional Forwarding Detection (BFD) sessions might continuously switch from on to off for several minutes after an in-service software upgrade (ISSU). [PR980476](#)

---

### QFabric Systems

- On a QFabric system, if the fabric control Routing Engines are not load balanced, when you request a "component all" style software upgrade, the upgrade fails. [PR892310](#)
- On a QFabric system, when you configure an alias for a Node device or an Interconnect device, use that alias when you configure a flow group. [PR1032693](#)
- In a QFabric system, traffic loss is expected if the master of an RSNG is rebooted. Issue a switchover to make the device the backup of the RSNG before rebooting the device. [PR1229949](#)
- On QFabric systems, during the reboot of the CPE primary device, there could be packet loss of control packets for a few seconds. The control packets have provisioning for retransmission, so protocols may not be impacted as long as the protocol **dead-interval** or **hold-interval** is more than 10 seconds during the CPE primary reboot. The data traffic is not impacted by the CPE primary reboot. [PR1252908](#)
- SNMP traps sent out from a QFabric system have an additional object **jnxQFabricEventSource** (that is, 1.3.6.1.4.1.2636.3.42.1.1.1) to identify the component from which the SNMP trap was generated such as node-group, interconnect device name, and so on. This object is not part of the MIB definitions. As a workaround, program NMS Tools to handle this additional object and correctly identify the source of the trap. [PR1269928](#)
- On QFabric systems, an incorrect alarm is displayed when the fan tray is removed—an incorrect FPC slot value is displayed in the alarm. [PR1273894](#)

---

## Routing Policy and Firewall Filters

---

- On QFX Series Virtual Chassis, packets that are generated in the CPU and exit from a non-master FPC port might be subjected to an egress port-based firewall filter (PACL) and be egress filtered, while packets that exit from a master FPC port might not be egress filtered. [PR923659](#)

---

## Routing Protocols

---

- The device does not properly process a Neighbor Solicitation sent to its Subnet-Router Anycast address. [PR693235](#)

---

## Security

---

- The following control packets share the same policer (burst and bandwidth) in hardware, so changing one in the DDoS protection CLI also changes the DDoS parameter for other protocols:
  - STP, PVSTP, and LLDP share DDoS parameters
  - l3mtu-fail, TTL, and ip-opt share DDoS parameters
  - RSVP, LDP, and BGP share DDoS parameters
  - unknown-l2mc, RIP, and OSPF share DDoS parameters

[PR1211911](#)

---

## Software Installation and Upgrade

---

- On a QFX5100 switch, system logs might not be retained after a unified in-service software upgrade (ISSU), due to the data disk being reformatted during the ISSU. [PR964950](#)
- On QFX5100 switches, if a port mirroring analyzer is configured with a VLAN input and you perform an ISSU, the analyzer state is restored after the upgrade. If you later delete the analyzer configuration, mirroring stops but there might be harmless stale entries in the hardware. [PR970011](#)
- On QFX3500 and QFX5100 switches, the amount of time that it takes for Zero Touch Provisioning to complete might be lengthy because TFTP might take a long time to fetch required data. [PR980530](#)
- On a QFabric system, during an NSSU upgrade from Junos OS Release 13.2X52 to 14.1X53-D40, traffic loss may be observed during RSNG upgrade. [PR1207804](#)
- On QFX5100, ISSU from Junos OS Release 14.1X53-D30 to 14.1X53-D40 is not supported. As a workaround, perform a regular software upgrade by downloading the new software version and rebooting the switch during a maintenance window. [PR12209272](#)

## System Management

---

- On the QFX Series, using the **set license keys key** and **request system license add terminal** commands together does not work for installing and deleting licenses. For example, if you install the license using the **set license keys key**, use the **delete license keys** command to delete the license. This is true for the **request system license** commands. [PR1023672](#)

## Storage and Fibre Channel

---

- Each Fibre Channel fabric on an FCoE-FC gateway supports a maximum of four Fibre Channel over Ethernet (FCoE) VLAN interfaces.
- The maximum number of logins for each FCoE node (ENode) is in the range of 32 through 2500. (Each ENode can log in to a particular fabric up to the maximum number of configured times. The maximum number of logins is per fabric, so an ENode can log in to more than one fabric and have its configured maximum number of logins on each fabric.)
- The maximum number of FCoE sessions for the switch, which equals the total number of fabric login (FLOGI) sessions plus the total number of fabric discovery (FDISC) sessions, is 2500.
- The maximum number of FIP snooping sessions per QFX3500 switch is 2500.
- When you configure FIP snooping filters, if the filters consume more space than is available in the ternary content-addressable memory (TCAM), the configuration commit operation succeeds even though the filters are not actually implemented in the configuration. Because the commit operation checks syntax but does not check available resources, it appears as if the FIP snooping filters are configured, but they are not. The only indication of this issue is that the switch generates a system log message that the TCAM is full. You must check the system log to find out if a TCAM full message has been logged if you suspect that the filters have not been implemented.
- You cannot use a fixed classifier to map FCoE traffic to an Ethernet interface. The FCoE application type, length, and value (TLV) carries the FCoE priority-based flow control (PFC) information when you use an explicit IEEE 802.1p classifier to map FCoE traffic to an Ethernet interface. You cannot use a fixed classifier to map FCoE traffic to an Ethernet interface because untagged traffic is classified in the FCoE forwarding class, but FCoE traffic must have a priority tag (FCoE traffic cannot be untagged).

For example, the following behavior aggregate classifier configuration is supported:

```
[edit class-of-service]
```

```
user@switch# set congestion notification profile fcoe-cnp input ieee-802.1 code-point 011 pfc
```

```
user@switch# set interfaces xe-0/0/24 unit 0 classifiers ieee-802.1 fcoe
```

For example, the following fixed classifier configuration is not supported:

```
[edit class-of-service]
```

---

```
user@switch# set interfaces xe-0/0/24 unit 0 forwarding-class fcoe
```

- On a QFX Series device, a DCBX interoperability issue between 10-Gigabit Ethernet interfaces on QFX Series devices and 10-Gigabit Ethernet interfaces on another vendor's devices can prevent the two interfaces from performing DCBX negotiation successfully in the following scenario:
  1. On a QFX Series 10-Gigabit Ethernet interface, LLDP is running, but DCBX is disabled.
  2. On another vendor's device 10-Gigabit Ethernet interface, both LLDP and DCBX are running, but the interface is administratively down.
  3. When you bring another vendor's 10-Gigabit Ethernet interface up by issuing the **no shutdown** command, the device sends DCBX 1.01 (CEE) TLVs, but receives no acknowledge (ACK) message from the QFX Series device, because DCBX is not enabled on the QFX Series device. After a few tries, another vendor's device sends DCBX 1.00 (CIN) TLVs, and again receive no ACK messages from the QFX Series device.
  4. Enable DCBX on the QFX Series 10-Gigabit Ethernet interface. The interface sends DCBX 1.01 (CEE) TLVs, but the other vendor's device ignores them and replies with DCBX 1.00 (CIN) TLVs. The other vendor's device does not attempt to send or acknowledge DCBX 1.01 TLVs, only DCBX 1.00 TLVs.

In this case, the QFX Series device ignores the DCBX 1.00 (CIN) TLVs because the QFX Series does not support DCBX 1.00 (the QFX Series supports DCBX 1.01 and IEEE DCBX). The result is that the DCBX capabilities negotiation between the two interfaces fails.

