

Release Notes: Junos[®] OS Release 14.1X53-D48 for the EX Series and QFX Series

Release 14.1X53-D48
November 13, 2018
Revision 1

Contents

Junos OS Release Notes for EX Series Switches	5
New and Changed Features	5
New Features in Release 14.1X53-D40	6
New Features in Release 14.1X53-D35	7
New Features in Release 14.1X53-D26	9
New Features in Release 14.1X53-D25	10
New Features in Release 14.1X53-D15	10
New Features in Release 14.1X53-D10	12
Changes in Behavior and Syntax	22
Authentication and Access Control	22
Dynamic Host Configuration Protocol	23
Interfaces and Chassis	24
Network Management and Monitoring	24
Routing Policy and Firewall Filters	25
Virtual Chassis and Virtual Chassis Fabric	25
Known Behavior	27
General Routing	28
High Availability	28
Infrastructure	28
J-Web	28
Platform and Infrastructure	29
User Interface and Configuration	29
Known Issues	29
General Routing	30
High Availability (HA) and Resiliency	32
Infrastructure	32

Layer 2 Features	33
Layer 3 Features	33
Platform and Infrastructure	33
Routing Protocols	34
User Interface and Configuration	34
Virtual Chassis	34
Resolved Issues	35
Resolved Issues: Release 14.1X53-D48	35
Resolved Issues: Release 14.1X53-D47	37
Resolved Issues: Release 14.1X53-D46	44
Resolved Issues: Release 14.1X53-D45	46
Resolved Issues: Release 14.1X53-D44	48
Resolved Issues: Release 14.1X53-D43	49
Resolved Issues: Release 14.1X53-D42	52
Resolved Issues: Release 14.1X53-D40	54
Resolved Issues: Release 14.1X53-D35	65
Resolved Issues: Release 14.1X53-D30	72
Resolved Issues: Release 14.1X53-D27	79
Resolved Issues: Release 14.1X53-D26	80
Resolved Issues: Release 14.1X53-D25	81
Resolved Issues: Release 14.1X53-D16	85
Resolved Issues: Release 14.1X53-D10	86
Documentation Updates	88
Bridging and Learning	88
Interfaces and Chassis	89
Security	89
Migration, Upgrade, and Downgrade Instructions	89
Upgrade and Downgrade Support Policy for Junos OS Releases	89
Product Compatibility	90
Hardware Compatibility	90
Junos OS Release Notes for the QFX Series	91
New and Changed Features	91
New Features in Release 14.1X53-D46	92
New Features in Release 14.1X53-D40	92
New Features in Release 14.1X53-D35	99
New Features in Release 14.1X53-D30	100
New Features in Release 14.1X53-D27	103
New Features in Release 14.1X53-D26	104
New Features in Release 14.1X53-D25	105
New Features in Release 14.1X53-D15	106
New Features in Release 14.1X53-D10	111
Changes in Behavior and Syntax	122
Authentication and Access Control	123
Ethernet Switching	123
Interfaces and Chassis	123
MPLS	124
Network Management and Monitoring	124
Open vSwitch Database (OVSDB)	125
SNMP	125

Software Upgrade	125
Virtual Chassis and Virtual Chassis Fabric	125
Known Behavior	127
EVPN	127
High Availability	127
Interfaces and Chassis	128
Layer 2 Features	128
MPLS	128
OVSDB	129
Platform and Infrastructure	129
Routing Protocols	129
Security	130
Software Installation and Upgrade	130
Storage and Fibre Channel	130
Traffic Management	132
VXLAN	135
Known Issues	136
Class of Service (CoS)	137
EVPN	137
General Routing	138
High Availability (HA) and Resiliency	140
Interfaces and Chassis	140
Layer 2 Ethernet Services	140
Layer 2 Features	140
MPLS	141
Platform and Infrastructure	142
Routing Protocols	143
User Interface and Configuration	144
Virtual Chassis	144
VXLAN	144
Resolved Issues	144
Resolved Issues: Release 14.1X53-D48	145
Resolved Issues: Release 14.1X53-D47	147
Resolved Issues: Release 14.1X53-D46	152
Resolved Issues: Release 14.1X53-D45	154
Resolved Issues: Release 14.1X53-D44	156
Resolved Issues: Release 14.1X53-D43	158
Resolved Issues: Release 14.1X53-D42	162
Resolved Issues: Release 14.1X53-D40	163
Resolved Issues: Release 14.1X53-D35	176
Resolved Issues: Release 14.1X53-D30	182
Resolved Issues: Release 14.1X53-D27	194
Resolved Issues: Release 14.1X53-D26	196
Resolved Issues: Release 14.1X53-D25	198
Resolved Issues: Release 14.1X53-D16	202
Resolved Issues: Resolved Before Release 14.1X53-D16	205
Documentation Updates	206
Bridging and Learning	207
Network Management and Monitoring	207

Virtual Chassis and Virtual Chassis Fabric (VCF)	207
Migration, Upgrade, and Downgrade Instructions	207
Upgrading to a Controlled Version of Junos OS	208
Upgrading Software on QFX5100 Standalone Switches	208
Performing an In-Service Software Upgrade (ISSU)	211
Preparing the Switch for Software Installation	211
Upgrading the Software Using ISSU	211
Product Compatibility	213
Hardware Compatibility	213
Third-Party Components	214
Finding More Information	214
Documentation Feedback	214
Requesting Technical Support	214
Self-Help Online Tools and Resources	215
Opening a Case with JTAC	215
Revision History	216

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 14.1X53-D48 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

The following EX Series switches are supported in Junos OS Release 14.1X53-D48: EX3200, EX3300, EX4200, EX4300, EX4500, EX4550, EX4600, EX6210, EX8208, and EX8216.



NOTE: These release notes include information on all Junos OS Release 14.1X53 releases.

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 22](#)
- [Known Behavior on page 27](#)
- [Known Issues on page 29](#)
- [Resolved Issues on page 35](#)
- [Documentation Updates on page 88](#)
- [Migration, Upgrade, and Downgrade Instructions on page 89](#)
- [Product Compatibility on page 90](#)

New and Changed Features

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

- [New Features in Release 14.1X53-D40 on page 6](#)
- [New Features in Release 14.1X53-D35 on page 7](#)
- [New Features in Release 14.1X53-D26 on page 9](#)
- [New Features in Release 14.1X53-D25 on page 10](#)
- [New Features in Release 14.1X53-D15 on page 10](#)
- [New Features in Release 14.1X53-D10 on page 12](#)

New Features in Release 14.1X53-D40

Authentication and Access Control

- **Voice VLAN fallback (EX Series)**—Starting in Junos OS Release 14.1X53-D40, you can configure authentication fallback options to specify how VoIP clients sending voice traffic are supported if the RADIUS authentication server becomes unavailable. When you configure the server fail fallback feature you must specify an action that the switch applies to end devices when the authentication servers are unavailable. The switch can accept or deny access to supplicants or maintain the access already granted to supplicants before the RADIUS timeout occurred. You can also configure the switch to move the supplicants to a specific VLAN.

[See [Configuring RADIUS Server Fail Fallback \(CLI Procedure\)](#).]

Interfaces and Chassis

- **Half-duplex link support (EX4300 switches)**—Starting with Junos OS 14.1X53-D40, half-duplex communication is supported on all built-in network copper ports on EX4300 switches. *Half-duplex* is bidirectional communication, but signals can flow in only one direction at a time. *Full-duplex* communication means that both ends of the communication can send and receive signals at the same time. Half-duplex is configured by default on EX4300 switches. If the link partner is set to autonegotiate the link, then the link is autonegotiated to full duplex or half duplex. If the link is not set to autonegotiation, then the EX4300 link defaults to half-duplex unless the interface is explicitly configured for full duplex.

To explicitly configure full duplex:

```
[edit]
user@switch# set interfaces interface-name speed 10m-or-100m
[edit]
user@switch# set interfaces interface-name ether-options no-auto-negotiate
```

To verify a half-duplex setting:

```
user@switch> show interfaces interface-name extensive
```

[See [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#).]

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.](#)]

New Features in Release 14.1X53-D35

Hardware

- **Revert EX2200 and EX2200-C switches to the factory-default configuration using the Factory reset/Mode button on the switch**—Starting with Junos OS Release 14.1X53-D35, you can transition EX2200 and EX2200-C switches to the factory-default configuration by pressing the Factory reset/Mode button located below the LED labeled **POE** on the far right side of the front panel of the switches for 10 seconds. You can transition the switches to the initial setup mode by pressing the button for 10 seconds more.

Interfaces

- **GRE tunneling (EX4300 switches)**—Starting with Junos OS Release 14.1X53-D35, generic routing encapsulation (GRE) tunneling is supported on EX4300 switches. Tunneling provides a private, secure path for transporting packets through an otherwise public network by encapsulating packets inside a transport protocol known as an IP encapsulation protocol. GRE is an IP encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel. GRE tunneling is accomplished through routable tunnel endpoints that operate on top of existing physical and other logical endpoints. GRE tunnels connect one endpoint to another and provide a clear data path between the endpoints.

Configure tunnels to use GRE:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number family inet address
user@switch# set gr-0/0/0 unit number tunnel source source-address
user@switch# set gr-0/0/0 unit number tunnel destination destination-address
```



NOTE: The switch supports IPv4 as the tunneling (delivery) protocol. It supports IPv4 and IPv6 as the payload protocol.

J-Web Interface

- **J-Web (EX4600 switches)**—Starting with Junos OS Release 14.1X53-D35, you can configure and monitor software features on EX4600 switches by using the J-Web interface.

The following limitations apply to using J-Web on EX4600 switches:

- 802.1X authentication configuration is not supported.
- Power over Ethernet (PoE) configuration and monitoring is not supported.
- Class-of-service (CoS) configuration is not supported.



NOTE: On EX4600 switches, the maximum number of LAG devices that you can configure is 1000.

For more information, see [J-Web for EX Series Ethernet Switches](#).

Platform and Infrastructure

- **Workaround for sudden shutdowns while crossing negative temperature thresholds (EX2200 switches)**—Starting with Junos OS Release 14.1X53-D35, you can configure a time interval in seconds for the switch to remain powered on after crossing the temperature-shutdown limit.

Configure the time interval:

```
[edit]
user@switch# set chassis shutdown-delay-period seconds
```

You can configure an operating-temperature range and a time interval in seconds for raising an alarm once the temperature crosses either end of the operating range. The alarm will be raised periodically at each time interval that passes while the switch remains out of operating-temperature range.

Configure the operating-temperature range and time interval:

```
[edit]
user@switch# set chassis operating-temperature temperature-range low-value
high-value alarm-interval seconds
```


Port Security

- **DHCP snooping table update for changed MAC address (EX4300 and EX4600 switches)**—Starting with Junos OS Release 14.1X53-D35, the DHCP snooping table is updated in the event of a change to a client's MAC address. If a client requests for an IP address that matches an IP address in the DHCP snooping table, but has a MAC address that does not match the one bound to that IP address in the DHCP snooping table, then a placeholder binding is created using the client IP address and the new MAC address. When the switch receives a DHCPACK message from the DHCP server, this binding is added to the DHCP snooping table, replacing the original binding. This new feature requires no configuration changes to be made by the user.

Routing Policy and Firewall Filters

- **Firewall filter with policer action as forwarding-class and loss priority (PLP) (EX4300 switches)**—Starting with Junos OS Release 14.1X53-D35, on EX4300 switches you can configure the firewall with policer action as forwarding-class and loss priority (PLP). When the traffic hits the policer, PLP changes as per the action rule. The supported PLP designations are low, high, and medium-high. You configure policer actions at the `[edit firewall]` hierarchy level.

New Features in Release 14.1X53-D26

Hardware

- **New optical transceivers support on EX4300 switches**—Starting with Junos OS Release 14.1X53-D26, EX4300 switches support the following optical transceivers:
 - EX-SFP-GE10KT13R14 (1000BASE-BX-U, 10 km)
 - EX-SFP-GE10KT14R13 (1000BASE-BX-D, 10 km)

- EX-SFP-GE10KT13R15 (1000BASE-BX-U, 10 km)
- EX-SFP-GE10KT15R13 (1000BASE-BX-D, 10 km)

New Features in Release 14.1X53-D25

Authentication and Access Control

- **Access control (mixed EX4300 and EX4600 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D25, EX4600 switches operating in a mixed Virtual Chassis with EX4300 switches support controlling access to your network by using several different authentication methods: 802.1X authentication, MAC RADIUS authentication, or **captive portal**. You enable the **authentication-whitelist** statement at the **[edit switching-options]** hierarchy level instead of at the **[edit ethernet-switching-options]** hierarchy level.

Access control features in a mixed EX4300 and EX4600 Virtual Chassis are supported only on EX4300 switch interfaces.

[See [Access Control on a Mixed EX4300-EX4600 Virtual Chassis](#).]

MPLS

- **MPLS stitching for virtual machine connections (EX4600)**—By using MPLS, the stitching feature of Junos OS 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 link 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](#).]

New Features in Release 14.1X53-D15

Interfaces and Chassis

- **Default logging for Ethernet ring protection switching (ERPS) (EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, EX8200 standalone switches; EX2200, EX3300, EX4200, EX4500, EX4550, EX8200 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D15, the listed EX Series switches automatically log basic state transitions for the ERPS protocol. No configuration is required to initiate this logging. Basic state transitions include ERPS interface transitions from up to down, and down to up; and ERPS state transitions from idle to protection, and protection to idle.

The basic state transitions are logged in a single file named **erp-default**, which resides in the **/var/log** directory of the switch. The maximum size of this file is 15 MB.

Default logging for ERPS can capture initial ERPS interface and state transitions, which can help you troubleshoot issues that occur early in the ERPS protocol startup process. However, if more robust logging is needed, you can enable traceoptions for ERPS by entering the **traceoptions** statement in the **[edit protocols protection-group]** hierarchy level.

Be aware that for ERPS, only default logging or traceoptions can be active at a time on the switch. That is, default logging for ERPS is automatically enabled and if you enable traceoptions for ERPS, the switch automatically disables default logging. Conversely, if you disable traceoptions for ERPS, the switch automatically enables default logging.

- **Power over Ethernet (EX4600 switches)**—Starting with Junos OS Release 14.1X53-D15, EX4600 switches support Power over Ethernet (PoE) when operating in a mixed-mode Virtual Chassis with an EX4300 switch. You can enable PoE configuration statements and run PoE operational commands on an EX4600 switch only when the switch is operating in a mixed-mode Virtual Chassis.

You can configure PoE at the `[edit poe]` hierarchy level.

[See [Understanding PoE on EX Series Switches](#).]

MPLS

- **MPLS enhancements (EX4600 switches)**—Starting with Junos OS Release 14.1X53-D15, a set of procedures is provided for augmenting network layer packets with label stacks, thereby turning them into labeled packets. MPLS has emerged as an elegant solution to meet the bandwidth-management and service requirements for next-generation IP-based backbone networks.

The following MPLS features have been added to EX4600:

- BGP L3 VPN Carrier-over-Carrier and Interprovider

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routing devices in different autonomous systems (ASs). Instead of using the label distribution protocols LDP or RSVP, MPLS can piggyback on routing protocols such as BGP and OSPF.

- Ethernet over MPLS pseudowire based on LDP (draft Martini / L2 Circuit)

Ethernet-over-MPLS supports sending Layer 2 Ethernet frames transparently over MPLS using a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. Pseudowire is a software mechanism for emulating various networking or telecommunications services across packet-switched networks that use Ethernet, IP, or MPLS. An Ethernet pseudowire is used to carry Ethernet or 802.3 PDUs over an MPLS network enabling service providers to offer emulated Ethernet services over existing MPLS networks. There are several label distribution protocols used such as Label Distribution Protocol (LDP) or RSVP; another technique is piggybacking on routing protocols such as BGP and OSPF.

- Static and dynamic Ethernet pseudowire over LDP and RSVP tunnels

Pseudowire is a software mechanism for emulating various networking or telecommunications services across packet-switched networks that use Ethernet, IP, or MPLS. Label Distribution Protocol (LDP) and RSVP are label distribution protocols used by MPLS.

- Pseudowire over aggregated Ethernet on core-facing interfaces

Pseudowire is a software mechanism for emulating various networking or telecommunications services across packet-switched networks that use Ethernet, IP, or MPLS.

- RSVP fast-reroute including link-protection and node-link-protection

One label distribution protocol used for MPLS data transmission is RSVP.

[See [MPLS Feature Support on the QFX Series and the EX4600 Switch](#).]

Security

- **Media Access Control Security (MACsec) support (EX4600 switches)**—Starting with Junos OS Release 14.1X53-D15, MACsec is supported on all built-in SFP+ interfaces on an EX4600 switch. MACsec is also supported on all eight SFP+ interfaces on the EX4600-EM-8F expansion module when it is installed in an EX4600 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 also *Documentation Updates*.

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

New Features in Release 14.1X53-D10

Authentication and Access Control

- **IPv6 for RADIUS AAA (EX3300, EX4200, EX4300, EX4500, and EX8200 switches and EX4300 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D10, EX3300, EX4200, EX4300, EX4500, and EX8200 switches and EX4300 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

- **RVI support for private VLANs (EX8200 switches and EX8200 Virtual Chassis)**—Starting with Junos OS Release 14.1X53-D10, you can configure a routed VLAN interface (RVI) on an EX8200 switch or EX8200 Virtual Chassis to handle the Layer 3 traffic of intersecondary VLANs (community VLANs and isolated VLANs) in a private VLAN (PVLAN). By using an RVI to handle the routing within the PVLAN, you eliminate the need for an external router with a promiscuous port connection to perform this function.

One RVI serves the entire PVLAN domain regardless of whether the domain consists of one or more switches. After you configure the RVI, Layer 3 packets received by the secondary VLAN interfaces are mapped to and routed by the RVI.

[See [Configuring a Routed VLAN Interface in a Private VLAN \(CLI Procedure\)](#).]

- **Support for private VLANs (EX4300)**—Starting with Junos OS Release 14.1X53-D10, EX4300 switches support private VLANs (PVLANS). PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the known communication between known hosts. PVLANS can be used to help ensure the security of service providers sharing a server farm, or to provide security to subscribers of various service providers sharing a common metropolitan area network.



NOTE: An interface can belong to only one PVLAN domain.

[See [Understanding Private VLANs on EX Series Switches](#).]

- **Support for Layer 2 protocol tunneling (EX4300)**—Starting with Junos OS Release 14.1X53-D10, EX4300 switches support Layer 2 protocol tunneling (L2PT). L2PT enables service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain. This feature is useful when you want to run Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network. For example, it can help you provide transparent LAN services over a metropolitan Ethernet infrastructure. L2PT operates under the Q-in-Q tunneling configuration; therefore, you must enable Q-in-Q tunneling before you can configure L2PT.

The Layer 2 protocol to be tunneled can be one of the following: 802.3AH, CDP, LACP, LLDP, MVRP, STP, VTP, GVRP, or VSTP.



NOTE: L2PT does not support the following on EX4300 switches:

- drop-threshold or shutdown-threshold statements
- The all option for setting the Layer 2 protocol
- 802.1X authentication

[See [Understanding Layer 2 Protocol Tunneling on EX Series Switches](#).]

- **MAC notification (EX4300 and EX4600)**—Starting with Junos OS Release 14.1X53-D10, MAC notification is supported on EX4300 and EX4600 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 206](#).

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

- **Default VLAN and multiple VLAN range support (EX4300)**—Starting with Junos OS Release 14.1X53-D10, the default VLAN and multiple VLAN range are supported on EX4300 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 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 enabled for VoIP or 802.1X 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] hierarchy level.

Class of Service

- **Explicit congestion notification (ECN) support (EX4300)**—Starting with Junos OS Release 14.1X53-D10, ECN marking is supported on EX4300 switches—you enable it for packets in scheduler queues. Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.

ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets.

To enable ECN, issue the **set class-of-service schedulers *name* explicit-congestion-notification** command.

Infrastructure

- **Licensing enhancements (EX Series)**—Starting with Junos OS Release 14.1X53-D10, licensing enhancements on EX 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:

```

	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 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

- **Support for aggregated Ethernet link protection enhancements (EX4500)**—Starting with Junos OS Release 14.1X53-D10, aggregated Ethernet link protection is enhanced on EX4500 switches to support a collection of Ethernet links within a LAG bundle. Link protection could earlier be used to protect a single link within a LAG bundle only. The ability to provide link protection for a collection of links in a LAG bundle is provided using link protection subgroups, which are introduced as part of this feature.

[See [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\)](#).]

J-Web

- **J-Web interface available in two packages (EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, EX4550, EX6200)**—Prior to this release, the J-Web interface was available as a single package as part of Junos OS. Starting with Junos OS Release 14.1X53-D10, the J-Web interface is available in two packages:
 - The Platform package is installed as part of Junos OS, which provides basic functionalities of J-Web. You can use the Platform package to create a basic configuration and maintain your EX Series switch.
 - The Application package is an optionally installable package, which provides complete functionalities of J-Web that enable you to configure, monitor, maintain, and troubleshoot your switch. You must download the Application package and install it over the Platform package on your switch.

For detailed information about the J-Web packages, see [Release Notes: J-Web Application Package Release 14.1X53-A1 for Juniper Networks EX Series Ethernet Switches](#).

- **Browser support enhancements for the J-Web interface (EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, EX4550, EX6200)**—Starting with Junos OS Release 14.1X53-D10, the J-Web interface supports the following browsers:
 - Microsoft Internet Explorer versions 9 and 10
 - Mozilla Firefox versions 24 through 30
 - Google Chrome versions 27 through 36



TIP: For best viewing of the J-Web application, set the screen resolution to 1440 X 900.

Layer 3 Protocols

- **IS-IS protocol (EX3300)**—EX3300 switches now support the Intermediate System-to-Intermediate System (IS-IS) protocol. On EX3300 switches, the IS-IS configuration is available at the **[edit protocols]** hierarchy level.

[See [Layer 3 Protocols Supported on EX Series Switches.](#)]

MPLS

- **Ethernet-over-MPLS (L2 circuit) (EX4600)**—Starting with Junos OS Release 14.1X53-D10, Ethernet-over-MPLS is supported on EX4600 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-based Layer 3 VPNs (EX4600)**—Starting with Junos OS Release 14.1X53-D10, MPLS-based Layer 3 VPNs are supported on EX4600 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 device 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.
 - **MPLS LSP protection (EX4600)**—Starting with Junos OS Release 14.1X53-D10, the following types of MPLS LSP protection are supported on EX4600 switches:
 - Fast reroute (FRR)

- Link protection
- Node link protection

[See [MPLS Overview](#).]

Network Management and Monitoring

- **Chef for Junos OS (EX4300)**—Starting with Junos OS Release 14.1X53-D10, Chef for Junos OS is supported on EX4300 switches.
- **Puppet for Junos OS (EX4300)**—Starting with Junos OS Release 14.1X53-D10, Puppet for Junos OS is supported on EX4300 switches.
- **Network analytics (EX4300)**—Starting with Junos OS Release 14.1X53-D10, EX4300 switches support the network analytics feature. The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data by using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed. The analytics manager (analyticsm) in the Packet Forwarding Engine collects traffic and queue statistics, and the analytics daemon (analyticsd) in the Routing Engine analyzes the data and generates reports. You can enable network analytics by configuring microburst monitoring and high-frequency traffic statistics monitoring.

[See [Network Analytics Overview](#).]

- **Ethernet frame delay measurement (EX2200)**—Starting with Junos OS Release 14.1X53-D10, you can obtain Ethernet frame delay measurements (ETH-DM) on an EX2200 switch. You can configure Operation, Administration, and Maintenance (OAM) statements for connectivity fault management (IEEE 802.1ag) to provide on-demand measurements of frame delay and frame delay variation (jitter). You configure the feature under the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.
- **Support for native analyzers and remote port-mirroring capabilities (EX4300)**—Starting with Junos OS Release 14.1X53-D10, native analyzers and remote port mirroring are supported on EX4300 switches. A native analyzer configuration contains both an input stanza and an output stanza in the analyzer hierarchy for mirroring packets. In remote port mirroring, the mirrored traffic is flooded into a remote mirroring VLAN that can be specifically created for the purpose of receiving mirrored traffic. On EX4300 switches, the analyzer configuration is available under the **[edit forwarding-options]** hierarchy level.

Port Security

- **IPv6 access security (EX2200 and EX3300)**—Starting with Junos OS Release 14.1X53-D10, the following IPv6 access security features are supported on EX2200 and EX3300 switches: DHCPv6 snooping, IPv6 Neighbor Discovery Inspection, IPv6 source guard, and RA guard. DHCPv6 snooping enables a switch to process DHCPv6 messages between a client and a server and build a database of the IPv6 addresses assigned to the DHCPv6 clients. The switch can use this database, also known as the binding table, to stop malicious traffic. DHCPv6 includes the relay agent Remote-ID option, also known as Option 37, to optionally append additional information to the messages sent by the client towards the server. This information can be used by the server to assign addresses and configuration parameters to the client. IPv6 Neighbor Discovery inspection analyzes neighbor discovery messages sent between IPv6 nodes on the same link and verifies them against the DHCPv6 binding table. IPv6 source guard inspects all IPv6 traffic from the client and verifies the source IPv6 address and source MAC address against the entries in the DHCPv6 binding table. If no match is found, the traffic is dropped. RA guard examines incoming Router Advertisement (RA) messages and decides whether to forward or block them based on statically configured IPv6/MAC address bindings. If the content of the RA message does not match the bindings, the message is dropped.

Starting with this release, Remote-ID (Option-37) is not added by default on when you enable **dhcpv6-snooping**.

You configure DHCPv6 snooping, IPv6 Neighbor Discovery Inspection, and IPv6 source guard at the **[edit ethernet-switching-options secure-access-port vlan *vlan-name*]** hierarchy level. You configure RA guard at the **[edit ethernet-switching-options secure-access-port interface *interface-name*]** hierarchy level.

[See [Port Security Overview](#).]

- **IPv6 access security (EX4300)**—Starting with Junos OS Release 14.1X53-D10, DHCPv6 snooping supports a configuration to optionally append the relay agent Remote ID (Option-37), Interface-ID (Option-18), and Vendor-Class (Option-16) to the DHCPv6 packets sent by a client. You can configure these options under the **[edit vlans *vlan-name* forwarding-options dhcp-security dhcpv6-options]** hierarchy level.
- **Media Access Control Security (MACsec) support for switch to host connections (EX4200, EX4300, and EX4550)**—Starting with Junos OS Release 14.1X53-D10, MACsec is supported on links connecting EX4200, EX4300, and EX4550 switches to host devices, such as phones, servers, personal computers, or other endpoint devices. This feature also introduces MACsec dynamic mode and the ability to retrieve MACsec Key Agreement (MKA) keys from a RADIUS server, which are required to enable MACsec on a switch to host link.

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

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 enables 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](#).]

- See Also**
- [Changes in Behavior and Syntax on page 22](#)
 - [Known Behavior on page 27](#)
 - [Known Issues on page 29](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 88](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)
 - [Product Compatibility on page 90](#)

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 the EX Series.

- [Authentication and Access Control](#)
- [Dynamic Host Configuration Protocol](#)
- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)
- [Routing Policy and Firewall Filters](#)
- [Virtual Chassis and Virtual Chassis Fabric](#)

Authentication and Access Control

- **LLDP neighbor port info display (EX4300 switches)**—Starting with Junos OS Release 14.1X53-D40 for EX Series switches, the **neighbor-port-info-display** CLI statement is supported at the **[edit protocols lldp]** hierarchy level. You can use this statement to configure the type of LLDP neighbor port information that the switch displays in the **Port info** field in the output of the **show lldp neighbors** CLI command. By default, the **Port info** field in the output of the **show lldp neighbors** CLI command displays the port description TLV.

[See [neighbor-port-info-display](#).]

- **Increase in TACACS message length (EX 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.

- **Support for the accounting-port statement (EX Series)**—Starting with Junos OS Release 14.1X53-D25, the **accounting-port** CLI statement is now supported at the [**edit access radius-server server-address**] hierarchy level on all EX Series switches. This command was supported only on EX4300, EX4600, and EX9200 switches in earlier Junos OS releases. The **accounting-port** statement enables you to specify the port on which to contact the RADIUS accounting server. The default port number is 1813, as specified in RFC 2866.

Dynamic Host Configuration Protocol

- **Format change for VLAN ID in DHCP Option 18**—On EX4300 and EX4600 switches with DHCP snooping configured, when the VLAN ID is appended to the prefix of DHCP option 18, it will appear in decimal format instead of hexadecimal format.

Interfaces and Chassis

- On EX4300 switches, when you configure the DHCP **relay-option-82** option, the circuit ID is added by default. In the case of an IRB interface, the DHCP relay option 82 will contain a description or name of the physical layer interface instead of the name of the IRB interface. To include the name of the IRB interface, you can use the **include-irb-and-l2** statement. To display only the IRB interface without the names of the Layer 2 interface and VLAN, use the **no-vlan-interface-name** statement.

default	VLAN-tagged interface	ge-1/2/3:10
	Dual-tagged interface	ge-1/2/3:10-20
	Pure Layer 3 interface	ge-1/2/3:0
	IRB interface	ge-1/2/3:0:v10
use-vlan-id		ge-1/2/3:0:10
include-irb-and-l2		ge-1/2/3:0:v10+irb.10
include-irb-and-l2 and use-vlan-id		ge-1/2/3:0:10+irb.10
no-vlan-interface-name		irb.10
no-vlan-interface-name and use-vlan-id		Mutually exclusive
no-vlan-interface-name and include-irb-and-l2		ge-1/2/3:0+irb.10
use-interface-description		l2_descr:v10
	If no description found	ge-1/2/3:0:v10
use-interface-description and use-vlan-id		Mutually exclusive
use-interface-description and include-irb-and-l2		l2_descr:v10+irb.10
	If no description found	ge-1/2/3:0:v10+irb.10
use-interface-description and no-vlan-interface-name		irb_descr
	If no description found	irb.10
use-interface-description, no-vlan-interface-name, and include-irb-and-l2		l2_descr+irb.10
	If no description found	ge-1/2/3:0+irb.10

Network Management and Monitoring

- **Juniper MIBs loading errors fixed (EX Series)**—Starting with Junos OS Release 14.1X53-D48, duplicated entries and errors while loading MIBs on the ManageEngine MIB browser are fixed for the following MIB files:

- jnx-chas-defines.mib
- jnx-ifotn.mib

[See [MIB Explorer](#).]

Routing Policy and Firewall Filters

- **Support for enhanced mask length on IPv6 destination-address match conditions for loopback filters (EX Series switches)**—Starting with Junos OS Release 14.1X53-D15, the maximum mask length of IPv6 destination-address match conditions in loopback (lo0) filters on EX Series switches is /128.

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).

- **Increased time to rejoin Virtual Chassis after a member is rebooted (EX2200, EX3300, EX4200, EX6200, and EX8200 Virtual Chassis)**—Starting in Junos OS Release 14.1X53-D35, when one or more member switches in an EX2200, EX3300, EX4200, EX6200, or EX8200 Virtual Chassis are rebooted, the Virtual Chassis master's delay time before reinstating the rebooted switch as a member in the Virtual Chassis is increased from two minutes to ten minutes. As a result, after rebooting a Virtual Chassis member, up to 15 or 20 minutes total elapsed time might be required for the member to completely rejoin the Virtual Chassis. The increased delay time allows the Virtual Chassis to correctly rebuild its Virtual Chassis port (VCP) adjacency information, and avoids unexpected mastership election contention or failure of the Virtual Chassis to re-form.

- **Automatic software update (EX2200 Virtual Chassis)**—Starting in Junos OS Release 14.1X53-D44, the automatic software update feature can be used to automatically update Junos software on members of an EX2200 Virtual Chassis running Junos OS Release 12.3R12 and later. Automatic software update is not supported on an EX2200 Virtual Chassis in releases prior to 14.1X53-D44.
- **Automatic Virtual Chassis port conversion disabled by default (EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis)**—Starting in Junos OS Release 14.1X53-D47, automatic Virtual Chassis port (VCP) conversion is disabled by default in an EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis. Previously, automatic VCP conversion was always enabled by default on these switches in a Virtual Chassis.

When automatic VCP conversion is enabled, if you add a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

Automatic VCP conversion is not needed when using dedicated VCPs or default-configured VCPs on both sides of the link to interconnect two members. You can also manually configure network or uplink ports that are supported as VCPs on both ends of the link, instead of using the automatic VCP conversion feature.



NOTE: When automatic VCP conversion is enabled in a Virtual Chassis with switches that have dedicated VCPs (EX4200, EX4500, or EX4550 Virtual Chassis), if network or uplink ports are automatically converted into VCPs to create a redundant link with a dedicated VCP connection, you must reboot the Virtual Chassis to avoid creating a traffic loop within the Virtual Chassis. This step is also required if the ports for the redundant link are manually configured into VCPs.

To enable automatic VCP conversion in an EX2200, EX3300, EX4200, EX4500, and EX4550 Virtual Chassis, configure the **auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level on the Virtual Chassis. Subsequently deleting the **auto-conversion** statement returns the Virtual Chassis to the default behavior, in which automatic VCP conversion is disabled.

- **New configuration option to disable automatic Virtual Chassis port conversion (EX4300 and EX4600 Virtual Chassis)**—Starting in Junos OS Release 14.1X53-D47, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in an EX4300 or

EX4600 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

Automatic VCP conversion is not needed when using default-configured VCPs on both sides of the link to interconnect two members. On both ends of the link, you can also manually configure network or uplink ports that are supported as VCPs, whether or not the automatic VCP conversion feature is enabled.

Deleting the **no-auto-conversion** statement from the configuration returns the Virtual Chassis to the default behavior, which reenables automatic VCP conversion.

- See Also**
- [New and Changed Features on page 5](#)
 - [Known Behavior on page 27](#)
 - [Known Issues on page 29](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 88](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)
 - [Product Compatibility on page 90](#)

Known Behavior

The following are changes in known behavior in Junos OS Releases 14.1X53 for the EX Series.

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

- [General Routing](#)
- [High Availability](#)
- [Infrastructure](#)
- [J-Web](#)
- [Platform and Infrastructure](#)
- [User Interface and Configuration](#)

General Routing

- On an EX4600 switch, high ICMP delays are experienced while pinging directly connected IRB interfaces. This is because of a hardware limitation, and the transit traffic is not affected. [PR1164135](#)

High Availability

- During a nonstop software upgrade (NSSU) on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)

Infrastructure

- On EX4550 switches, if you configure IGMP on all interfaces and create a large number of multicast groups, the maximum scale for IGMP can be achieved on some interfaces but not on all interfaces. [PR1025169](#)
- When link protection, node-link protection, or fast reroute is configured on high-traffic MPLS label-switched paths (LSPs), a traffic convergence delay of 680 ms to 1.5 seconds can occur. The link protection provides protection against a link failure along an RSVP LSP. The node-link protection establishes a bypass LSP through a different device. Fast reroute provides redundancy for an LSP path. [PR1039717](#)

J-Web

- In the J-Web interface, you cannot commit some configuration changes in the Ports Configuration page or the VLAN Configuration page because of the following limitations for port-mirroring ports and port-mirroring VLANs:
 - A port configured as the output port for an analyzer cannot be a member of any VLAN other than the default VLAN.
 - A VLAN configured to receive analyzer output can be associated with only one interface.