### Traffic Management

---

- On a QFX5100 switch, running **tcpdump** on the console might cause system instability or cause protocols such as STP or LACP to fail. [PR932592](#)
- CoS on Virtual Chassis access interfaces is the same as CoS on QFX Series access interfaces with the exception of shared buffer settings. All of the documentation for QFX Series CoS on access interfaces applies to Virtual Chassis access interfaces.

Virtual Chassis access interfaces support the following CoS features:

- Forwarding classes—The default forwarding classes, queue mapping, and packet drop attributes are the same as on QFX Series access interfaces:

Default Forwarding Class	Default Queue Mapping	Default Packet Drop Attribute
best-effort (be)	0	drop
fcoe	3	no-loss
no-loss	4	no-loss
network-control (nc)	7	drop
mcast	8	drop

- Packet classification—Classifier default settings and configuration are the same as on QFX Series access interfaces. Support for behavior aggregate, multifield, multidestination, and fixed classifiers is the same as on QFX Series access interfaces.
- Enhanced transmission selection (ETS)—This data center bridging (DCB) feature that supports hierarchical scheduling has the same defaults and user configuration as on QFX Series access interfaces, including forwarding class set (priority group) and traffic control profile configuration.
- Priority-based flow control (PFC)—This DCB feature that supports lossless transport has the same defaults and user configuration as on QFX Series access interfaces, including support for six lossless priorities (forwarding classes).
- Ethernet PAUSE—Same defaults and configuration as on QFX Series access interfaces.
- Queue scheduling—Same defaults, configuration, and scheduler-to-forwarding-class mapping as on QFX Series access interfaces. Queue scheduling is a subset of hierarchical scheduling.
- Priority group (forwarding class set) scheduling—Same defaults and configuration as on QFX Series access interfaces. Priority group scheduling is a subset of hierarchical scheduling.
- Tail-drop profiles—Same defaults and configuration as on QFX Series access interfaces.
- Code-point aliases—Same defaults and configuration as on QFX Series access interfaces.
- Rewrite rules—As on the QFX Series access interfaces, there are no default rewrite rules applied to egress traffic.
- Host outbound traffic—Same defaults and configuration as on QFX Series access interfaces.



The default shared buffer settings and shared buffer configuration are also the same as on QFX Series access interfaces, except that the shared buffer configuration is global and applies to all access ports on all members of the Virtual Chassis. You cannot configure different shared buffer settings for different Virtual Chassis members.

- **Similarities in CoS support on VCP interfaces and QFabric system Node device fabric interfaces**—VCP interfaces support full hierarchical scheduling (ETS). ETS includes:
  - Creating forwarding class sets (priority groups) and mapping forwarding classes to forwarding class sets.
  - Scheduling for individual output queues. The scheduler defaults and configuration are the same as the scheduler on access interfaces.
  - Scheduling for priority groups (forwarding class sets) using a traffic control profile. The defaults and configuration are the same as on access interfaces.
  - No other CoS features are supported on VCP interfaces.



**NOTE:** You cannot attach classifiers, congestion notification profiles, or rewrite rules to VCP interfaces. Also, you cannot configure buffer settings on VCP interfaces. Similar to QFabric system Node device fabric interfaces, you can only attach forwarding class sets and traffic control profiles to VCP interfaces.

The behavior of lossless traffic across 40-Gigabit VCP interfaces is the same as the behavior of lossless traffic across QFabric system Node device fabric ports. Flow control for lossless forwarding classes (priorities) is enabled automatically. The system dynamically calculates buffer headroom that is allocated from the global lossless headroom buffer for the lossless forwarding classes on each 40-Gigabit VCP interface. If there is not enough global headroom buffer space to support the number of lossless flows on a 40-Gigabit VCP interface, the system generates a syslog message.



**NOTE:** After you configure lossless transport on a Virtual Chassis, check the syslog messages to ensure that there is sufficient buffer space to support the configuration.



**NOTE:** If you break out a 40-Gigabit VCP interface into 10-Gigabit VCP interfaces, lossless transport is not supported on the 10-Gigabit VCP interfaces. Lossless transport is supported only on 40-Gigabit VCP interfaces.

- **Differences in CoS support on VCP interfaces and QFabric system Node device fabric interfaces**—Although most of the CoS behavior on VCP interfaces is similar to CoS behavior on QFabric system Node device fabric ports, there are some important differences:

- Hierarchical scheduling (queue and priority group scheduling)—On QFabric system Node device fabric interfaces, you can apply a different hierarchical scheduler (traffic control profile) to different priority groups (forwarding class sets) on different interfaces. However, on VCP interfaces, the schedulers you apply to priority groups are global to all VCP interfaces. One hierarchical scheduler controls scheduling for a priority group on all VCP interfaces.

You attach a scheduler to VCP interfaces using the global identifier (*vcp-\**) for VCP interfaces. For example, if you want to apply a traffic control profile (which contains both queue and priority group scheduling configuration) named *vcp-fcoe-tcp* to a forwarding class set named *vcp-fcoe-fcset*, you include the following statement in the configuration:

```
[edit]
user@switch# set class-of-service interfaces vcp-* forwarding-class-set vcp-fcoe-fcset
output-traffic-control-profile vcp-fcoe-tcp
```

The system applies the hierarchical scheduler *vcp-fcoe-tcp* to the traffic mapped to the priority group *vcp-fcoe-fcset* on all VCP interfaces.

- You cannot attach classifiers, congestion notification profiles, or rewrite rules to VCP interfaces. Also, you cannot configure buffer settings on VCP interfaces. Similar to QFabric system Node device fabric interfaces, you can only attach forwarding class sets and traffic control profiles to VCP interfaces.
- Lossless transport is supported only on 40-Gigabit VCP interfaces. If you break out a 40-Gigabit VCP interface into 10-Gigabit VCP interfaces, lossless transport is not supported on the 10-Gigabit VCP interfaces.
- 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 interface 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.

### Virtual Chassis and Virtual Chassis Fabric

- On a mixed-mode Virtual Chassis Fabric, during a Routing Engine switchover, the system might experience a 200-300 millisecond loss of traffic. [PR964987](#)
- On a mixed Virtual Chassis Fabric (VCF), control plane packets, including control packets for OSPF or PIM, are not mirrored by the native analyzer when the output port belongs to another member switch. [PR969542](#)
- If a VCF is connected to a Juniper Networks router with a flexible PIC concentrator (FPC) and an xSTP bridge protocol data unit is distributed to the FPC, there might be traffic loss when the FPC is rebooted. [PR990247](#)
- When a Virtual Chassis port (VCP) is added between two QFX5100 member switches that are already interconnected using a VCP, a VCP link aggregation group (LAG) is formed and some multicast packets between the two member switches might be duplicated. [PR1007204](#)
- On a mixed-mode VCF, if you perform a nonstop software upgrade (NSSU) and a MAC address is present on the ingress or egress Packet Forwarding Engine, in some cases known Layer 2 unicast traffic might still be flooded over the VLAN. [PR1013416](#)
- On QFX5100 mixed-mode Virtual Chassis or Virtual Chassis Fabric (VCF) systems that include QFX3500 or QFX3600 switches, a MACsec configuration cannot be committed because MACsec is not supported on QFX3500 and QFX3600. [PR1024921](#)
- On a mixed Virtual Chassis Fabric (VCF), a VCP link between two members disappears after you perform a nonstop software upgrade. The **show virtual-chassis protocol adjacency member** command output shows the state of the VCP link as **initializing**. [PR1031296](#)
- In a mixed-mode Virtual Chassis with QFX3500 switches, if multicast packets are sent to the Routing Engine at a very high rate, the Virtual Chassis might become unresponsive. [PR1117133](#)
- In a large-scale Virtual Chassis Fabric (VCF), you might see timeout errors when running the command **request system reboot all-members at now** to immediately reboot all members of the VCF. As a workaround, omit the **at now** option, and run the command simply as **request system reboot all-members**. [PR1215130](#)

## VXLAN

---

- On a QFX5100 switch with a VXLAN configured, (S,G) interface entries downstream from a VXLAN interface might be missing from the multicast routing table but be present in the kernel and packet forwarding engine. In this circumstance, traffic is forwarded as expected. [PR1027119](#)
- VXLANs with the VLAN IDs of 1 and 2 are configured on a QFX5100 switch. The replicated packets for these VXLANs should include the VLAN tags of 1 or 2, respectively. Instead, the replicated packets for these VXLANs are untagged, which might result in the packets being dropped by a Juniper Networks device that receives the packets. To avoid this situation, when configuring a VXLAN on a QFX5100 switch, we recommend using a VLAN ID of 3 or higher. [PR1072090](#)
- QFX5100 switches do not support ingress VLAN firewall flood filters. If you configure such a filter by issuing the **set vlans forwarding-options flood input** command on a QFX5100 switch, the filter is implemented on egress traffic instead of on ingress traffic, which causes unexpected results. The unexpected results especially impact packets in which a VLAN header is added or removed in egress traffic, for example, IRB traffic and VXLAN traffic. As a workaround for these types of traffic, we recommend applying a filter policy on the ingress VLAN traffic and not using the **flood** keyword in the command that you issue. [PR1166200](#)
- QFX5100 switches do not support ingress VLAN firewall flood filters. If you configure such as a filter by issuing the **set vlans forwarding-options flood input** command on a QFX5100 switch, the filter is implemented on egress traffic instead of on ingress traffic, which causes unexpected results. Further, if the filter includes **policer** as the action, the rate at which traffic is flooded to egress interfaces is reduced by a factor of the number of egress interfaces with respect to the committed information rate (CIR). For example, if the CIR is 4g and the number of egress interfaces in the VLAN is 2, the amount of traffic flooded to each interface is reduced to approximately half of the CIR traffic ( $4g/2 = 2g$ ). Similarly, if the CIR is 4g and the number of egress interfaces in the VLAN is 4, the amount of traffic flooded to each interface is reduced to approximately a quarter of the CIR traffic ( $4g/4 = 1g$ ). [PR1166439](#)
- QFX5100 switches do not support ingress VLAN firewall flood filters. If you configure such as a filter by issuing the **set vlans forwarding-options flood input** command and specify **policer** as the action on a QFX5100 switch, the filter is implemented on egress traffic instead of on ingress traffic, which causes unexpected results, especially for integrated routing and bridging (IRB) traffic or VXLAN traffic. For example, in the case of Layer 2 traffic intended for VLAN 101 and temporarily encapsulated with a VLAN header (VLAN 100), such a filter applied to VLAN 100 might result in the ingress interfaces in VLAN 101 being flooded by traffic intended for VLAN 100. Further, in the case of routing traffic between VLANs, traffic intended for VLAN 101 might be routed to the IRB interface associated with VLAN 100, or in the case of VXLAN traffic, to a virtual tunnel endpoint (VTEP) on which VLAN 100 is configured. [PR1168777](#)
- While setting up multihoming active-active mode on a link aggregation group (LAG) interface on a QFX5100 standalone switch or QFX5100 Virtual Chassis in an EVPN-VXLAN topology, check to see if the LAG interface on which you are configuring an Ethernet segment identifier (ESI) is already configured for flexible VLAN tagging. If

so, all logical subinterfaces associated with this interface must be activated. If one logical subinterface is deactivated, unexpected behavior related to the ESI and designated forwarder (DF) election might occur. [PR1189830](#)

#### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 34](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

## Known Issues

The following issues are outstanding in Junos OS Release 14.1X53 for the QFX Series. 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>.

- [High Availability](#)
- [Interfaces and Chassis](#)
- [MPLS](#)
- [Multicast Protocols](#)
- [Routing Protocols](#)
- [QFabric Systems](#)
- [VXLAN](#)

### High Availability

- On a Virtual Chassis Fabric, NSSU might fail in the following situations: there is routed traffic with no end hosts, or hosts are missing because of an NSSU node reboot. [PR1199744](#)
- On a QFX5100 switch, you cannot perform a unified in-service software upgrade (ISSU) from Junos OS Release 14.1X53-D30 to Junos OS Release 14.1X53-D40. As a workaround, during a maintenance window, download the new software version, perform a regular software upgrade, and reboot the switch. [PR1229272](#)

## Interfaces and Chassis

---

- On a QFX5100 switch, with a fully meshed MC-LAG topology configured, sometimes there is more traffic loss when the ICL interface goes down and then back up compared to when you have Junos OS Release 14.1X53-D35 software installed. The root cause has been identified, and this issue does not affect MC-LAG functionality. [PR1209322](#)

## MPLS

---

- On QFX5100 switches, the **traceroute mpls ldp** command output shows incorrect information when using an IRB interface between the ingress provider edge (PE) switch and the provider (P) switch. This occurs when running LDP over RSVP over an MPLS core network. [PR1217132](#)

## Multicast Protocols

---

- When a static multicast route with a next-table nexthop is changed from a table that cannot forward the traffic to one that can and then reverted back to the original table, the traffic might continue to flow out the downstream interface even though the static route is no longer pointing to the table that allowed for the traffic increase. [PR1217958](#)
- On QFX5100 switches, multicast route leaking does not support a Layer 3 interface (IPv4) as an upstream port. As a workaround, use an integrated routing and bridging (IRB) interface. [PR1250430](#)

## Routing Protocols

---

- On QFX5100, TCP packets with dest ipv6 as link-local address and dest port 179 are dropped in the Packet Forwarding Engine. [PR1267565](#)

## QFabric Systems

---

- On QFabric systems, excess mount points might be created on the DRE that could result in multiple TCP connections being opened between the DRE and the NFS server. As a workaround, reboot the DRE. [PR1259008](#)
- The QFabric director retrieves syslogs and SNMP traps from different components—such as node-groups, node-devices, and interconnects—and logs them in the **/tmp/sfc-captures/misc** directory. Over a period of time, this can consume a large amount of disk space, as these logs are not purged. [PR1272190](#)
- On QFabric systems, if you delete a Layer 3 interface without deleting the corresponding VLAN, then there will be multicast traffic loss on that VLAN. [PR1272596](#)
- The ccif service in the QFabric directors is responsible for spawning VMs (like Fabric Manager, Network Node Group, and so on) and load-balancing them between the two Director Groups. The ccif service uses a GFS-based shared-disk file system to share data across the two Director Groups. The ccif service is managed by the RedHat Cluster Management. Due to hard-disk errors, the ccif service could encounter I/O errors while

accessing the shared file system. When this happens, the Cluster Manager moves the ccif service to the Failed state. [PR1277504](#)

- On QFabric systems, if IGMP snooping is enabled, unknown multicast traffic might be lost after an NNG switchover during a DG upgrade, due to misprogramming of mrouter ports in VLANs. As a workaround, issue the **clear igmp-snooping membership** command. [PR1280110](#)

## VXLAN

- On QFX5100 switches, the Layer 3 routes that form virtual extensible LAN (VXLAN) tunnels use per-packet load balancing by default, which means that load balancing is implemented if there are ECMP paths to the remote tunnel endpoint. This is different from normal routing behavior in which per-packet load balancing is not used by default. (Normal routing uses per-prefix load balancing by default.) [PR1018814](#)
- On QFX5100 switches when VXLAN is configured, a VXLAN table is created to resolve routes to remote VTEPs. If the underlay is OSPF IS-IS or EBGp, the routes can distribute the traffic over multiple paths if load balancing is configured. However, if the underlay is IBGP, the route will select one of the available paths rather than using all the available paths. [PR1154961](#)

### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 34](#)
- [Known Behavior on page 38](#)
- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

## Resolved Issues

This section lists the issues fixed in the Junos OS Release 14.1X53 for the QFX Series.