[PR400814](#)

- In the J-Web interface for EX4500 switches, the **Ports Configuration** page (Configure>Interfaces>Ports), the **Port Security Configuration** page (Configure>Security>Port Security), and the **Filters Configuration** page (Configure>Security>Filters) display features that are not supported on EX4500 switches. [PR525671](#)
- If a Virtual Chassis contains more than six members, the Support Information page (Maintain>Customer Support>Support information) might not load. [PR777372](#)
- While committing an EZsetup, if the system gets disconnected, J-Web reflects the status as success irrespective of the original commit status. [PR866976](#)

Platform and Infrastructure

- On the EX4300 platform, a firewall filter applied on the lo0 interface might not work as expected. The OSPF adjacencies form irrespective of the source address. The same configuration works on the EX4200 as expected. [PR1164711](#)
- On EX4300 switches, support for UNI and NNI functionality on the same physical interface is not available. According to current behavior, UNI and NNI functionality might need to be configured on different physical interfaces. If you attempt to switch the role of the port between UNI and NNI, it is recommended that the physical IFL configuration under the interface be deleted, reconfigured completely (as UNI or NNI), and then the new configuration committed. [PR1214960](#)
- If a packet received at standby bdbrepd process is an invalid or corrupted packet, then it is designed to assert while the packet is validated. Possible causes for this could be an intermittent connectivity issue between the Virtual Chassis links that connects the master and standby Virtual Chassis members. [PR1380772](#)

User Interface and Configuration

- On an EX3300 Virtual Chassis, if you use the J-Web user interface to request support information for all members at the same time, the switch might not be able to retrieve the information. As a workaround, request support information for each member one at a time. [PR911551](#)
- In the Monitor>Interface page, the background color of the graph changes after the page is refreshed. [PR994915](#)
- If you uninstall the J-Web platform package by using the CLI, reinstalling the application package might not restore J-Web. [PR1026308](#)
- In J-Web, Maintain>Update J-Web page, Select Application package>Update J-Web>local file does not work in Internet Explorer (IE) version 9 and later. This issue occurs because of the default security settings in IE version 9 and later. [PR1029736](#)

- See Also**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 22](#)
 - [Known Issues on page 29](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 88](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)
 - [Product Compatibility on page 90](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1X53 for the EX Series.

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

- [General Routing](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Layer 2 Features](#)
- [Layer 3 Features on page 33](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)
- [User Interface and Configuration](#)
- [Virtual Chassis](#)

[General Routing](#)

- When an EX4300 switch connects to any other switch through a 40-gigabit DAC connection, the link might not come up. This issue occurs because the EX4300 switches have autonegotiation enabled on the 40-gigabit DAC interfaces by default, and other switches have auto-negotiation disabled by default. As a workaround, disable autonegotiation on the EX4300 switch, this recovers the connection. When the 40-gigabit interface works as a Virtual Chassis port (VCP) on both sides in Virtual Chassis/VCF scenario, it does not have this issue, and disabling autonegotiation is not required. [PR935197](#)
- On EX4300 switches, a maximum of 5000 supplicants are supported for the dot1xd process. [PR962292](#)
- On EX4600 switches, the Zero Touch Provisioning might take more time to complete because TFTP might take longer time to fetch the required data. [PR980530](#)
- On an EX4300 switch, a MAC address that is specified as part of a MAC-based VLAN is authenticated on an interface—for example, xe-1/1/—on which 802.1X authentication in multiple supplicant mode is configured. However, the same MAC address might not be authenticated on another interface—for example, xe-2/1/1—if the MAC address moves to interface xe-2/1/1 from interface xe-1/1/1. [PR1007589](#)
- On EX3300 and EX4200 switches, captive portal authentication is used to redirect Web browser requests to a login page. After the client is successfully authenticated, there might be a delay of 1–3 minutes before the captive portal redirects the browser to the login page, and sometimes the redirection might fail. [PR1026305](#)
- On an EX3300 switch, if there are multiple Telnet or SSH sessions open, the switch might become unresponsive. [PR1029340](#)
- On a mixed EX4200 and EX4500 Virtual Chassis, an EX3300 Virtual Chassis, an EX6200 switch, or EX8200 switch, during an NSSU, duplicate packets might be seen. [PR1062944](#)
- Substantial traffic losses might occur when executing an NSSU on a mixed EX4200 and EX4500 Virtual Chassis or on an EX3300 Virtual Chassis, an EX6200 switch, an EX8200 switch, or an EX8200 Virtual Chassis. [PR1062960](#)

- When LACP is configured together with MACsec, the links in the bundle might not work. Rebooting the switch might solve the problematic links, but might also create the same issue on other child interfaces. [PR1093295](#)
- The dot1x process authentication might fail on EX3300 switches as the **NAS-Port-Type** attribute in the access-request message is sent as an unknown value. [PR1111863](#)
- On Enhanced Layer 2 Software (ELS) platforms (EX4300 and EX4600), if Q-in-Q tunneling is enabled and if you configure a redundant trunk group (RTG) on a Q-in-Q interface, the RTG configuration cannot be applied—there is a commit check error. [PR1134126](#)
- On EX4300 and EX4600 switches with 802.1X authentication and VoIP configured, the initially VoIP phone authenticates on data VLAN using EAPOL authentication method and then tagged traffic comes for voice VLAN, later on, the dot1x phone does not send any packet in data VLAN. Because of high dot1x reauthentication timer value (default: 3600 seconds), the MAC aging (default: 300 seconds) happens for data VLAN that disconnects the dot1x session and it flushes out both data VLAN and voice VLAN. As a workaround, the timeout interval for the MAC entries should be higher than the re-authentication interval. [PR1146457](#)
- On EX4300 Virtual Chassis, NSSU from Junos OS Release 14.1X53-D35 to Junos OS Release 15.1 is not supported. [PR1148760](#)
- On EX8200 Virtual Chassis, traffic might be lost for multicast and Layer 3 protocols (such as RIP, OSPF, BGP, and VRRP) during NSSU. [PR1185456](#)
- On EX Series switches, eswd scheduler slips might occur when the switch cannot reach the TFTP server to store the DHCP snooping database file. The eswd scheduler slip might affect Layer 2 switching features, such as MAC address learning and spanning-tree protocols, resulting in service impacts. [PR1201060](#)
- Support to include-irb-and-l2 an enhanced DHCP relay for EX Series switches (non-ELS). [PR1203507](#)
- On EX4600 switches, MACsec statistics are collected as part of the Ethernet periodic thread (for example, collected every 1 second), causing significant utilization of the CPU with MACsec sessions. The utilization is also proportional to the number of MACsec sessions, which might result in some problems such as high CPU and MACsec session drop. [PR1247479](#)
- Currently, other than on the EX4600, PIC 1 is not supported on any other platforms inline with QFX5100. The following command cannot be used on PIC 1: **set chassis fpc 0 pic 1 port <x> channel-speed disable-auto-speed-detection**. This results in a commit error, with the following message: **[edit chassis fpc 0 pic 1 port 2 channel-speed] 'channel-speed disable-auto-speed-detection' PIC:1 not valid for Auto-speed-disable mode. error: configuration check-out failed** So if you want to disable auto-channelisation on PIC1, you have to disable auto-speed-detection for whole FPC. **set chassis fpc 0 auto-speed-detection disable**. [PR1362647](#)
- The EX4300 Virtual Chassis systems might fail to register some jnxOperating SNMP OIDs related to the Routing Engines. This behavior is more likely if Virtual Chassis members 0 and 1 (FPC0 and FPC1) are not selected as Routing Engines. [PR1368845](#)

High Availability (HA) and Resiliency

- During a NSSU on an EX4300 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)
- In a rare scenario, GRES might not reach the ready state and might fail to start, because the Routing Engine does not receive the state ACK message from the Packet Forwarding Engine after performing GRES. This is a timing issue. It might also stop Routing Engine resource releasing and then cause resource exhausting. As a workaround, reboot the system if this problem occurs. [PR1236882](#)
- On an EX4300 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)

Infrastructure

- On EX4550 switches, 128-byte packets are dropped if the CPU is at 97 percent load or greater. Packets of different sizes are not dropped under these conditions. [PR862767](#)
- The **SNMP_RTSLIB_FAILURE: mdb_conv_update_seq_err: mdb conv: socket seq** error is thrown while performing SNMP walk continuously. [PR986613](#)
- On EX4500 and EX4200 Virtual Chassis, if you configure an IPv4 GRE interface on an IPv6 interface, the GRE tunnel might not work properly. In this case, traffic is not forwarded through the tunnel. [PR1008157](#)
- On EX4550 switches, if you configure IGMP on all interfaces and create a large number of multicast groups, the maximum scale for IGMP can be achieved on some interfaces but not all interfaces. [PR1025169](#)
- Ping does not work on an interface after it is disabled and reenabled. [PR1039743](#)
- On EX4200 switches, when MPLS is configured on an interface and if you configure CoS behavior aggregate (BA) classifiers on the same interface, the BA classifiers might not work. As a workaround, use the multifield classifier instead of BA. [PR1044470](#)
- On EX4300 switches, traffic might be lost for Layer 3 protocols (such as RIP, OSPF, BGP, and VRRP) during a NSSU. [PR1065405](#)
- On EX Series switches except EX4300/EX4600/EX9200, after eswd process restart with Ethernet ring protection switching (ERPS) configured on trunk interface, traffic drop might be seen on the default VLAN. [PR1207047](#)
- On EX8200 switches functioning as designated routers, if the source multicast interface is a VLAN and contains two 10-Gbps active link aggregated interfaces as VLAN members then some PIM groups might not be able to send out native multicast traffic because of an incorrectly programmed Packet Forwarding Engine. [PR1209585](#)
- On an EX4300 Virtual Chassis, during an upgrade, failover, or switchover operation on the backup Routing Engine member, VM core files and ksyncd core files generated might be generated and the following log message: **/kernel: Nexthop index allocation failed: regular index space exhausted**. [PR1212075](#)

Layer 2 Features

- On EX4200 switches, after a switch reboot, a Q-in-Q tunneling interface might not function as expected. The problem occurs when the interface is a member of a P-VLAN with mapping set to swap and is also a member of a non-private VLAN. The PVID of the access interface does not get set when the P-VLAN is configured before the non-private VLAN. The problem does not occur when the non-private VLAN is configured before the P-VLAN. [PR937927](#)

Layer 3 Features

- When the same MAC address is learnt across different VLAN on different interfaces by issuing the **clear ethernet-switching table** command, l2ald core files are generated. [PR1381739](#)

Platform and Infrastructure

- On EX4300 and EX4600 switches, if a remote analyzer has an output IP address that is reachable through a route learned by BGP, the analyzer might be in down state. [PR1007963](#)
- When you delete the Virtual Chassis port (VCP) connecting an EX4300 switch to the Virtual Chassis, the EX4300 switch splits from the Virtual Chassis. To add the EX4300 switch back into the Virtual Chassis, enter the **request virtual-chassis reactive** command to take the switch out of linecard role and then enter the **request virtual-chassis vc-port set pic-slot <slot-number> port <port-number>** command to create the VCP. [PR1013386](#)
- On EX4300 switches, Layer 2 traffic is dropped in some cases. [PR1157058](#)
- On an EX4300 switch, packets received on a Layer 2 interface might be dropped if their destination MAC address matches the MAC address of the destination Layer 3 interface. [PR1162277](#)
- On EX4300 switches, when a loopback filter has a term that matches the destination IP address of the transit multicast packets, the transit multicast packets might hit this rule. [PR1163745](#)
- On an EX4300 Virtual Chassis switch with the Ethernet port configured as a VCP, if you issue a command to take the PIC offline and then bring it back online, then the VCP is still down. [PR1184981](#)
- On EX4300 switches and EX4300 Virtual Chassis, the Hot Standby Router Protocol (HSRP) packets might be dropped in a VLAN if IGMP snooping is configured. As a workaround, configure the switch to flood the multicast address 224.0.0.2. [PR1211440](#)
- On an EX4300 switch, if you install a firewall filter with filter-based forwarding rules to multiple bind points, available TCAM might be exhausted. In this case, the filter is deleted from all the bind points. As a workaround, apply the filter to the bind points with a series of commits, and apply the filter to some of the bind points with each commit. [PR1214151](#)
- On EX4300 switches with redundant trunk groups (RTG) configured, Layer 3 protocol packets such as OSPF or RIP packets might not be sent. [PR1226976](#)

- On EX4300 switches, when a policer with the action of loss of priority is applied to the lo0 interface, all ICMP packets might be dropped. [PR1243666](#)
- A cosmetic alarm might be seen in an EX4300 Virtual Chassis when the `/var` partition file on the backup member is full. The alarm is seen for more than 10 minutes (default) even when the `/var` file has a lot of space available. [PR1354007](#)

Routing Protocols

- On EX4500 switches, if you initiate a BGP session with a peer that is not configured and the peer autonomous system is a member of a confederation group, the rpd process generates a core file. As a workaround, configure a peer for each peer in the confederation autonomous systems. [PR963565](#)

User Interface and Configuration

- On an EX Series switch using the J-Web interface, the J-Web interface might pause indefinitely after STP, RSTP, or MSTP is selected from the Configure>Switching>Spanning tree menu. [PR1046051](#)
- When you install Junos OS Release 14.1X53-D47, the J-Web auto-upgrade notification feature does not work in the J-Web platform package. As a workaround, use the manual upgrade procedure: 1) Go to Maintain>Update J-Web>Download application package. Junos OS Release 14.1X53-A3 package download begins. 2) After the download completes, browse and select the package file and click the **update** button. [PR1372705](#)

Virtual Chassis

- On EX4200 and EX4500 Virtual Chassis, if the Virtual Chassis contains more than two members and the Virtual Chassis members connect as ring topology with redundancy dedicated VCP links, the Virtual Chassis might flood burst BUM traffic out of interface when the backup Routing Engine reboots. [PR1064483](#)
- On EX4300 Virtual Chassis, the message `/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 process that is not active). As a workaround, you can use a syslog filter to mask the messages. [PR1209847](#)

- See Also**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 22](#)
 - [Known Behavior on page 27](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 88](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)
 - [Product Compatibility on page 90](#)

Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

- [Resolved Issues: Release 14.1X53-D48 on page 35](#)
- [Resolved Issues: Release 14.1X53-D47 on page 37](#)
- [Resolved Issues: Release 14.1X53-D46 on page 44](#)
- [Resolved Issues: Release 14.1X53-D45 on page 46](#)
- [Resolved Issues: Release 14.1X53-D44 on page 48](#)
- [Resolved Issues: Release 14.1X53-D43 on page 49](#)
- [Resolved Issues: Release 14.1X53-D42 on page 52](#)
- [Resolved Issues: Release 14.1X53-D40 on page 54](#)
- [Resolved Issues: Release 14.1X53-D35 on page 65](#)
- [Resolved Issues: Release 14.1X53-D30 on page 72](#)
- [Resolved Issues: Release 14.1X53-D27 on page 79](#)
- [Resolved Issues: Release 14.1X53-D26 on page 80](#)
- [Resolved Issues: Release 14.1X53-D25 on page 81](#)
- [Resolved Issues: Release 14.1X53-D16 on page 85](#)
- [Resolved Issues: Release 14.1X53-D10 on page 86](#)

Resolved Issues: Release 14.1X53-D48

General Routing

- On EX4500 switches, the **interface disable** command might not work for an SFP+ interface if it is connected to another vendor's network device. [PR1354673](#)
- On EX3200, EX3300, EX4200, and EX4500 platforms, the output of the **show interface ge-x/x/x extensive** command shows **Duplex: Half-duplex** in the output when **link-mode** is set as automatic or not set (default). This is a display issue and has no service impact. [PR1364659](#)
- On EX4600 Virtual Chassis running Junos OS Release 14.1X53-D43 to Junos OS Release 14.1X53-D47, the **show interfaces ae <interface name> extensive** command might display duplicate entries for member interfaces. [PR1369713](#)
- On EX Series platform, if redundant trunk group (RTG) is enabled with a large-scale MAC address, a MAC refresh frame might not be sent out from the new primary link after RTG failover by deactivating the former primary link on the peer side. [PR1372999](#)

- After the upgrade to Junos OS 15.1R6-x one of the port (ge-0/0/22) is not coming up on standalone EX4550 switch. The reason being MACSec bypass is not set on the port for unknown reason. [PR1380991](#)
- The SFP-T transceivers from the vendor are not recognized by Junos OS on EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, resulting in their invalidation. [PR1382825](#)

Infrastructure

- On EX4300 switches, the l2cpd process might crash when a nonexistent logical interface is being parsed. [PR1198334](#)

Interfaces and Chassis

- On EX4300 and EX4600 switches, MC-LAG peer might not send ARP request to the host. [PR1360216](#)

Layer 2 Features

- On EX4200, EX3300, and EX4500 Virtual Chassis, eswd core files might be observed on the backup Routing Engine. [PR1346566](#)
- On EX3300, EX3200, EX4200, EX4500, EX4550, EX6200, and EX8200 switches, if the firewall filter is configured with the **dot1q-tag** match condition on certain values, the filter might not work correctly. [PR1369592](#)
- On EX4550 switches, the **eswd[1200]: ESWD_MAC_SMAC_BRIDGE_MAC_IDENTICAL: Bridge Address Add: XX:XX:db:2b:26:81 SMAC is equal to bridge mac hence don't learn** error message is seen in the syslog every few minutes on the ERPS owner. Because the log is caused by an ERPS PDU in the ERPS setup, you can ignore the message. [PR1372422](#)

Network Management and Monitoring

- On EX4300 switches, the event-policy-generated traps are sent with UTC, even though **timezone** is defined under system hierarchy. [PR1380777](#)

Platform and Infrastructure

- On EX4300-32F fiber with SFP-T transceiver installed, the corresponding ports might be still up after system halt. This is specific to the case where the SFP-T transceiver is installed in one of the first 32 ports (built-in ports). [PR1354857](#)
- On EX4300 switches, in a rare situation the remote interface starts flapping unexpectedly. [PR1361483](#)
- On EX4300 and EX4600 switches, during any client's dynamic VLAN membership creation in a dot1x scenario, the l2ald process might crash. [PR1363964](#)
- On an Enhanced Layer 2 Software (ELS) platform, if **storm-control** is configured under **ether-option**, after BUM traffic is sent, which exceeds the storm control limit, the interface might physically shut down and might not be brought back even though **recover-timeout** has been configured. [PR1364654](#)

- On an EX4300 platform, the Packet Forwarding Engine might crash after frequent MAC moves or when the sequence of MAC learning and deleting is continuously performed, which eventually causes memory exhaustion. [PR1367141](#)
- On an EX4300 switch in an RSTP senario, if you set wrong bridge-id as RSTP bridge-id. It might cause loops in the networks. [PR1383356](#)
- On EX4300 and EX4600 Virtual Chassis, the IRB interface that is associated only with master chassis interface might not turn down when the master chassis is rebooted or halted. [PR1381272](#)

Virtual Chassis

- In a mixed Virtual Chassis (VC) configuration, the fast failover does not work properly if the link between EX4500 and EX4200 is restarted, and this might cause traffic loss. [PR1353908](#)

Resolved Issues: Release 14.1X53-D47

Authentication and Access Control

- On EX Series Virtual Chassis switches, the LLDP TLV for MAC/PHY configuration status displays the MAU (medium attachment unit) as unknown [PR1185137](#)
- In Power over Ethernet (PoE) using Link Layer Discovery Protocol (LLDP) scenario, the LLDP Power-via-MDI TLV and LLDP Media Endpoint Discovery (LLDP-MED) TLV transmits the wrong Power Class type. [PR1296547](#)
- If dynamic assignment of VoIP VLAN is used, the switch might not send correct VoIP VLAN information in LLDP MED packets after any configuration change and commit. [PR1311635](#)

General Routing

- On dual Routing Engine platforms, if changes occur on an aggregated Ethernet interface that results in marking ARP routes as down (for example, bringing down one of the member links) due to an interface state pending operation issue on the backup Routing Engine, in a race condition, the backup Routing Engine might crash and reboot with the following error message: panic:rn timer_index_alloc: nhindex XXX could not be allocated err=X. [PR1179732](#)
- If you issue the command "request system snapshot" on a Virtual Chassis, some Virtual Chassis members might go down if traceoption or syslog is enabled, due to the snapshot copy causing a CPU-busy condition with multiple kernel errors and also the Virtual Chassis Control Protocol (VCCP) adjacency going down. [PR1180386](#)
- In case EX4550-32T is configured with 100m fixed speed without auto-negotiation, sometimes an interface does not come up. At that time, the peer device which is supporting Auto-MDI does not detect correctly and causes link down. [PR1235868](#)
- On EX4300 switches, there is no option to enable DHCP snooping without having to enable other port security features such as IP source guard or DAI. [PR1245559](#)
- Junos OS: Short MacSec keys may allow man-in-the-middle attacks (CVE-2018-0021); Refer to <https://kb.juniper.net/JS10854> for more information. [PR1251909](#)

- On EX4200 and EX3200 platforms using PSU module EX-PWR3-930-AC, the PSU is not detected by the show chassis hardware command and is listed as "absent" in the show chassis environment command output. [PR1256980](#)
- DHCP request/discover duplication between L2 interfaces on Junos OS Release 15.1R5 [PR1268550](#)
- On EX4200 Virtual Chassis: Memory Leak for chassisd. [PR1285832](#)
- The jhdcpd process might create a core-dump if dhcpv6-security is configured and client is sending dhcpv6 packet with rapid commit. [PR1287074](#)
- In EX2200, EX3300, EX4200, EX4500, and EX4550 platforms with a Virtual Chassis environment, the Simple Network Management Protocol (SNMP) output for some SNMP values (for example, CPU, memory, temperature, and so on) might not be read any more if the member ID is changed from (0,1) to other IDs. [PR1299330](#)
- On EX2200 Series switches, when Redundant Power System (RPS) is connected and not powered on, the Small form-factor pluggable (SFP) interface might flap and this has impact on traffic forwarding. [PR1307748](#)
- On EX3300 platform, when a network port is used for a Virtual Chassis port, it does not work properly. Once it goes down, it does not come up even though it is physically correct. This issue has been seen by only using network port and this issue has service impact. [PR1310819](#)
- Traffic drop occurs on sending L3 traffic across MPLS LSP. [PR1311977](#)
- On EX2200/EX3300/EX4200/EX4500/EX4550/EX8200 Virtual Chassis platform, an interface MAC address might not be restored the configuration is deleted or rolled back. The issue might cause the hardware address and current address to not be the same. [PR1319234](#)
- EX Series switches do not send RADIUS requests after configuration of an interface-range change. This might cause some hosts to remain in a connecting state and not get authenticated. [PR1326442](#)
- All the multicast traffic must use only Queue 8 (mcast-be). In rare scenarios, multicast traffic may be forwarded via other multicast queues than mcast-be queue (Queue number 8) in EX4300. This has been fixed during device initialization. [PR1347232](#)
- On EX4600, analytics feature for queue-monitoring does not work properly and might generate the log though the current latency value does not cross the high latency threshold value. [PR1348749](#)

High Availability (HA) and Resiliency

- In a rare scenario, GRES might not reach the ready state and might fail to start, because the Routing Engine does not receive the state ack message from the Packet Forwarding Engine after performing GRES. This is a timing issue. It might also stop Routing Engine resource releasing and then cause resource exhausting. [PR1236882](#)

Infrastructure

- On EX4300 switches, if you configure a firewall filter policer with action forwarding-class on an egress filter, the software might allow the configuration to commit although that action is not supported. [PR1104868](#)
- On EX series switches except EX4300/EX4600/EX9200, bootp packets with dhcp options are dropped when DHCP snooping is enabled. [PR1173118](#)
- On an EX4300 switch or an EX4300 Virtual Chassis that has a generic routing encapsulation (GRE) tunnel configured on an integrated routing and bridging interface (IRB), the associated GRE statistical counters might not be updated after the GRE interface is deactivated and then reactivated. [PR1183521](#)
- A maximum of 8 elements within all the "vlan-id-list" can be accepted on a physical interface for EX4300 due to product limitation. This is regardless of the span of each element (a single VLAN or a range spanning 500 VLANs) and whether 8 elements are on a single logical unit or over 8 logical units (or somewhere in between). If there are more than 8 elements configured on a physical interface then the configuration is accepted but only the first 8 elements will be acted upon. The behavior remains the same only a warning message is displayed. However, the commit check succeeds. [PR1225281](#)
- On EX2200/EX3300 Series switches, if configuring one IPv6 firewall filter with source-address/destination-address on lo0 (loopback 0) interface, but Packet Forwarding Engine result in the IPv6 firewall filter not working correctly for combination of some IPv6 address and subnet mask. [PR1247149](#)
- On EX4300 switches, AE interface with LACP enabled might be down if the interface member VLAN is PVLAN. [PR1264268](#)
- On EX4600 Series switches in Virtual Chassis scenario, when em0 interface of FPC member is connected to another FPC Layer 2 (L2) interface of the same Virtual Chassis, it can be seen that no connectivity for management is provided by em0 interface. [PR1299385](#)
- On standalone EX2200/EX3200/EX3300/EX4200/EX4500/EX4550/EX6200/EX8200 or Virtual Chassis (Virtual Chassis) with these switches, when Ternary Content-Addressable Memory (TCAM) is in "out of memory space" condition, pfem might generate a core file when adding a new route entry in TCAM. [PR1304299](#)
- On EX Series switches except EX4300/EX4600/EX9200, if Spanning Tree Protocol (STP/RSTP/MSTP/VSTP) is configured, the topology change (for example: connecting one downstream device) might cause the pfem process to crash. [PR1312042](#)
- On EX2200/EX3300/EX3200/EX4200/EX4500/EX4550/EX6200/EX8200 platforms, file system corruption might happen if bad blocks are in the flash/filesystem. The

upgrade might fail. In order to avoid this, additional steps during installation are added to improve the existing behavior. [PR1317628](#)

- On EX4600, priority-based flow control (PFC) frames might not work. [PR1322439](#)
- ifinfo core files can be created on EX4600 Virtual Chassis [PR1324326](#)
- On EX2200/EX3300/EX3200/EX4200/EX4500/EX4550 platforms, high CPU load for the sfd process might be seen if the high rate of ARP packets are received (e.g. 500pps) and IGMP Snooping is enabled for that VLAN. [PR1325026](#)
- Support for archiving dmesg file, as of today only the Last reboot logs are recorded. [PR1327021](#)
- In a Virtual Chassis composed of EX4200, EX4500, or EX4550 switches, if two member switches are already connected with a dedicated VCP link and a redundant VCP link is added between the two members using uplink ports converted into VCPs, traffic might loop in the Virtual Chassis. The issue can occur whether the redundant link is added intentionally or inadvertently due to miscabling, and whether the link is converted into a VCP link manually or by the VCP automatic conversion feature. As a workaround to stop the looping behavior, reboot the Virtual Chassis after adding the additional VCP link, or reboot the Virtual Chassis after correcting the miscabling and removing unintentional VCP settings. NOTE: When enabled, VCP automatic conversion is invoked if the Virtual Chassis is preprovisioned, LLDP is enabled on the ports on both sides of the link, and the ports on both sides of the link are network ports that are not already converted into VCPs. [PR1346438](#)

Interfaces and Chassis

- On EX4300-VC platforms, the MAC address assigned to an AE member interface is not the same as that of its parent AE interface upon master RE halt. [PR1333734](#)
- On EX4600/QFX5100 platform, if the ICL link is configured on a single interface (such as GE-0/0/0, without LAG) and one member of MC-LAG is down, and both MC-LAG peers are rebooted, packets might drop on ICL of MC-LAG peer where MC-LAG is up. [PR1345316](#)

Layer 2 Features

- The eswd process might crash after a Routing Engine switchover in an EX Series Virtual Chassis scenario. The crash happens due to disordered processing of a VLAN/vmember by eswd and L2PT modules. As the order of processing does not remain the same every time, the crash is random across the switchover. [PR1275468](#)
- In x Spanning Tree Protocol (xSTP) scenario on EX4500/EX4550, some ports might not come up on PIC 1 or PIC 2 when the third PIC is inserted. [PR1298155](#)
- On EX2200/EX3200/EX3300/EX4200/EX4500/EX6200/EX8200/QFX3500/QFX3600 Series switches, ERPS route update fails during moving of a new non-ERPS member interface associated with a VLAN between the Ethernet ring protection switching (ERPS) groups, and it can result in the traffic stop. [PR1301595](#)

MPLS

- On QFX5100 switches, unified ISSU is not supported with MPLS configuration. [PR1264786](#)

Network Management and Monitoring

- On EX2200/EX3300/EX4200/EX4500/EX4550/EX8200/XRE200 platform configured with sFlow and mac-radius authentication, MAC authentication requests might incorrectly be sent because transit DHCPv6 traffic is picked up by the sFlow agent. [PR1298646](#)

Platform and Infrastructure

- On ungraceful exit of telnet (quit/shell logout), perm and env files created by pam were not deleted. [PR1142436](#)
- Starting with 14.1X53-D47, on EX4300, MSTP can be configured on the VLAN running on RTG as long as MSTP is not enabled on the RTG interfaces. [PR1176352](#)
- On EX4300 switches and EX4300 Virtual Chassis, Hot Standby Router Protocol (HSRP) packets might be dropped in a VLAN if IGMP snooping is configured. As a workaround, configure the switch to flood multicast 224.0.0.2. [PR1211440](#)
- On an EX4300, EX4600, EX9200, or QFX5100 standalone switch or its Virtual Chassis or VCF, with a port configured in access mode and with dot1x enabled, if this port is converted to trunk mode, then this port might not be able to learn a MAC address or might drop packets silently. [PR1239252](#)
- On EX4300 platform with power redundancy N+N mode, PoE interfaces flap when any side power supply unit (PSU) is removed and only left one PSU. [PR1258107](#)
- On all EX Series platforms, the configuration of speed and auto-negotiation properties might not be committed successfully if it is applied on a group of interfaces [PR1258851](#)
- On EX4300 Series switches with flexible-vlan-tagging and extended-vlan-bridge configured a traffic blackhole might be observed if vlan-id is not matched between a logical interface and VLAN configuration. [PR1259310](#)
- On mixed Virtual Chassis (Virtual Chassis) / Virtual Chassis Fabric (VCF), QFX5100 works as RE (Route Engine) and EX4300 works as Line Card. The knob "interface-mac-limit" configured for interfaces on EX4300 does not work. [PR1259634](#)
- On EX4300 Virtual Chassis, a 10-gigabit VCP might not get a neighbor after a system reboot [PR1261363](#)
- On EX4300 Series switches with MLD (Multicast Listener Discovery) snooping enabled, NS (Neighbor Solicitation) messages might be dropped, which can result in IPv6 Ping not work. [PR1263535](#)
- In Virtual Chassis scenario, when the master member FPC reboots and the interface on which the ARP is learned goes down along with the master FPC, traffic loss might be observed for about 10 seconds. At that time, the ARP entry cannot be learned from the remaining FPC. [PR1283702](#)

- On EX4300 Series switches, the Filter-Based Forwarding (FBF) might not work properly after deactivating/activating. [PR1293581](#)
- On EX4300 switches, packets whose size is larger than 1452 bytes will be dropped after generic routing encapsulation (GRE). [PR1293787](#)
- On EX4300 switches, when unknown unicast ICMP packets are received by interface, packets are getting routed, so TTL is getting decremented. [PR1302070](#)
- On EX4300 platform, OSPF packets with IEEE P-bit 6 might change to 0 while being received if OSPF is configured on VLAN-tagged L3 interfaces or IRB interfaces. [PR1306750](#)
- On EX4300 platform with PIM and IGMP-Snooping enabled on IRB interface, if an IGMPv2 report which creates (*G) entry is sent first, then multicast data traffic for the same group is sent, the multicast receiver connected to EX4300 might not be able to get the multicast streaming. [PR1308269](#)
- Autonegotiation is not working as expected between EX4300 and SRX5800. [PR1311458](#)
- On EX4300 Virtual Chassis, IGMP snooping might not learn a multicast router interface dynamically if PIM hello messages are received on the interface where igmp-snooping is configured. [PR1312128](#)
- On EX4300 switch, if the interface with 1G SFP port is configured with the no-auto-negotiation option, the interface might stay down after reboot. [PR1315668](#)
- IGMPv3 neighborhood information is now in sync with the kernel entries. [PR1317141](#)
- On EX4300 Virtual-Chassis platform, high latency might be observed between Master RE and other FPC if traffic burst is received on Master RE every 3 to 4 seconds. [PR1319795](#)
- On EX4300 platform, multicast traffic might not be forwarded to one of the receivers if IGMPv3 and IGMPv2 reports are received for the same group on the same VLAN. [PR1323499](#)
- On EX4300 platform with IGMP snooping configured for VLAN all, MAC learning issues might occur and new VLANs creation might fail when loading a large VLAN configuration with different port combinations under each VLAN. [PR1325816](#)
- On all Junos platforms with a LAG enabled, l2cpd might create a core-dump if 'set protocols layer2-control mac-rewrite' or 'set protocols layer2-control bpdu-block' is configured on any of child members of a LAG. [PR1325917](#)
- When exhausting the TCAM (by adding, for example, new prefixes to prefix-list) , warning messages appears and filter is just not programming and doesn't work. But when we deactivate-activate it after that, warning messages also appears, but filter is programming. [PR1330148](#)
- On EX4300 platforms, L2ALD storm control logs will not be generated if the interface is on RTG configuration. [PR1335256](#)
- On EX4300 Series switches, when configuring an interface as Redundant Trunk Group (RTG) backup interface and configuring multicast-router-interface for same interface under igmp-snooping, the loop is generated between RTG interfaces and causes

Internet Group Management Protocol (IGMP) packets going out of RTG backup interface. [PR1335733](#)

- On EX series platforms, when an interface with LLDP (Link Layer Discovery Protocol) and VoIP (Voice over IP) CoS (Class of Service) configuration flaps, memory leak happens in l2cpd (Layer 2 Control Protocol process) due to memory allocated not being released. [PR1337347](#)
- The "show spanning-tree statistics bridge" command output gives 0 for all VLAN instance IDs. [PR1337891](#)
- On EX Series and QFX Series switches, when a media access control (MAC) source address filter is configured with "accept-source-mac", if a MAC move limit is also configured, then the filter will not work as expected. [PR1341520](#)
- In an MSTP scenario with the character size of MSTP region name exceeds 31 characters, any commit might trigger MSTP work abnormally even the configuration change does not relate to MSTP. [PR1342900](#)
- On EX4300 platform, the firewall filter might not be programmed in PFE even though PFE shows enough TCAM entries are available. The system might go into the error state and the failed filter might not work. [PR1345296](#)
- On EX Series switches, when VLAN translation is configured on an interface, the VLAN tag of the traffic, which is destined to the interface itself, will not be rewritten, and then the traffic will be dropped. [PR1348094](#)
- On EX4300 platform, traffic drop might happen if LLC packets are sent with DSAP and SSAP as 0x88 and 0x8e respectively. On other platforms, traffic drop might happen when dot1x is configured. [PR1348618](#)