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

- [Resolved Issues: Release 14.1X53-D43 on page 56](#)
- [Resolved Issues: Release 14.1X53-D42 on page 58](#)
- [Resolved Issues: Release 14.1X53-D40 on page 62](#)
- [Resolved Issues: Release 14.1X53-D35 on page 63](#)
- [Resolved Issues: Release 14.1X53-D30 on page 65](#)
- [Resolved Issues: Release 14.1X53-D27 on page 65](#)
- [Resolved Issues: Release 14.1X53-D26 on page 66](#)
- [Resolved Issues: Release 14.1X53-D25 on page 66](#)

- [Resolved Issues: Release 14.1X53-D16 on page 67](#)
- [Resolved Issues: Resolved Before Release 14.1X53-D16 on page 69](#)

#### **Resolved Issues: Release 14.1X53-D43**

---

- [Class of Service \(CoS\)](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [MPLS](#)
- [Multicast Protocols](#)
- [Network Management and Monitoring](#)
- [Platform and Chassis](#)
- [Port Security](#)
- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Spanning-Tree Protocols](#)
- [Virtual Chassis and Virtual Chassis Fabric](#)
- [VLAN Infrastructure](#)

##### ***Class of Service (CoS)***

- On QFX5100, EX4300, or EX4600, traffic might be dropped when there is more than one **forwarding-classes** under **forwarding-class-sets**. [PR1255077](#)

##### ***Infrastructure***

- On QFX5100 and EX4600, password required for user **root** even after SSH public key authentication is enabled. [PR1234100](#)

##### ***Interfaces and Chassis***

- The AE interface might be down after NSSU is done on QFX5100 or EX4600 switches. [PR1227522](#)
- Incorrect statistics might be shown for an AE interface after rebooting device or clearing interface statistics. [PR1228042](#)
- Wrong value recorded in Resource error counter of QFX Series switch interface. [PR1245748](#)
- SFP-T in QFX5100-48S-6Q does not work at 100M full duplex in Releases 14.1X53-D35 and later (works in 14.1X53-D30). [PR1250453](#)
- On QFX5100, **show interface** incorrectly displays an interface as **Link-mode: Auto Speed: Auto** even though the interface is configured for, and up at, 100M/Full. [PR1260986](#)



### **Layer 2 Features**

- On QFX5100, MAC learning will be very slow when clearing MAC addresses in cases of scale MAC learning (128k). [PR1240114](#)

### **MPLS**

- On QFX3500 Virtual Chassis, traffic loss may be observed when MPLS explicit-null is configured on core routers. [PR1234441](#)
- VC/VCF-l2ckt: FXPC core is seen when deactivating core interface on MPLS l2ckt configuration using IRB interface [PR1242203](#)

### **Multicast Protocols**

- Layer 3 interface (inet family) is not supported as upstream port in multicast route leaking. [PR1250430](#)

### **Network Management and Monitoring**

- 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. [PR1234119](#) , [PR1159544](#)

### **Platform and Chassis**

- Incorrect signedness comparison in the ioctl(2) handler allows a malicious local user to overwrite a portion of the kernel memory. Refer to JSA10784 for more information, <https://kb.juniper.net/JSA10784>. [PR1184592](#)
- QFX Series switches:CSIM:major alarm - 'Management Ethernet 1 Link Down' [PR1228577](#)
- 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](#)

### ***Port Security***

- On QFX3500, DHCP binding may not work when untrusted ARP inspection is enabled in the snooping device. [PR1229399](#)

### ***Routing Policy and Firewall Filters***

- Egress firewall filters containing policers might not process packets correctly if TCAM entries are programmed over multiple slices of TCAM memory space. [PR1232926](#)

### ***Routing Protocols***

- VRRPv2 for IPv4 not working correctly. Router with physical interface higher IPv4 address preempts for mastership in case of a priority tie. [PR1204969](#)

### ***Spanning-Tree Protocols***

- QFX5100 does not transfer BPDU packets though xSTP is disabled. [PR1262847](#)

### ***Virtual Chassis and Virtual Chassis Fabric***

- VCF not communicating properly with backup spine. [PR1141965](#)
- VCF does not forward BUM traffic after **fabric-tree-root** is configured. [PR1257984](#)

### ***VLAN Infrastructure***

- The traffic might not be transmitted correctly after a logical interface is deleted from one VLAN and added to another VLAN on EX9200, EX4300, QFX Series. [PR1228526](#)

## **Resolved Issues: Release 14.1X53-D42**

---

- [Authentication and Access Control](#)
- [Firewall Filters](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [MPLS](#)
- [Multicast Protocols](#)
- [Network Management and Monitoring](#)
- [Virtual Chassis and Virtual Chassis Fabric](#)
- [VXLAN](#)

### ***Authentication and Access Control***

- On QFX Series switches, LLDP does not work on management and internal Ethernet (em) interfaces. [PR1224832](#)

### ***Firewall Filters***

- On QFX5100 switches running 14.1X53-D30.3, when you apply an IPv6 firewall filter, the system might crash with a PFE panic. [PR1234729](#)

### ***High Availability (HA) and Resiliency***

- In a QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF), if the master Routing Engine crashes when nonstop active routing (NSR) is configured and the **[edit system] switchover-on-routing-crash** statement is set, the Virtual Chassis or VCF fails to perform the switchover to the backup Routing Engine. The **switchover-on-routing-crash** statement helps to prevent loss of traffic during a Routing Engine switchover when NSR is enabled by switching over immediately to the backup Routing Engine. [PR1220811](#)
- On QFabric systems during an RSNG NSSU from Junos OS Release 12.2 to Release 14.1X53-D40, traffic might be lost after the mastership switchover and when the old master is rebooted. This happens when the old master has a lower MAC address than that of the backup. [PR1228413](#)
- If you are performing a topology-independent in-service software upgrade (TISSU) from one version of Junos OS Release 14.1X53 to another on a QFX5100 switch and the network analytics feature **[edit service analytics]** is configured, the upgrade might not succeed. In addition, the fxpc process might stop working, and you might notice that a core file is generated. As a workaround, before performing the TISSU, delete the network analytics configuration and commit the configuration. After the TISSU is completed, reconfigure the network analytics feature and commit the configuration. [PR1234945](#)

### ***Infrastructure***

- On EX4300 switches, a message such as **/kernel: %KERN-5: tcp\_timer\_keep: Dropping socket connection due to keepalive timer expiration** might be seen repeatedly. There is no service impact from the condition that causes the message (a Packet Forwarding Engine timeout trying to connect to a daemon that is not active). [PR1209847](#)
- On an EX Series or QFX Series Virtual Chassis, during an upgrade, failover, or switchover operation on the backup Routing Engine member, you might see vmcore and ksyncd core files created and see the log message **/kernel: Nexthop index allocation failed: regular index space exhausted**. [PR1212075](#)
- In an EX4600 Virtual Chassis or an QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF), when using scp on the management interfaces to copy files greater than about 150 MB, you might see protocol flapping and Routing Engine TCP connections dropping. [PR1213286](#)
- On QFX Series switches, the BIOS event log is monitored every 10 minutes, which causes unnecessary overhead and might result in a scheduler slip, which can cause a

delay of approximately 4 to 6 seconds, leading to issues such as LACP or BGP flapping. As a workaround, contact JTAC to disable the scripts manually. [PR1221420](#)

- On QFX5100 switches, an fxpc process might generate a core file. [PR1231071](#)
- In a QFabric system, after a Node device is replaced in a Node group, you might observe issues when running the **file copy** and **request routing-engine login other-routing-engine** commands between redundant server Node group members. As a restoration-only workaround:
  - If the Node with the bad TNP entry is the backup, reboot the backup Node.
  - If the Node with the bad TNP entry is the master, do a switchover, then reboot the new backup member that has the bad TNP entry.