Routing Protocols

- In a rare condition after a BGP session flaps, BGP updates might not be sent completely, resulting in BGP routes being shown in the advertising-protocol table on the local end but not shown in the receive-protocol table on the remote end. [PR1231707](#)

Spanning Tree Protocols

- On EX8200 platform with dual Routing Engines, rebooting both Routing Engines at the same time with any STP protocol configured, the port might continue to stay in a blocking state if it continues to receive BPDUs from the peer end. As a workaround, please restart the eswd daemon. [PR1305954](#)

User Interface and Configuration

- On an EX Series switch that is supporting the zeroize feature, after the switch is booted up from "request system zeroize" and then a configuration is saved, the saved configuration won't be restored after the switch is rebooted. [PR1228274](#)

Virtual Chassis

- On QFX5100/EX4600 Virtual Chassis or VCF topology, it takes 10 minutes to obtain the Routing Engine role if you reboot the chassis. The issue is seen only when there is offline chassis in Virtual Chassis/VCF topology. [PR1225696](#)

- On EX-Series switches except for EX4300/EX4600 packet drop might be seen during the failover or switchover from the master switch to backup switch in a virtual-chassis (Virtual Chassis). It is because of the delay in ARP update during the failover or switchover of the master Routing Engine (RE). [PR1278214](#)
- On EX2200 Virtual Chassis, system-mac gets changed after mastership RE changes. Due to this MAC inconsistency, unexpected packet loss would be seen for a long time. This issue can be restored by the method: ping from the peer and reboot the master EX2200. [PR1347213](#)

Resolved Issues: Release 14.1X53-D46

General Routing

- In some scenario's, the 36th port in captive portal is not redirecting to the URL as configured. This problem is seen with the 'set system services web-management https system-generated-certificate' configured. [PR1217743](#)
- On EX4300 switch, if igmp-snooping is configured, the IGMP leave packet might be flooded to all ports (including the receive port) in the VLAN. Note: IGMP control packets on snooping enabled device are not supposed to be flooded to all ports in VLAN. [PR1228912](#)
- On EX/MX/QFX-Series platform where MC-LAG with IPv6 is supported, the l2ald memory might leak for every IPv6 ND (Neighbor Discovery) message it receives from a peer MC-LAG and it does not free the memory allocated, causing l2ald memory exhaustion and an l2ald process crash. [PR1277203](#)
- Starting from Junos OS 15.1R3, the 40G link with SR4 transceivers on EX4550 device will fail to come up after a PIC offline/online event or a link UP and DOWN event. [PR1281983](#)
- On EX4600 switches, if interface is configured with 100m speed explicitly and no-auto-negotiation, the interface might be down after reboot. [PR1283531](#)
- MACsec issue: The "show security macsec statistics" command does not show expected results. Statistics are incorrectly cleared for each physical interface (IFD) under eth periodic (1 second). [PR1283544](#)

Infrastructure

- EX-Series Switches, fsck is run with '-C' option which skips the file system corruption check if the partition has been marked clean during the boot 'nand-media' check. Due to this there have been multiple instances where the partition has had file system issues even when cleanly shutdown. This change is to enforce fsck during the boot cycle to strengthen the file system check during boot time. HOW TO RECOVER: * The switch will repair the corruption during the boot cycle when the file system check (fsck) is run. [PR1191072](#)
- A vulnerability in IPv6 processing has been discovered that may allow a specially crafted IPv6 Neighbor Discovery (ND) packet to be accepted by the router rather than discarded. The crafted packet, destined to the router, will then be processed by the routing engine (RE). A malicious network-based packet flood, sourced from beyond the local broadcast domain, can cause the RE CPU to spike, or cause the DDoS protection ARP protocol

group policer to engage. When this happens, the DDoS policer may start dropping legitimate IPv6 neighbors as legitimate ND times out. Refer to JSA10749 for more information. [PR1220211](#)

- When EX4550-32T is booting up, 1G interface would be up for 60 sec and turned down. Then turn up again few sec later. While unexpected link up is seen, peer device would send traffic to that port and cause traffic black hole. [PR1257932](#)
- No space in an EX8200 line card to save pfem core dumps [PR1263024](#)

Interfaces and Chassis

- For EX Series switches, in a rare condition (for example: reboot or reloading configuration), the MAC address of an AE interface and its member links might be inconsistent, causing unexpected behavior for some routing protocols. [PR1272973](#)

Layer 2 Features

- On EX-Series platforms except for EX4300/EX4600/EX9200, the Multiple Spanning Tree Protocol (MSTP) might not be able to detect the topology changes after Nonstop software upgrade (NSSU) process, which might lead to packet loop. The topology change count is shown as 0 after that. user@switch> show spanning-tree bridge msti 2 ++++++Output Snipped+++++ STP bridge parameters for MSTI 2 MSTI regional root : 8194.78:fe:3d:b1:e4:01 Hello time : 2 seconds Maximum age : 20 seconds Forward delay : 15 seconds Number of topology changes : 0 >>>>> showing 0 Topology change last recvd. from : 88:a2:5e:35:70:04 Local parameters Bridge ID : 8194.78:fe:3d:b1:e4:01 Extended system ID : 0 Internal instance ID : 2 [PR1284415](#)

Platform and Infrastructure

- The PFE might coredump on EX4300. When hitting this issue, all the interfaces flap, this causes service impact and traffic drop. [PR1214727](#)
- On EX4300 switches, certain multicast traffic might impact the network, for example, cause OSPF to flap. Issues might occur when multicast packets use the same interface queue as certain network protocol packets (for example, OSPF, RIP, PIM, and VRRP). [PR1244351](#)
- DHCP option 2 is not working if configuring switch as DHCP server. [PR1252437](#)
- On EX4300 Series switches, in virtual-chassis (VC) scenario, pfex might restart while doing master reboot or during Nonstop software upgrade (NSSU) if the old master reboots at the end of NSSU phases. [PR1258863](#)
- The packets with certain UDP destination port might be dropped on EX Series Virtual Chassis except EX4300 or EX4600 Virtual Chassis. [PR1262969](#)
- On EX4300/EX4600/QFX5200/QFX5100/QFX3500/QFX3600 platform, with DHCP relay traffic flowing, CPU usage of pfex_junos might go high. The issue might be seen if DHCP relay function is on and DHCP relay packets are received continually. [PR1276995](#)

Security

- The Juniper Networks enhanced jdhcpd process might experience high CPU utilization, or crash and restart upon receipt of an invalid IPv6 UDP packet. Both high CPU utilization and repeated crashes of the jdhcpd process might result in a denial of service as DHCP service is interrupted. Refer to JSA10800 for further details. [PR1119019](#)
- A buffer overflow vulnerability in Junos OS CLI might allow a local authenticated user with read only privileges and access to Junos CLI, to execute code with root privileges. Refer to JSA10803 for further details. [PR1149652](#)
- Two vulnerabilities in telnetd service on Juniper Networks Junos OS might allow a remote unauthenticated attacker to cause a denial of service through memory and/or CPU consumption. Please refer to JSA10817 for more information. [PR1159841](#)
- Junos: Potential remote code execution vulnerability in PAM (CVE-2017-10615); Refer to <https://kb.juniper.net/JSA10818> for more information. [PR1192119](#)
- Junos: EX Series PFE and MX MPC7E/8E/9E PFE crash when fetching interface stats with extended-statistics enabled (CVE-2017-10611); Refer to <https://kb.juniper.net/JSA10814> for more information. [PR1247026](#)
- On Junos OS devices with SNMP enabled, a network-based attacker with unfiltered access to the Routing Engine can cause the Junos OS snmpd process (daemon) to crash and restart by sending a crafted SNMP packet. Repeated crashes of snmpd process can result in a partial denial-of-service condition. Additionally, it might be possible to craft a malicious SNMP packet in a way that can result in remote code execution. Refer to <https://kb.juniper.net/JSA10793> for more information. [PR1282772](#)

Software Installation and Upgrade

- EX4300 Virtual Chassis: More than expected traffic loss during NSSU. [PR1115398](#)
- New switch added in EX2200 Virtual Chassis is not getting automatic software update from master switch. [PR1270412](#)

System Management

- Netconf syntax error reported if the resync character is split in multiple streams. [PR1161167](#)

Virtual Chassis

- On EX4300 FRU removal/insertion trap not generated for non-master (backup/line card) FPCs. [PR1293820](#)

Resolved Issues: Release 14.1X53-D45

General Routing

- On EX9200 and EX4300 switches, 802.1X supplicants might not be reauthenticated by server fail fallback authentication after the server becomes reachable. [PR1157032](#)
- In 802.1X (dot1x) single-supplicant mode, after username and password were configured on interfaces and dot1x supplicants were started, the users were

authenticated with the Radius_DataVlan VLAN, but the Ethernet-switching table was not updated for one of the interfaces. [PR1283880](#)

Infrastructure

- Some error messages will be seen on the PDB-unsupported platforms. [PR1103035](#)
- On an EX4300 egress VLAN-based firewall filter on a Q-in-Q interface, after a switch reboot, firewall counters might not increment as expected. [PR1165450](#)
- On EX4200VC or EX4500/4550VC, PFE does not update learned MAC to RTG active interface after RTG failover. This issue is seen with RTG which is configured across FPCs in a VC setup. [PR1209323](#)
- When run CLI command "request system snapshot slice alternate" on EX series switches, a timeout error might occur and the CLI command would not output as expected. [PR1229520](#)

Platform and Infrastructure

- On EX4300/4600 and QFX5100 serial platforms, the GRE traffic failed to pass through the GRE tunnel if tunnel destination resolved by BGP which have indirect next-hops. [PR1211189](#)
- On EX4300 switches, problems with connectivity might arise on 100M interfaces set to full duplex and half duplex or on 10M interfaces set to full duplex or half duplex. The links appear, but connectivity to end devices might not work. The port does not transmit packets even though port statistics show packets as transmitted. [PR1212093](#)
- On EX4300 Series switches, Dynamic Host Configuration Protocol (DHCP) with PXE boot server is not working as expected due to PXE unicast ACK packet dropped. The communication between the DHCP client and PXE server might be affected. [PR1230096](#)
- On EX4300 switches, traffic is not forwarded through the GRE tunnel in some cases. [PR1254638](#)
- On EX4300 Series switches with flexible-vlan-tagging and extended-vlan-bridge configured, traffic blackhole might be observed if vlan-id is not matched between an interface IFL and VLAN configuration. [PR1259310](#)
- The jdhcpd might core due to memory leak if Dynamic Host Configuration Protocol (DHCP) security is enabled, and then DHCP relay might stop working. As result, DHCP client might not get IP address from DHCP server. [PR1273452](#)
- On EX4300-VC, when persistent learning with a mac-limit of 1 is enabled on the interface, then the switch might not forward the Internet Group Management Protocol (IGMP) report upstream to the router or any Layer2 device connected via the interface. [PR1285807](#)

Resolved Issues: Release 14.1X53-D44

General Routing

- Dot1x authentication may fail in Ex switches as the NAS-Port-Type attribute in the access-request message is sent as "unknown" value. [PR1111863](#)
- When clients are authenticated with dynamic VLAN assignment on an 802.1X-enabled interface, if they are connected/disconnected within a short time (within sub-seconds), the logical interface and the bridge domain or VLAN might remain in a problematic state, thus cause the clients to be denied when accessing the network. As a workaround, restart the l2-learning process to recover the port/interface from the problematic state. [PR1230073](#)
- On EX Series switches except EX4300, EX4600, and EX9200, the switch cannot send DHCP option 2 when extended DHCP local server is configured. The switch sends DHCP option 2 incorrectly when traditional DHCP server is configured. [PR1252437](#)
- After the MACsec session flaps, data traffic sent over the MACsec-enabled link might not be properly received and the receiving device might report the received frames as "framing errors" in the output of show interfaces command. [PR1269229](#)
- Starting in Junos OS Release 14.1X53-D44, the automatic software update feature can be used to automatically update Junos software on members of an EX2200 Virtual Chassis running Junos OS Release 12.3R12 and later. Automatic software update is not supported on an EX2200 Virtual Chassis in releases prior to 14.1X53-D44. [PR1270412](#)

Infrastructure

- On EX8200 Series switches, if a layer 3 interface is configured vlan-tagging, then the switch might put wrong source mac address when it routes traffic to this layer 3 interface. [PR1262928](#)
- From Junos OS release 13.2X50-D15, for EX Virtual Chassis (VC) switches except EX4300/EX4600/EX9200-VC, when small UDP (<80 bytes) packets are forwarded between endpoints across Virtual Chassis port (VCP) link, a certain User Datagram Protocol (UDP) destination port gets black holed. [PR1262969](#)

Platform and Infrastructure

- On an EX4300, if you install a firewall filter with filter-based forwarding rules to multiple bind points, it might exhaust the available TCAM. In this case, the filter is deleted from all the bind points. You can work around this issue by applying the filter to the bind points with a series of commits, applying the filter to some of the bind points with each commit. [PR1214151](#)
- On a EX4300 switch that has dhcp-service dhcp-snooping file <file name> write-interval 600 configured, there is a possibility that you will see the following logs in file messages, but you wont see any core dumps associated with it: jdhcpd: LIBJNX_EXEC_EXITED: Command stopped: PID 8081, signal='Unknown signal: -1', core dumped, command '/usr/bin/tftp' {master:0}[edit] root# run show system core-dumps fpc0: -----
/var/crash/*core*: No such file or directory /var/tmp/*core*: No such file or directory

/var/tmp/pics/*core*: No such file or directory /var/crash/kernel.*: No such file or directory The same logs won't be seen if the file is found on the local device. [PR1257975](#)

- On Enhanced Layer 2 Software (ELS) and MX platforms, due to a memory leak issue, the l2ald process might crash when many dot1x clients are being reauthenticated, for example, 150 clients with transmit-period set to 5. It is around 40–60 bytes memory leak per reauthentication for one dot1x client. Here the leak is due to the interaction between dot1x and the l2ald process; with more frequent reauthentication and more clients, the crash will be observed more often. [PR1269945](#)
- On Virtual Chassis (VC) based on EX4300/EX4600/EX9200/QFX3500/QFX3600/QFX5100, the irb interfaces which only associated with physical interfaces on master chassis do not turn down when master chassis rebooted or halted. [PR1273176](#)

Routing Protocols

- On EX4600/QFX3500/QFX3600/QFX5000 series switches, when new Filter-based Forwarding (FBF) firewall filter is applied on Integrated Routing and Bridging (IRB) interface which is not L3 interface, or while binding/unbinding the FBF filter on L3 interfaces, the FXPC might hit 100% CPU usage. [PR1263896](#)

Security

- NTP.org and FreeBSD have published security advisories for vulnerabilities resolved in ntpd (NTP daemon). Server-side vulnerabilities are only exploitable on systems where NTP server is enabled within the [edit system ntp] hierarchy level. A summary of the vulnerabilities that might impact Junos OS is in JSA10776. Refer to JSA10776 for more information. [PR1234119](#)

Virtual Chassis

- When you add an EX4300 switch to the VCF, the following error message is seen: ?ch_opus_map_alarm_id alarm ignored: object 0x7e reason?. [PR1234780](#)
- When the linecard role FPC is removed and rejoined to the Virtual Chassis immediately, the LAG interface on the master/backup would not be reprogrammed in the rejoined FPC. [PR1255302](#)
- On an EX4550 switch in a Virtual Chassis configuration, fast-failover function for VCP will work properly when you initially add this configuration. However, if the device is rebooted, the function would not take effect next time. [PR1267633](#)

Resolved Issues: Release 14.1X53-D43

General Routing

- According to the IEEE, edge port feature has been supported from 802.1w and later xstp, so it should be removed under "protocol stp" hierarchy. [PR1028009](#)
- On EX4300-VC/EX4600-VC/QFX5100-VC and other unconfirmed EX/QFX Virtual Chassis, when a switchover with GRES enabled is performed, this warning might appear: All Packet Forwarding Engines are not ready for RE switchover and may be reset. [PR1158881](#)

- On EX4300 Series switches with DHCP snooping enabled, if ND inspection (neighbor-discovery-inspection) is configured together, the routing table entry for IRB (on which a connection with a DHCPv6 server is configured) might be removed. [PR1201628](#)
- During boot up, EX4200, EX4550, and EX4300 switches might have no display or might display gibberish on the LCD. It is an LCD corruption issue. [PR1233580](#)
- On EX4600 or QFX5100 standalone switches and Virtual Chassis, MACsec connections are deleted randomly after a switch reboot, optics removal, deactivation or activation of a MACsec configuration, or fxpc process restart. [PR1234447](#)
- On QFX5100 and EX4600 switches, if traceroute is used between endpoints and the path travels through a GRE tunnel, hops in the tunnel are displayed by an asterisk in the traceroute output. [PR1236343](#)
- On EX4600/QFX3500/QFX3600/QFX5100 Series switches, MACsec statistics are collected as part of the Ethernet periodic thread (for example, collected every 1 second), causing significant utilization of the CPU with MACsec sessions. The utilization is also proportional to the number of MACsec sessions, which can result in some problems such as high CPU and MACsec session drop. [PR1247479](#)
- On an EX Series switch or a Virtual Chassis with 802.1X (dot1x) enabled, in a scenario with more than 254 clients (suplicants), plenty of clients might be going to the server-reject VLAN and have limited access to the server-reject VLAN although the clients have correct credentials. For a few authenticated clients, the authentication method might be displayed as "Server-Reject" although the client was authenticated in the correct VLAN---that is, the data VLAN. [PR1251530](#)
- Dot1x EAP clients not getting authenticated when there is a high number of authentication requests sent from switch. [PR1259241](#)

Class of Service (CoS)

- On QFX5100/EX4600/EX4300 Series switches, if forwarding-class-sets with more than one forwarding-classes is applied to interface, and the scheduler for these forwarding-classes under this forwarding-class-sets are not configured with shaping-rate, then it might cause traffic to be dropped for this interface. [PR1255077](#)

Infrastructure

- On EX Series switches except EX4300/EX4600/EX9200, if there are a large scale of IFLs (logical interface) configured, eswd (ethernet-switching) daemon may not learn MAC addresses after switch rebooted. [PR1248051](#)
- QFX5k/EX4600/EX4300 series, when the system received traffic when the TTL is 1 and DF bit been set, (eg, reply for a tracerouter), the system will reply with " ICMP Destination Unreachable (Fragment needed) " and "MTU 0"

```
user@root> ping 20.20.20.2 ttl 1 do-not-fragment PING 20.20.20.2 (20.20.20.2): 56 data bytes 36 bytes from 20.20.20.2: frag needed and DF set (MTU 0) Vr HL TOS Len ID Flg off TTL Pro cks Src Dst 4 5 00 0054 3c0d 2 0000 01 01 ed71 20.20.20.1 20.20.20.2 ..
```

[PR1251523](#)

Interfaces and Chassis

- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating incorrect configuration. This issue is fixed and it is now not allowed to configure the same IP address (the length of the mask does not matter) on different logical interfaces. [PR1221993](#)

Multiprotocol Label Switching (MPLS)

- In MPLS layer 2 or layer 3 VPN scenario, QFX5100/EX4600 Series switches work as PE router and the core interface of PE using IRB interface. When deactivating/disabling/deleting underline member interface of the IRB, and if the (parent) IPv4 nexthop is uninstalled first before cleaning up the (child) MPLS nexthop, the fxpc process might crash and restart. And the FXPC core will be seen. [PR1242203](#)

Network Management and Monitoring

- On EX Series switches except EX4300/EX4600/EX9200, when RTG (Redundant Trunk Group) switchovers are done, then the /var/log/shadow.log or /var/log/shadow_debug.log is rotated. And it might cause PFE process to crash. [PR1233050](#)
- 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 might impact Junos OS is in JSA10776. Refer to JSA10776 for more information. [PR1234119](#) , [PR1159544](#)
- After the reboot of the EX4600 Virtual Chassis, authentication of SNMPv3 users fails due to the change of the local engine ID. [PR1256166](#)

Platform and Infrastructure

- On EX4300 switches, Layer 2 traffic is dropped in some cases. [PR1157058](#)
- 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](#)
- The interface which is configured with 1G and **no-auto-negotiation** might be down after reboot on EX4300 switch. [PR1223234](#)
- On QFX5100 and EX4600, password required for user **root** even after SSH public key authentication is enabled. [PR1234100](#)
- On EX4300 switches, when a policer with the action of loss of priority is applied to the lo0 interface, all ICMP packets might be dropped. [PR1243666](#)
- An SFP+ might not be recognized after an EX4300 reboots. [PR1247172](#)

- The egress PE device (EX4300) sends out LLDP frames toward the CE device with the destination MAC address of 01:00:0c:cd:cd:d0, which is a duplicated frame and is rewritten by the ingress (PE) device. [PR1251391](#)
- Mastership switchover from member0 to member1 doesn't generate SNMP trap or log message. All other combination does generate. [PR1253408](#)
- On EX4300 switches, traffic is not forwarded through the GRE tunnel in some cases. [PR1254638](#)
- On EX4300 switches, if a filter is configured with a policer action (for example, with action **then loss-priority low** or **discard**) and is applied to the lo0 interface, BGP transit packets might hit the TTL0 entry of the loopback filter and might be dropped. [PR1258038](#)

Subscriber Access Management

- On EX2200/EX3300 switches, the authd core file is created during the authenticating. And the authentication might be failed, also it can lost all session data previously maintained. [PR1241326](#)

Resolved Issues: Release 14.1X53-D42

Firewall Filters

- On EX Series switches, the dfwc (daemon that performs as a firewall compiler) might fail to get filter information from the kernel in COMMIT_CHECK (configuration validation) mode. As a result, the filter index is regenerated starting from index 1. This will create the mismatch of filter index as compared to the existing filters in the system. [PR1107139](#)

General Routing

- On EX Series switches, Ethernet switching process (eswd) scheduler slips might occur when the switch cannot reach the TFTP server to store the DHCP snooping database file. The eswd scheduler slips might affect Layer 2 switching features, such as MAC address learning and spanning-tree protocols, resulting in service impacts. [PR1201060](#)
- On EX Series Virtual Chassis that support PoE, when the master Routing Engine member is rebooted, PoE devices connected to the master might not come back online after the reboot. As a workaround to avoid this issue, when configuring PoE interfaces, use the set poe interface all configuration command instead of configuring specific interfaces individually. To recover connections after seeing this issue, disable and reenabale the ports affected by the issue. [PR1203880](#)
- IGMP-Snooping is for IPv4 and should not affect IPv6 multicast traffic. On EX4300, EX4600, and QFX5100 switches in a Virtual Chassis configuration, IPv6 multicast packets might be affected and 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 EX4300 switches with DHCP relay configured, DHCP return packets, e.g. DHCPREPLY and DHCPOFFER, that are received across a GRE tunnel might not be forwarded to clients, which can impact DHCP services. [PR1226868](#)
- On an EX4300 switch or a Virtual Chassis with 802.1X (dot1x) enabled, in a scenario with more than 254 clients (supplicants), plenty of clients might be going to the server-reject VLAN and have limited access to the server-reject VLAN although the clients have correct credentials. For a few authenticated clients, the authentication method might be displayed as "Server-Reject" although the client was authenticated in the correct VLAN---that is, the data VLAN. [PR1251530](#)

Infrastructure

- 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 generated and see the log message `"/kernel: Nexthop index allocation failed: regular index space exhausted"`. [PR1212075](#)
- When you load and commit a configuration on an EX2200 or EX3300 switch, the system might automatically go into db mode. As a result, you might not be able to access the switch through SSH, and a vmcore file is generated. [PR1237559](#)
- On EX Series switches, CoS (for example: rewrite rules) are unbound from the IFL (logical interface) when deleting ISIS interface and hence device is not marking traffic correctly. [PR1239827](#)

Network Management and Monitoring

- On EX4600 switches, when temperatures for FPCs are polled, the temperatures might not be polled for all SNMP members. [PR1232911](#)

Platform and Infrastructure

- On EX4300 virtual chassis switch, the ethernet port is configured as VCP port (virtual-chassis port), if issuing PIC offline then online command, the VCP port is still down. [PR1184981](#)
- On EX4300 switches, if a Layer 3 interface receives a frame with the CFI/DEI bit set to 1, this frame might be dropped and not be processed further. [PR1237945](#)
- On EX4300 Series Switches, when a VLAN is mirrored, the mirrored packets may contain 38 additional bytes. The IP address in this packet is randomly generated and may appear as one of many existing, valid IP addresses on the Internet. It may appear as ERSPAN as well, which is a proprietary non-Juniper protocol. These addresses and packet types can be ignored. They may appear as alerts in certain IDP / IDS's and packet analyzer applications which can be ignored. [PR1239635](#)
- On EX4300 switches, problems of connectivity might arise on 100 Mbps interfaces set to full/half duplex or on 10 Mbps interfaces set to full/half duplex. The interface will show up, but connectivity to end devices might not work. The interface does not transmit packets even though interface statistics show packets are transmitted. [PR1249170](#)

Routing Protocols

- On EX4500 switches, if you initiate a BGP session with a peer that is not configured and the peer autonomous system is a member of a confederation group, the routing protocol process (rpd) generates a core file. As a workaround, configure a peer for each peer in the confederation autonomous systems. [PR963565](#)
- On EX4300 switches, with redundant trunk groups (RTGs) configured, Layer 3 protocol packets such as OSPF or RIP packets might not be sent. [PR1226976](#)

Virtual Chassis

- On EX4300 Virtual Chassis, 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 process that is not active). As a workaround, you can use a system-logging (syslog) filter to mask the messages. [PR1209847](#)
- On member switches in an EX Series Virtual Chassis, the `"request virtual-chassis vc-port set"` CLI command allows specifying an invalid or non-existent Virtual Chassis port (VCP) interface name. An entry with the invalid VCP interface name is added to the database, and the CLI command `"show virtual-chassis vc-port"` displays these entries with the invalid VCP interface names, but these entries cannot subsequently be removed. [PR1215004](#)

Resolved Issues: Release 14.1X53-D40

General Routing

- On EX/QFX Series switches, when dot1x is configured, the dot1xd process may crash while authenticating a large number of mac users. [PR984182](#)
- On an EX Series or an MX Series platform acting as a DHCPv6 server, the server does not send a Reply packet after receiving a Confirm packet from the client; the behavior is not compliant with the RFC3315 standard. [PR1025019](#)
- Certain QFX and EX Series devices do not pad Ethernet packets with zeros, and thus some packets can contain fragments of system memory or data from previous packets. This issue is also known as 'Etherleak' and often detected as CVE-2003-0001. Refer to JSA10773 for more information. [PR1063645](#)
- On EX2200 switches, if you issue the CLI command `"request system services dhcp release interface-name"`, an IP address release message DHCP packet is sent from the client and processed at the server. At the same time, the client clears the IP address on the same interface, and the clearance of the IP address on the interface leads to acquisition of a new IP address from the server. If you then issue the CLI command `"show system services dhcp client interface-name"`, the output of that command indicates that the issued operational command had no impact. [PR1072319](#)
- On EX4500, EX4550, EX6200, and EX8200 switches, if you replace an SFP (1G) optic with an SFP+ (10G) optic on one port, the adjacent port might go down. For example, install an SFP optic on port-0/0/36 and port-0/0/37. If you replace the SFP optic with

an SFP+ optic on port-0/0/36 and port-0/0/37, then port-0/0/36 might go down during insertion of the SFP+ optic on port-0/0/37. [PR1073184](#)

- By enabling this configurations, it Drops spanning-tree protocol BPDUs (for STP, MSTP, and RSTP) entering any or a specified interface The BPDU drop feature can be specified only on interfaces on which no spanning-tree protocol is configured. This behavior is same as EX platforms. [PR1084116](#)
- When LACP is configured together with MACsec, the links in the bundle might not all work. Rebooting the switch might solve the problematic links, but could also create the same issue on other child interfaces. [PR1093295](#)
- If MAC move limit is configured to drop traffic, QFX and EX Series switches might forward traffic instead of dropping traffic when the MAC move limit is exceeded. [PR1105372](#)
- On EX4300 Virtual Chassis, after you have run an NSSU, the master might detect the backup coming up after an upgrade and reprograms the trunk, even though the backup member links are down. The traffic might drop when the master tries to push the traffic through trunk members that are not up yet, and the traffic resumes once the links are up. [PR1115398](#)
- On EX4300 switches, when storm-control or storm-control-profiles with action-shutdown is configured, if the storm-triggered traffic is control traffic such as LACP, the physical interface will be put into an STP blocking state rather than turned down, so valid control traffic might be trapped to the control plane and unrelated interfaces might be set down as an LACP timeout. [PR1130099](#)
- On the EX8200/EX8200-VC platform, restart chassis-control (chassisd) several times on master Routing Engine in EX8200 after reboot, or restart chassis-control on master Routing Engine in line card chassis (LCC) of EX8200-VC after LLC reboot, then executing 'show snmp mib walk .1.3.6.1.4.1.2636.3.1.6' might returns nothing. This issue only impacts SNMP for specific MIB, it does not have any service/ traffic impact. [PR1140495](#)
- On EX/QFX virtual chassis, NSSU from a pre-14.1X53-D30 release to the 14.1X53-D30~D34, or from 14.1X53-D30~D34 to 14.1X53-D35 and later, the upgrade might hang and do not complete. [PR1142275](#)
- On EX Series Switches except EX4300/EX4600, configuring custom MAC address for VLAN interface via "set interface vlan mac x". When changing the family for any interfaces to inet, the MAC address for VLAN interface might be get changed unexpectedly. If this issue happens, the source MAC address of routed traffic (forwarded traffic) will get changed with a different address. [PR1143299](#)
- In a scenario which an EX4200 switch fails to communicate with a dot1x server. If the switch receives EAPOL packets from clients. The switch restarts the authentication process. [PR1147894](#)
- On EX4300/EX4600/EX9200/QFX5100/MX Series platforms configured for 802.1X authentication, if the VLAN assigned to an access port is changed, then the supplicants authenticated are disconnected and the users are not able to authenticate anymore. [PR1148486](#)

- On an EX3300 switch, if you configure IGMP snooping with a VLAN that is not on the switch, the configuration does not commit. [PR1149509](#)
- On QFX/EX4600 Series switches, in a rare timing condition, if there was already a request to gather some info from the QSFP and remove it at the same time, the packet forwarding engine manager (fxpc) might crash. [PR1151295](#)
- On EX2200/3300, negative temp value may be displayed incorrectly in "show chassis environment ". At that time, you may see the temperature value over 200 degrees C.
root@EX2200> show chassis environment 2 Class Item Status Measurement Power
FPC 0 Power Supply 0 OK Temp FPC 0 CPU OK 254 degrees C / 489 degrees F FPC
0 Exhaust Area OK 7 degrees C / 44 degrees F FPC 0 EX-PFE1 OK 254 degrees C / 489
degrees F FPC 0 EX-PFE2 OK 1 degrees C / 33 degrees F FPC 0 Local Intake OK 247
degrees C / 476 degrees F FPC 0 Remote Intake OK 248 degrees C / 478 degrees F
[PR1157692](#)
- During dot1x process, the access reject from authenticated the server-reject-vlan is assigned to the voip vlan. This memory is not freed during the Port-based Network Access Control client data cleanup results in the memory leak, and it seems can't learn mac addresses and created a dot1x core file. [PR1160059](#)
- An insufficient authentication vulnerability on platforms where Junos OS instances are run in a virtualized environment, may allow unprivileged users on the Junos OS instance to gain access to the host operating environment, and thus escalate privileges. [PR1161762](#)
- On EX8200 Virtual Chassis with a DC power supply for the external Routing Engine (XRE200), when one DC power supply fails, no logs or SNMP traps are generated. [PR1162165](#)
- On EX2200/EX3300/EX3200/EX4200/EX4500/EX4550/EX6200/EX8200 series switches configured dhcp-relay and dhcp-snooping under specific routing-instance, the DHCP client requests his IP address via a broadcasted DHCP discover and receives an ip address correctly. When the lease time is running half over, the DHCP client sends a Unicast DHCP request to the DHCP server that gave him his IP address. The request works and both the DHCP server and client show that the lease time is renewed, but the switch acting as relay-agent does not update the dhcp snooping binding table with the renewed lease time. This causes the lease time to run out in the dhcp snooping binding table, even though the lease time has not run out on the DHCP server, then the switch starts dropping all of the DHCP clients packets, so the end user loses connectivity to the network. [PR1162941](#)
- On a QFX5100 switch with an integrated routing and bridging (IRB) interface configured as a Layer 3 interface and with two hosts (Host A and Host B) connected to the switch, if you deactivate the IP address on Host A and then configure the same IP address on Host B, the outgoing interface of the IP address might not be changed in the ARP table. [PR1166400](#)
- On EX4600/QFX Series switches, if vlan-rewrite is configured on the ingress interface (e.g. vlan-rewrite 21 12), VLAN translation might not work in the reverse direction, this causes source device to drop traffic due to VLAN mismatch. [PR1168525](#)
- If the configuration is pushed to an EX Series switch using ZTP (Zero Touch Provisioning), then after a subsequent reboot, the configuration might be deleted. [PR1170165](#)

- On EX4600/QFX5100 switches, when a VLAN is mirrored, the mirrored packets may contain 38 additional bytes. The IP address in this packet is randomly generated and may appear as one of many existing, valid IP addresses on the Internet. It may appear as ERSPAN as well, which is a proprietary non-Juniper protocol. These addresses and packet types can be ignored. They may appear as alerts in certain IDP / IDS's and packet analyzer applications which can be ignored. [PR1170589](#)
- Multiple ports cannot be selected from Jweb at a time. [PR1170640](#)
- 1G fiber link ports might be down with MACsec configured on EX4300 Series switches when EX4300 is rebooted. [PR1172833](#)
- On Virtual Chassis with a fixed-configuration switch as members, traffic loss might occur for about 10-15 seconds when the master leaves the Virtual Chassis for an upgrade during NSSU. [PR1173754](#)
- On EX series switch except EX9200/EX4300/EX4600, if POE is configured, when one IP phone (Aastra 6725ip phone) is connected with POE interface, the phone cannot receive PoE power from switch. [PR1174025](#)
- On EX4300, EX4600, and EX9200 switches, when root guard is in effect or cleared, there is no appropriate syslog message. [PR1176240](#)
- On EX3300 and EX4200 switches, after the "request system zeroize media" command has been executed, J-Web might not work. [PR1177214](#)
- If you issue the command "request system snapshot" on a Virtual Chassis, some Virtual Chassis members might go down if traceoption or syslog is enabled, due to the snapshot copy causing a CPU-busy condition with multiple kernel errors and also the Virtual Chassis Control Protocol (VCCP) adjacency going down. [PR1180386](#)
- On EX2200-C switches, during a software upgrade to Junos OS Release 14.1X53-D35 or 15.1R3, the error messages "Triggering freezing circuitry" or "Triggering overheat circuitry" might be generated after rebooting, and then the switch shuts down. [PR1183631](#)
- On EX Series switches except for EX4300/4600/9200, while processing xstp-disabled interface with BPDU block configuration, current code flow sets the bpdud_control flag for RSTP enabled interfaces as well. This might result in RSTP enabled port getting blocked on receipt of a BPDU. [PR1185402](#)
- The issue was caused by the inconsistency resulting between the HAL and HALP layers. When storm-control is programmed on an IFL and if the IFL is Link-Down, further programming (HALP programming) is thwarted until the IFL comes up. Any change applied on the storm-control instance attached to that IFL at this stage results in this inconsistency. "By default the storm control configuration is enabled on the switch on all interfaces, let's say if the link is down and if we unbind the default storm control configuration when the link stays down, then can we go to this state" [PR1187271](#)
- On EX4200 Virtual Chassis, when the an interface flaps and it has "hold-time up" configured over a long period of time (for example, 16 days), a chassis manager (chassism) process memory leak might occur due to the incorrectly accumulated task timer. About 128 bytes of the process leak every time the memory leak is triggered. [PR1188403](#)

- Customer may see JDHCP core dump on EX4300-VC running on JUNOS 14.1X53-D35 and DHCP relay service. The issue can be seen after the ungraceful reboot of the VC. The DHCP service shall work normal after the DHCP process is running. [PR1190258](#)
- In case EX4300 handles DHCP packet which needs to relay, dhcpd might be crashed with core dump. [PR1192735](#)
- When try to add more than 24 POE devices into EX4300 in a mixed EX4600-EX4300 or QFX5100-EX4300 VC, the new device is powered on, and then another port goes down. [PR1195946](#)
- On EX4300, EX4600, QFX3500, QFX3600, QFX5100 platforms, when any type of spanning tree (STP, RSTP, MSTP, or VSTP) is configured, the MAC address part of the bridge ID might be set to all zeros (for example, 4096.00:00:00:00:00:00) after you power cycle the device without issuing the "request system halt" command. As a workaround, issue the restart l2-learning command. [PR1201293](#)
- On EX4300/EX4600/EX9200/QFX3500/QFX3600/QFX5100 platform, if the platform itself acts as DHCP (Dynamic Host Configuration Protocol) relay agent, and DAI (Dynamic ARP Inspection) is enabled as well. The jdhcpd might crash and restart if DHCP packets with invalid hardware address are received. It has service impact. As a workaround, please disable all DHCP-security features or remove the client which is sending DHCP packets with the invalid MAC address. [PR1206241](#)
- A vulnerability in IPv6 processing has been discovered that may allow a specially crafted IPv6 Neighbor Discovery (ND) packet destined to an EX Series Ethernet Switches to cause a slow memory leak. A malicious network-based packet flood of these crafted IPv6 NDP packets may eventually lead to resource exhaustion and a denial of service. [PR1206593](#)
- On an EX4600 switch, when you remove the 40GBASE-ER4 QSFP+ module, the show chassis hardware command still shows that the module is inserted. [PR1208805](#)
- If a QFX5100 switch or VCF is configured with IGMP snooping but not with any PIM-related configuration, a mcsnoopd memory leak might occur when the device receives PIM Hello packets that need to be forwarded further. When PIM hellos are arriving on the device, 12 bytes are allocated for every PIM hello packet, increasing the amount of memory consumed by the mcsnoopd process. As a workaround, either restart the mcsnoopd process or apply a firewall filter that discards PIM packets on the loopback (lo0) interface of the device in the input direction. [PR1209773](#)

Authentication and Access Control

- LLDP packets not received on ae interface. [PR1168970](#)

Infrastructure

- On EX8200 Series switches, in rare condition, the Packet Forwarding Engine manager (pfem) might crash and generate the core file when changing Protocol Independent Multicast (PIM) mode from sparse mode to dense mode or vice versa. [PR1087730](#)
- On EX8200 switches, the pfem process might crash and generate a core file. This might impact traffic. [PR1138059](#)

- On EX Series switches except EX4300/EX4600/EX9200 with ERPS (Ethernet ring protection switching) configured, no VLAN will be included in data-channel if data-channel has not been explicitly configured, then MAC flush might not happen for any data VLAN while receiving an SF (signal failure). This might lead to a traffic issue before the MAC address has aged out. [PR1152188](#)
- On EX8200-VC, if an interface has L3 config no matter with or without vlan-tagging, the switch uses incorrect source mac address to send traffic, the incorrect mac that is being used is not the mac for the chassis, just a mac from a different interface. [PR1153858](#)
- On an EX2200/EX-3300/EX-4200/EX-4500/EX4550 Virtual Chassis with its ingress on the backup and egress on the master (vice versa), the connectivity might be lost after the new VLAN deletion/addition process where the "vlan-pruning" is already enabled for all VLANs due to the VCP member programming in hardware is modified. [PR1167170](#)
- On EX Series switches with ERPS (Ethernet ring protection switching) configured (except EX4300/4600/9200), many SF (signal failure) packets might appear in a link-end ring node during a link failure that occurred for a short time. [PR1169372](#)
- If port firewall filters are configured locally for the interface, then VSAs (vendor-specific attribute) take precedence when they conflict with the filters. If the VSAs and the local port firewall filters do not conflict, they are merged. For Filter term merge involving dot1x Filters, 'Implicit Discard' terms are not merged with other terms and are appended towards the end of the final merged firewall filter. For Ex8200 this applies for both normal Filter merges involving VACL /PACL and also dot1x Filter term merges. [PR1171441](#)
- EX4300 may generate a Core dump if the QFSP's are inserted [PR1172350](#)
- On EX4300 switches with firewall loopback rule "ip-options", only "any" is available for ip option match. [PR1173347](#)
- On EX4300 switch, if configuring one filter to discard MSTP BPDU, it cannot drop MSTP packet. [PR1184343](#)
- On an EX8200 Virtual Chassis, doing Routing Engine failovers before booting up the line cards might cause the VLAN interface MAC address to be automatically and incorrectly set to "00:00:00:00:00:01". [PR1185678](#)
- On EX4600 and QFX5100 switches that are configured with native-vlan-id, the switch sends untagged traffic. But if you delete native-vlan-id, the switch keeps sending untagged traffic. [PR1186436](#)
- On EX4300/EX4600/QFX3500/QFX3600/QFX5100 switches with vlan-rewrite configured on an AE interface, a VLAN rewrite might fail and result in traffic loss. [PR1186821](#)
- On EX4300 Virtual Chassis, when upgrading from Junos OS Release 15.x to Release 16.x via NSSU, the backup or any line card upgrades first to a new image, and then the old master might have an upgrade failure, and keep rebooting. [PR1190164](#)
- On EX8200 Series switches with redundant trunk group (RTG) configured, when the RTG active link fails over to the standby link, during which it triggers the MAC addresses

movement from old active link to new active link, this might result in packets drop on peer device due to MAC lookup failure. [PR1194318](#)

- On EX4300 Series switch, if vlan-rewrite is configured on aggregated links, it does not work. [PR1194585](#)
- On EX4600/QFX3500/QFX3600/QFX5k series switches, when traffic enters a MPLS interface and is destined to the loopback interface in a routing instance, the firewall filter might not work properly. [PR1205626](#)
- On EX and QFX platforms, firewall filters with syslog might not work, because as part of packet processing, packets were incorrectly mapped to the pcmd queue instead of the DFW queue. [PR1208491](#)
- On EX4200VC or EX4500/4550VC, PFE does not update learned MAC to RTG active interface after RTG failover. This issue is seen with RTG which is configured across FPCs in a VC setup. [PR1209323](#)
- When a VLAN is configured on all switches in a ring topology, traffic loop might occur after removing the VLAN from one of the switches. [PR1229744](#)

Interfaces and Chassis

- In the bridge domain configuration with IRB interface environment, the IRB interface INET/ISO MTU is set to 1500. When the MTU on IRB interface is deleted, the MTU wouldn't be changed. [PR990018](#)
- Fixed in 14.1X53-D40 [PR996005](#)
- If an EX4550-32F switch in a Virtual Chassis reboots and comes online, LACP interfaces on any member of the Virtual Chassis might go down and not come up. [PR1035280](#)
- On an EX4300 switch or Virtual Chassis, when the I2C bus gets stuck, the chassisd daemon might get stuck and become unresponsive. If you issue a chassisd-related show command, the command returns an error message: "error: the chassis-control subsystem is not responding to management requests". [PR1038830](#)
- You might be unable to commit your configuration if you modify the subnet of an IP address on an IRB interface by using the "replace pattern" command. [PR1119713](#)
- The 40G copper QSFP link between EX4300 and QFX10002 with 3M/5M DAC has issue, interface remains down on QFX10002 side but interface keep flapping on EX4300 side. [PR1177888](#)
- On Junos based platforms, if static MAC is configured on an interface, after device reboot, chassisd might overwrite this MAC address to default MAC, this behavior has impact on the service/traffic. [PR1186478](#)

Layer 2 Features

- The Packet Forwarding Engine manager daemon (FXPC) might crash on a QFX5100 switch if multiple processes attempt to access the Ethernet-switching table/database at the same time. [PR1146937](#)
- On QFX5100 and EX4600 switches, every time a MAC address is learned, some messages might be output to syslog and be repeated frequently. The logged messages have no impact on service traffic. [PR1171523](#)
- On EX Series switches, if the device is configured as a part of an Ethernet ring protection switching (ERPS) ring, deactivating or deleting the ERPS configuration might cause traffic to stop forwarding through one or more VLANs. [PR1189585](#)
- An EX Series Switch might not process ERPS PDUs that are received from other nodes. This could lead to ERPS ring not operating correctly. [PR1190007](#)

Layer 3 Features

- On a switch that have secure-access-port configured, when you change MTU size of interfaces and commit, there is a high probability that VRRP session flaps between VRRP Master & Backup. No traffic loss is observed during vrrp mastership switchover [PR1163652](#)
- On EX2200/EX3300/EX3200/EX4200/EX4500/EX4550/EX6200/EX8200 series switches, when VRRP configuration changes from ethernet-switching to inet family and vice-versa, then the local IP of master VRRP switch can not be reached on backup VRRP switch and vice-versa. But virtual IP is always reached on both switches. [PR1171220](#)
- On EX Series Switches, configuring MPLS might result in the following messages being logged. These messages do not imply operational impact and can safely be ignored.
Jun 15 15:11:04 Switch rpd[1420]: task_get_socket: domain AF_UNIX type SOCK_STREAM protocol Otask_connect: addr /var/run/mplssoamd_control: No such file or directory
[PR1192238](#)
- GARPs were being sent whenever there is a mac (fdb) operation (add or delete). This is now updated to send GARP when interface is UP & l3 interface attached to the vlan.
[PR1192520](#)

Network Management and Monitoring

- On EX Series switch except EX4300/EX4600/EX92xx, SNMP walking the VLAN info with long name might cause the eswd (ethernet switching process) process to crash. [PR1157338](#)

Platform and Infrastructure

- On EX4300 switches with sFlow configured, some harmless log messages regarding sFlow might be seen continuously. [PR1116568](#)
- On a EX4300-VC platform, if a Q-in-Q S-VLAN interface with MC-LAG is configured, when the backup EX4300 is acting as master, you might lose connection to the management IP address through the interface. As a result, management traffic will be dropped. [PR1131755](#)

- On EX4300 switches and EX4300 Virtual Chassis, PIM register messages are not forwarded to a rendezvous point (RP) when the RP is not directly connected to the first-hop router of the multicast source. [PR1134235](#)
- This is an enhancement of existed feature to display Static Egress and Untagged ports. In the Q-BRIDGE-MIB, we represent VLAN port membership in ASCII string containing comma-separated ASCII-encoded numbers that are indexes into the dot1dBase table. But as per RFC4363b dot1qVlanStaticEgressPorts, dot1qVlanStaticUntaggedPorts is supposed to return a PortList for each VLAN which is defined as a bit vector octet string where each bit represents a single port on the switch. Each octet stores the status of 8 ports. Bit 1 = first port, Bit 2 = second port, etc. up to the total number of ports on the switch. If the bit value is 1, then that port is a member of the VLAN. If the bit value is 0, then that port is not a member of the VLAN. By Default, dot1qVlanStaticEgressPorts, dot1qVlanStaticUntaggedPorts will be displaying the vlan details in ASCII format (older behaviour). CLI syntax: set switch-options mib dot1q-mib port-list bit-map/string. And, if customer wants in RFC4363b way, the above CLI needs to be configured so that MIBs (dot1qVlanStaticEgressPorts, dot1qVlanStaticUntaggedPorts) will return the values in bit format. Please find Description of PR-1149118 for more clarification. [PR1149118](#)
- On ARM platforms such as EX3300, configuring internal IPsec security associations containing authentication hmac-sha2-256 can cause a kernel alignment exception. [PR1149565](#)
- On EX4300 Series switches, the filter is configured under loopback interface with the parameter of "destination-address" and action of "count" or "log", if the destination address of transit packets is matched with filter destination-address, the filter will count or log these packets. [PR1149670](#)
- On an EX4300 Virtual Chassis with Q-in-Q enabled, when "vlan-id-list" is configured on a C-VLAN interface and, for example, if the VLAN range vlist element is in [1-3] or [5-50], C-VLAN traffic is not sent properly across the Q-in-Q network from the C-VLAN interface. [PR1159854](#)
- On EX4300 switches, when xSTP is configured, if you unplug a loopback cable between ports of different FPCs and then plug it back in, the interface might go down and a BPDU error might be detected on this port, causing traffic to drop on another egress port. As a restoration workaround, clear the Ethernet-switching table. [PR1160114](#)
- On EX4300 Series switches, few SSM Multicast (PIM source-specific multicast) streams doesn't egress out downstream port even though port is present in igmp-snooping membership. [PR1162054](#)
- On EX4300 switches, if IGMP snooping is enabled, packets with destination 224.0.0.0/24 might be dropped, except for well-known addresses (for example, 224.0.0.5/6 for OSPF). [PR1167859](#)
- On EX4300 Series switches, ICMP-tagged packets might be seen on the egress interface of a PVLAN access port. The correct behavior is that traffic is sent untagged. [PR1169116](#)
- On EX4300 Series switches, when dhcp-security is enabled on the VLAN, the Unicast packets (e.g. DHCP Offers and ACKs) might be forwarded to all ports in the VLAN. [PR1172730](#)

- On EX4300 switches with IGMP snooping enabled with flexible-vlan-tagging configured on ingress and egress interfaces for passthrough multicast traffic, IGMPv2 membership report messages might not be forwarded from the receiver to the sender. [PR1175954](#)
- On EX4300 series switches, if an ethernet port receives a frame with CFI/DEI bit set to 1, then this frame would not be bridged to an untagged (access) port but it could be bridged to a trunk port. [PR1176770](#)
- When IGMP snooping and storm control are enabled, EX/QFX is supposed to forward traffic with destination IP address 224.0.0.0/24 to all the ports on the VLAN. But for EX4300 except for the well-known addresses in this range, for example, 224.0.0.5/6 for OSPF, 224.0.0.20 for VRRP, all other multicast traffic with a destination in 224.0.0.0/24 will be dropped. [PR1176802](#)
- The commit synchronize command fails because the kernel socket gets stuck. [PR1177692](#)
- On EX4300 or EX4300-VC platforms, if VLAN Spanning Tree Protocol (VSTP) is configured, when some operations about VSTP (for example: Deactivating/activating VSPT interface, Deactivating/activating VSPT VLAN and so on) are done, it might cause pfex process crash. [PR1178539](#)
- If you upgrade the Power over Ethernet (PoE) firmware on a member of an EX4300 Virtual Chassis, the PoE firmware upgrade process might fail or get interrupted on that member switch. You can recognize that this problem has occurred if the member switch is not listed in the command output when you issue the "show poe controller" command. The problem is also indicated if you issue the ?show chassis firmware detail? command and the ?PoE firmware? version field is not shown in the output or has a value of 0.0.0.0. As a workaround, upgrade the Junos software to a release marked as fixed in this PR, and then upgrade the PoE firmware on the affected member switch. To confirm PoE firmware has been successfully upgraded and to check the version, issue the command "show chassis firmware detail". [PR1178780](#)
- An unauthenticated root login may allow upon reboot when a commit script is used. A commit script allows a device administrator to execute certain instructions during commit, which is configured under the [system scripts commit] stanza. Please Refer to <https://kb.juniper.net/JSA10835> for more information. [PR1179601](#)
- On EX4300 switches, if there is a mismatch in the speed configuration between two interfaces, the link might be autonegotiated to half-duplex mode instead of full-duplex mode. [PR1183043](#)
- On EX4300 series switches configured dscp and 802.1p rewrite-rules on an interface, if deleting 802.1p rewrite-rules from this interface, but still 802.1p rewrite is happening along with dscp rewrite. [PR1187175](#)
- On EX4300/EX4600 and QFX Series switches with VSTP enabled for multiple VLANs and participated in a VSTP topology. When the BPDU packets are received on Packet Forwarding Engine (PFE) from other switches, the switch will send BPDU packets to Routing Engine (RE) for further VSTP computing. But, in rare cases, the switch might not send VSTP packets for all VLANs to RE. For example, for an uncertain VLAN, BPDU packets are not reaching RE, even though VSTP is enabled for that VLAN. This will result in this VLAN consider itself as the root bridge. And further, advertising itself as

the root bridge and sending BPDUs to other VSTP switches, other switches might block related port. This might not follow the network design. [PR1187499](#)

- When EX4300 switch receives VRRP advertisement packet on trunk port, but it has different vlan-tag with that port, EX4300 switch may transfer it along with vlan-tag instead of dropping it. [PR1192800](#)
- On EX4300 switches, if you configure a policer on the loopback filter, host-bound traffic might drop even though the traffic does not exceed the specified limit. [PR1196822](#)
- When you install an SFP in an operating EX4300 switch, the SFP might be recognized as either unsupported or as an SFP+-10G. As a workaround, reboot the switch. [PR1202730](#)
- On EX4300 switches, a firewall filter might not be programmed correctly when multiple action modifiers (such as forwarding-class, priority, loss-priority) are performed in the same firewall filter term. [PR1203251](#)
- On EX4300/4600 and QFX5100 serial platforms, the GRE traffic failed to pass through the GRE tunnel if tunnel destination resolved by BGP which have indirect next-hops. [PR1211189](#)
- On EX4300 switches, problems with connectivity might arise on 100M interfaces set to full duplex and half duplex or on 10M interfaces set to full duplex or half duplex. The links appear, but connectivity to end devices might not work. The port does not transmit packets even though port statistics show packets as transmitted. [PR1212093](#)
- On EX4300/EX4600/EX9200/QFX Series switches, if you activate dhcp-security features for IPv6, a JDHCPD core file might be generated. [PR1212239](#)
- On EX4300 switches, EBGp packets with ttl=1 and non-EBGP packets with ttl=1, whether destined for the device or even transit traffic, go to the same queue. In the event of a heavy inflow of non-EBGP ttl=1 packets, occasionally valid EBGp packets might be dropped, causing EBGp to flap. [PR1215863](#)
- On an EX4300 switch, a loopback policer might not work. [PR1219946](#)

Routing Protocols

- For devices populated with master and backup Routing Engines (RE) and configured for nonstop active routing (NSR) and Protocol Independent Multicast (PIM) configuration, the routing protocol process (rpd) might crash on the backup Routing Engine due to a memory leak. This leak occurs when the backup Routing Engine handling mirror updates about PIM received from the master Routing Engine deletes information about a PIM session from its database. But because of a software defect, a leak of 2 memory blocks (8 or 16 bytes) might occur for every PIM leave. If the memory is exhausted, the rpd may crash on the backup Routing Engine. There is no impact seen on the master Routing Engine when the rpd crashes on the backup Routing Engine. Use the "show system processes extensive" command to check the memory. [PR1155778](#)
- ALL traffic destined for leaked route are forwarded to CPU. The traffic expected to be treated as transit. Ping between the default routing instance and routing-instance which leaking routes via the rib-group takes around 40-90 ms. juniper@abc:~\$ ping 10.1.1.1 PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data. 64 bytes from 10.1.1.1: icmp_seq=1

tll=62 time=17.9 ms 64 bytes from 10.1.1.1: icmp_seq=2 tll=62 time=42.3 ms 64 bytes from 10.1.1.1: icmp_seq=3 tll=62 time=15.3 ms 64 bytes from 10.1.1.1: icmp_seq=4 tll=62 time=18.0 ms 64 bytes from 10.1.1.1: icmp_seq=5 tll=62 time=21.1 ms [PR1167156](#)

- On EX4300-VC platform with IGMP-snooping enabled, when IGMP-hosts subscribe for the same group, IGMP queries might not go through between the line card and the and master. [PR1200008](#)

Virtual Chassis

- SDK can raise false alarms for parity error messages like "soc_mem_array_sbusrdma_read" & "soc_ser_correction: mem write" on QFX5100. [PR1161821](#)
- On EX3300 Virtual Chassis, the vcp-snmp-statistics configuration statement is not visible in the [edit virtual-chassis] hierarchy. [PR1178467](#)

Resolved Issues: Release 14.1X53-D35

General Routing

- Setting link speed to 100 Mbps does not work in the following situations: - When network interfaces are used on an EX4600 switch - When an EX4600-EM-8F expansion module is installed in a QFX5100-24Q switch or an EX4600 switch [PR1032257](#)
- If precision-timers and traceoptions are enabled for BGP then both main-thread and precision-timers pthread try to rotate the same tracefile without taking any locks. As a result all the status commands for rpd & krt may timed-out. [PR1044141](#)
- On EX4300 switches, if a Gigabit Ethernet interface is directly connected to an MX104 management interface (fxp0), the physical link will be down. [PR1069198](#)
- On QFX series switches, the wrong source IP address is being used when the switch initiates traffic when em0 is configured with a 192.168.1.XXX/16 subnet. [PR1071517](#)
- On an EX2200 or EX3300 switch on which Dynamic Host Configuration Protocol (DHCP) relay is enabled, when a client requests an IP address, the system might generate a harmless warning message such as: /kernel: Unaligned memory access by pid 19514 [jdhcpd] at 46c906 PC[104de0] . [PR1076494](#)
- On an EX4300 Virtual Chassis or a mixed mode Virtual Chassis with an EX4300 as a member, if you disable root login connections to the console port by issuing the "set system ports console insecure" command, users can still log in as root from the backup and linecard members of the Virtual Chassis. [PR1096018](#)
- On EX4600 switches, the EX4600-EM-8F expansion module interfaces might not come up when the module is removed and re-inserted or when the PIC is taken offline and then brought online. [PR1100470](#)
- On EX Series switches, even though not sending join to specific ports traffic initiated, the traffic might be received in other ports which IGMP-snooping is enabled. It seems that affected interfaces are dynamically recognized as multicast router interfaces for other VLANs. [PR1115300](#)

- On EX4300/EX4600/EX9200/QFX5100/QFX3500/QFX3600, after altering VoIP interface vlan membership configuration by vlan-id, the l2ald core files will be created, and it might have impact on traffic forwarding. [PR1118634](#)
- On EX8200 switches, a nonstop software upgrade (NSSU) might fail during the master Routing Engine upgrade step, and an NSSU process might abort with this message: "mgd: unable to execute /var/etc/reboot.ex: Authentication error". [PR1122628](#)
- On EX Series platform with xSTP enabled, while Edge delay is running, and if Forward delay expires, it might cause interface not be transitioned to FORWARDING immediately, then to be LEARNING state, after the expiry of the second Forward delay timer, it is moved to FORWARDING. Sometimes, it is treated as TC change, and TC is propagated on interface (TCN is observed on this interface). Flapping edge port might cause this issue, and it might cause edge port to be transitioned from non-edge to edge state. [PR1124853](#)
- On EX Series switches with dual Routing Engines or on an EX Series Virtual Chassis, the switch might send multiple proposal BPDUs on an alternate port after a Routing Engine switchover or a nonstop software upgrade (NSSU), resulting in the peer device receiving multiple proposal BPDUs and triggering a dispute condition. The peer port states constantly alternate between "FORWARDING" and "BLOCKING". [PR1126677](#)
- On EX Series switches, if 802.1X authentication (dot1x) is configured on all interfaces, an 802.1X-enabled interface might get stuck in the "Initialize" state after the interface goes down and comes back up, and 802.1X authentication fails. Also, if 802.1X authentication (dot1x) is configured on all interfaces and the no-mac-table-binding configuration statement is configured under the [edit protocols dot1x authenticator] hierarchy level, the dot1x process (dot1xd) might generate core files after it is deactivated and then reactivated, and 802.1X authentication might be temporarily impacted until the process restarts automatically. [PR1127566](#)
- On EX Series switches with bridge protocol data unit (BPDU) protection configured on all edge ports, edge ports might not work correctly and might revert to the unblocking state when the drop option is configured under the [edit ethernet-switching-options bpd-block interface xstp-disabled] hierarchy. [PR1128258](#)
- On EX4300 switches, when there is a redundant trunk group (RTG) link failover, media access control (MAC) refresh packets might be sent out from a non-RTG interface that is in the same VLAN as the RTG interface, and a traffic drop might occur because of MAC flapping. [PR1133431](#)
- On EX Series switches, in rare scenarios, changing the VLAN ID might generate stale or null values in the connectivity fault management daemon (cfmd)'s data structure, and the cfmd might crash. The cfmd is restarted automatically, but the crash might cause a delay in re-establishment of the cfmd functionality. [PR1137453](#)
- On EX Series switches, if multiple source MAC addresses are flooded into a port on which MAC authentication is enabled, a dot1xd process core file might be created. [PR1140634](#)
- On EX Series switches, an interface with a non-Juniper Networks 1000BASE-EX SFP Module-40km might not come up because register values are not set to correct values. This issue only occurs at initial deployment of the switched or when the switch is

upgraded to Junos OS Release 12.3R8, 13.2X51-D30, 14.1X53-D10, 15.1R2 onwards.

[PR1142175](#)

- On EX/QFX virtual chassis, NSSU from a pre-14.1X53-D30 release to the 14.1X53-D30~D34, or from 14.1X53-D30~D34 to 14.1X53-D35 and later, the upgrade might hang and do not complete. [PR1142275](#)
- The Configuration for Generating an Alarm on crossing temperature threshold? was not played to all the members of the VC. Due to which the members of VC were unaware of the configuration. This configuration is now propagated to all members of VC and alarm is raised once the temperature threshold is crossed. [PR1142904](#)
- On EX Series Switches except EX4300/EX4600, configuring custom MAC address for VLAN interface via "set interface vlan mac x". When changing the family for any interfaces to inet, the MAC address for VLAN interface might be get changed unexpectedly. If this issue happens, the source MAC address of routed traffic (forwarded traffic) will get changed with a different address. [PR1143299](#)
- On EX4300/EX4600/EX9200/QFX5100/MX Series platforms configured for 802.1X authentication, if the VLAN assigned to an access port is changed, then the supplicants authenticated are disconnected and the users are not able to authenticate anymore. [PR1148486](#)
- On EX ELS switches, if you change the server-fail vlan, all authenticated supplicants are disconnected. They are then authenticated again, and during this disconnection and reconnection, there is a service impact of 3-4 seconds. [PR1151234](#)
- On EX2200/3300, negative temp value may be displayed incorrectly in "show chassis environment ". At that time, you may see the temperature value over 200 degrees C.
root@EX2200> show chassis environment 2 Class Item Status Measurement Power
FPC 0 Power Supply 0 OK Temp FPC 0 CPU OK 254 degrees C / 489 degrees F FPC
0 Exhaust Area OK 7 degrees C / 44 degrees F FPC 0 EX-PFE1 OK 254 degrees C / 489
degrees F FPC 0 EX-PFE2 OK 1 degrees C / 33 degrees F FPC 0 Local Intake OK 247
degrees C / 476 degrees F FPC 0 Remote Intake OK 248 degrees C / 478 degrees F
[PR1157692](#)

Infrastructure

- On EX2200 and EX300 Virtual Chassis, the Internal state in ERPS is not updated properly in certain conditions . As a workaround, check the interface state and update the ERPS engine accordingly so that they are always in sync. [PR975104](#)
- On an EX Virtual Chassis scenario, the Packet Forwarding Engine manager (pfem) might crash and generate the core files on all VC members due to a NULL pointer check bug. This is a corner issue and it was only seen when configuring 252 OSPFv3 virtual routing and forwarding (VRF) instances sessions via script. The issue was not seen when tried to replicate manually. [PR987682](#)
- On EX4300 switches, traffic sampling is not supported. If you configure traffic sampling, the sampling process (sampled) might generate a core file. [PR1091826](#)
- On EX8200 switches with multicast protocols configured, when a multicast-related (non-aggregated Ethernet) interface goes down and comes back up, ARP installation

for certain hosts might fail because stale entries have not been cleared, and traffic might be lost as well. [PR1105025](#)

- On EX4600 and QFX5100 switches, when the Virtual Router Redundancy Protocol (VRRP) priority is modified to change the VRRP mastership after cosd restart (or device restart), packets might be dropped on interfaces that have both inet and inet6 families enabled. [PR1105963](#)
- On EX4200 and EX4550 switches, the xe- interfaces in a 10-Gigabit SFP+ expansion module (EX4550-EM-8XSFP) or an SFP+ MACsec uplink module (EX-UM-2X4SFP-M) might stop forwarding traffic when the module is removed and re-inserted or the PIC goes offline and comes back online. The issue is resolved in 14.1X53-D35 and 15.1R3 [PR1113375](#)
- On EX4500 switches, if MPLS and CoS behavior aggregate (BA) classifiers are configured on the same interface, the BA classifiers might not work. As a workaround, use multifield (MF) classifiers instead of BA classifiers. [PR1116462](#)
- On EX Series switches, if you deactivate an output interface that is configured with family mpls, a nondefault CoS classifier configured on the interface might be deleted, putting traffic in the wrong queue. [PR1123191](#)
- On EX4300 switches, if the switch works as part of a target subnet, while receiving the targeted broadcast traffic, the packets might be forwarded to the destination with the switch's MAC address as the destination MAC address, rather than the Layer 2 broadcast frame with destination MAC address FFFF.FFFF.FFFF. [PR1127852](#)
- 14.1X53-D30.3 [PR1129896](#)
- On EX8200-VC, if an interface has L3 config no matter with or without vlan-tagging, the switch uses incorrect source mac address to send traffic, the incorrect mac that is being used is not the mac for the chassis, just a mac from a different interface. [PR1153858](#)

Layer 2 Features

- On EX Series switches, the the Ethernet switching process (eswd) might crash when base config is loaded with minimal Ethernet ring protection switching (ERPS) configuration. This is due to a NULL access during Interface Unit (IFL) processing. [PR1066379](#)
- On EX4600 and QFX5100 switches, the VSTP BPDUs might be reinjected to the Packet Forwarding Engine and not be sent out of an interface when the interface has been added to the VSTP configuration and is configured with flexible-vlan-tagging. [PR1117540](#)
- On QFX5100 switches, if you configure a PVLAN inter-switch-link on an existing working trunk port, normal VLAN traffic might break. [PR1118728](#)
- On EX4300, EX4600, and QFX Series switches, traffic received on the backup redundant trunk group (RTG) link might get forwarded to other interfaces following an RTG link failover. [PR1119654](#)
- On EX Series switches, if you configure Ethernet ring protection (ERP) with interfaces configured with vlan members all and commit the changes, then add a new VLAN and

commit the configuration again, the Ethernet switching process (eswd) might crash when a non-ERP interface goes down and then back up. [PR1129309](#)

- On EX Series switches, the Ethernet Switching Process (eswd) is getting multiple MAC learn notification from the Software Infrastructure Daemon (sfid) for same the MAC addresses might cause the stale MACs entry to stop the ageing process, which results in the Ethernet switching table reaches the max limit. [PR1147854](#)
- On EX Series switches except EX4300, EX4600, and EX9200, the Ethernet switching process (eswd) might crash if you delete a VLAN's tag and then add the VLAN's name to the configuration under the [edit ethernet-switching-options unknown-unicast-forwarding] hierarchy during the same commit. [PR1152343](#)

Multiprotocol Label Switching (MPLS)

- On QFX/EX4600 Series switches, while receiving an IPv6 packet whose destination IPv6 address does not have an entry in the IPv6 neighbor table, they would fail to send out an IPv6 neighbor discovery packet and traffic to these IPv6 hosts might be dropped. [PR1134599](#)

Network Management and Monitoring

- On EX Series switches (except EX4300, EX4600, and EX9200), when syslog is enabled and an RPM probe is set to greater than 8000 bytes, the message ?PING_RTT_THRESHOLD_EXCEEDED? is not displayed, although it should be. [PR1072059](#)
- On EX Series switches, there are two issues regarding SNMP MIB walks: A private interface---for example, pime.32769---must have an ifIndex value of less than 500. If you do not add the private interface to a static list of rendezvous point (RP) addresses, the mib2d process assigns an ifIndex value from the public pool (with ifIndex values greater than 500) to the interface, which then will have an incorrect ifIndex allocation. A random "Request failed: OID not increasing" error might occur when you issue the "show snmp mib walk" command, because the kernel response for a 10-gigabit interface during an SNMP walk might take more than 1 second, and the mib2d process receives duplicate SNMP queries from the snmpd process. [PR1121625](#)

Platform and Infrastructure

- On an EX4300 switch with Bidirectional Forwarding Detection (BFD) configured, the BFD packets are getting forward to best effort queue (Queue 0), instead of network control queue (Queue 3). When queue 0 is under congestion, the BFD session might flap continuously. [PR1032137](#)
- On EX4300 switches with redundant trunk groups (RTGs) configured, after an RTG primary link comes online from the offline state, it becomes the active link and the other link becomes the backup link. After that occurrence, the Layer 2 address learning daemon (l2ald) sends a MAC refresh packet out of the new active RTG logical interface, which is not yet programmed in the Packet Forwarding Engine, causing the primary link to incorrectly update the MAC entry and also causing traffic loss. [PR1095133](#)
- When Firewall filter is configured with action syslog, firewall log does not show information for the interfaces on backup member. Example, filter with action syslog

applied to the interfaces xe-0/0/34 and xe-1/0/34 but the firewall logs show the information for xe-0/0/34 only. root@W01-TUS# run show log firewall Jun 22 22:00:06 W01-TUS clear-log[16806]: logfile cleared Jun 22 22:00:10 W01-TUS pfex: PFE_FW_SYSLOG_IP: FW: xe-0/0/34.0 A icmp 213.164.95.73 213.164.95.74 8 0 (1 packets) Jun 22 22:00:14 W01-TUS last message repeated 4 times [PR1098476](#)