#### [PR1236898](#)

- On a QFX5100 switch, Gratuitous Address Resolution Protocol (GARP) reply packets are not updating the Address Resolution Protocol (ARP) table. GARP request packets, however, are updating the ARP table as expected. [PR1246988](#)

### **Interfaces and Chassis**

- On QFX5100, when resilient hashing is enabled on ECMP paths, flows on other paths should not be rehashed when one path goes down. But for host routes (/32 routes), rehashing might happen in some cases. [PR1137998](#)
- On a QFX3600 Node device in a QFabric system, configuring ports Q0 through Q7 as 40-Gbps data plane uplink (fte) ports in a network Node group might result in misconfiguration of some 10 Gbps (xe) and 40 Gbps (xle) network ports on the Node device, causing those network ports to go down. [PR1208137](#)

### **MPLS**

- On the QFX5100 switch, single label tagged MPLS traffic received from the MPLS core at the PHP node will not get forwarded to the egress ECMP IPv4 next hop. [PR1212519](#)
- On EX Series and QFX Series switches, if you change a Layer 2 circuit configuration from Ethernet CCC encapsulation to VLAN CCC encapsulation, traffic losses might occur at the pseudowire tunnel initiation point. As a workaround, restart the Packet Forwarding Engine on which the problem occurs. [PR1222888](#)

### **Multicast Protocols**

- On QFX5100 switches, multicast statistics are not supported on channelized interfaces and aggregated Ethernet interfaces. [PR1197915](#)
- On EX4300, EX4600, and QFX5100 switches in a Virtual Chassis configuration, IPv6 multicast packets might not be flooded in a VLAN if IGMP snooping is enabled and the ingress interface is on a different FPC than the egress interface. [PR1205416](#)
- On QFabric systems, when you disable and then reenables PIM and graceful restart is configured in a routing instance, the following message occurs: **GR restart pending 1**. After this message is triggered, a PIM neighbor reset or the start of a new multicast

group might result in a PIM join problem state. As a workaround, clear the PIM join.  
[PR1242218](#)

### ***Network Management and Monitoring***

- On a QFabric system, log messages from all components are stored in Director Group devices, which increases the usage of cluster disk space. System logging (syslog) messages are generated by the monitoring cluster area (pbData), but the system does not generate SNMP notification alerts and traps when the pbData area reaches a critical level. [PR1221104](#)
- On QFX5100 switches, an SNMP script that requests a high number of OIDs, for example, 374 or more, might time out. As a workaround, set the timeout value higher than 1 second. [PR1223775](#)

### ***Virtual Chassis and Virtual Chassis Fabric***

- If you change the VNIs of all the VXLANs in a QFX5100 Virtual Chassis Fabric, VXLAN traffic does not converge for some of the VXLANs. [PR1028588](#)
- On a non-mixed Virtual Chassis Fabric (VCF), LACP flaps when the switch in the master Routing Engine role is rebooted using the CLI or because of a power cycle. This issue is not experienced after a Routing Engine switchover. As a workaround, configure a slow LACP timeout. [PR1034377](#)
- When a nonstop software upgrade (NSSU) is performed on 20 members of a Virtual Chassis Fabric (VCF) with a multidimensional scale configuration, members might fail to upgrade, and the message **watchdog: scheduling fairness gone for <nnnn> seconds now** might appear continuously on the console of the master Routing Engine. [PR1142361](#)
- On a Virtual Chassis Fabric (VCF), if three or more members are rebooted simultaneously, there could be traffic loss on the VCF for up to 5 minutes. [PR1146687](#)
- Reconfiguring the longest prefix match (LPM) profile through a CLI command in a scaled Virtual Chassis Fabric (VCF) with a lot of host-bound traffic might cause the VCF to split. To avoid this situation, configure the desired LPM profile at the beginning or restart the VCF manually after the configuration. [PR1152102](#)

### ***VXLAN***

- If you configure a QFX5100 switch to be a VXLAN virtual tunnel endpoint and also configure it to be a PIM RP, the multicast tree does not successfully converge and multicast traffic is dropped. [PR1027159](#)
- If you change the VNIs of all the VXLANs in a QFX5100 Virtual Chassis Fabric, VXLAN traffic does not converge for some of the VXLANs. [PR1028588](#)
- In a large-scale VXLAN configuration, if an aggregated Ethernet Layer 2 interface configured on a QFX5100 switch goes down and then comes back up repeatedly, the VTEP interfaces associated with the MAC addresses learned on the aggregated Ethernet interface could lose traffic for up to 18 seconds. [PR1151456](#)

## Resolved Issues: Release 14.1X53-D40

---

- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)
- [System Management](#)
- [VXLAN](#)

### *High Availability (HA) and Resiliency*

- On a QFX Virtual Chassis, an NSSU from a release between 14.1X53-D30 and 14.1X53-D34 to 14.1X53-D35 might hang and not complete for a Virtual Chassis with five or more devices. [PR1142275](#)

### *Infrastructure*

- On QFX5100 switches in a Virtual Chassis Fabric (VCF), the **clear arp** command does not clear ARP entries for interfaces defined in a routing instance. As a workaround, explicitly specify the interface name for which to clear ARP entries, as follows: **clear arp interface *interface-name***. [PR1159447](#)
- On a QFabric system, system logging (syslog) messages from all components are stored in the MySQL database on the Director. When syslog messages are generated at a high rate, a continuous deadlock might occur from the MySQL server side. Eventually, all incoming syslog insert transactions are kept waiting in the database queue to acquire a lock and expire after 50 seconds, so the syslog messages are not inserted in the database. New syslog messages might not be displayed when you issue the **show log messages** command on the Director. After some time has passed, when the lock is released, the new logs might be seen, even logs that were missing. As a restoration workaround, restart the mysql service on the master DG and wait 15 minutes. Then restart the sfc service on both DGs. [PR1174011](#)

### *Interfaces and Chassis*

- In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. You can, however, enable auto-negotiation by issuing the **set interface *name* ether-options auto-negotiation** command on the interface for which you want to change the interface speed. With autonegotiation enabled, the interface auto-detects the speed in which to operate. [PR1170909](#)
- On a QFabric system, traffic might drop if there is a mismatch in the ordering of the fabric (fte) interface numbers between the Network Node Group (NNG) and the Interconnect (IC) device or if there is a new node addition or an interface ID change caused by any configuration change on the fte interface. As a workaround, correct the ordering of the FTE links between the node and the IC (lower to higher on the node and corresponding lower to higher on the IC). [PR1188574](#)

### **Network Management and Monitoring**

- On a QFX3000-G or QFX3000-M QFabric system, in rare cases, the MySQL DB might be locked, with the result that MySQL and the sfcsnmpd service do not run on Director and any request directed to them does not get a response. SNMP traps and MIB walks might not work as expected. In this problematic situation, the QFabric stops sending SNMP traps to a network management system (NMS), and the NMS cannot get SNMP information from the QFabric. As a restoration workaround, restart the sfcsnmpd process from the Director. [PR1165565](#)