- On EX4300 switches with Virtual Router Redundancy Protocol (VRRP) configured on an integrated routing and bridging (IRB) logical interface, when the IRB logical interface is disabled or deleted, the kernel does not send VRRP dest-mac-filter delete messages to the Packet Forwarding Engine, possibly causing loss of traffic that comes from another device's same VRRP group master VIP to the backup (or backup to master). [PR1103265](#)
- On EX4300 and QFX Series switches, the analytics daemon (analyticd) runs on the devices even if there is no analytics configuration, which might cause system instability due to the high number of files opened by analyticd. [PR111613](#)
- On EX4300 switches, if a policer ICMP filter is applied on the loopback interface, while an ICMP packet comes in, it might be dropped on the ingress Packet Forwarding Engine and the ARP request might not be generated. [PR1121067](#)
- On EX4300 switches, configuring "set groups group_name interfaces <*> unit 0 family ethernet-switching" and committing the configuration might cause the Layer 2 address learning process (l2ald) to generate a core file.. [PR1121406](#)
- On EX4300 switches with a Q-in-Q configuration, when Layer2 Protocol Tunneling (L2PT) for VLAN Spanning-Tree Protocol (VSTP) is enabled, the C-VLAN (inner VLAN or customer VLAN) might not be encapsulated in the PDUs that go out through the trunk port. [PR1121737](#)
- On EX4300, EX4600, EX9200, and QFX Series switches, the lldp-med-bypass feature does not work. [PR1124537](#)
- On EX4500 switches running TACACS, you might notice the following log message: mgd[65984]: %DAEMON-5-UI_TACPLUS_ERROR: TACACS+ failure: . [PR1126386](#)
- On a EX4300 Virtual Chassis, if a redundant trunk group (RTG) interface flaps, when control packets originating from the switch are going over that RTG interface, the core device become nonresponsive and you would have to reload the device to restore connectivity. [PR1130419](#)
- On EX4300 Virtual Chassis, traffic from or to a Routing Engine through an AE member interface that is not in the master might be dropped, but traffic transmitted through the switch (that is, hardware switched) is not affected. [PR1130975](#)
- On a EX4300-VC platform, if a Q-in-Q S-VLAN interface with MC-LAG is configured, when the backup EX4300 is acting as master, you might lose connection to the management IP address through the interface. As a result, management traffic will be dropped. [PR1131755](#)
- On an EX4300 switch, when an SNMP walk is performed to query the native VLAN, the query might return a value of 0 for most of the trunk interfaces instead of the configured native VLAN ID. [PR1132752](#)

- On EX4300 series switches in Ethernet Ring Protection Switching (ERPS) scenario, control plane might assign more than one STP instance to a VLAN on ERPS ring after system reboot, this will cause Ping packets forwarding issue. [PR1132770](#)
- On EX4200, EX4300, EX4550, EX4600, and QFX5100 switches with Media Access Control Security (MACsec) enabled on an AE subinterface, MACsec might not work because the MACsec Key Agreement (MKA) session is not established with a peer after flexible-vlan-tagging is configured on the AE interface. [PR1133528](#)
- On EX4300 switches, a filter might not work as expected when you commit a filter-based forwarding (FBF) configuration for the first time after rebooting the switch. [PR1135751](#)
- On EX4300/EX4600/QFX series switches, and ACX5000 line of routers when dhcp client configuration deleted on the interface, if there is no dhcp group configuration for dhcp clients, DHCP-relay might not work. [PR1136236](#)
- DHCP client receives IP assignment from DHCP server, the switch updated its snooping database. DHCP client, after T1 (half of the DHCP lease time) time, will renew it's IP address by sending DHCP REQUEST toward the DHCP server. During this time, the switch changes the binding state from BOUND to RENEWING. In this scenario, DHCP server for some reason is out of service, therefore no DHCP ACK sent back from DHCP server. After roughly 2 minutes later, the switch deletes this binding entry from DHCP snooping database. Subsequently, DHCP client with valid IP address loses its connectivity to the network with due to arp inspection. [PR1138118](#)
- On EX4300 Series switches with et interface (40G/100G interface), when interface is rebooted (for example: rebooting switch), et interface might be in incorrect status, so it might cause some impacts (for example: OSPF stuck in INIT state). [PR1143601](#)
- On EX Series switches, the following DEBUG messages might be incorrectly output with logging level INFO: %USER-6: [EX-BCM PIC] ex_bcm_pic_eth_an_config %USER-6: [EX-BCM PIC] ex_bcm_pic_check_an_config_change [PR1143904](#)
- On EX4300 switches, if an IPv6 firewall filter term exceeds the maximum, the Packet Forwarding Engine manager (pfex) might crash continuously. [PR1145432](#)
- On EX4300 switches with redundant trunk groups (RTGs) configured, VSTP BPDUs coming into an RTG backup interface might be incorrectly forwarded out of interfaces other than the RTG primary interface. [PR1151113](#)
- On EX/QFX/MX Series platforms, if "interface-mac-limit" is configured on an interface range, the commit might fail. [PR1154699](#)

Routing Protocols

- On QFX5100/EX4600 Series switches, when eACL (Egress routing ACL filter) is applied to more than 64 interfaces, a memory corruption issue might occur, resulting in the Packet Forwarding Engine manager (fxpc) process to crash. [PR1123374](#)
- On QFX5100 switches, you might see the "soc_mem_read: invalid index -1 for memory EGR_L3_INTF" log message. You can ignore the message; there is no functional impact on the switch. [PR1126035](#)

- Configuring analyzers might lead to sub-optimal use of allocated TCAM space. When this happens, the following logs might be displayed: [Sat Nov 21 08:45:18 2015 LOG: Err] PFE: Unknown next-hop (nh_id 2532) for sampling [Sat Nov 21 08:45:19 2015 LOG: Err] PFE: Unknown next-hop (nh_id 2532) for sampling [PR1136837](#)

Virtual Chassis

- On a two-member EX Series Virtual Chassis with the same mastership priority configured on both members, if there are more than 34 SFPs present in the current master and if a reboot is issued in the current master, then the backup becomes the master. When the original master rejoins the Virtual Chassis, it regains mastership. [PR111669](#)

Resolved Issues: Release 14.1X53-D30

General Routing

- On EX Series switches with integrated routing and bridging (IRB) interface configured, if the JSRV interface is created prior to the IRB interface after restarting the device or chassis daemon (chassisd), it might cause all IRB interfaces to be disappeared. [PR965097](#)
- On EX4500 and EX4550 switches, if you disable an interface on an EX-SFP-10GE-LR uplink module by issuing the CLI command "set interface interface-name disable", and then the interface through which a peer device is connected to the interface on the uplink module goes down, the CPU utilization of the chassis manager process (chassism) might spike, causing chassism to create a core file. [PR1032818](#)
- On an EX8216 switch, if the switch interface board (SIB) or the Switch Fabric (SF) module fails, there are no spare fabric planes available for switchover, which might cause a traffic outage. Depending on the nature of the SIB failure, the plane might need to be taken offline to resolve the issue. [PR1037646](#)
- If precision-timers and traceoptions are enabled for BGP then both main-thread and precision-timers pthread try to rotate the same tracefile without taking any locks. As a result all the status commands for rpd & krt may timed-out. [PR1044141](#)
- Inconsistent/Incorrect AE IFD stats because of incorrect handling the child IFD stat flags. As AE stats is an aggregate of the child IFD stats, these requests are processed differently as compared to stand alone interfaces thereby introducing inconsistencies in the next poll cycle. [PR1048276](#)
- On EX4200 Series switches and EX4550 Series switches, if you transit traffic through Ethernet over Sonet and also configure Media Access Control Security (MACsec) for switch to switch connections, packets might be dropped. [PR1056790](#)
- On EX4300 switches, if a Gigabit Ethernet interface is directly connected to an MX104 management interface (fxp0), the physical link will be down. [PR1069198](#)
- On EX Series switches with DHCPv6 snooping configured, when the VLAN ID is appended to the prefix of DHCPv6 Option 18, it will appear in decimal format instead of hexadecimal format. For second issue, if configuring "option-18 use-string" and "option-16 use-string" might truncate first 4 bytes. [PR1070488](#)

- On a mixed Virtual Chassis with EX4300 and EX4600 switches, if deactivating dot1x on the interface, the MAC address does not get learned on the interface. As a workaround, after deactivating dot1x, deactivate and activate the interface on which dot1x was configured. [PR1070885](#)
- On a QFX5100 Virtual Chassis, the log messages as "fpc0 vccpd irt socket connect failed (no route to host)" are seen continuously, it is harmless. [PR1075437](#)
- On EX4300/EX4600/QFX Series switches, if a Redundant Trunk Groups(RTG) interface is created and changed the primary/backup interface type(access to trunk or vice versa) during the same commit, the RTG interface might not work correctly. [PR1076601](#)
- On EX9200 and QFX5100 switches, if you configure DHCP relay with the DHCP server and the DHCP client in separate routing instances, unicast DHCP reply packets (for example, a DHCP ACK in response to a DHCP RENEW) might be dropped. [PR1079980](#)
- On EX Series switches except EX9200, EX4300, and EX4600, the configuration of options for the circuit-id CLI statement at the [edit forwarding-options dhcp-relay group group-name relay-option-82] hierarchy level does not work as expected. The DHCP option 82 Circuit ID should be formatted as "SWITCH-NAME: <physical interface name>: vlan-name", but instead, the string is "SWITCH-NAME:vlan-name". [PR1081246](#)
- On EX4600 and QFX Series switches, IGMP snooping might not be enabled after you reboot the switch. You might see the same issue might be seen after you run a nonstop software upgrade (NSSU) on the switch. [PR1082453](#)
- EX switch in subscriber management environment, it might respond a Disconnect-NAK message after receiving Disconnect-request message from RADIUS server. This would result that the user can not be disconnected. [PR1087008](#)
- On EX2200 Series switches, when connectivity-fault-management (CFM) is configured on LAG interface, MAC learning might not work after the ethernet switching table is cleared from CLI. [PR1087886](#)
- On EX4600 or QFX5100 switches with Junos OS release 14.1X53-D25, when configuring 40G interface on slot1 and slot2 convert to 10G interface, it might throw an error message mentioning config not supported for these restricted ports channelization. [PR1094071](#)
- On EX Series and QFX Series switches, RADIUS authentication might fail when the switch receives an access-accept message containing another vendor's vendor specific attribute (VSA). [PR1095197](#)
- On EX/QFX Series switches with Junos OS release 14.1X53-D10 onwards, when DHCP relay is configured on the IRB interface for BOOTP relay, if the client is connected to the physical interface that belongs to the same VLAN as the IRB interface, and sends BOOTP request packets to the server, BOOTP reply packets from the server might be dropped on the IRB interface. [PR1096560](#)
- On QFX/EX Series switches, after one AE member failure of MC-LAG, IGMP-reports might not be forwarded over Interchassis Link (ICL), however the Unicast works well. [PR1097526](#)
- From 14.1X53-D30, NAS-Port-Type is included in the Access-Request for Dot1x / Captive portal authentication, for Ex platforms. The CLI command 'set access profile

<profile-name> radius options' is enabled for all Ex platforms to enable custom configuration. [PR1097865](#)

- On EX4300, EX4600, and QFX Series switches, if a trunk port is deleted and then reconfigured as an access port in the same commit, the Layer 2 address learning daemon (l2ald) might generate a core file. [PR1105255](#)
- On a QFX Series Virtual Chassis Fabric (VCF) or Virtual Chassis with GRES enabled, the backup Routing Engine might continuously reboot after you configure "forward-and-send-to-re" or "forward-only" under the [edit interface interface-name unit unit-number family inet targeted-broadcast] hierarchy. [PR1106151](#)
- On EX Series switches, an interface with an EX-SFP-1GE-LH transceiver might not come up and the transceiver might be detected as an SFP-EX transceiver. [PR1109377](#)
- The following debug logs may be seen on an EX4300 if the optic does not support diagnostics. "ex_bcm_pic_sfpp_get_optics_info: Diagnostics is not supported in Optics 0 of pic 2" This is a debug message and does not impact traffic. [PR1117479](#)
- On EX8200 switches, a nonstop software upgrade (NSSU) might fail during the master Routing Engine upgrade step, and an NSSU process might abort with this message: "mgd: unable to execute /var/etc/reboot.ex: Authentication error". [PR1122628](#)

Authentication and Access Control

- On EX4500/EX8200 Series switches with the Link Layer Discovery Protocol (LLDP) enabled and Edge Virtual Bridging (EVB) configured, when a switch is connected to a Virtual Machine (VM) server using Virtual Ethernet Port Aggregator (VEPA) technology, the EVB TLV in LLDP packets might be sent to the wrong multicast MAC address(01:80:c2:00:00:0e) instead of correctly one(01:80:c2:00:00:00). [PR1022279](#)

Class of Service (CoS)

- On QFX5100 and EX4600 switches, if an interface that is enabled for flow control is connected to an EX Series switch (except EX9200), even low-rate traffic (host-bound traffic) received might cause a MAC pause frame to be sent from the interface to the peer device, and other transmitting traffic from the interfaces might be affected (for example, LACP flapping might occur). [PR1113937](#)

Infrastructure

- On EX4200, EX4500, EX6200, and EX8200 switches that are configured with distributed periodic packet management (PPM) mode, if you configure the Bidirectional Forwarding Detection (BFD) minimum-receive-interval value to the custom interval, BFD packets might be sent to a remote neighbor at a rate that exceeds the remote minimum-receive-interval value. As a workaround, configure PPM in centralized mode. [PR1055830](#)
- On an EX8200 Virtual Chassis, if you configure vlan-tagging on an interface without configuring a family for the interface, the Packet Forwarding Engine might program an improper MAC (the local chassis MAC) instead of the router MAC, which is used for routing. As a workaround, configure family inet on the interface. [PR1060148](#)

- On EX4600 switches, the VLAN translation is not working. If configured this feature, the incoming VLAN tag will be not translated correctly. [PR1070637](#)
- On EX4500 and EX4550 Virtual Chassis, UDP (i.e. NFS) fragmented packets might be dropped if these packets ingress over an aggregated bundle and traverse VCP links. [PR1074105](#)
- On EX4600 Series switches, when the device install Junos OS release 14.1X53-D25, it might cause VCP ports to down and the PFE manager (fxpc) process crash with a core file generated. The second scenario is when performing in-service software upgrade (ISSU) to Junos OS releases 14.1X53-D25, then performing a manual reboot of the device, it might cause device into the unusable state. [PR1074930](#)
- On EX Series switches, the bidirectional forwarding detection (BFD) liveness detection timer might not be updated correctly after modifying the minimum interval. [PR1084316](#)
- On EX8200 virtual chassis, local ECMP hashing changes when a remote (non-local) interface flaps if the number of local interfaces does not equal the number of remote interfaces, which might impact ECMP load balancing. [PR1084982](#)
- On an EX4300 Virtual Chassis, when Layer 3 traffic transits across different members, the source MAC address of the egressing packets might use the hardware address instead of the correct current address. This results in a service impact. [PR1087541](#)
- On EX4300 switches, if you configure a firewall filter on a loopback (lo0) interface to accept BGP flow and an OTHER term with the discard action, and the receiving host-inbound traffic with a designated TCP port 179 to the Routing Engine, existing BGP sessions might go down. [PR1090033](#)
- On Virtual Chassis with GRES enabled, if an IPv6 Neighbor discovery next hop is out of sync between RE's, and when this next hop is re-allocated on the master RE again, the kernel on the backup RE may crash, which may cause a temporary traffic loss. [PR1096005](#)
- On EX Series switches, the Packet Forwarding Engine Manager (pfem) process might crash and generate a core dump when the TCAM is fully utilized. In that particular case, utilization of TCAM was caused by flapping links and multiple dot1x logging out/in events. [PR1107305](#)
- On EX3200 and EX4200 switches with multiple member interfaces on an aggregated Ethernet (AE) interface and with a large-scale CoS configuration enabled on the AE interface, a Packet Forwarding Engine limitation might be exceeded, the Packet Forwarding Engine might return an invalid ID, and the Packet Forwarding Engine manager (pfem) process might generate core files. [PR1109022](#)
- On EX4500 or EX4550 Virtual Chassis, if an NFS/UDP fragmented packet enters the Virtual Chassis through a LAG and traverses a Virtual Chassis port (VCP) link, CPU utilization might become high, and the software forwarding infrastructure (sfid) process might generate a core file. [PR1109312](#)

Interfaces and Chassis

- If an EX4550-32F switch in a Virtual Chassis reboots and comes online, LACP interfaces on any member of the Virtual Chassis might go down and not come up. [PR1035280](#)

- On a two-member EX8200 Virtual Chassis, if the Link Aggregation Control Protocol (LACP) child interfaces span different Virtual Chassis members, the MUX state in the LAG member interfaces might remain in the "attached/detached" state after you have disabled and reenabled the AE interface. [PR1102866](#)

Layer 2 Features

- In a mixed QFX3500 and EX4300 Virtual Chassis that has configured persistent MAC and MAC limiting, traffic is not received on Aggregated Ethernet (AE) interfaces on EX4300 switches when the EX4300 switches are acting as the linecard role. [PR1033618](#)
- On ELS platforms, when the multicast VLAN Registration Protocol (MVRP) is configured on an aggregated Ethernet (AE) interface and the CLI command "no-attribute-length-in-pdu" is configured, MVRP can become unstable. [PR1053664](#)
- On all EX Series switches except EX4300, EX4600, and EX9200, configuring more than 1000 IPv4 addresses can prevent gratuitous ARP packets from being sent to peers. [PR1062460](#)
- Occasionally, on an EX Series switch Virtual Chassis (VC) with spanning tree protocol (STP) enabled when each member's mastership priority values are different, rebooting some or all VC members can cause a traffic failure, even after bootup is complete. [PR1066897](#)
- On EX4300 Series switches after mac-move limit is configured, the initial one or two packets are always lost. [PR1079043](#)
- On all EX Series switches except EX4300, EX4600, and EX9200 switches, when Multiple Spanning-Tree Protocol (MSTP) is configured, the Ethernet switching process (ESWD) might crash and generate multiple types of core files in the large-scale VLANs that are associated with Multiple Spanning-Tree Instances (MSTIs). [PR1083395](#)
- On EX4600/QFX5100 Series switches with L3VPN scenario, ping packets sent from CE to remote CE may not work for back to back PE connection. [PR1096698](#)
- On an EX3300 switch, in a broadcast storm situation in which DHCP snooping is enabled and there are repeated DHCP requests and acknowledgements arriving on the switch as a result of IP addresses not being accepted by clients, the eswd process might create a core file. [PR1105709](#)
- On EX4200 switches with DHCP snooping configured, when a host moves from one interface to another interface and then renews its DHCP lease, the DHCP snooping database might not get updated, and thus the host might not connect on the new interface. [PR1112811](#)
- On EX4300, EX4600, and QFX Series switches, traffic received on the backup redundant trunk group (RTG) link might get forwarded to other interfaces following an RTG link failover. [PR1119654](#)

Multiprotocol Label Switching (MPLS)

- On EX4600 switches, when configuring "revert time" for redundant Layer 2 circuits, if the Layer 2 circuits pseudowire(PW) switched over from hot-standby backup PW to primary PW, the PFE manager(fxpc) process might crash with a core file generate. [PR1050534](#)

Network Management and Monitoring

- On all Junos based platforms, when configuring an invalid SNMP source address, SNMP traps might not be sent even changed the SNMP source address to a valid interface address. [PR1099802](#)

Platform and Infrastructure

- On EX4300 Series switches, if you configure the policer on the loopback interface (lo0) to limit the host-inbound traffic, the policer might not work as expected. [PR1037554](#)
- On EX4300 switches, EX4600 switches, and QFX Series switches, the firewall filter can no longer be programmed in the hardware after a kernel handles system initialization with extraordinary high activity such as a commit or reboot. [PR1062604](#)
- Some qualified quad small form-factor pluggable plus (QSFP+) direct-attach copper cables (DACs) might not be recognized by EX4300 switches. The DAC links come up, but the DACs are not listed in output for the CLI command "show chassis hardware". [PR1063056](#)
- On EX4300 switches, interfaces might stop forwarding traffic when you reboot the switch or apply a CoS scheduler configuration. [PR1072200](#)
- A specific device configuration can result in a commit failure condition. When this occurs, a user is logged in without being prompted for a password while trying to login through console, ssh, ftp, telnet or su etc., This issue relies upon a device configuration precondition to occur. Typically, device configurations are the result of a trusted administrative change to the system's running configuration. Refer to JSA10802 for further details. [PR1075580](#)
- On ELS switches (like EX4300s), if you try to commit a configuration that includes the access-ports statement, the configuration commit operation fails, and you might see this message: "error: 'access-ports' is not a valid interface-range or alias name." As a workaround, remove the access-ports configuration from your switch before you upgrade the switch to Release 14.1X53-D20 or later or use an interface-range. [PR1080873](#)
- On EX4300 Series switches, filters with the "is-fragment" match condition does not work properly when used as a VLAN filter. As a workaround, apply the filters to the inet interface or IRB interface. [PR1081617](#)
- On EX4300 Series switches, FBF (Filer-Based-Forwarding) does not work if a source VLAN has a LAG interface configured. [PR1082093](#)
- On EX4300 Series switches, when the device configured "set system processes dhcp-service accept-any-source-port" and the "dhcp-relay" does not configure for

the subnets which the traffic is routed, if the non-DHCP traffic send to the device with destination port is 67 or 68, the traffic might be dropped. [PR1088087](#)

- On EX4300 standalone or EX4300-VC switches with VSTP enabled, if VSTP packets are received at high rate (2000 pps) during VSTP convergence, then VSTP might not get convergence, this convergence issue might cause high CPU utilization and result in l2cpd crash. [PR1090591](#)
- On EX4300 switches with redundant trunk groups (RTGs) configured, after an RTG primary link comes online from the offline state, it becomes the active link and the other link becomes the backup link. After that occurrence, the Layer 2 address learning daemon (l2ald) sends a MAC refresh packet out of the new active RTG logical interface, which is not yet programmed in the Packet Forwarding Engine, causing the primary link to incorrectly update the MAC entry and also causing traffic loss. [PR1095133](#)
- On EX4300 Series switches, when VRRP is configured on different routing-instances, and if the hosts connected to different instances are set VRRP virtual IP (VIP) as default gateway(DGW), packets forwarding might not work well between the instances. [PR1098160](#)
- On EX4300 Series switches, the received LACPDUs are reaching lacpd in kernel instead of pppman in the linecard, which result in continuous logs indicating LACP timeout on the member interfaces. [PR1100588](#)
- On EX4300/EX4600/QFX Series switches, when configuring preemptive-cutover timer for a redundant trunk group (RTG), when the primary goes down, is replaced by the secondary link, if the secondary link goes down within the preemptive cutover time (by default, it is 120 seconds), even at this moment the primary link is up, the primary link is still in the blocked state. [PR1101678](#)
- On EX4300 switches, VSTP BPDUS are not flooded in the VLAN when VSTP is not configured on the switches. [PR1104488](#)
- On EX4300 Series switches, when an XE interface is re-configured to a GE interface by using "replace-pattern" command and vice-versa, if the interface is configured to rewrite packets with Differentiated Services code point (DSCP) markings, the packets will still be forwarded, but packets will no longer be dscp marked after the change. [PR1106140](#)
- On EX4300 switches, port vector corruption on a physical port might be caused by the interface's flapping multiple times, causing a Packet Forwarding Engine manager (pfem) crash and a Routing Engine reboot. [PR1121493](#)
- On an EX4300 switch, when an SNMP walk is performed to query the native VLAN, the query might return a value of 0 for most of the trunk interfaces instead of the configured native VLAN ID. [PR1132752](#)

Routing Protocols

- On EX4600 and QFX Series switches, you might not be able to commit the configuration when the arp-type match condition is configured in a firewall filter. [PR1084579](#)

- On QFX5100/EX4600 Series switches, when eRACL (Egress routing ACL filter) is applied to more than 64 interfaces, a memory corruption issue might occur, resulting in the Packet Forwarding Engine manager (fxpc) process to crash. [PR1123374](#)

User Interface and Configuration

- On EX4300/EX4600/EX9200/QFX Series switches, when configuring an interface range, if the interface range includes large-scale physical interfaces, and is configured with the "family" option set to "ethernet-switching", committing the configuration might take a long time to complete. [PR1072147](#)

Virtual Chassis

- On a EX8200 Virtual Chassis, when configuring multiple LAG interfaces which LAG contains more than 12 members physical interface, if the LAG interface flapping. It might cause the Packet Forwarding Engine Manager (pfem) crashes and generates the core file. [PR1065682](#)
- On EX4600/QFX based VC or mixed VC scenario, when saving the recently committed configuration as the rescue configuration may fail with error messages. As a workaround, manually copy current configuration to a rescue configuration and then copy rescue configuration from master to backup routing engine. [PR1074772](#)
- A non-master member of a Virtual Chassis or a Virtual Chassis Fabric might go down and the status "linecard / inactive / NotPrsnt" might be displayed in output from the show virtual-chassis CLI command. This event might occur in either of these scenarios:
1. On an EX4600 Virtual Chassis or on a QFX Virtual Chassis or Virtual Chassis Fabric (VCF), during a nonstop system software upgrade (NSSU) or standard upgrade of the non-master member to Junos OS Release 13.2X51-D35 or a later release
2. On a QFX Virtual Chassis or VCF on which automatic software update is configured, and prospective member switches are running Junos OS Release 13.2X51-D35 or a later release and are joining the Virtual Chassis or the VCF [PR1096895](#)

Resolved Issues: Release 14.1X53-D27

General Routing

- On EX Series switches with integrated routing and bridging (IRB) interface configured, if the JSRV interface is created prior to the IRB interface after restarting the device or chassis daemon (chassisd), it might cause all IRB interfaces to be disappeared. [PR965097](#)
- On EX4300 switches, if a Gigabit Ethernet interface is directly connected to an MX104 management interface (fxp0), the physical link will be down. [PR1069198](#)

Platform and Infrastructure

- On EX4300 Series switches, if you configure the policer on the loopback interface (lo0) to limit the host-inbound traffic, the policer might not work as expected. [PR1037554](#)

Resolved Issues: Release 14.1X53-D26

- [General Routing](#)
- [Infrastructure](#)
- [Multiprotocol Label Switching \(MPLS\)](#)
- [Platform and Infrastructure](#)
- [User Interface and Configuration](#)

General Routing

- On 8200-VC , NSSU is not supported [PR894369](#)
- On a mixed Virtual Chassis with EX4300 and EX4600 switches, if deactivating dot1x on the interface, the MAC address does not get learned on the interface. As a workaround, after deactivating dot1x, deactivate and activate the interface on which dot1x was configured. [PR1070885](#)
- On a QFX5100 Virtual Chassis, the log messages as "fpc0 vccpd irt socket connect failed (no route to host)" are seen continuously, it is harmless. [PR1075437](#)
- On EX9200 and QFX5100 switches, if you configure DHCP relay with the DHCP server and the DHCP client in separate routing instances, unicast DHCP reply packets (for example, a DHCP ACK in response to a DHCP RENEW) might be dropped. [PR1079980](#)

Infrastructure

- On EX4600 switches, the VLAN translation is not working. If configured this feature, the incoming VLAN tag will be not translated correctly. [PR1070637](#)
- On EX4600 Series switches, when the device install Junos OS release 14.1X53-D25, it might cause VCP ports to down and the PFE manager (fxpc) process crash with a core file generated. The second scenario is when performing in-service software upgrade (ISSU) to Junos OS releases 14.1X53-D25, then performing a manual reboot of the device, it might cause device into the unusable state. [PR1074930](#)

Multiprotocol Label Switching (MPLS)

- On EX4600 switches, when configuring "revert time" for redundant Layer 2 circuits, if the Layer 2 circuits pseudowire(PW) switched over from hot-standby backup PW to primary PW, the PFE manager(fxpc) process might crash with a core file generate. [PR1050534](#)

Platform and Infrastructure

- A certain type of legitimate traffic which can be maliciously crafted or may be valid traffic for networking requirements in customer sites received from ingress interfaces through the EX-PFE leak out blocked Spanning Tree Protocol egress interfaces creating an artificial loop in network topology. This can lead to a high bandwidth usage denial of service condition. Continued packets can create a persistent denial of service condition. See <http://kb.juniper.net/JSA10719>. [PR1069179](#)
- On EX4300 switches, interfaces might stop forwarding traffic when you reboot the switch or apply a CoS scheduler configuration. [PR1072200](#)
- On ELS switches (like EX4300s), if you try to commit a configuration that includes the access-ports statement, the configuration commit operation fails, and you might see this message: "error: 'access-ports' is not a valid interface-range or alias name." As a workaround, remove the access-ports configuration from your switch before you upgrade the switch to Release 14.1X53-D20 or later or use an interface-range. [PR1080873](#)
- On EX4300 Series switches, FBF (Filer-Based-Forwarding) does not work if a source VLAN has a LAG interface configured. [PR1082093](#)

User Interface and Configuration

- On EX4300/EX4600/EX9200/QFX Series switches, when configuring an interface range, if the interface range includes large-scale physical interfaces, and is configured with the "family" option set to "ethernet-switching", committing the configuration might take a long time to complete. [PR1072147](#)

Resolved Issues: Release 14.1X53-D25

- [General Routing](#)
- [Authentication and Access Control](#)
- [Class of Service \(CoS\)](#)
- [Infrastructure](#)
- [Layer 2 Features](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)

General Routing

- On 8200-VC , NSSU is not supported [PR894369](#)
- On a EX Virtual Chassis, one of the VC member is reboot or after switch failover might cause Connectivity-fault Management Process (cfmd) keep send next hop add request to kernel, which result in traffic drop because of next hop space index exhaustion. [PR1016587](#)
- 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](#)
- With this JUNOS release onwards, when the Juniper Enterprise ID (2636) is included in the prefix of DHCPv6 option 37 (Remote ID) or DHCPv6 option 16 (Vendor Class ID), it will be encoded in binary. These options are appended to DHCPv6 packets when DHCPv6 snooping is configured on the switch. [PR1052956](#)
- On MX Series, EX Series and QFX Series platforms, during the jdhcpd initialization phase, jdhcpd attempts to open the dhcpv4 socket without checking whether dhcp-security is configured or not, which might cause the helper bootp for DHCPv4 and DHCP-Relay DHCPv6 not to work together. [PR1053807](#)
- On EX3200/EX4200 Series switches with flow-control enabled (which is enabled by default), pause frames are continuous and frequent received on the interfaces might cause continuous soft-reset of the Packet Forwarding Engine (PFE), which result in the Flexible PIC Concentrator (FPC) to detach from the Virtual Chassis (VC), to cause other traffic related issues. [PR1056787](#)
- When you use the SNMP GET request to poll jnxOperatingState for FPCs that are not present on a Virtual Chassis Fabric (VCF) or an EX Series Virtual Chassis, incorrect results are displayed. Non-existent FPCs might be reported to be UP and RUNNING. This issue does not affect SNMP walks. [PR1061960](#)
- On EX4300/EX9200/MX240/MX480/MX960 platforms, when clients are authenticated with dynamic VLAN assignment on an 802.1X-enabled interfacen, disabling 802.1X authentication on the interface might cause the Layer 2 Address Learning daemon (l2ald) to generate a core file. [PR1064491](#)
- On EX4300 switches with VLAN Spanning Tree Protocol (VSTP) running on aggregated Ethernet (AE) interfaces, the root port might receive VSTP BPDUs that are intended for other interfaces (port IDs). This issue can cause the root bridge to flap. The issue can also cause the root bridge to dispute the BPDUs and not converge. [PR1066137](#)

Authentication and Access Control

- On EX series switch, the ptopoConnLastVerifyTime MIB might return an incorrect value. [PR1049860](#)
- On all Junos platforms, the output for the ptopoConnRemotePort MIB might display an incorrect value for portIDMacAddr. [PR1061073](#)

Class of Service (CoS)

- An SNMP MIB walk for the OID under jnxCos might disappear and reappear for each iteration. The output might not return for all interfaces. [PR1001092](#)

Infrastructure

- On a EX4550 Virtual Chassis with enable VLAN prune, if the Link Aggregation Control Protocol (LACP) child interfaces cross different VC members, then switching routing engine mastership and rebooting new backup routing engine might cause LACP abnormal, which result in LAG interfaces down. [PR1021554](#)
- On EX Series switches, if multiple L3/non-L2 subinterfaces are enabled on a physical interface(IFD), and the family is deleted on a sub-interface or a subinterface itself is deleted, then an impact on traffic might be seen that the traffic gets software switched (sent to RE) instead of hardware switched by the Packet Forwarding Engine (PFE). [PR1032503](#)
- On EX4500 and EX8200 switches, if the switch is configured as a P router for MPLS, MPLS labels might be seen on the P router where the packets transit the RE on both input and output MPLS interfaces. This might result in high CPU usage and can impact service. [PR1038618](#)
- On EX4600/QFX Series switches, after disabling and re-enabling an AE interface will cause traffic failure. This will require a reboot to have the interface to recover. [PR1044580](#)
- On EX4200/EX3300/EX3200 Series switches, high levels of traffic bound for the Routing Engine (RE) might cause the watchdog timer to expire, causing the switch to reboot. This issue is mainly seen with Protocol Independent Multicast (PIM) configurations when the multicast route is not present in the Packet Forwarding Engine (PFE) for some amount of time, during which the multicast traffic for that route is routed to the CPU. However the issue can be seen with any traffic going to CPU at high rate. [PR1047142](#)
- When EX4300 Virtual Chassis(VC) configured with Class of Service(CoS) on all interface, broadcast packets may not pass through VC via Virtual Chassis Port(VCP) links. [PR1053978](#)
- On EX4300/EX4600/QFX Series switches, in rare conditions, Packet Forwarding Engine(PFE) crash might cause Inter-Process Communication(IPC) communication drop, and interface flap. [PR1056816](#)

Layer 2 Features

- On EX Series switches except EX4300/EX4600/EX9200, if removed or disconnected interface and cause it goes down, it might not generate the Simple Network Management Protocol(SNMP) mac-notification trap to server, but it might send the mac-removal trap along with mac-add trap when interface come up. [PR1070638](#)

Platform and Infrastructure

- On an EX4300 platform, when traffic destination to a layer 3 physical interface coming, the traffic will be layer 2 traffic with its destination MAC as the layer 3 physical interface; if the traffic receive from an ingress layer 2 interface, which with access mode configuration and without a layer 3 interface(e.g irb/vlan.x) associated to its VLAN, the traffic might be dropped. [PR1033796](#)
- On EX4300/EX4600/QFX Series switches with nonstop bridging (NSB) and VLAN Spanning Tree Protocol (VSTP) enabled, traffic is flowing through non VSTP configured VLAN might stop forwarding after perform graceful Routing Engine switchover (GRES) when port have multiple VLANs and part of VLANs associated with VSTP. [PR1041887](#)
- On EX4300 Series switches, when Layer 3 interface transit traffic with UDP destination port of 520, these traffic might get dropped. [PR1050473](#)
- On EX series which was running with 13.2X51, mcsnoopd core generated while EX was booting up. [PR1051706](#)
- On EX/QFX Series switches which is running on 14.1X53, could not enter configuration mode after issued "file list" command. [PR1054796](#)
- egress input analyzer would not forward traffic to the analyzer interface when the size of the Packets are more than 8182 bytes [PR1057640](#)
- On QFX/EX4600 Series switches, when Dynamic Host Configuration Protocol(DHCP) packets with double tag going through the trunk interface which configured Virtual Local Area Network(VLAN) members was bound Layer 3 (IRB) interface, it might be cause the DHCP packets dropped. [PR1059557](#)
- On EX4300 switches, EX4600 switches, and QFX Series switches, the firewall filter can no longer be programmed in the hardware after a kernel handles system initialization with extraordinary high activity such as a commit or reboot. [PR1062604](#)
- On a Virtual Chassis (VC} or a Virtual Chassis Fabric (VCF) with EX4300 switch members, when one member splits from the VC or VCF, the physical interfaces (IFDs) remain in an Up state which results in a traffic outage. [PR1065451](#)
- A certain type of legitimate traffic which can be maliciously crafted or may be valid traffic for networking requirements in customer sites received from ingress interfaces through the EX-PFE leak out blocked Spanning Tree Protocol egress interfaces creating an artificial loop in network topology. This can lead to a high bandwidth usage denial of service condition. Continued packets can create a persistent denial of service condition. See <http://kb.juniper.net/JSA10719>. [PR1069179](#)
- On EX4300 switches, the Dynamic Host Configuration Protocol(DHCP) option-82 flag in a DHCP message cannot be applied to a trunk interface because the trunk interface is always set to the default trust interface. [PR1071644](#)

Routing Protocols

- In a rare condition, the routing protocol daemon (rpd) might crash and create a core file if there is internal BGP (IBGP) route churn and BGP next hop fails to update.
[PR1060133](#)

Resolved Issues: Release 14.1X53-D16

- [General Routing](#)
- [Infrastructure](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)

General Routing

- On EX4300 Series switches with 802.1X authentication configured, when an 802.1X-enabled interface flaps, the dot1x daemon (dot1xd) might generate frequent core files due to a memory leak. [PR1049635](#)
- 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](#)
- When you use the SNMP GET request to poll jnxOperatingState for FPCs that are not present on a Virtual Chassis Fabric (VCF) or an EX Series Virtual Chassis, incorrect results are displayed. Non-existent FPCs might be reported to be UP and RUNNING. This issue does not affect SNMP walks. [PR1061960](#)
- On EX4300 switches with VLAN Spanning Tree Protocol (VSTP) running on aggregated Ethernet (AE) interfaces, the root port might receive VSTP BPDUs that are intended for other interfaces (port IDs). This issue can cause the root bridge to flap. The issue can also cause the root bridge to dispute the BPDUs and not converge. [PR1066137](#)

Infrastructure

- On EX4600/QFX Series switches, after disabling and re-enabling an AE interface will cause traffic failure. This will require a reboot to have the interface to recover.
[PR1044580](#)
- When EX4300 Virtual Chassis(VC) configured with Class of Service(CoS) on all interface, broadcast packets may not pass through VC via Virtual Chassis Port(VCP) links. [PR1053978](#)
- On EX4300/EX4600/QFX Series switches, in rare conditions, Packet Forwarding Engine(PFE) crash might cause Inter-Process Communication(IPC) communication drop, and interface flap. [PR1056816](#)

Platform and Infrastructure

- When apply-groups are used in the configuration, the expansion of interfaces <*> apply-groups will be done against all interfaces during the configuration validation

process, even if the apply-group is configured only under a specific interface stanza.

[PR967233](#)

- On QFX/EX4600 Series switches, when Dynamic Host Configuration Protocol(DHCP) packets with double tag going through the trunk interface which configured Virtual Local Area Network(VLAN) members was bound Layer 3 (IRB) interface, it might be cause the DHCP packets dropped. [PR1059557](#)

Routing Protocols

- In a rare condition, the routing protocol daemon (rpd) might crash and create a core file if there is internal BGP (IBGP) route churn and BGP next hop fails to update. [PR1060133](#)

Resolved Issues: Release 14.1X53-D10

General Routing

- If vcp interfaces are monitored using show or monitor cli options, log message "/kernel: vcp-255/0/0: invalid PFE queue counter pairs to copy" will be observed in syslogs. Occasionally these cosmetic messages can occur without user interaction when the system does internal periodic checks on the VCP interface counters. These messages could be seen on VC / VCF systems that use VCP interfaces and are built with QFX5100 / QFX3x00 / EX4300 members. [PR869043](#)
- On EX4500 switches with an uplink module installed, if the uplink module is removed and then installed in less than 10 seconds, the chassis manager (chassism) may create a core file. [PR941499](#)
- On EX4300 switches running a release earlier than Junos OS Release 13.2X51-D26, when 802.1X authentication is configured on the switch, the RADIUS Access-Request packets sent from the switch to the RADIUS server do not include the Tunnel-Private-Group-ID attribute. [PR1017594](#)
- Static IPv6 snooping can be configured but does not take effect. [PR1028544](#)

Infrastructure

- In case of the traffic flowing through one of the interfaces on the master switch of virtual chassis, and if the master switch is rebooted or halted, the FDB entry remains in incomplete / discard state about 30 seconds. At that time, the traffic which uses the FDB entry will not flow. [PR1007672](#)
- On EX Series switches with IP source guard (IPSG) enabled, traffic might be dropped due to an IPSG drop rule after you delete the IPSG configuration and if you enable or disable 802.1X authentication on an interface that belongs to an IPSG-enabled VLAN and you change the interface to another VLAN that does not have IPSG enabled. [PR1011279](#)

- On EX Series Virtual Chassis, if VLAN pruning is enabled on a VLAN, traffic on that VLAN might be dropped on the Virtual Chassis Port (VCP) if the link is changed from trunk to access mode and then back to trunk mode. [PR1012049](#)
- When changing configuration, PFEM might be crashed and generates core file. It would be happened while deleting route entry in PFE. [PR1029908](#)

Layer 2 Features

- On EX Series switches, an Ethernet switching daemon (eswd) memory leak might occur if the following conditions are met: 1. If a VLAN has the VLAN index 0, and the VLAN is deleted but the memory is not freed accordingly. 2. In a Multiple VLAN Registration Protocol (MVRP) scenario, when a VLAN map entry is deleted but the memory is not freed accordingly. [PR956754](#)

Layer 3 Features

- On EX Series switches, when GRE tunnel is configured, and the tunnel source address is on the VLAN tagging interface, ARP might not be resolved for the GRE tunnel with data traffic. [PR974434](#)

Platform and Infrastructure

- On EX4300 switches, when there is a limit on the number of MAC addresses that can be learned on an aggregated Ethernet interface, and the action configured on the interface is to shut down after reaching the MAC limit, the aggregated Ethernet interface might not shut down. [PR933168](#)
- On EX4300 switches, two-way Ethernet frame delay measurement(OAM CFM) does not work in centralized mode. [PR960168](#)
- On EX4300 switches, CPU starvation may be observed after executing CLI command "request system software add < >", causing scheduler slips and protocol flap. [PR1015214](#)
- When Rapid Spanning Tree Protocol (RSTP)/VLAN Spanning Tree Protocol (VSTP) converging between EX4300/EX4600/QFX switch and other vendors' systems (eg. Aruba wireless controllers), the proposal Bridge Protocol Data Unit (BPDU) which are 68 bytes in length sent by switch might be dropped by other vendors' systems and this will cause the switch to go into a BLK state due to WLCs's frame length check fail. Because the expected frames length should be 64 bytes. [PR1015220](#)
- On an EX4300 switch running Junos OS Release 13.2X51-D20 and later, the switch might drop traffic after the MAC move limiting feature is enabled. [PR1019668](#)
- On an EX4300 switch that has both multiple VLANs and integrated routing and bridging (IRB) interfaces configured, if the 802.1p tag of a VLAN is the equivalent of a VLAN hardware token and egress firewall filters (access control lists (ACLs)) are configured, both VLANs could match a firewall filter term because of a duplication rule wherein tagged and untagged traffic could be identified. [PR1020327](#)
- On an EX4300 switch, when Aggregated Ethernet (AE) interface is configured as members of 3500 VLANs, and the child members of the AE interface is in trunk mode,

then packet forwarding engine manager (pfex) may crash with a core file generated and the AE interface stay down. [PR1023861](#)

- On EX4300 switches ports might stay up even when they have no cable connections. [PR1027025](#)

Routing Protocols

- On a platform with an IGMP configuration in which two receivers are joined to the same (S,G) and IGMP immediate-leave is configured, when one of the receivers sends a leave message for the (S,G), the other receiver might not receive traffic for one through two minutes. [PR979936](#)

User Interface and Configuration

- On the EX Series switches, the J-Web service might respond slowly or become unresponsive. [PR1017811](#)

Virtual Chassis

- In a three-member EX8200 Virtual Chassis (VC) scenario with Link Aggregation Group interfaces configured, traffic may be dropped on LAG interfaces after rebooting one member of the VC. [PR1016698](#)

- See Also**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 22](#)
 - [Known Behavior on page 27](#)
 - [Known Issues on page 29](#)
 - [Documentation Updates on page 88](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)
 - [Product Compatibility on page 90](#)

Documentation Updates

This section lists changes and errata in Junos OS Release 14.1X53 for the EX Series switches documentation.

- [Bridging and Learning on page 88](#)
- [Interfaces and Chassis on page 89](#)
- [Security on page 89](#)

Bridging and Learning

- Two new MIBs related to MAC notification are provided at Junos OS Release 14.1X53-D10:
 - `jnxL2aldMacHistoryEntry`
 - `jnxL2aldMacNotificationMIBGlobalObjects`

These MIBs are not yet described in the documentation.

Interfaces and Chassis

- On an EX4300 switch, if you disable autonegotiation on an interface, auto-MDIX is disabled at the same time. This information does not currently appear in the documentation.

Security

- Media Access Control Security (MACsec) support (EX4600 switches) was added in Junos OS Release 14.1X53-D15, but that feature was not listed in the first versions of the Junos OS Release 14.1X53-D15 release notes. We have added the feature listing in revision 3 of the release notes. See *New and Changed Features*.

- See Also**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 22](#)
 - [Known Behavior on page 27](#)
 - [Known Issues on page 29](#)
 - [Resolved Issues on page 35](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)
 - [Product Compatibility on page 90](#)

Migration, Upgrade, and Downgrade Instructions

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 89](#)

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release, even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases

ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

For information on software installation and upgrade, see the [Installation and Upgrade Guide](#).

- See Also**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 22](#)
 - [Known Behavior on page 27](#)
 - [Known Issues on page 29](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 88](#)
 - [Product Compatibility on page 90](#)

Product Compatibility

- [Hardware Compatibility on page 90](#)

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

- See Also**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 22](#)
 - [Known Behavior on page 27](#)
 - [Known Issues on page 29](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 88](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 89](#)

Junos OS Release Notes for the QFX Series

These release notes accompany Junos OS Release 14.1X53-D48 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-D48: QFX3500, QFX3600, and QFX5100.



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-D46 but are supported in other Junos OS Release 14.1X53 releases are included in these release notes.

- [New and Changed Features on page 91](#)
- [Changes in Behavior and Syntax on page 122](#)
- [Known Behavior on page 127](#)
- [Known Issues on page 136](#)
- [Resolved Issues on page 144](#)
- [Documentation Updates on page 206](#)
- [Migration, Upgrade, and Downgrade Instructions on page 207](#)
- [Product Compatibility on page 213](#)

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-D46 on page 92](#)
- [New Features in Release 14.1X53-D40 on page 92](#)
- [New Features in Release 14.1X53-D35 on page 99](#)
- [New Features in Release 14.1X53-D30 on page 100](#)
- [New Features in Release 14.1X53-D27 on page 103](#)
- [New Features in Release 14.1X53-D26 on page 104](#)
- [New Features in Release 14.1X53-D25 on page 105](#)
- [New Features in Release 14.1X53-D15 on page 106](#)
- [New Features in Release 14.1X53-D10 on page 111](#)

New Features in Release 14.1X53-D46

Interfaces and Chassis

- **Link Aggregation Control Protocol (LACP) force-up enhancements (QFX5100 switches and QFX5100 Virtual Chassis)**—Starting in Junos OS Release 14.1X53-D46, if an aggregated Ethernet interface (AE) on a switch has multiple member links and one member link in that AE is in the force-up state with its peer's LACP down, and then if LACP comes up partially—that is, if LACP is established with a non-force-up member link—force-up is disabled on the member link on which **force-up** has been set, and that member link is ready for connection establishment through LACP. Force-up is eligible only if the server-side interface has LACP issues.

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, QFX5100 Virtual Chassis)**—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 206.

[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:

                                Licenses    Licenses    Licenses    Expiry
                                used      installed  needed
Feature name
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](#).]

- See Also**
- [Changes in Behavior and Syntax on page 122](#)
 - [Known Behavior on page 127](#)
 - [Known Issues on page 136](#)
 - [Resolved Issues on page 144](#)
 - [Documentation Updates on page 206](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 207](#)
 - [Product Compatibility on page 213](#)

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](#)
- [Ethernet Switching](#)
- [Interfaces and Chassis](#)
- [MPLS](#)
- [Network Management and Monitoring](#)
- [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.

Ethernet Switching

- **L2 Learning protocol**—On QFX5100 Switches, the new parameter **exclusive-mac mac** is added to enhance the MAC move feature. This feature is used to track MAC addresses when they appear on a different physical interface or within a different unit of the same physical interface. When you configure the **exclusive-mac mac** parameter at the **[edit protocols l2-learning global-mac-move]** hierarchy level, the specified MAC address is excluded from the MAC move limit algorithm. The MAC address will not be tracked.

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](#).]
- Starting with Junos OS Release 14.1X53-D47, on QFX5100 switches, the configuration statement **source-destination-only-loadbalancing** under the **[edit forwarding-options enhanced-hash-key]** hierarchy is not visible in the CLI. The statement is not supported on QFX5100.

MPLS

- On QFX5100 PE switches with Layer 2 circuit configured, enabling VLAN bridge encapsulation on a CE interface drops packets if flexible Ethernet services and VLAN CCC encapsulation are configured on the same logical interface. You can configure only one encapsulation type: either **set interfaces xe-0/0/18 encapsulation flexible-ethernet-services** or **set interfaces xe-0/0/18 encapsulation vlan-ccc**.

Network Management and Monitoring

- **Juniper MIBs loading errors fixed (QFX Series)**—Starting with Junos OS Release 14.1X53-D48, duplicated entries and errors while loading MIBs on the Manage Engine MIB browser are fixed for the following MIB files:
 - jnx-chas-defines.mib
 - jnx-ifotn.mib

[See [MIB Explorer](#).]

Open vSwitch Database (OVSDB)

- **Automatic configuration of trunk interfaces that handle untagged packets in OVSDB-managed VXLANs (QFX5100, QFX5100 Virtual Chassis)**—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.

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](#).]

- **Adaptive load balancing (ALB) feature (Virtual Chassis Fabric)**—Starting in Junos OS Release 14.1X53-D46, the adaptive load balancing (ALB) feature for Virtual Chassis Fabric (VCF) is being deprecated to avoid potential VCF instability. The **fabric-load-balance** configuration statement in the **[edit forwarding-options enhanced-hash-key]** hierarchy is no longer available to enable and configure ALB in a VCF. When upgrading a VCF to a Junos OS release where ALB is deprecated, if the configuration has ALB enabled, you should delete the **fabric-load-balance** configuration item before initiating the upgrade.

See [Understanding Traffic Flow Through a Virtual Chassis Fabric](#) and [fabric-load-balance](#).

- **New configuration option to disable automatic Virtual Chassis port conversion (QFX3500, QFX3600, and QFX5100 Virtual Chassis)**—Starting in Junos OS Release 14.1X53-D47, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in a QFX3500, QFX3600, or QFX5100 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:
 - LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
 - The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.

- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

Automatic VCP conversion is not needed when using default-configured VCPs on both sides of the link to interconnect two members. On both ends of the link, you can also manually configure network or uplink ports that are supported as VCPs, whether or not the automatic VCP conversion feature is enabled.

Deleting the **no-auto-conversion** statement from the configuration returns the Virtual Chassis to the default behavior, which reenables automatic VCP conversion.

- See Also**
- [New and Changed Features on page 91](#)
 - [Known Behavior on page 127](#)
 - [Known Issues on page 136](#)
 - [Resolved Issues on page 144](#)
 - [Documentation Updates on page 206](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 207](#)
 - [Product Compatibility on page 213](#)

Known Behavior

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

EVPN

- During unified ISSU, layer 2 unicast traffic drop might be seen for approximately 10 seconds and then it converges. [PR1355066](#)

High Availability

- During a nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a traffic loop or loss might occur if the Junos OS software version that you are upgrading and the Junos OS software version that you are upgrading to use different internal message formats. [PR1123764](#)
- On a QFX5100 Virtual Chassis, when you perform an NSSU, there might be more than five seconds of traffic loss for multicast traffic. [PR1125155](#)
- 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](#)

Interfaces and Chassis

- 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](#)
- 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.
- Configuring link aggregation group (LAG) hashing with the **[edit forwarding-options enhanced-hash-key] inet vlan-id** statement uses the VLAN ID in the hashing algorithm calculation. On some switching platforms, when this option is configured for a LAG that spans FPCs, such as in a Virtual Chassis or Virtual Chassis Fabric (VCF), packets are dropped due to an issue with using an incorrect VLAN ID in the hashing algorithm. As a result, the **vlan-id** hashing option is not supported in a Virtual Chassis or VCF containing any of the following switches as members: EX4300, EX4600, QFX3500, QFX3600, QFX5100 switches. Under these conditions, use any of the other supported **enhanced-hash-key** hashing configuration options instead. [PR1293920](#)
- If minimum links and FUP are configured, then an aggregated Ethernet interface will be down even when fewer than 2 interfaces are down, i.e., when only single interface is UP, AE will be DOWN. [PR1313873](#)
- If ICCP and ICL links are disabled and subsequently enabled in an MC-LAG, there could be a traffic loss of around 6 seconds. [PR1122509](#)

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](#)
- L2TP is not supported on QFX5100 switches. [PR1212269](#)
- 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 S-VLAN 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. As a workaround, always include **vlan-id-list** in the Q-in-Q configuration. [PR1216724](#)

MPLS

- On QFX5100 switches, MPLS ECMP with penultimate hop popping (PHP) does not work with single labels. [PR1212113](#)
- On the QFX5100 switches in a Layer 3 VPN setup, when traceroute is run on an ingress PE device to a laptop, only Phop, and laptop are displayed. Topology: PE1-----P-----PE2-----Laptop. The following output shows that the egress PE2 device

is missing: {master:0} user@host> traceroute routing-instance VPN1 200.200.200.25
 traceroute to 200.200.200.25 (200.200.200.25), 30 hops max, 40 byte packets 1 10.1.50.2
 (10.1.50.2) 10.973 ms 22.267 ms 21.850 ms MPLS Label=299920 CoS=0 TTL=1 S=0 MPLS
 Label=16 CoS=0 TTL=1 S=1 2 200.200.200.25 (200.200.200.25) 10.820 ms 10.686 ms
 10.697 ms. This is a hardware limitation. [PR1188551](#)

OVSDB

- 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](#)
- The TSN receives BUM packets from the originating TOR (which sends BUM traffic to TSN) and replicates them to other TORs or SW VTEPs. Upon GRES, BUM traffic loss for 2 minutes might happen while doing GRES on QFX5100 Virtual Chassis with a scale of 60K MAC and 2K remote VTEPs. [PR1268529](#)

Platform and Infrastructure

- On QFX5100 Series switches, the Link Aggregation Control Protocol (LACP) in fast mode can go down and then come back up. This causes a timeout and a service outage during a unified ISSU or an NSSU. In addition, after rebooting the master Routing Engine is rebooted, switches can experience intermittent traffic loss on non-LAG interfaces, and redundant trunk groups (RTG) convergence time can be too long. [PR1116923](#)
- On QFX5100 Virtual Chassis, generic routing encapsulation (GRE) counters might not increment with a firewall filter and PIM configured. [PR1124170](#)
- When you try to configure VLAN on an interface as well as through apply-groups, the configuration might fail to commit. This type of configuration is not supported. [PR1186657](#)

Routing Protocols

- On QFX Series switches, the output of the **show route multicast extensive** command does not display correct statistics because the Packet Forwarding Engine hardware does not support multicast stream-specific statistics. [PR607228](#)
- On a QFX5100 Series switch, if one firewall filter is configured with **source-port-range-optimize** or **destination-port-range-optimize** and multiple noncontiguous source-port or destination-port match conditions, it fails. [PR1163523](#)
- If the L3_DEFIP table in the Packet Forwarding Engine is full, then does not install any more active routes from the Routing Engine. If those active routes are deleted from the Packet Forwarding Engine, then it will program the rest of the routes from Routing Engine to Packet Forwarding Engine. [PR1231774](#)
- In QFX5100, ECMP traffic is unevenly distributed over an aggregated interface that has an even number of child interfaces. As a workaround, you can use the following CLI/VTY commands to improve the traffic load-balancing: **root@host# set forwarding-options enhanced-hash-key ? Possible completions: + apply-groups Groups from which to inherit**

configuration data + apply-groups-except Don't inherit configuration data from these groups ecmp-resilient-hash Set resilient hashing for ECMP > hash-mode Hashing mode > inet Configure inet4 fields > inet6 Configure inet6 fields > layer2 Configure layer2 fields VTY COMMANDS: user@host)# set hash-params ecmp-rx-fte Set hashing parameters for ECMP traffic rcvd on fte ecmp-rx-xe Set hashing parameters for ECMP traffic rcvd on xe hglag-rx-bfte Set hashing parameters for HGLAG traffic rcvd on bfte hglag-rx-fte Set hashing parameters for HGLAG traffic rcvd on fte hglag-rx-xe Set hashing parameters for HGLAG traffic rcvd on xe lag-rx-fte Set hashing parameters for LAG traffic rcvd on fte lag-rx-xe Set hashing parameters for LAG traffic rcvd on xe macroflow-hash Set macro flow-based hash parameters. For example, to improve IPv4 traffic with SRC address changing and DST address being constant)

set forwarding-options enhanced-hash-key hash-mode layer2-payload, set forwarding-options enhanced-hash-key inet no-incoming-port, set forwarding-options enhanced-hash-key inet no-incoming-device, set forwarding-options enhanced-hash-key inet no-ipv4-destination-address, set hash-params lag-rx-xe offset 0xa, set hash-params macroflow-hash seed 9999, set hash-params macroflow-hash algorithm 1. When no-ipv4-destination-address is enabled, as only source IP address is changed and not destination IP address. On enabling no-ipv4-destination-address, unnecessary hashing on invalid parameter is avoided. [PR1346350](#)

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 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 might be observed during RSNG upgrade. [PR1207804](#)

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

- 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.

VXLAN

- 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](#)

- See Also**
- [New and Changed Features on page 91](#)
 - [Changes in Behavior and Syntax on page 122](#)
 - [Known Issues on page 136](#)
 - [Resolved Issues on page 144](#)
 - [Documentation Updates on page 206](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 207](#)

- [Product Compatibility on page 213](#)

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 <https://www.juniper.net/prsearch>.

- [Class of Service \(CoS\) on page 137](#)
- [EVPN on page 137](#)
- [General Routing on page 138](#)
- [High Availability \(HA\) and Resiliency on page 140](#)
- [Interfaces and Chassis on page 140](#)
- [Layer 2 Ethernet Services on page 140](#)
- [Layer 2 Features on page 140](#)
- [MPLS on page 141](#)
- [Platform and Infrastructure on page 142](#)
- [Routing Protocols on page 143](#)
- [User Interface and Configuration on page 144](#)
- [Virtual Chassis on page 144](#)
- [VXLAN on page 144](#)

Class of Service (CoS)

- On QFX5100 switches, with the **CoS traffic-control-profiles** configuration (without the **guaranteed-rate**), the CoS configuration is not actually pushed to the Packet Forwarding Engine. CoS configuration is validated in 2 stages - **commit-check** and **commit-sync**. In this case, **commit-check** passes but **commit-sync** fails. Hence, there is syslog error and configuration is not pushed down. But, once the commit check is passed, CoS removes the default scheduler associated with the physical interface which should ideally happen after the commit synchronization. Without default scheduler attached to the physical interface, the control packets are not prioritized and hence you see link flap and mastership change. The link can be used as reference, it has to be tested before implementing in production.

https://www.juniper.net/techpubs/en_US/junos15/topics/example/cos-hierarchical-port-scheduling-ets-configuring.html
PR1183139

EVPN

- VXLAN ping and traceroute overlay do not follow the same path as the data packets over VXLAN tunnel when ECMP uplinks on first-hop TOR. [PR1106169](#)
- On QFX5100 switches, EVPN routes from compute nodes can be withdrawn when no change has taken place on either the compute node or the QFX5100 switch. [PR1106510](#)
- During a unified ISSU, the following syslogs might be seen: **Nov 5 11:45:11 st-pdt-opus06 l2ald[10183]: vgd:pipe is unblocked Nov 5 11:45:11 st-pdt-opus06 l2ald[10183]: vgd: pipe is now flow blocked Nov 5 11:45:12 st-pdt-opus06 l2ald[10183]: vgd:pipe is unblocked Nov 5 11:45:12 st-pdt-opus06 l2ald[10183]: vgd: pipe is now flow blocked Nov 5 11:45:12 st-pdt-opus06 l2ald[10183]: vgd: pipe is unblocked Nov 5 11:45:12 st-pdt-opus06 l2ald[10183]: vgd: pipe is now flow blocked Nov 5 11:45:13 st-pdt-opus06 l2ald[10183]: vgd:pipe is unblocked Nov 5 11:45:13 st-pdt-opus06 l2ald[10183]: vgd: pipe is now flow blocked.** These logs indicates that the pipe between l2ald and vgd is blocked when the pipe buffer is full. It is blocked until sufficient data has been read from the pipe to allow the write to complete. When the pipe is unblocked, it is notified and the queued data is flushed. This is normal when the communication traffic is heavy. [PR1136533](#)
- QFX5100 switches do not support ingress VLAN access control list (IVACL) flood filters. If you configure such as a filter by issuing the **set vlans <vlan-name> forwarding-options flood input <filter-name>** 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](#)
- When the **clear ethernet switching** command is used on an OVSDB-enabled switch causes delay in relearning MAC addresses, and depending on the scale of VNs, this delay can increase. The **clear ethernet-switching recovery-timeout** command is used

to recover the switch once the logical interfaces are down because of a MAC move limit shutdown. [PR1275025](#)

General Routing

- On a QFX5100 switch, running **tcpdump** on the console might cause system instability or cause protocols such as STP or LACP to fail. [PR932592](#)
- On Juniper switches, when an QFX5100 connect to any other Juniper switches through a 40G DAC connection, the link might not come up. This is because QFX5100 has auto-negotiation enabled on 40G DAC interface by default, any other Juniper switches have auto-negotiation disable by default. As a workaround, disable auto-negotiation on the QFX5100 will recover the connection. When 40G interface works as virtual chassis port (VCP) on both side in Virtual Chassis or Virtual Chassis Fabric (VCF) scenario, it does not have this issue, and auto-negotiation disable is not required. [PR935197](#)
- On the QFX5100 switch, a MAC address that is specified as part of a MAC-based VLAN is authenticated on an interface, for example, xe-1/1/1, on which 802.1X authentication in multiple supplicant mode is configured. However, the same MAC address might not be authenticated on another interface, for example, xe-2/1/1, if the MAC address moves to interface xe-2/1/1 from interface xe-1/1/1. [PR1007589](#)
- On a QFX5100 switch with VXLAN configured, adding or deleting an interface to/from the VLAN to which the VXLAN is associated, the switch might drop traffic for devices connected to other interfaces in the same VLAN. [PR1019378](#)
- QFX5100 1G sfp link will not come up with some devices. [PR1021260](#)
- In a QFX5100 Virtual Chassis Fabric (VCF) setup, a kernel synchronization process crashes and generates a core file after NSSU. [PR1023140](#)
- LFM adjacencies on the vcf drops when back-up and line card members are rebooted during NSSU, resulting in a state of "Active Send Local" until NSSU is completed. [PR1023831](#)
- On a mixed Virtual Chassis or Virtual Chassis Fabric (VCF) that contains at least one QFX3500 or QFX3600 member switch, MACsec configuration cannot be committed. [PR1024921](#)
- Traffic convergence delay time for link protection, node-link protection, and fast reroute is more than 50ms for the QFX5100-48T switch. [PR1026957](#)
- 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](#)
- When a transceiver on a QFX5100, QFX3500, and QFX3600 switch is removed and reinserted into an interface within 30 seconds after issuing the **set virtual-chassis vc-port set** command to convert the interface into a Virtual Chassis port (VCP), the VCP is not created. [PR1029829](#)
- In rare cases, a race condition might occur, in which a duplicate SNMP index might be assigned to the same interface. As a result, the mib2d daemon might crash. This issue should not cause any service impact. [PR1033249](#)

- On QFX5100 Series switches, when the device connects to EX4550 Series switch by 40G interface, when EX4550 switch is rebooted, the 40G interface on the QFX5100 switch might come up as channelized 10G ports. As a workaround, configure **set chassis fpc <fpc-slot> pic <pic-slot> port <port-num> channel-speed disable-auto-speed-detection** on QFX5100. [PR1049314](#)
- a) DHCP Relay in forwarding mode does not maintain any binding of the DHCP Client. In this case, if the DHCP server responds the INFORM-ACK directly to the relay agent address, the relay looks for a matching binding entry. In the absence of a binding entry, this ACK is dropped. b) DHCP server looks for an existing binding entry for the DHCP client. And if present, fetches the relay agent address from the entry and uses it to send the unicast INFORM-ACK. [PR1066679](#)
- On QFX5100 Series switches, with default factory settings, if adding an interface to OVSDB configuration and port bindings are pushed from NSX, the transaction fails and moves to failed queue. [PR1082218](#)
- On QFX Series Switches, nonstop software upgrade (NSSU) cannot be used to upgrade from a Junos OS Release 14.1X53 image to a Junos OS Release 15.1 or later image. [PR1087893](#)
- In a mixed mode Virtual Chassis with QFX3500 switches, if multicast packets are sent to the Routing Engine at a high rate, the Virtual Chassis might become unresponsive. [PR1117133](#)
- On QFX5100 Series switches, the auto-negotiation must disable on both ends of a 40 Gigabit Ethernet interface in order for the interface to remain up. For example, on each switch, issue the **set interface et-x/y/z ether-options no-auto-negotiation** command. [PR1118318](#)
- In a large scale VXLAN and OVSDB setup (for example, 100K MAC/1K VNI), Routing Engine switchover causes secure sockets layer (SSL) connection to controller break around 4 minutes. And no new MAC entries are learnt during this time. Existing and programmed MAC entries will remain and the switch will continue to forward traffic for those MACs. [PR1136123](#)
- On QFX Series switches, if VSTP is not configured on a switch, VSTP or PVST+ BPDUs might flood on the VLAN because the scenario wishes to be supported, where two adjacent switches are configured with VSTP and the intermediate switch is not and it can act as a transparent bridge for VSTP or PVST+. [PR1199367](#)
- Output from **show chassis environment** says fan tray testing or absent in QFX3500 Virtual Chassis with EX4300. [PR1200638](#)
- In a large scale QFX Virtual Chassis Fabric (VCF), the timeout errors might be seen when running the command **request system reboot all-members at now** to immediately reboot all members of the VCF. [PR1215130](#)
- The smid process might crash on QFX3500, this has been fixed through internal PRs and contains a fix from Junos OS Release 15.1R5 and later. [PR1245772](#)
- Currently other than QFX5100-24q PIC 1 is not supported on any other platforms inline with QFX5100. The command below cannot be used on PIC 1. **set chassis fpc 0 pic 1 port <x> channel-speed disable-auto-speed-detection**. This will result in a commit error

[edit chassis fpc 0 pic 1 port 2 channel-speed] channel-speed

disable-auto-speed-detection PIC:1 not valid for Auto-speed-disable mode. **error:**

configuration check-out failed. So if you want to disable auto-channelisation on PIC1, you have to disable auto-speed-detection for whole FPC using **set chassis fpc 0**

auto-speed-detection disable command. [PR1362647](#)

- BGP session bounce might sometime cause, NOT to flood BUM traffic to all remote VTEPs. [PR1373093](#)

High Availability (HA) and Resiliency

- 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](#)

Interfaces and Chassis

- BPDU guard gets cleared for sometime during Network node-group switchover. [PR856614](#)
- The priority of a VRRP group can be tied to the operational state of an interface, using the "track interface" keywords. If the interface goes down, the priority of the VRRP group is reduced by a specified amount, possibly triggering a VRRP mastership change. The priority of a VRRP group, tracking an aggregated Ethernet interface on an RSNG/SNG, is not being reduced (as expected) when the aggregated Ethernet interface is operationally down. [PR882628](#)
- On QFX5100 switches, if you configure **MC-LAG**, **RB 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](#)

Layer 2 Ethernet Services

- From dcd, when two configurations as mentioned below are committed together force-up value is not set from dcd. The force-up value is pushed to kernel and further read from lacpd to move the interfaces state directly to CD without lacpd's from peer. This is a configuration that is missed from dcd and whenever there is a change from earlier versions to Junos OS Releases 14.1X53-D40 and later. This causes mismatch in kernel and dcd configuration database **set groups force-up interfaces ge-0/0/1 ether-options 802.3ad lacp force-up** and **set groups force-up interfaces ge-0/0/1 ether-options 802.3ad ae0**. [PR1331361](#)

Layer 2 Features

- 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](#)
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric switches, when an xSTP bridge protocol data unit is distributed to the FPC (one member switch), there might be traffic loss if 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 Virtual Chassis Fabric (VCF), if you perform a nonstop software upgrade 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 switches, the Layer 3 routes that form 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 a mixed-mode Virtual Chassis Fabric (VCF) with **interface-mac-limit** configured, if you remove the complete **mac-limit** configuration, the mac-limit behavior might remain. As a workaround, reboot the device. [PR1044460](#)
- In a QFX5100 Virtual Chassis or Virtual Chassis Fabric, an NSSU to Junos OS Release 14.1X53-D35 might cause a traffic loss for a few seconds for BUM traffic. [PR1128208](#)
- If the QFX5100 has multiple ae interfaces with child members (1 Gigabit or 10 Gigabit respectively). If some ae interfaces are configured for MSTP and some are not, for example, ae1 and ae2. ae1 is part of MSTP but ae2 is not. Then, ae2 child members do not forward transit or CPU originated traffic. [PR1163227](#)
- Packets might be dropped when using egress UNI VLAN-ID without **vlan-id-list** configuration. Packets are dropped when egress UNI VLAN-ID is not matched with customer inner tag-id. [PR1216732](#)
- On random initialization of QFX5100 the programming of storm control profile is missed within hardware on random interfaces. This is not visible over CLI and the configuration still shows intact. This happens as a result of interface speed not properly getting detected within the hardware. [PR1354889](#)