### **System Management**

- On a QFX3500 switch, when you upgrade from Junos OS Release 14.1X53-D35 to Junos OS Release 14.1X53-D40, the ZTP configuration fails. [PR1228814](#)

### **VXLAN**

- On a QFX5100 switch configured with a VXLAN and PIM, the (S,G) route for the VXLAN multicast group can get stuck pointing to the pime interface even though the RP has joined the multicast group. This does not affect traffic forwarding for multicast traffic as the forwarding state for the (S,G) route points correctly to the uplink interface to the RP. [PR1023447](#)

## **Resolved Issues: Release 14.1X53-D35**

- [High Availability](#)
- [Interfaces and Chassis](#)
- [Virtual Chassis and Virtual Chassis Fabric](#)
- [VPNs](#)
- [VXLAN](#)

### **High Availability**

- On a QFX Series switch or Virtual Chassis that is performing a nonstop software upgrade (NSSU) and that has aggregated Ethernet link bundles with member links on multiple switches or line cards, traffic traversing the aggregated Ethernet interface might be lost when the backup Routing Engine reboots as part of the NSSU.  
As a workaround, enter the **clear ethernet-switching table interface aggregated-ethernet-interface-name** command to clear the learned MAC addresses from the Ethernet switching table for the aggregated Ethernet interface. [PR1126855](#)
- On a Virtual Chassis Fabric (VCF), an in-service software upgrade (ISSU) from a release between 14.1X53-D30 and 14.1X53-D34 to 14.1X53-D35 might show traffic loss on ECMP links. As a workaround, follow these two steps:
  - Do an ISSU to a 14.1X53-D32 image.
  - From the 14.1X53-D32 image, do an ISSU to 14.1X53-D35.

[PR1129004](#)

### **Interfaces and Chassis**

- An EX4600-EM-8F expansion module installed in a QFX5100-24Q switch or an EX4600 switch does not support the 100 Mbps speed on the 10-Gigabit Ethernet interfaces. [PR1032257](#)

### **Virtual Chassis and Virtual Chassis Fabric**

- In a mixed QFX3500 and EX4300 Virtual Chassis with a QFX3500 switch acting in the master role, the Virtual Chassis mastership might change when the Virtual Chassis receives multicast traffic. A mixed QFX3500 and EX4300 Virtual Chassis with a QFX3500 switch acting in the master role is not a supported configuration in this release of Junos OS because of this issue. [PR1126216](#)
- On a QFX Series Virtual Chassis Fabric (VCF), rebooting a leaf node might change the size of the VCF, resulting in a flood loop of the unicast or multicast traffic. To fix the issue, use the new CLI statement **fabric-tree-root**. See “Changes in Behavior and Syntax” on page 34 for details. [PR1093988](#)

### **VPNs**

- On a QFX5100 switch using EVPN with VXLAN, the Ethernet segment identifier (ESI) value of the most significant octet (type byte) must be 00 when you are manually configuring an ESI even though the switch accepts other configuration values. [PR1085837](#)
- On a QFX5100 switch using EVPN with VXLAN, the QFX5100 switch encodes the route target extended community value from the second byte of the type 4 ESI value. The IETF standard (draft-ietf-l2vpn-evpn-11) states that the route target extended community value be encoded from the high-order 6-byte portion of the type 4 ESI value. This behavior is also seen on MX Series routers, and does not cause incompatibility problems between QFX Series switches and MX Series routers using EVPN with VXLAN. [PR1085872](#)

### **VXLAN**

- On a QFX5100 switch configured with a VXLAN and PIM, the (S,G) route for the VXLAN multicast group can get stuck pointing to the pime interface even though the rendezvous point (RP) has joined the multicast group. This does not affect traffic forwarding for multicast traffic as the forwarding state for the (S,G) route points correctly to the uplink interface to the RP. [PR1023447](#)
- If you perform a graceful Routing Engine switchover in a Virtual Chassis Fabric acting as a VXLAN virtual tunnel endpoint, known unicast traffic might be dropped from the VXLAN. [PR1026408](#)
- When a Layer 2 interface on a QFX5100 switch is deactivated and reactivated again, an encapsulated VXLAN packet received on a Layer 3 interface on another networking device is sent to the Routing Engine kernel instead of being de-encapsulated and forwarded to the Layer 2 interface. [PR1049752](#)



- On a QFX5100 switch using BGP with VXLAN that has recently deactivated a VRF import policy, all entries in the BGP VXLAN routing table might disappear when routing is reenabled. [PR1084108](#)
- On an aggregated Ethernet OVSDb interface with member links connecting to multiple member switches on a QFX5100 Virtual Chassis, a reboot of one member switch might impact VXLAN traffic encapsulation traversing the member links on other FPCs. As a workaround, disable and reenab the aggregated Ethernet interface to restore the working state. [PR1126915](#)

---

## Resolved Issues: Release 14.1X53-D30

---

- [Interfaces and Chassis](#)
- [OVSDb](#)
- [System Management](#)

### *Interfaces and Chassis*

- In a mixed QFX3500 and EX4300 Virtual Chassis configured for persistent MAC and MAC limiting, traffic is not received on aggregated Ethernet interfaces on EX4300 switches when the EX4300 switches are operating in the linecard role. [PR1033618](#)
- On a QFX5100 switch, if you configure MC-LAG with the **force-up option**, the FXP might create a core file and generate the following error message: **0x0806b7b8 in panic (format\_string=0x9c72614 "Memory corruption in block %p\n" at ../../src/pfe/platform/fxpc/fxpc\_panic.c:93.** [PR1024354](#)
- On a QFX5100 switch, the C0 SFP management port (em0) sometimes fails to come online. [PR1075001](#)

### *OVSDb*

- In an IP fabric using EBGp and VTEP, convergence issues might cause traffic to be dropped for several minutes after BGP is deactivated on a spine device. [PR1091007](#)
- When QFX5100 switches are part of an OVSDb-managed VXLAN, traffic load balancing on a link aggregation group (LAG) does not work over ECMP for Layer 3 VXLAN interfaces. [PR1090791](#)
- On a QFX5100 switch, deleting a VLAN and an interface configuration on an OVSDb-managed interface might cause the switch to reboot. A "Core was generated by `fxpc'" message appears on the console, and a core dump file is created. [PR1091446](#)

### *System Management*

- On EX Series and QFX Series switches that are configured with the **include-option-82 nak** option so that DHCP servers include option 82 information in NAK messages, two copies of option-82 are sometimes appended to DHCP ACK packets. [PR1064969](#)

---

## Resolved Issues: Release 14.1X53-D27

---

No issues that were previously reported in any version of the Junos OS Release 14.1X53 release notes have been resolved in Junos OS Release 14.1X53-D27 for the QFX Series.

## Resolved Issues: Release 14.1X53-D26

---

- [Interfaces and Chassis](#)
- [VXLAN](#)

### *Interfaces and Chassis*

- On EX4600 and QFX5100 switches, the 100Mbps LED functionality is not working. The LED does not glow when 100Mbps traffic is sent or received on the switch, and no output is displayed when the **show chassis led** command is entered to gather information on the 100Mbps interface. [PR1025359](#)

### *VXLAN*

- On a QFX5100 Virtual Chassis, when approximately 3000 Virtual Extensible LANs (VXLANs) are configured and associated with logical interfaces for the same OVSDB-managed interface, a high level of memory usage might occur. As a workaround, disable the 802.1X and multicast snooping processes using the **set system processes dot1x-protocol disable** and **set system processes multicast-snooping disable** statements. [PR1073677](#)

## Resolved Issues: Release 14.1X53-D25

---

- [Class of Service](#)
- [MPLS](#)
- [Network Management and Monitoring](#)
- [Software Installation and Upgrade](#)
- [QFabric System](#)

### *Class of Service*

- On the QFX Series, applying a class-of-service (CoS) configuration globally (using the \* wildcard) to all interfaces on a device can cause inconsistency in the packet forwarding state if the device has interfaces that are members of a link aggregation (LAG) interface bundle and also interfaces that are not members. [PR1001605](#)

### *MPLS*

- Pseudowire emulates a service over packet-switched network (PSN) using only virtual wire. You can emulate any circuit end to end using pseudowire. In the event of a failure, the switchover from the active pseudowire to the standby/backup pseudowire takes longer than expected. [PR1025899](#)
- Pseudowire emulates a service over packet-switched network (PSN) using only virtual wire. You can emulate any circuit end to end using pseudowire. However, the switchover-delay timer for an Active/Standby pseudowire topology might not work as expected. [PR1026336](#)

### ***Network Management and Monitoring***

- This issue applies to the Cloud Analytics Engine feature. On QFX5100-48T-6Q switches, when error counters have been incremented in ERROR\_CNTR TLV, a probe does not recognize those errors. ERROR\_CNTR TLV is not populated in Cloud Analytics Engine probe responses. [PR1034928](#)
- This issue applies to the Cloud Analytics Engine feature. The default aging timer value is 60 seconds for a flow to get aged out. The flow will be marked for aging in the range of 60 to 120 seconds for the default aging timer value from the latest probe time stamp. It will be deleted in next 60 seconds (because the default aging time is 60 seconds). Total time to delete the flow will be a maximum of 3 minutes from the latest flow time stamp for default aging timer value. [PR1037738](#)
- This issue applies to the Cloud Analytics Engine feature. If an invalid Compute Agent IP address is specified in the CA-Discovery file uploaded to the Data Learning Engine, the Data Learning Engine will fail. [PR1041925](#)
- This issue applies to the Cloud Analytics Engine feature. Some application flows might not get registered in the signature database when high rate flow is on. [PR1048083](#)
- This issue applies to the Cloud Analytics Engine feature. Depending on how CentOS installation and configuration is performed, the loopback address might not be configured. This will cause Data Learning Engine components to be inaccessible. There must be at least one entry in the /etc/hosts file for a default loopback IP address mapping to local host. [PR1048890](#)

### ***Software Installation and Upgrade***

- On QFX5100 switches with a large number of firewall terms configured, firewall filters might stop working after you perform an in-service software upgrade (ISSU). [PR966445](#)

### ***QFabric System***

- On a QFabric system, if you include a QFX5100 switch as a Node device in a redundant server Node group, the Master LED light might not turn on or be displayed in the output of the **show chassis led** command. [PR1048853](#)

---

### **Resolved Issues: Release 14.1X53-D16**

- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)
- [Port Security](#)
- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Software Installation and Upgrade](#)
- [Virtual Chassis and Virtual Chassis Fabric](#)
- [VLAN Infrastructure](#)

### ***Interfaces and Chassis***

- On QFX5100 switches with a VXLAN configured, if you add an interface to the VLAN that the VXLAN is associated with or delete an interface from that VLAN, the switch might drop traffic for devices connected to other interfaces in the same VLAN. [PR1019378](#)
- On a QFX Series switch, when you reboot the switch with an enabled 40-Gigabit Ethernet interface, the interface might be disabled after the reboot. As a workaround, remove and then reinsert the attached cable. [PR1014139](#)
- On QFX5100 switches, disabling a member link of an AE interface might cause packets to be sent to a port that is down, which results in traffic loss. As a workaround, to restore service, bring the port that is down back up again. [PR1050260](#)
- On a QFX5100 switch, issuing the **request system reboot** command might not shut down SFP-T interfaces. [PR1050650](#)
- On a QFX5100-48T switch that uses QSFP+ transceivers (QSFP-40G-SR4), if you upgrade the switch software to Junos OS Release 14.1X53-D15, the QSFP+ transceivers might not be detected after the upgrade. [PR1051903](#)
- On a QFX5100 switch, DHCP offer packets with double tags on trunk interfaces might be dropped. [PR1059557](#)
- On a QFX5100-48T-6Q switch, when you configure a Gigabit Ethernet interface on a fiber SFP transceiver with the speed set to 1g and the settings *full duplex* and *no autonegotiation*, the interface goes down. [PR1063118](#)

### ***Network Management and Monitoring***

- This issue applies to the Cloud Analytics Engine feature. The CPU utilization value is incorrect in the Cloud Analytics Engine probe response statistics. [PR1024840](#)
- This issue applies to the Cloud Analytics Engine feature. Packets are not mirrored when a mirror IP address is configured on a remote device. [PR1052028](#)

### **Port Security**

- In a mixed-mode Virtual Chassis Fabric with storm control enabled, if autonegotiation is enabled on a 1-gigabit interface (the default setting), the storm-control value for allowed bandwidth might be set to 0, which would cause traffic to be dropped. As a workaround, manually configure the link speed instead of using autonegotiation. [PR1051756](#)

### **Routing Policy and Firewall Filters**

- If you configure a QFX5100 switch with a firewall filter that redirects traffic to a different interface (by using the **interface** action modifier, rebooting the switch might cause the Packet Forwarding Engine daemon (fxpc) to crash and generate core files. [PR1037563](#)

### **Routing Protocols**

- On a QFX5100 switch, the routing protocol daemon (rpd) might crash and create a core file if there is interior BGP (IBGP) route churn while IBGP multipath is configured and there are multiple levels of IBGP next-hop recursion. [PR1060133](#)

### **Software Installation and Upgrade**

- On a QFX5100 switch, if a port mirroring analyzer is configured with a VLAN input and you perform an in-service software upgrade (ISSU), the analyzer state is restored after the upgrade. If you later delete the analyzer configuration, mirroring stops but there might be residual harmless stale entries in the hardware. [PR970011](#)
- On a QFX5100 switch, if you perform an ISSU, there might be approximately 2 seconds of IPv4 or IPv6 traffic loss during the em0 handoff. [PR985462](#)

### **Virtual Chassis and Virtual Chassis Fabric**

- On a QFX5100 Virtual Chassis Fabric (VCF), routing protocols (for example, OSPF or ISIS) might flap after a master Routing Engine (RE) role switch powers off, causing traffic loss. This issue is not seen on leaf switches. [PR1029066](#)

### **VLAN Infrastructure**

- On QFX Series platforms, naming a VLAN *vlan-rewrite* causes an error when you commit the configuration. [PR1054996](#)

---

## **Resolved Issues: Resolved Before Release 14.1X53-D16**

- [Interfaces and Chassis](#)
- [Layer 3 Protocols](#)
- [OVSDB](#)
- [Software Installation and Upgrade](#)
- [VXLAN](#)

### ***Interfaces and Chassis***

- On QFX5100 switches, traffic might be dropped on a 40G channelized port. [PR1015221](#)
- On a QFX5100 switch, after performing an in-service software upgrade (ISSU), Layer 3 traffic might be interrupted on a configured VLAN or IRB interface. [PR1014130](#)

### ***Layer 3 Protocols***

- On a QFX5100 switch, if you perform an in-service software upgrade on a QFX5100 switch with the virtual routing redundancy protocol (VRRP) configured and there are a large number of VRRP groups or there are many VRRP transitions, you might see duplicate VRRP my\_station\_tcam entries. [PR1028607](#)