MPLS

- In the event of link failure when multiple LSPs are using a link-protected and fast-rerouted link, the convergence time is proportional to the number of LSPs sharing the protected link. [PR1015806](#)
- In the event of link failure when multiple LSPs are using a link-protected and fast-rerouted link, the convergence time is proportional to the number of LSPs sharing the protected link. [PR1016146](#)
- When a link fails on a transit router that hosts a Layer 2 circuit over an RSVP tunnel, the traffic convergence time is approximately 350 ms 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 might drop. [PR1027755](#)

- On QFX5100 using the Ethernet tagged mode of operation on a pseudowire, Layer 2 control protocols might 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](#)
- On QFX5100 switches using the IS-IS routing protocol as an interior gateway protocol between customer edge (CE) switches for an Layer 2 circuit, the CEs might fail to form an IS-IS adjacency over a pseudowire. As a workaround, use an alternative IGP protocol such as OSPF. Both IS-IS and OSPF link state protocols use the same algorithm for computing the best path through the network. [PR1032007](#)
- On a QFX5100 switch, the enhanced hash key does not work for MPLS-IP packets. [PR1095136](#)
- 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](#)
- On QFX Series switches, if you change a Layer 2 circuit configuration from Ethernet CCC encapsulation to VLAN CCC encapsulation, traffic loss might occur at the pseudowire tunnel initiation point. [PR1222888](#)

Platform and Infrastructure

- 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 QFX5100 switches with a large number of firewall terms configured, firewall filters might stop working after you perform a unified ISSU. [PR966445](#)
- In a mixed-mode Virtual Chassis Fabric (VCF), control plane packets such as OSPF or PIM might not be mirrored by the native analyzer when the output port belongs to another member in the Virtual Chassis. [PR969542](#)
- On a Virtual Chassis Fabric, if you issue the **show interfaces gr-0/0/0 extensive** command, GRE statistics for logical interface gr-0/0/0.0 are not updated properly and it takes a long time for the CLI to respond. [PR979629](#)
- 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 QFX5100 switches acting as a VXLAN virtual tunnel endpoint (VTEP), known unicast traffic might be dropped from the VXLAN after GRES (for example, NSSU, ISSU). [PR1026408](#)
- On QFX5100 Virtual Chassis, generic routing encapsulation (GRE) counters might not increment with a firewall filter and PIM configured. [PR1124170](#)

- On QFX3500 and QFX3600 switches with ECMP enabled, if you add or delete routes continuously, the Packet Forwarding Engine might stop forwarding traffic, causing a traffic blackhole. [PR1137890](#)
- On a QFX5100 Virtual Chassis, when you perform a non-stop software upgrade 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](#)

Routing Protocols

- On a mixed-mode Virtual Chassis Fabric (VCF), when you add a new member to an existing VCF, routing protocols might transit down and up. [PR957292](#)
- When a static multicast route with a next-table next hop is changed from a table that cannot forward the traffic to one that can and then revert 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. For example, 1) starting state: output rate is 100k pps show configuration groups vrf1 routing-instances r1 routing-options static route 233.252.0.1/32 next-table r4.inet.0; route 233.252.0.2/32 next-table r4.inet.0; 2) change the route such that one of the routes now has a next-table of inet.0 and outbound traffic rate increases to 101k pps show configuration groups vrf1 routing-instances r1 routing-options static route 233.252.0.1/32 next-table inet.0; route 233.252.0.2/32 next-table r4.inet.0; 3) revert the change to return to the original configuration (traffic rate stays at 101k pps) show configuration groups vrf1 routing-instances r1 routing-options static route 233.252.0.1/32 next-table r4.inet.0; route 233.252.0.2/32 next-table r4.inet.0; [PR1217958](#)
- If ddos-protection is disabled on QFX5100 VC and multicast traffic is being sent, Virtual Chassis might become unstable, with high CPU usage and it might crash eventually, creating FXPC core files. Disabling ddos protection disables rate limiting for all host bound traffic. This is not a recommended setting on the device since high amount of control traffic can overwhelm the system, causing system instability. [PR1238875](#)
- If a QFX5100 switch has a host route with ECMP next hops and receives a better path with single next hop then next hop in hardware will not be changed. 1) Route 10.244.1.1/32, ECMP next-hops. Note: These commands must be entered without typos/mistakes. Verify each command before typing. **user@host# show route ip hw host Rtt VRF Destination Type NH-ID Interface HW NH-ID Src-MAC-Address Port Dst-MAC-Address -----**
----- 0 110.244.1.1 Unicast 131071 200011
(ECMP) 100012 64:64:9b:52:d9:07 5 64:64:9b:54:0a:07 100013 64:64:9b:52:d9:08 6
64:64:9b:54:0a:08 100014 64:64:9b:52:d9:09 7 64:64:9b:54:0a:09. 2) A better path
has been learned for 10.244.1.1/32 (uses single next-hop). HW NH-ID stays unchanged
user@host# show route ip hw host Rtt VRF Destination Type NH-ID Interface HW NH-ID
Src-MAC-Address Port Dst-MAC-Address -----
----- 0 110.244.1.1 Unicast 1730 irb.89
ifl 549 200011 (ECMP) 100011 64:64:9b:52:d9:07 5 64:64:9b:54:0a:07 100013
64:64:9b:52:d9:08 6 64:64:9b:54:0a:08 100014 64:64:9b:52:d9:09 7 64:64:9b:54:0a:09.
[PR1387713](#)

User Interface and Configuration

- If a configuration file contains groups related configuration is loaded by command **load replace**, a **commit confirmed** operation might fail. When this issue occurs, the new configuration is committed even if you do not confirm it within the specified time limit. [PR925512](#)

Virtual Chassis

- On a mixed Virtual Chassis Fabric (VCF), a Virtual Chassis port (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](#)

VXLAN

- 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](#)

- See Also**
- [New and Changed Features on page 91](#)
 - [Changes in Behavior and Syntax on page 122](#)
 - [Known Behavior on page 127](#)
 - [Resolved Issues on page 144](#)
 - [Documentation Updates on page 206](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 207](#)
 - [Product Compatibility on page 213](#)

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-D48 on page 145](#)
- [Resolved Issues: Release 14.1X53-D47 on page 147](#)
- [Resolved Issues: Release 14.1X53-D46 on page 152](#)
- [Resolved Issues: Release 14.1X53-D45 on page 154](#)
- [Resolved Issues: Release 14.1X53-D44 on page 156](#)
- [Resolved Issues: Release 14.1X53-D43 on page 158](#)
- [Resolved Issues: Release 14.1X53-D42 on page 162](#)

- [Resolved Issues: Release 14.1X53-D40 on page 163](#)
- [Resolved Issues: Release 14.1X53-D35 on page 176](#)
- [Resolved Issues: Release 14.1X53-D30 on page 182](#)
- [Resolved Issues: Release 14.1X53-D27 on page 194](#)
- [Resolved Issues: Release 14.1X53-D26 on page 196](#)
- [Resolved Issues: Release 14.1X53-D25 on page 198](#)
- [Resolved Issues: Release 14.1X53-D16 on page 202](#)
- [Resolved Issues: Resolved Before Release 14.1X53-D16 on page 205](#)

Resolved Issues: Release 14.1X53-D48

Class of Service (CoS)

- Firewall filter cannot filter packets with DST IP as 224/4 and DST MAC = QFX_intf_mac on loopback interface using a single match condition for source address 224.0.0.0/4. [PR1354377](#)

General Routing

- IPv6 firewall syslog action shows source, destination address wrongly correct address: 2001:DB8:4:0:0:0:0:2 2001:DB8:4:0:0:0:0:1 PFE_FW_SYSLOG_IP6_TCP_UDP: FW: .local..0 A tcp SA 120:b80d:400:0:0:0:0:200 DA 120:b80d:400:0:0:0:0:100 sport: 0 dport: 0 (258 packets) ^M PFE_FW_SYSLOG_IP6_TCP_UDP: FW: .local..0 A tcp SA 120:b80d:400:0:0:0:0:200 DA 120:b80d:400:0:0:0:0:100 sport: 0 dport: 0 (252 packets) ^M. [PR1104378](#)
- The initial implementation of auto-channelization relied upon the success or failure of certain timing related state machines. In some instances such as when an upstream device is rebooting, or in the process of initializing interfaces this can result in incorrectly (auto) channelizing a native 40G link. Once channelized the port must be manually reconfigured to restore native 40G connectivity which can impact some ZTP boot scenarios. This change modifies the decision tree to include reading of the applicable EEPROM register of the inserted qSFP to determine if the cable is capable of breakout before performing auto-channelization. [PR1317872](#)
- On QFX5100 switches, well known ports are used as source port in the VxLAN scenario. Per RFC, it is recommended the dynamic or private port which can range 49152-65535. [PR1335227](#)
- Fan RPM spikes every time the temperature sensor reaches its threshold level and revert to normal level when the temperature decreases. There is no functional impact to fan control software because of this fluctuation. [PR1345181](#)
- On QFX5100 switches, the Packet Forwarding Engine might drop the ARP reply packets after changing the interface MAC address. [PR1353241](#)
- On QFX5100-VC, VME interface might be unreachable after link flap of em0 on master FPC. [PR1362437](#)
- On QFX3500 and QFX3600 platforms, OSPF might remain in init status after loading the Junos OS Release 14.1X53-D47.4 image. [PR1362996](#)

- On QFX5100 switches in VC/VCF scenario, the chassisd might crash after issuing the CLI **show chassis hardware**. This can result in VCP down and traffic drop. [PR1366746](#)
- On QFX Series switches except QFX10000, if IS-IS packet is received with DMAC as 09:00:2b:00:00:05 (ISO 9542, All Intermediate System Network Entities Address) and Jumbo frame with EtherType as 0x8870 (non-standard, used by Cisco), such packet will be dropped, resulting in failure in the adjacency. [PR1368913](#)
- On QFX5100-VC running Junos OS Release 14.1X53-D43 through Junos OS Release 14.1X53-D47, command **show interfaces ae<interface-name> extensive** might display duplicate entries for member interfaces. [PR1369713](#)
- On QFX5100, IPv6 routed packet is transmitted over VRRP virtual IP address though its VRRP state is in transition to master. [PR1372163](#)
- On QFX5100 Virtual Chassis platform with GRES configured, if the backup member has file of **/var/run/consoleredirect.pid**, then reboot the master member of the Routing Engine switchover, the backup cannot become the master member. [PR1372521](#)
- On QFX Series platform, if RTG redundant trunking group (RTG) is enabled with a large-scale MAC address, MAC refresh frame might not be sent out from the new primary link after RTG failover by deactivating the former primary link on peer side. [PR1372999](#)
- A QFX5xxx Packet Forwarding Engine might show DISCARD next hop for overlay-bgp-lo0-ip when the QFX5xxxx is the leave in a leave-spine topology. [PR1380795](#)
- In Open vSwitch Database (OVSDb) environment, Virtual Chassis master copies **/var/db/ovsdatabase** to backup every 10 seconds and Virtual Chassis backup writes the whole OVS database to SSD frequently. This causes a high write I/O and shortens the SSD lifetime. [PR1381888](#)

Infrastructure

- On QFX5100 platform, a complete packet loss is experienced if **mac-move-limit** is enabled on an interface which has encapsulation **flexible-vlan-tagging** configured and has a port which has Layer 2 and Layer 3 VLAN. [PR1357742](#)

Interfaces and Chassis

- On QFX3500, QFX3600, and QFX5000 Series switches, MC-LAG peer might not send ARP request to the host. [PR1360216](#)

Layer 2 Features

- After rebooting one unit in Virtual chassis, the unit cannot establish the LAG because lacp packet drops. [PR1361054](#)
- On QFX5100 switches, IPv6 traffic over VxLAN tunnel does not hash, this might result in some unexpected issue in ECMP scenario. [PR1368258](#)
- On QFX5100 switches, if changing an interface from VXLAN to a member of an aggregated Ethernet interface, the DHCP relay might not work and the DHCP client might not get IP address normally. [PR1377521](#)

MPLS

- On all QFX5000 platforms, if the P/PE router is configured with no-decrement-ttl, the routing protocol process (rpd) sends the **NO_PROPAGATE_TTL** flag even for the tunnel transit case. [PR1366804](#)

Routing Protocols

- On QFX5100 platforms, the switch might get into an improper state where it is unable to correct parity errors in the Packet Forwarding Engine memory. Traffic might get silently dropped and get discarded for specific destination IPs. [PR1364657](#)

Resolved Issues: Release 14.1X53-D47**EVPN**

- With VXLAN configured for 30 VXLAN VNIs, L3 Unicast traffic loss may be observed on deleting and adding back all the VXLAN VNI's. [PR1318045](#)
- Given three leaf VTEPs: two remote VTEPs and one local VTEP, the programming for a MAC address might become mis-programmed on the local VTEP. This might happen when a MAC address in the EVPN database moves from remote VTEP (VTEP #1) to a local VTEP (VTEP #2) and then to a different remote VTEP (VTEP #3), the programming for the MAC address on the device with VTEP #2 is still point to remote VTEP #1. It will not be updated with the correct VTEP where the MAC address has moved (VTEP #3). [PR1335431](#)

General Routing

- Memory leak in JDHCP during dhcp session RELEASE/BIND [PR1181723](#)
- On EX Series or QFX Series Virtual Chassis, if new members are not zeroized prior to being added to the Virtual Chassis, and then one of the new members splits from the Virtual Chassis, then whenever you run "commit" or "commit check", the commit might hang for a long time and then report a timeout error on the FPC that split from the Virtual Chassis. [PR1211753](#)
- During the last stage of NSSU, before rebooting the master, NSSU state is set to idle and reboot is issued around 10 seconds after. The traffic drop is observed for these 10 seconds. [PR1219693](#)
- A QFX5100-48S or QFX5100-96S might incorrectly show the media type of an SFP-T copper module as fiber in the output of the 'show interface' command. [PR1240681](#)
- On EX/QFX Series switches, if Dynamic Host Configuration Protocol (DHCP) server uses boot file name option, when doing ZTP (Zero Touch Provisioning), the device cannot receive the image with error info of "Image File Not Set", causing image and configuration upgrade failure. [PR1247648](#)
- Junos OS: Short MacSec keys may allow man-in-the-middle attacks (CVE-2018-0021); Refer to <https://kb.juniper.net/JS10854> for more information. [PR1251909](#)
- On QFX5100 Series Switches, the following errors might get displayed with multicast configuration/traffic. The messages do not indicate traffic impact, however multicast

statistics might not work due to these messages. Feb 15 07:28:49 switch fpc0
brmc_ipmc_get_multicast_stats:3947 brmc_ipmc_stat_get failure Feb 15 07:28:49
switch fpc0 brmc_rt_stats:1906 brmc_ipmc_get_multicast_stats failure err=-7.
[PR1255497](#)

- On EX4600/QFX5100 Series switches, when an Integrated Routing and Bridging interface (IRB) is configured with the underlying layer 2 interfaces, if an Address Resolution Protocol (ARP) reply is received whose destination Media Access Control (MAC) is the same with IRB's MAC, the packet is consumed and also flooded in the Virtual Local Area Network (VLAN) as the ARP reply's MAC address received on the underlying layer 2 interface is not the interface's MAC. [PR1294530](#)
- Network Analytics process may be incorrect instantiated leading to traffic statistics not being transmitted. When this occurs the 'Sent' value for 'show analytics collector' will display as zero and 'show analytics traffic-statistics' will be empty: root@QFX5100> show analytics collector Address Port Transport Stream format State Sent 10.10.10.72 50020 udp json n/a 0 10.10.10.167 50020 udp json n/a 0 root@QFX5100> show analytics traffic-statistics CLI issued at 2018-03-26 22:15:56.411671 [PR1297535](#)
- On Enhanced Layer 2 Software (ELS) platform, if an interface is configured under a VLAN "A" but the same VLAN "A" is not configured in the chassis, there won't be any commit error being generated after performing committing configuration, which might lead to software upgrade failure. [PR1302904](#)
- On QFX5100 platform, for a subinterface of AE interface, the run-time pps statistics value is zero. This is a cosmetic issue. It does not have any service/traffic impact. [PR1309485](#)
- On QFX5100, QFX3500, and QFX3600 platforms, traffic loss might be seen if sending traffic via the 40G interface which is connected with peers through DWDM and the CRC errors of the interface may also keep on increasing after flapping the interface on QFX side. [PR1309613](#)
- On QFabric, a core file creation of cosd can be observed on RSNG/NNG some times if the configuration includes FCset configurations applied with non-wild card on all interfaces that includes AE interfaces. [PR1311158](#)
- Traffic drop occurs on sending L3 traffic across MPLS LSP. [PR1311977](#)
- On QFX5100 platform, transit traffic over GRE tunnels might hit CPU and trigger a DDoS violation on L3NHOP in below cases 1. When Unicast routes are formed to reach the Tunnel destination. As a workaround, With ECMP configuration removed, delete and reprogramming of GRE interface will resolve the issue 2. If deleting specific route for GRE tunnel destination IP. As a workaround, restart PFE process. [PR1315773](#)
- On QFX3500, QFX3600, or QFX5100 with Simple Network Management Protocol (SNMP) protocol enabled, if an interface connected to VoIP product, has Link Layer Discovery Protocol (LLDP) and LLDP-MED enabled, l2cpd might drop core files repeatedly. [PR1317114](#)
- On QFX5100 switches, if openflow is configured with interfaces and controller options, then the openflow session might flap constantly. This issue is caused by a malformed Openflow response packet. [PR1323273](#)

- On Enhanced Layer 2 Software (ELS) platform, VLAN or VLAN bridge might not be added or deleted if there is an interface bridge domain (IFBD) hardware token limit exhaustion. It might cause new IFBDs not be created or old IFBDs not be deleted. [PR1325217](#)
- On QFX5100 Series switches with Ethernet Virtual Private Network with Virtual Extensible Local Area Network (EVPN/VxLAN) multi-homing configuration, if the aggregation interface (AE) is configured with Service Provider style, then deleting one VxLAN might cause traffic loop for multi-homing scenario. [PR1327978](#)
- In Virtual Chassis (VC) or Virtual Chassis Fabric (VCF) scenario using QFX5100, if VXLAN is configured on access ports which are in the same VLAN, it might interfere another independent/unrelated port, a Virtual Chassis port (VCP) or a network port. As a result, members of the Virtual Chassis or the VCF are split. [PR1330132](#)
- After adding new leaf node to VCF, spine fpc loop sent back frame via ingress AE port issue has been fixed from 14.1X53-D47 [PR1335909](#)
- The analyzer status might show as down when port mirroring is configured to mirror packets from an AE member. [PR1338564](#)
- The interfaces with SFP-T transceivers are detected by RSTP as LAN interface type instead of point to point. The problem appears because of an incorrect duplex variable assignment for the link partner. [PR1341640](#)
- FXPC process might generate a core file when removing VXLAN configuration. [PR1345231](#)
- QFX5100-48T 10G interface might be auto-negotiated at 100M speed instead of 10G after peer device reboot. [PR1347144](#)
- On QFX Series switches with AE interface configured, the GTP (GPRS Tunnel Protocol) traffic cannot be hashed correctly when transmitted through the AE interface. [PR1351518](#)
- On QFX3500/QFX3600 platform, OSPF might remain in init status after loading the 14.1X53-D47.4 image. [PR1362996](#)

Interfaces and Chassis

- On QFX3500, QFX3600, or QFX5100 Series switches with MC-AE configured, when local and peer MC-AE are both down and then local MC-AE is up and peer MC-AE is still down, ARP reply might be dropped in this scenario. [PR1282349](#)
- On EX/QFX platform with MC-LAG enabled, if "redundancy-group-id-list" isn't configured under ICCP, upgrading might encounter commit failure during bootstrap. [PR1311009](#)
- On EX4600/QFX5100 platform, if the ICL link is configured on a single interface (such as GE-0/0/0, without LAG) and one member of MC-LAG is down, and both MC-LAG peers are rebooted, packets might drop on ICL of MC-LAG peer where MC-LAG is up. [PR1345316](#)
- If CVLAN(customer virtual local area network) range 16(e.g., vlan-id-list 30-45)is configured in a Q-in-Q(i.e., 802.1ad) scenario, all the 16 VLANs might not pass traffic. [PR1345994](#)

Layer 2 Features

- 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](#)
- When l2ald daemon (l2-learning) is restarted there might be l2ald core file generated. [PR1229838](#)
- When a VTEP interface is flapping frequently, a core dump may be seen which causes traffic forwarding to stop until the pfe is recovered from the core dump. [PR1230198](#)
- On QFX3500/QFX3600/QFX5100 Series switches, if RTG and xSTP are configured on the same VLAN, RTG interface might go to blocked state and packets cannot be forwarded as expected over the RTG interfaces. [PR1230750](#)
- On QFX5100 platform, ARP entry might be learned on STP blocking ports if GARP reply packets or broadcast ARP reply packets are received on spanning tree blocking ports. As a result, traffic loss might be seen. [PR1324245](#)
- When there are multiple logical units on a lag (ae) interface, ingress pop might not work when the configuration is changed on the interface and rolled back. [PR1331722](#)
- When there are multiple logical units on a lag (ae) interface, ingress pop might not work when the configuration is changed on the interface and rolled back. [PR1331722](#)
- On QFX5100 Series platforms, the DHCP packet might be forwarded by the MSTP blocked port if the "dhcp-security group * overrides no-option82" is enabled, which might lead to MAC flapping and form a loop. [PR1345610](#)

MPLS

- On QFX5100 switches, unified ISSU is not supported with MPLS configuration. [PR1264786](#)
- On QFX3500/QFX3600/QFX5100 Series switches with Dynamic Host Configuration Protocol (DHCP) relay configured under Border Gateway Protocol (BGP)-Layer 3 Virtual Private Network (VPN), DHCP clients connect to the switch can not get IP address over BGP-L3VPN. [PR1303442](#)
- On QFX5100 platforms with hot-standby for the l2circuit scenario, the device might not forward traffic if the primary path fails over to standby circuit. [PR1329720](#)

Platform and Infrastructure

- On mixed Virtual Chassis (Virtual Chassis) / Virtual Chassis Fabric (VCF), QFX5100 works as RE (Route Engine) and EX4300 works as Line Card. The knob "interface-mac-limit" configured for interfaces on EX4300 does not work. [PR1259634](#)
- The autonegotiation is not working as expected between EX4300 and SRX5800. [PR1311458](#)
- In Virtual Chassis scenario, when the master member FPC reboots and the interface on which the ARP is learned goes down along with the master FPC, traffic loss might

be observed for about 10 seconds. At that time, the ARP entry cannot be learned from the remaining FPC. [PR1283702](#)

Routing Policy and Firewall Filters

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

Routing Protocols

- In a rare condition, an mt tunnel interface flap cause a backup Routing Engine core file to be created. The exact root cause is not known. [PR1156701](#)
- In situations where BGP multipath is used and there is a large route scale, the lead route (the one selected as active) may not been deleted right away and remains as the active route. The router doesn't consider routes as multipath feasible if they are received from a peer that has gone down. Because of this, the active route will not be feasible for multipath and the router will not able to find a lead route for BGP multipath (lead route has to be the active route). This causes the tearing down of BGP multipath and re-creation of BGP multipath later when the active route is deleted and the new active route became feasible for BGP multipath. [PR1156831](#)
- On QFX5100 Series switches, if Protocol Independent Multicast (PIM) source-specific multicast (SSM) is used, IPv6 multicast traffic from the source might be 100% dropped. [PR1292519](#)
- On QFX5100 platforms, some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update during normal operations, and as a result, multicast traffic from impacted groups traversing through the device might be silently discarded. [PR1320723](#)
- Consistent load balancing minimizes flow remapping in an equal-cost multipath (ECMP) group. Previously on QFX5100 switches, the CLI command 'set policy-options policy-statement ECMP term 2 then load-balance consistent-hash' hid the 'consistent-hash' attribute from the load-balance object. [PR1322299](#)
- On EX4600 or QFX5100 platform, Intermediate System to Intermediate System (IS-IS) Level 2 (L2) Hello packets are dropped when they come from a Brocade device, then ISIS L2 adjacency will fail. The issue is seen only for Jumbo ISIS L2 packets. [PR1325436](#)
- On QFX5100, if an Integrated Routing and Bridging (IRB) interface is loopbacked with a physical interface in another VLAN on the switch, then the IRB interface is not be accessible to remote networks. [PR1333019](#)

Virtual Chassis

- On QFX5100 Virtual Chassis or VCF topology, it takes 10 minutes to obtain RE role if you reboot chassis. The issue is seen only when there is offline chassis in Virtual Chassis/VCF topology. [PR1225696](#)
- On QFX5100 Switches Virtual-Chassis, traffic loop might be seen during network port to VCP (Virtual Chassis Port) conversion. Once those interfaces are removed from Virtual-Chassis, VLAN programming might be affected. [PR1346851](#)

- On QFX5100 Switches platforms, performing vulnerability test using NMAP application might cause fxpc process to crash resulting in traffic loss and coredump. This issue is seen only in Virtual Chassis (VC) environment. Example of Nmap command which causes problem "nmap -v -sO 192.168.101.1". [PR1351411](#)

Resolved Issues: Release 14.1X53-D46

General Routing

- In a data center interconnect (DCI) scenario, when two QFX5100-24Qs in different data centers are interconnected using a 40G link and when DWDM is used in the connection especially with ADVA and single mode fiber (SMF) on one side and multi mode fiber (MMF) on the other, the 40G connection between the two QFX5100-24Qs may not be stable. Sometimes the link will come up and sometimes not. Frame errors might be seen constantly. [PR1178799](#)
- On QFX5100, receiving malformed PIM Hello packets can cause 24-byte memory leaks. [PR1224397](#)
- MACsec issue: The "show security macsec statistics" command does not show expected results. Statistics are incorrectly cleared for each physical interface (IFD) under eth periodic (1 second). [PR1283544](#)
- On QFX5100-48T switch with AE interface configured, if there is a speed setting to 1G on AE member xe interface, the AE link flap might be seen every time when changing configuration and no matter what config is changed. [PR1284495](#)
- If storm control is enabled with the shutdown action on QFX3500, QFX3600, QFX5100, EX4300, or EX4600, the interface with DN and SCTL flags will lose the SCTL flag and will remain permanently down after GRES. [PR1290246](#)
- In QFX5100 if a fan module is released, a major alarm is raised instead of a minor alarm. [PR1291622](#)
- QFX5100 FXPC coredump when a large number of routes is pushed to program in the hardware. [PR1294033](#)
- On QFX5100 switches, the 40-gigabit interface might not come up if a specific vendor direct attach copper (DAC) cable is used. [PR1296011](#)
- On QFX Series platforms with the ZTP feature enabled, the DHCP clients are not getting an IP address if the DHCP pool with /31 subnet is configured. However, if the DHCP pool with /30 or /24 is configured, it works fine. With /31 configured, the DHCP client state remains as "requesting":
user@host> show dhcp client binding IP address
Hardware address Expires State Interface 0.0.0.0 00:00:5E:00:53:00 0 SELECTING
irb.0 0.0.0.0 00:00:5E:00:53:01 0 SELECTING vme.0 10.160.136.65 00:00:5E:00:53:03
0 REQUESTING et-0/0/0.0 [PR1298234](#)
- qfx5100 crash and fxcp core during normal operation [PR1306768](#)
- On all QFX Series platform, all the Internet Control Message Protocol (ICMP) requests that are sending to the Integrated Routing and Bridging (IRB) interface might be dropped for 4-60 seconds if an IRB interface is configured as its gateway in a failover scenario for Virtual-Chassis. [PR1319146](#)

EVPN

- In an EVPN VXLAN scenario, a previous learned MAC address from a remote Ethernet segment Identifier (ESI) cannot be changed to local even it is connected directly. The MAC address of the host might remain as learned from ESI instead of local interface until the MAC address is aged out. [PR1303202](#)

Platform and Infrastructure

- SDK can raise false alarms for parity error messages like `soc_mem_read: invalid index` and `soc_mem_pipe_select_read: invalid index 7289 for memory L3_DEFIP acc_type 1` on QFX5100. [PR1212682](#)
- Dropping the TCP RST packet incorrectly on PFE might cause traffic drop. [PR1269202](#)

Interfaces and Chassis

- QFX5100: Packets are getting dropped if outer TPID is set with 0x9100. [PR1267178](#)

Multiprotocol Label Switching (MPLS)

- QFX5100/EX4600: Stale MPLS label entries might exist on MPLS table in PFE after deleting or disabling the underlying interface of IRB/AE interface. [PR1243276](#)
- On QFX5100/QFX3500/QFX3600/EX4600 Series switches in Multiprotocol Label Switching (MPLS) penultimate-hop popping (PHP) scenario, after MPLS next-hop changed and then back, traffic might stop passing LSP. [PR1309058](#)

Multicast Protocols

- Multicast traffic is black-holed when the master reboot is done on a QFX5100 or EX4600 Virtual Chassis. [PR1164357](#)

Routing Protocols

- On EX4600/QFX Series switches with unicast-in-lpm configured, EBGp packets with `ttl=1` and non-EBGP packets with `ttl=1`, whether destined for the device or even transit traffic, go to the same queue. This might result in valid EBGp packets drop which can cause EBGp flap. [PR1227314](#)
- QFX5100 might log "Cannot program filter "xxx" (type VFP FBF)" but the VFP entries did not reached `max_count 512`. [PR1229375](#)
- If the number of 'Ref count' entries used by firewall filter applied on loopback interface is more than 255, log 'dc-pfe: list_destroy(): non-empty list (1)' is printed after commit the firewall filter configuration. [PR1286209](#)

Security

- The Juniper Networks enhanced jdhcpd process might experience high CPU utilization, or crash and restart upon receipt of an invalid IPv6 UDP packet. Both high CPU utilization and repeated crashes of the jdhcpd process might result in a denial of service as DHCP service is interrupted. Refer to JSA10800 for further details. [PR1119019](#)
- A buffer overflow vulnerability in Junos OS CLI might allow a local authenticated user with read only privileges and access to Junos CLI, to execute code with root privileges. Refer to JSA10803 for further details. [PR1149652](#)
- Two vulnerabilities in telnetd service on Juniper Networks Junos OS might allow a remote unauthenticated attacker to cause a denial of service through memory and/or CPU consumption. Please refer to JSA10817 for more information. [PR1159841](#)
- Junos: Potential remote code execution vulnerability in PAM (CVE-2017-10615); Refer to <https://kb.juniper.net/JSA10818> for more information. [PR1192119](#)
- On Junos OS devices with SNMP enabled, a network-based attacker with unfiltered access to the Routing Engine can cause the Junos OS snmpd process (daemon) to crash and restart by sending a crafted SNMP packet. Repeated crashes of snmpd process can result in a partial denial-of-service condition. Additionally, it might be possible to craft a malicious SNMP packet in a way that can result in remote code execution. Refer to <https://kb.juniper.net/JSA10793> for more information. [PR1282772](#)

Virtual Chassis

- On QFX platform with non stop routing configured on performing a routing engine switch, there is possibility of drops for a short duration. Commit marker sequence has been modified to check state of commit, and only if a valid entry is seen then warning is prompted. [PR1225829](#)

VXLAN

- AE interface cannot forward traffic in VXLAN configuration. [PR1213701](#)

Resolved Issues: Release 14.1X53-D45

General Routing

- In a Layer 3 VPN, if IRB is used between the penultimate hop and the PE node, if checking VRF connectivity using PE to PE ping, then pinging to the PE loopback address or interface IP address from the remote PE does not work. [PR1211462](#)
- Due to some register values at PHY for tuning the cable is not optimal, the interface might experience continuous flapping. [PR1273861](#)
- Previous PR 1169106, changed the behavior for 'rxbps' to report bits per second for streaming data, instead of bytes. The output for "show analytics traffic-statistics interface" was changed from 'Octets per second' to 'Bits per second' as seen below, but the actual value reported remained in bytes: QFX5100> show analytics traffic-statistics interface ge-0/0/2 Time: 00:00:00.363490 ago, Physical interface: ge-0/0/2 Traffic Statistics: Receive Transmit Total octets: 87926097472 11412 Total packets: 1373845261 41 Unicast packet: 1373845261 3 Multicast packets: 0 34 Broadcast

packets: 0 4 Bits per second: 762063768 1584 <<<<<<<< Display shows Bits per second, but 762063768 BYTES are reported Packets per second: 1488405 0 CRC/Align errors: 0 0 Packets dropped: 0 0 Code changes in this PR completed the changes to reflect the correct value in bits. Additionally, documentation has since been changed to reflect 'rxbps' represents 'Total bits received per second':
https://www.juniper.net/documentation/en_US/junos/topics/concept/analytics-streaming-statistics-remote-understanding.html [PR1285434](#)

- When ovsdatabase in QFX5100 was corrupted accidentally, ovsdb-server daemon cannot launch properly even though rebooting QFX5100. [PR1288052](#)
- On QFX5100 Series Switches with EVPN/VxLAN deployed, VLAN flood index might not be programed correctly on PFE. Due to this, the ARP requests to the virtual gateway are dropped, and traffic forwarding is affected. [PR1293163](#)

Class of Service (CoS)

- In current design, in order for the knob "transmit-rate" applied within the "forwarding-class-set" to work properly, must configure the knob "guaranteed-rate" for "forwarding-class-set", this is mandatory. Without "guaranteed-rate" configured, if configuring "transmit-rate" in value then "transmit-rate" applied within the "forwarding-class-set" does not work, this is as per design, but if configuring "transmit-rate" in percent then "transmit-rate" applied within the "forwarding-class-set" still works, this is not as per design. [PR1277497](#)

EVPN

- On QFX5100 deleting a vxlan causes traffic disruption in all other vxlans. Rolling back the vxlan deletion alone would not resolve the issue. Rollback and l2-learning restart thereafter is needed to recover. Issue is fixed in JunOS 14.1X53-D40 onwards. [PR1215883](#)
- In a VxLAN scenario, the Packet Forwarding Engine manager daemon (fxpc) and kernel crash might be observed after adding MTU configuration on QFX5000-VC platform. [PR1283966](#)
- A new option exclusive-mac is added under protocols l2-learning global-mac-move set protocols l2-learning global-mac-move exclusive-mac <mac>. [PR1285749](#)