### ***OVSDB***

- If you enter a **show configuration** command after installing the OVSDB software package (jsdn-i386-release) on a QFX5100 Virtual Chassis or VCF, you see the warning **ddl\_sequence\_number\_match: sequence numbers don't match**. [PR1019087](#)

### ***Software Installation and Upgrade***

- ISSU does not work with VXLANs on QFX5100 switches. [PR1024457](#)

### ***VXLAN***

- On a QFX5100 switch with a VXLAN configured, (S,G) interface entries downstream from a VXLAN interface might be missing from the multicast routing table but be present in the kernel and Packet Forwarding Engine. In this circumstance, traffic is forwarded as expected. [PR1027119](#)
- If a 32-member VCF loads the MDconfig without any routes and traffic and receives the **nh\_comp\_msg\_parse** message, the FXPC might create a core file. [PR1029884](#)
- The **interface-mac-limit** statement is not supported with VXLANs. If you configure this statement with a VXLAN, MAC learning might not occur and traffic might not be forwarded. In this circumstance, delete the **interface-mac-limit** statement and the VXLAN configuration, then reconfigure the VXLAN. [PR1032552](#)

### **Related Documentation**

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 34](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

## Documentation Updates

This section lists the errata or changes in Junos OS Release 14.1X53 documentation for QFX Series.

### Bridging and Learning

- Two new MIBs related to MAC notification are provided with Junos OS Release 14.1X53-D10:
  - [jnxL2aldMacHistoryEntry](#)
  - [jnxL2aldMacNotificationMIBGlobalObjects](#)

These MIBs are not yet described in the documentation.

### Network Management and Monitoring

- The Network Management and Monitoring on the QFX Series feature guide at Junos OS Release 14.1X53-D10 erroneously contained topics that applied to QFabric systems but not to QFX Series standalone switches. Those QFabric systems topics have been removed from the guide.

### Virtual Chassis and Virtual Chassis Fabric (VCF)

- The support plan for the maximum number of member devices in a Virtual Chassis Fabric (VCF) has been revised to support for a maximum of 20 devices for all platforms that support VCF. The announcement for 32-device support has been removed from New Features in Junos OS Release 14.1X53-D15 in these release notes.

#### **Related Documentation**

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 34](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 55](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)
- [Product Compatibility on page 77](#)

## 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 to a Controlled Version of Junos OS on page 72](#)
- [Upgrading Software on QFX5100 Standalone Switches on page 72](#)
- [Performing an In-Service Software Upgrade \(ISSU\) on page 75](#)

- [Preparing the Switch for Software Installation on page 75](#)
- [Upgrading the Software Using ISSU on page 75](#)

---

### Upgrading to a Controlled Version of Junos OS

---

Starting in Junos OS Release 14.1X53-D15, you can install a controlled version of Junos OS software on a QFX Series switch. The controlled version of Junos OS software is required to enable Media Access Control security (MACsec).

If you are upgrading your switch between a domestic version of Junos OS and a controlled version of Junos OS, keep the following issues in mind:

- You cannot use NSSU to upgrade or downgrade from a controlled version of Junos OS to a domestic version of Junos OS.
- In a Virtual Chassis, all member switches must be running the same release of Junos OS. A Virtual Chassis with member switches that are running domestic and export versions of the same Junos OS release does form.
- In a Virtual Chassis, all member switches must be running the same release of Junos OS.

To support MACsec, however, all member switches in the Virtual Chassis must be running the controlled version of Junos OS.

The upgrade or downgrade procedure from a domestic version of Junos OS to a controlled version of Junos OS is, otherwise, identical to any other Junos OS upgrade. See *Upgrading Software* for more information.

---

### 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.



**NOTE:** On Junos Release 14.1X53-D35.3, autonegotiation is disabled by default.

---





**NOTE:** On QFX5100 and EX4600 switches, the Host OS is not upgraded automatically, so you must use the **force-host** option if you want the Junos OS and Host OS versions to be the same.

However, pay attention to these notes regarding Junos OS and Host OS versions:

- The Junos OS and Host OS versions do not need to be the same.
- During an ISSU, the Host OS cannot be upgraded.
- Upgrading the Host OS is not required for every software upgrade, as noted above.



**NOTE:** On QFX5100 and EX4600 switches, you must use the **force-host** option if you are downgrading from Junos OS Release 14.1X53-D40 to any release earlier than 14.1X53-D40 otherwise the switch will issue core dumps.

The download and installation process for Junos OS Release 14.1X53-D10 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 14.1 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 14.1 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-5-14.1X53-D25-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 14.1 jinstall package, you can issue the **request system software rollback** command to return to the previously installed software.

## Performing an In-Service Software Upgrade (ISSU)

You can use an in-service software upgrade to upgrade the software running on the switch with minimal traffic disruption during the upgrade.



**NOTE:** ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 75](#)
- [Upgrading the Software Using ISSU on page 75](#)

## Preparing the Switch for Software Installation

Before you begin software installation using ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:



**NOTE:** If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the **request system snapshot** command.

## Upgrading the Software Using ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Upgrading Software*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
  - On the switch, enter:

```
user@switch> request system software in-service-upgrade
/var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, *jinstall-132\_x51\_vjunos.domestic.tgz*.



**NOTE:** During the upgrade, you will not be able to access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-qfx-5-13.2X51-D15.4-domestic ...
Install jinstall-qfx-5-13.2X51-D15.4-domestic completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
```

ISSU: IDLE  
Initiate em0 device handoff



**NOTE:** An ISSU might stop instead of abort if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).



**NOTE:** If the ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

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

#### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 34](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Product Compatibility on page 77](#)

## Product Compatibility

- [Hardware Compatibility on page 77](#)

### 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 34](#)
- [Known Behavior on page 38](#)
- [Known Issues on page 53](#)

- [Resolved Issues on page 55](#)
- [Documentation Updates on page 71](#)
- [Migration, Upgrade, and Downgrade Instructions on page 71](#)

## Third-Party Components

---

This product includes third-party components. To obtain a complete list of third-party components, see [Overview for Routing Devices](#).

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 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://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, see <http://www.juniper.net/support/requesting-support.html>.

## Revision History

---

1 June, 2017—Revision 2, Junos OS for QFabric Systems, Release 14.1X53-D121—Added item to Known Issues.

31 May, 2017—Revision 1, Junos OS for QFabric Systems, Release 14.1X53-D121

10 May, 2017—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D43

14 March, 2017—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D42

24 February, 2017—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D42

24 January, 2017—Revision 4, Junos OS for the EX Series and QFX Series, Release 14.1X53-D40

13 December, 2016—Revision 3, Junos OS for the EX Series and QFX Series, Release 14.1X53-D40

16 November, 2016—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D40

14 November, 2016—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D40

29 July 2016—Revision 5, Junos OS for the EX Series, OCX Series, and QFX Series, Release 14.1X53-D35—Removed item from QFX Series New Features, added item to Documentation Updates.

9 May 2016—Revision 4, Junos OS for the EX Series, OCX Series, and QFX Series, Release 14.1X53-D35—Item added to QFX Series “Changes in Behavior and Syntax” and “Resolved Issues”.

25 March 2016—Revision 3, Junos OS for the EX Series, OCX Series, and QFX Series, Release 14.1X53-D35—Item added to EX Series “Known Behavior” and “Documentation Updates”.

9 March 2016—Revision 2, Junos OS for the EX Series, OCX Series, and QFX Series, Release 14.1X53-D35—Moved PR966905 to “Known Behavior”.

2 March 2016—Revision 1, Junos OS for the EX Series, OCX Series, and QFX Series, Release 14.1X53-D35

Copyright © 2017, Juniper Networks, Inc. All rights reserved.

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

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