Interfaces and Chassis

- On a QFX5100 VC/VCF when you upgrade the firmware via normal upgrade or NSSU from 13.2X51-D30.4 to 14.1X53-D35 or from 14.1X53-D35 to 14.1X53-D40, you can sometimes encounter and DCD core. [PR1276745](#)
- On a QFabric system, unexpected behaviour or crash may be observed if make-before-break timeout is configured less than 30 seconds. [PR1286613](#)

Layer 2 Ethernet Services

- A new static MAC is configured under AE interface, but the MAC of the LACP PDUs sent out is not changed. [PR1204895](#)

Multiprotocol Label Switching (MPLS)

- On QFX5100/EX4600 Series switches, when deleting an IRB/AE interface with MPLS enabled, it might not delete related entries from MPLS routing table in PFE. Which leading stale MPLS routes in PFE. The stale entries in the MPLS forwarding table(PFE) will impact the scale scenarios. [PR1243276](#)

QFabric Systems

- 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](#)

Routing Protocols

- 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 Virtual Chassis Fabric, you might see an error such as MMU ERR Type: 1B error, Addr: 0x001052cf, module: 42, which indicates that there was an ECC error in the PFE MMU counter memory. ECC errors are corrected by the hardware without software intervention and are corrected only when a packet hits that memory. Reading an ECC-errored entry always generates an interrupt; however, the error will only be corrected when the packet hits the memory. Because this is a counter memory, the counter thread reads this memory continuously, and hence you see continuous error messages. [PR1198162](#)
- On QFX3500/QFX3600/QFX5100/EX4300/EX4600 Series switches, Border Gateway Protocol (BGP) packets with IPv6 link local address as destination address are not punted to CPU, so it results in BGP session failing to establish. [PR1267565](#)
- On QFX5100 and EX4600 switches, when you are adding or deleting routes on a system with a large number of routes, in rare cases, the fxpc process might access an already freed-up memory space, causing the fxpc process to crash and restart with a core file generated. [PR1271825](#)
- On QFX5100-24Q and QFX5100-48S, if IPv6 link local packets are from members other than the first member of a channelized interface (for example, xe-0/1/2:1, xe-0/1/2:2, or xe-0/1/2:3), IPv6 packets are dropped. [PR1283065](#)

Resolved Issues: Release 14.1X53-D44

General Routing

- If Media Access Control Security (MACsec) session flaps, dot1x might crash and core dump, then MACsec session would fail to be established. [PR1251508](#)

- 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](#)

Interfaces and Chassis

- On QFX3500/QFX3600/QFX5100/EX4600/EX4300 Series switches with MC-LAG configuration, if ARPs are resolved across VRFs by route leaking, it might cause traffic to be dropped in scaling ARP entries (for example: 3K). [PR1241297](#)

Layer 2 Features

- On QFX5100, MAC learning will be very slow when clearing MAC addresses in cases of scale MAC learning (128k). [PR1240114](#)
- On QFX5100/EX4600 Series switches, if one filter last term is configured with reject action and applied on lo0 (loopback interface) interface, then it might cause media access control address (MAC) learning flap when the Dynamic Host Configuration Protocol (IGMP)/Dynamic Host Configuration Protocol (DHCP) packets are received. [PR1245210](#)
- On QFX5100 switches, if you configure a Layer 3 interface with **vlan-tags outer 0x9100.xx**, then packets are dropped on this interface. [PR1267178](#)

Platform and Infrastructure

- On QFabric, high disk utilization might be seen on Master DG because the processes (e.g., sfcsmnpd, cnm, mgd, sfctrap handler, dgsnmpd etc.) keep opening these log files and keep updating incoming immediate logs. [PR1245817](#)
- In rare cases, the Packet Forwarding Engine might drop the TCP RST (reset) packet from the Routing Engine side while doing GRES or flapping an interface, and traffic might be dropped. [PR1269202](#)

QFabric Systems

- 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](#)

Routing Protocols

- A filter attached to the lo0 interface with terms containing either destination-port-range-optimize or source-port-range-optimize statements will unexpectedly discard traffic. [PR1228335](#)
- When polling the SNMP MIB jnxFirewallsEntry and if more than one firewall filter was configured and attached to any logical interfaces on the QFX3500 platform, the counters for only one firewall filter would be returned. Now all filters and counters are returned when polling the MIB. [PR1250776](#)

Virtual Chassis

- A VCF might experience an outage for a while when a Virtual Chassis port (VCP) is flapping. [PR1158798](#)

Resolved Issues: Release 14.1X53-D43

General Routing

- [illegible]

the PIC online. This has been resolved by removing the dependency to wait for all FPC's to get initialized to bring the PIC's to online state. In the images going forward starting from this release, as soon as any of the FPC's gets initialized, PIC online will happen on QFX3008 Interconnect. [PR1261685](#)

- On standalone QFX5100 or on Virtual Chassis (VC) / Virtual Chassis Fabric (VCF) with QFX5100, Media Access Control Security (MACsec) licenses may not be added sometimes. [PR1269667](#)

Class of Service (CoS)

- On QFX5100/EX4600/EX4300 Series switches, if forwarding-class-sets with more than one forwarding-classes is applied to interface, and the scheduler for these forwarding-classes under this forwarding-class-sets are not configured with shaping-rate, then it might cause traffic to be dropped for this interface. [PR1255077](#)

EVPN

- On QFX5100 Series with VxLAN/EVPN configured, when multiple IP addresses are configured for VTEP source interface, traffic might be dropped on spines. [PR1248773](#)
- Removing force-up on an active link can cause programming issues on the QFX5100. Traffic returning from the destination will not be forwarded on an egress of the QFX5100. [PR1264650](#)

Interfaces and Chassis

- The AE interface might be down after NSSU is done on QFX5100 or EX4600 switches. [PR1227522](#)
- On an EX4300, EX4600, or a QFX Series switch, if one logical interface is configured in one VLAN and then is deleted and added to another VLAN, traffic might not be transmitted correctly. [PR1228526](#)
- On QFX3500/QFX3600/QFX5100/EX4600/EX4300 Series switches with MC-LAG configuration, if ARPs are resolved across VRFs by route leaking, it might cause traffic to be dropped in scaling ARP entries (for example: 3K). [PR1241297](#)
- [QFX5100-VC] / [14.1X53-D40 & 14.1X53-D42.3] IGMP general query packet destined to 224.0.0.1 are sent back on the received interface, breaking the unicast connectivity. [PR1262723](#)

Junos Fusion Satellite Software

- The following conditions must be met before a Junos switch can be converted to a satellite device when the action is initiated from the aggregation device: 1. The Junos switch must be in factory default settings OR it must have include 'set chassis auto-satellite-conversion' in its configuration 2. The package used to do the conversion must be one of SNOS 3.0, SNOS 1.0R5, SNOS 2.0R2 or higher. [PR1249877](#)

Layer 2 Features

- After rebooting or clearing interface statistics, excessive input/output statistics might be observed in "show interface aeX" command on QFX5100/EX4600/EX4300 Series switches. [PR1228042](#)
- On QFX5100, in cases of scale MAC learning (128k), MAC learning is very slow when MAC addresses are being cleared. [PR1240114](#)
- On QFX5100/EX4600 Series switches, VxLAN/EVPN is configured with multi-homing mode, the DF (Designated Forwarder) might forward BUM traffic received from ESI (Ethernet Segment Identifier) peer to CE facing interface after deleting/adding back VLAN. [PR1260533](#)

Multiprotocol Label Switching (MPLS)

- In MPLS layer 2 or layer 3 VPN scenario, QFX5100/EX4600 Series switches work as PE router and the core interface of PE using IRB interface. When deactivating/disabling/deleting underline member interface of the IRB, and if the (parent) IPv4 nexthop is uninstalled first before cleaning up the (child) MPLS nexthop, the fxpc process might crash and restart. And the FXPC core will be seen. [PR1242203](#)

Network Management and Monitoring

- On Qfabric, once active LAG goes down on CPE-1 (control plane ethernet) due to LACP timeout or some other reason and physical interface member interface does not go down, director group (DG) is not moving to standby links on CPE-2. This will cause all system protocols flap on the FM. [PR1253825](#)

Routing Policy and Firewall Filters

- On QFX Series switches and EX4300 switch, the command of showing policy which has parameter of "load-balance consistent-hash" might cause rpd to crash. [PR1200997](#)

Routing Protocols

- A vulnerability in IPv6 processing has been discovered that may allow a specially crafted IPv6 Neighbor Discovery (ND) packet to be accepted by the router rather than discarded. The crafted packet, destined to the router, will then be processed by the routing engine (RE). A malicious network-based packet flood, sourced from beyond the local broadcast domain, can cause the RE CPU to spike, or cause the DDoS protection ARP protocol group policer to engage. When this happens, the DDoS policer may start dropping legitimate IPv6 neighbors as legitimate ND times out. Refer to JSA10749 for more information. [PR1220209](#)
- On QFX5100/EX4600 Series switches, firewall filter that contains filter-specific policers might not process packets correctly after committing for the TCAM entries of filter are programmed over multiple slices of TCAM memory space. Note: Firewall filter terms are programmed as TCAM entries in the TCAM memory table. So in scenario with multiple of filter terms (for example: more than 256), this state might be seen easily. [PR1232926](#)
- QFX5100 and EX4600 switches might not send router advertisement (RA) packets to clients when **igmp-snooping** is configured on a user VLAN, and the end clients connected to the devices can lose IPv6 connectivity. [PR1238906](#)
- 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](#)
- In a VCF scenario that includes an EX4300 switch, if **fabric-tree-root** is configured, then the broadcast, unknown, and multicast (BUM) traffic might not be forwarded. [PR1257984](#)

Resolved Issues: Release 14.1X53-D42

General Routing

- Ports on the uplink module (QFX-EM-4Q) on QFX5100-24Q model alone do not get converted to VCP ports even after explicitly converting them to VCP. As a workaround, after converting the ports to VC ports, rebooting the QFX5100-24Q would complete the the VCP conversion successfully. Issue is fixed in JunOS 14.1X53-D42 onwards. The in-built 24 ports on the QFX5100-24Q do not have any such issues. [PR1158657](#)
- In an EX4600 or QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF), when using scp on the management interfaces to copy files greater than about 150MB, you might see protocol flapping and Routing Engine TCP connections dropping. [PR1213286](#)
- On QFX3500/QFX3600/QFX5100 Virtual Chassis (VC) or Virtual Chassis Fabric (VCF) with "nonstop-routing (NSR)" and "switchover-on-routing-crash" configured, if the rpd on master Routing Engine crashes, the VC or VCF fails to perform switchover to the backup Routing Engine. [PR1220811](#)
- On QFX5100 and EX4600 switches, during a nonstop software upgrade (NSSU), if an aggregated Ethernet (AE) interface is configured with multiple subinterfaces across multiple Flexible PIC Concentrators (FPCs), the AE interface might go down. [PR1227522](#)
- 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. [PR1234945](#)

EVPN

- When the master RE is rebooted and comes back up in QFX5100-VC with VxLAN UDP port configured, the vtep tunnel is created before it gets the updated udp-port information, and the encapsulated packets might be sent with the incorrect UDP destination port. [PR1214750](#)

Infrastructure

- On QFabric systems, no more PFE IDs can be allocated when the SOURCE_TRUNK_MAP_MODBASE table reaches its limit. [PR1236584](#)
- 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](#)

Layer 2 Features

- The fxpc process can generate a core file on QFX5100. [PR1231071](#)

Network Management and Monitoring

- On QFabric system, sfcsmmpd log messages are getting populated after upgrading to 14.1X53-D40 and D41. The sfcsmmpd is supposed to run only on master DG, hence sfcsmmpd is not running on backup DG as expected. These logs messages are harmless in backup DG only and do not impact any functionality. [PR1238939](#)

QFabric Systems

- While installing Junos AIS (JAIS) on a QFabric system, JAIS is not getting pushed to Network Node group (NW-NG0) node. JAIS gets installed on all other nodes(director, redundant server node group (RSNG) and server node group (SNG)) except NW-NG group. [PR1233166](#)
- On Qfabric, issue show fabric based multicast commands from Network Node Group (NNG) and show is dereferencing the freed memory (command example: root@NNG> show fabric multicast) during route delete. This may cause rpdf crash. [PR1242781](#)

Routing Protocols

- On QFX5100/EX4600 Series switches, if one filter is configured with match of "ipv6 tcp-established", then committing configuration might cause pfe process to crash. [PR1234729](#)
- 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](#)

Resolved Issues: Release 14.1X53-D40

General Routing

- In case you are using QFX5100-48T-6Q, "show chassis hardware" displays QFX5100-48C-6Q like below. ----- root@QFX5100-48T> show chassis hardware Hardware inventory: Item Version Part number Serial number Description Chassis TR0214999999 QFX5100-48C-6Q -----, [PR1006271](#)
- On QFX Series mixed Virtual Chassis Fabric (VCF), software rollback with the force option (request system software rollback force) might not work. [PR1028666](#)
- Certain QFX and EX Series devices do not pad Ethernet packets with zeros, and thus some packets can contain fragments of system memory or data from previous packets. This issue is also known as 'Etherleak' and often detected as CVE-2003-0001. Refer to JSA10773 for more information. [PR1063645](#)
- By enabling this configurations, it Drops spanning-tree protocol BPDUs (for STP, MSTP, and RSTP) entering any or a specified interface The BPDUs drop feature can be specified only on interfaces on which no spanning-tree protocol is configured. This behavior is same as EX platforms. [PR1084116](#)

- QFX5100-48S-6Q or QFX5100-96S-8Q might display incorrectly as "QFX5100-24Q-2P" in the output of "show chassis hardware" after the Flexible PIC Concentrator (FPC) restart or master role switchover. [PR1093677](#)
- If MAC move limit is configured to drop traffic, QFX and EX Series switches might forward traffic instead of dropping traffic when the MAC move limit is exceeded. [PR1105372](#)
- In a Virtual Chassis Fabric (VCF) with three or four spine devices, the spine devices operating in the linecard role cannot assume the Routing Engine role, including in cases where the master or backup Routing Engine fails. [PR1115323](#)
- On a QFabric system, when enabling "set fabric routing-options traceoptions" command on the initial Network Node Group (NNG) master for an extended period might cause the log files (approximately 1GB) exhausted the /var partition. As a result, an abnormal shutdown of the master Routing Engine (RE) was seen and mastership switched over to backup RE. In the corner case, the new master experienced a re-synchronization failure with the line cards PFEMAN thread. Under these circumstances, the Packet Forwarding Engine manager (pafxpc) restart is expected which will lead to interface flapping. [PR1133679](#)
- On QFX5100 switches, the openflowd process might generate a core file. [PR1142563](#)
- From 14.1X53-D36, there is a commit check added to prevent more than one IFL per physical interface(IFD) assigned to one single VLAN. [PR1144123](#)
- On EX4300/EX4600/QFX3500/QFX3600/QFX5100 series switches, if you insert bad SFP or SFP+ optic in a port and replace it with a good optic, then the good optic might not come up. [PR1144190](#)
- On EX4600/QFX5100 and QFX10000 series switches, after performing command "show version detail", an error message "Error: abnormal communication termination with app-engine-management-service daemon" might be seen at the end of the output. [PR1144234](#)
- On QFX5100/EX4600 Series switches, the switch might be in abnormal state after Junos try to write the large file on virtual-disk and reboot at the same time due to the Quick Emulator (QEMU) cache mode is set to "none". The recommended setting is "write through" mode. This is Quick Emulator (QEMU) issue which reporting an error when Junos try to disk IO or blocking disk IO access from Junos. There are two pieces of software in QFX5100: hypervisor (host software) and JUNOS (a VM running on top of hypervisor). All peripherals (e.g. Disks, network cards, etc) on JUNOS are simulated by a piece of software called QEMU. [PR1146353](#)
- After the number of DHCP server IPs in the dhcp-relay configuration is modified (increased or decreased), messages log file will be filled with following error messages and eventually cause DHCP process (jdhcpd) to crash. jdhcpd:
%USER-3-DH_SVC_SENDMSG_FAILURE: sendmsg() from 10.161.102.1 to port 67 at 0.0.0.0 via interface 615 and routing instance VR08_v881_900_office_system failed: Can't assign requested address [PR1147831](#)
- On EX4600/QFX Series switches, in corner cases, the PFE manager (fxpc) might crash when an SFP-T transceiver is removed/inserted too quickly or the interface is deleted. [PR1152097](#)

- On EX4600/QFX series platform, or its virtual chassis or virtual chassis fabric, the device automatically restarts (for the UFT configuration to take effect) when the Unified Forwarding Table (UFT) profile is reconfigured or modified. When this happens in a scaled Virtual Chassis or Virtual Chassis Fabric (VCF) environment, the VC/ VCF might become unstable and fail to recover, the VCF (all member devices) must be rebooted to reestablish stable VC/ VCF operation. To avoid this situation, configure the desired UFT profile when initially setting up the standalone/ VC/ VCF, rather than as a configuration update later during the standalone/ VC/ VCF operation. After the fix, for standalone and virtual chassis with a single member, it works as before. For VC and VCF with more than one member, the member does not restart anymore. And, it generates a syslog message to notify the user to restart the system manual when UFT config is changed. [PR1152102](#)
- On QFX5100 switches, a child member might drop the incoming Link Aggregation Control Protocol (LACP) frames when this child member is moved from an access-mode VXLAN LAG interface to a trunk-mode VXLAN LAG interface. [PR1153042](#)
- On QFX Series and EX Series switches, if you configure VRRP with an MC-LAG between the master and backup switches, both VRRP members of IRB interfaces might stay in the master state after a software upgrade. [PR1157075](#)
- On a VCF platform, the memory usage limitation for the vccpd daemon is 131Mbytes in memory. Any VCP port flapping will cause a small memory leak (256KB~1MB) in VCF and if the memory usage is reached 131Mbytes, the vccpd will crash with a core dumped and then restart. In the meantime, a member of the VCF will disconnect from VCF, this will have service impact for a while before vccpd comes up again. [PR1158798](#)
- In VC or VCF deployment, if a connection between members is made after all other VCP links have been auto-converted to VCP, the new connection may not successfully convert to VCP. [PR1159242](#)
- On an EX4600 Virtual Chassis or a QFX Series Virtual Chassis or Virtual Chassis Fabric (VCF), if you convert the Virtual Chassis port (VCP) to a network port by issuing the "request virtual-chassis vc-port delete" command, broadcast and multicast traffic might be dropped due to the port remaining programmed as a VCP in the hardware. [PR1159461](#)
- An insufficient authentication vulnerability on platforms where Junos OS instances are run in a virtualized environment, may allow unprivileged users on the Junos OS instance to gain access to the host operating environment, and thus escalate privileges. [PR1161762](#)
- On an EX4600/QFX Virtual Chassis with the members of the LAG are on the same VC member device, the multicast packets getting dropped (approximately 120 sec) during the master Routing Engine (RE) role switch reboot. [PR1164357](#)
- On QFX10000/EX4600/EX4300 Series switches, the Digital Optical Monitoring (DOM) update takes more than 25 sec to update when the interface goes down right away, and the issue is not seen on other platforms. [PR1165507](#)
- On a QFX5100 switch with an integrated routing and bridging (IRB) interface configured as a Layer 3 interface and with two hosts (Host A and Host B) connected to the switch, if you deactivate the IP address on Host A and then configure the same IP address on

Host B, the outgoing interface of the IP address might not be changed in the ARP table. [PR1166400](#)

- If a QFX5100 Virtual Chassis is created with a QFX5100-48S in the routing-engine role and a QFX5100-48T in the linecard role, ports of the QFX5100-48T might be shown as having media type Fiber. [PR1166810](#)
- On a QFX Virtual Chassis Fabric (VCF), when adding more members to VCF, since more members lead to more physical interfaces and more DEVBUF (device buffer) type of memory. When members of a VCF over 24 it will most likely trigger the DEVBUF type memory reach its limitation and the syslog messages like the following will be seen in syslog file. /kernel: %KERN-5: kmem type devbuf using 331293K, approaching limit 412876K [PR1167390](#)
- On QFX5100-48T, when issuing 'show interface <INT> xtensive' (or strictly 'show interface <INT> media') that the "Local resolution:" section of the "Autonegotiation information" section continues to show that flow control is enabled for both tx and rx even though flow control has been explicitly configured as disabled and additionally shows as disabled in the top portion of output. [PR1168511](#)
- On QFX Series switches, up to four port-mirroring analyzers can be configured, with maximum two of these used for mirroring ingress traffic and maximum two mirroring egress traffic. If a configuration with more than four analyzer sessions per QFX switch/Node is committed, the commit will fail and a relevant error message will be reported. The current code change introduced by this PR generates an error message which is sent to the user's console when such configuration with more than four analyzer sessions is attempted. This code change does not remove the limitation of four analyzer sessions per QFX switch/Node. [PR1168528](#)
- When enable LLDP and interface description is long (greater than 32 chars) on remote switch, the l2cpd (Layer 2 Control Protocol process) might crash with core dump if performing SNMP MIB walk since LLDP code is running within l2cpd. [PR1169252](#)
- On a Qfabric system, the syslog message "on /: filesystem full" is observed continually on Diagnostic Routing Engine (DRE), and clearing "Linking" file will not cause service impact. The following log message will be observed when this issue occurs: DRE-0 /kernel - - - pid 1299 (sh), uid 0 inumber 37 on /: filesystem full DRE-0 /kernel - - - pid 1299 (sh), uid 0 inumber 37 on /: filesystem full [PR1169760](#)
- On EX4600/QFX5100 switches, when a VLAN is mirrored, the mirrored packets may contain 38 additional bytes. The IP address in this packet is randomly generated and may appear as one of many existing, valid IP addresses on the Internet. It may appear as ERSPAN as well, which is a proprietary non-Juniper protocol. These addresses and packet types can be ignored. They may appear as alerts in certain IDP / IDS's and packet analyzer applications which can be ignored. [PR1170589](#)
- On QFX5100-48T switch with a release before 14.1X53-D35, copper interface with auto-negotiation (AN) enabled by default when an interface without explicit auto-negotiation configured. From 14.1X53-D35 and onwards, the default behavior has been changed, when a copper interface without explicit auto-negotiation configured, it comes up with 100M and AN disabled by default. If the interface connects to an interface on peer end that with AN enabled, the link may not come up for AN is disabled on the QFX5100-48T side. After 14.1X53-D37 14.1X53-D39 14.1X53-D40 and onwards,

the default behavior has been changed again, AN will be enabled for Copper ports 0-47 by default. [PR1170909](#)

- On QFX5100-48S switch, if insert or remove a SFP-T optic from one port, then fxpc core might occur and traffic forwarding might be interrupted. [PR1170941](#)
- If you enable aggregated Ethernet links by deleting the disable command, LACP core files might be generated. [PR1173562](#)
- On QFX5100 device, packet loss and framing errors might be observed on QSFP+40GE-LX4 transceiver. [PR1177499](#)
- On EX4600, QFX3500, QFX3600, and QFX5100 switches, some SFP-T modules could not be recognized due to low timeout for I2C read/write. [PR1180097](#)
- FCoE sessions/non-FCoE traffic might be affected when links to Interconnect are disconnected caused by Queue corruption [PR1182274](#)
- On QFX5100 Series switches, in VXLAN scenario with scaled AE interfaces or when AE child member is deactivated-activated back, flood next-hop is not getting updated with physical child interface when child list entries are not populated completely for all ae sub-ifls. This might result in traffic drop. [PR1182495](#)
- When the show chassis environment is executed the temp value shows correct readings but jnxOperatingTemp.7.2.0.0 = 0 might show "0" This is a bug. [PR1190186](#)
- On Junos based platforms, if they are configured as DHCP client, DHCP offer packet which giaddress is not zero might be dropped. [PR1191452](#)
- On QFabric System QFX3000-G and QFX3000-M, when the command "show chassis environment pem interconnect-device" is executed, a chassisd core might be generated. Though, this does not create any production impact. [PR1193597](#)
- On QFX3000-G or QFX3000-M QFabric System, when unknown unicast frames received at a high speed on a certain node. The sfid of this Node might be stuck after some time and the Node stops learning MAC address. No new MACs are learnt after this. This will cause unknown unicast traffic flooding, and further result in network connectivity issues, and network performance degradation of the customer network. [PR1200829](#)
- On QFX5100-96S with 850W AC power supply inserted, in certain environments, because of a software defect, there is a statistical probability of the QFX5100 850W AC power supply shutting itself down. [PR1203591](#)
- After you added or removed PEM on QFX5100, "show chassis environment pem" does not output correct Current(A) and Power(A) usage. [PR1204850](#)
- On QFX5100 switches, 'Rx power low warning set' messages might be logged continuously for channelization ports that are in the DOWN state with snmpwalk running in the background. [PR1204988](#)
- On QFX5100 Series switches, configuring MSTI not in incremental order might cause the existing MSTI fails to learn MAC address, and then traffic forwarding is affected. [PR1205074](#)

- On QFX5100 Series switches, flow modifying operations may cause openflow process core dump due to purge timer software issue. [PR1206127](#)
- There are basically three arguments (periodic, diagnostic, and tx) for the `lcmd -f 0 -d chassis -c` command. These top-level commands expect different numbers of arguments. If any one of the arguments is missing when the command is executed on a QFX3500 or QFX3600 switch, the chassisd process might crash. [PR1206328](#)
- The QFX5100 returns wrong information for flow registration to the openflow controller even though `show openflow ...` output from the CLI shows the correct output. [PR1206572](#)
- If a QFX5100 switch or VCF is configured with IGMP snooping but not with any PIM-related configuration, a mcsnoopd memory leak might occur when the device receives PIM Hello packets that need to be forwarded further. When PIM hellos are arriving on the device, 12 bytes are allocated for every PIM hello packet, increasing the amount of memory consumed by the mcsnoopd process. As a workaround, either restart the mcsnoopd process or apply a firewall filter that discards PIM packets on the loopback (lo0) interface of the device in the input direction. [PR1209773](#)
- In a Layer 3 VPN, if IRB is used between the penultimate hop and the PE node, if checking VRF connectivity using PE to PE ping, then pinging to the PE loopback address or interface IP address from the remote PE does not work. [PR1211462](#)
- On a QFX5100 with a JPSU-850W-AC-AFO power supply inserted. In rare conditions, when closely spaced I2C commands were executed by the software within 100us, the JPSU-850W-AC-AFO may be reset. If the QFX5100 with a single power supply inserted, the box will reboot. With this fix, a 100 microseconds delay were added between every i2c commands to prevent JPSU-850W-AC-AFO reset, and to prevent QFX5100 unexpected reboot. [PR1211736](#)
- In an EX4600 or QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF), when using `scp` on the management interfaces to copy files greater than about 150MB, you might see protocol flapping and Routing Engine TCP connections dropping. [PR1213286](#)

Class of Service (CoS)

- On QFX Series switches, the `cosd` might crash and generate the core file when DSCP classifier with both multicast Forwarding Classes (Queue 8-11) and Unicast Forwarding Classes (Queue 0-7) is applied on a layer 3 interface. [PR1137104](#)
- On QFX5100 and EX4600 switches, ICMP, SSH, and ARP traffic generated by the switch might be forwarded to queue 7 (network-control); the default behavior is that the traffic would be forwarded to queue 0 (best-effort). [PR1178188](#)
- On QFX5100 Series switches, if shaper is configured with shaping value in terms of percent of interface speed, after interface speed changed through CLI, this shaping value will not be changed with new interface speed, and interface flap might happen then. [PR1184505](#)
- In ETS configuration, if transmit-rate is configured at queue-level, guaranteed rate should be configured at the TCP level. If not, `commit` does to fail, but a syslog message is logged to inform the config failure. The config is not pushed to kernel/PFE. In case of VC, when a member joins, since the config check is already done on master, the config

is sent to members. Since the guaranteed rate configured is 0, the logic to calculate the transmit-rate fails. [PR1195498](#)

EVPN

- cli to allow user configuration to provide (EBGP) AS number for per vni auto-derived route target [PR1108613](#)
- On QFX5100-96S standalone/VC/VCF mode, the packet forwarding engine manager daemon (fxpc) process might crash continuously when the SFP/SFP+ transceivers is removed and then inserted in the specific 10-gigabit ethernet (xe) interface (xe-*/0/95) which has extended-vlan-bridge/flexible-vlan-tagging configuration. [PR1159156](#)
- On QFX5100 switches, if a trunk interface is a VXLAN port, tagged frames matching the native VLAN ID might be sent out with the native VLAN tagged. [PR1164850](#)
- On QFX5100 which is running with VxLAN Open vSwitch Database (OVSDb) feature, the packet forwarding engine manager (fxpc) might crash and generates the core file due to heap memory exhaustion for kernel. This is specific issue with OVSDb and does not affect multicast VxLAN. [PR1187299](#)
- On QFX5100 deleting a vxlan causes traffic disruption in all other vxlans. Rolling back the vxlan deletion alone would not resolve the issue. Rollback and I2-learning restart thereafter is needed to recover. Issue is fixed in JunOS 14.1X53-D40 onwards. [PR1215883](#)

High Availability (HA) and Resiliency

- After graceful switchover is triggered in the master VRRP router for the first time, the master state for all the VRRP instances are toggled to the backup and come back to the master immediately. During this time, all the traffic is dropped and comes back. [PR1142227](#)
- On QFX5100/EX4600 Series switches with MC-LAG is configured, the MC-LAG local state is not updated and cause forwarding traffic drop after performing in-service software upgrade (ISSU). This issue is not seen with PFE restart/system boot up. [PR1151658](#)

Infrastructure

- On EX4300 Virtual Chassis, when upgrading from Junos OS Release 15.x to Release 16.x via NSSU, the backup or any linecard upgrades first to a new image, and then the old master might have an upgrade failure, and keep rebooting. [PR1190164](#)

Interfaces and Chassis

- With multi-chassis lag configuration, the switch-options service-id configuration is required. If service-id configuration is missing, there should be a commit error. This commit error is missing. [PR989778](#)
- Fixed in 14.1X53-D40 [PR996005](#)

- You might be unable to commit your configuration if you modify the subnet of an IP address on an IRB interface by using the "replace pattern" command. [PR1119713](#)
- 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](#)

Layer 2 Features

- On QFX Series switches, when transmitting large packet which is more than MTU configured and not be fragmented on the IRB interface, ICMP error packet about type3 with code4 can not be generated. The large packets are getting silently dropped. [PR1089445](#)
- The Packet Forwarding Engine manager daemon (FXPC) might crash on a QFX5100 switch if multiple processes attempt to access the Ethernet-switching table/database at the same time. [PR1146937](#)
- On an EX4300 switch in a VCF, if a Layer 3 AE interface is looped back with a Layer 2 port in the same VLAN, then traffic with the same destination MAC to the AE interface is dropped (for example, the ping address of the AE interface). [PR1157283](#)
- On EX4600/QFX5100/NFX250 series switches, configure LFM (link-fault-management) and action-profile with action link-down on an interface. If this interface is down and up at the first time, LFM Discovery is success and interface is able to be up. But if this interface is down and up at the second time, then LFM Discovery is failure and interface never recover from Link-Down state. [PR1158110](#)
- On a QFX5100 switch, if you delete a VLAN and create a new VLAN with a different VLAN ID but use the same VNI, and you commit those changes within a single commit, a MAC learning failure might occur on the newly created VLAN. These system logging messages might be displayed: fpc0 BCM-VIRTUAL,brcm_vxlan_hw_add(),263:Failed to Program vxlan bd(22) token(0xf) status(-8) fpc0 BCM-VIRTUAL,brcm_virtual_bd_add(),626:Cannot create Virtual-BD for bd(22) fpc0 BCM-VIRTUAL,brcm_virtual_port_add(),101:Port(ge-0/1/2) add came before bd(22) add fpc0 LBCM-L2,pfe_bcm_l2_addr_delete_by_vlan(),52:delete L2 entries associated with bd 21(65535) failed(-4) [PR1161574](#)
- On QFX5100/QFX3500, buffer is corrupted on port 0 (*/*/0) and error message MACDRAINTIMEOUT and dc bcm_check_stuck_buffers are observed, which could eventually lead to port 0 (*/*/0) flapping. [PR1162947](#)
- On QFX5100 switches with a CoS classifier configured on an AE interface, if you add or delete a subinterface, traffic loss of approximately 10 packets might occur while you are committing the changes. [PR1162963](#)
- On QFX5100 switch, syslog may contain repeated messages like so: fpc12 Unit: 0 port 47 start error detected. [PR1164096](#)
- Repeated "nh_unilist_update_weight:" error messages when CCC L2VPN is configured. These are harmless [PR1167846](#)

- On QFX5100/EX4600 platform, when a Private VLAN is trunked and that interface is disconnected (cable removal or system reboot) there is a section of code that causes an issue in how the switch handles VLANs. This issue might cause all VLANs to be dropped. This issue is also present when a PVLAN is added to a working trunk. [PR1169601](#)
- On EX4600 and QFX5100 series switches configured "tag-protocol-id" and "flexible-vlan-tagging". If the switch receives traffic, whose outer tag protocol ID is not 0x8100 (e.g. 0x88a8, 0x9100, or 0x9200, which are usually used by double tag traffic) on a trunk interface, the switch always uses 0x8100 as outer protocol tag ID when it sends out the traffic. [PR1170939](#)
- On EX4600/QFX Series switches, after add and delete the fifth logical interface, the first 4 AE subinterfaces might be down and lose connectivity. [PR1171488](#)
- On QFX5100 and EX4600 switches, every time a MAC address is learned, some messages might be output to syslog and be repeated frequently. The logged messages have no impact on service traffic. [PR1171523](#)
- On QFX and EX4600 platform, in the scenario that MSTP/RSTP/VSTP is configured to prevent layer-2 network loop, there might be a chance that xSTP convergence may fail on the interface that configured with flexible-vlan-tagging and encapsulation extended-vlan-bridge. [PR1179167](#)
- PFE stats counters should be always incremental. In some cases, a user can observe lower stats values than the previously values given, and this will trigger following logs errors: "pfed: downward spike received from pfe for ipackets_reply" or "pfed: downward spike received from pfe for opackets_reply" The fix for this issue will give this logs "info" severity. [PR1183184](#)
- Ipv6 Linklocal filter entry to match on Unicast LinkLocal Address will not be Hit on the channelized interfaces, other than Lane 0 port [PR1193313](#)
- MAC move limit configuration is not supported in 16.1 for QFX-5100 [PR1194699](#)

Multiprotocol Label Switching (MPLS)

- In MPLS scenario, on EX4600/QFX Series switches with AE interface configured, after change the IGP metric and disable the AE interface, the fxpc crash might be observed when child nexthop of a UNILIST is pointing to NULL. [PR1168150](#)
- Ping over LSP shows different behavior in regards to HLIM. [PR1179518](#)
- For 2 label PUSH cases, both labels are consuming entries in the same label table. This might result in instabilities of MPLS tunnels and packets drop when add/delete routes. Correct behavior should be that tunnel label goes in one table and VRF label should go in another table. [PR1185550](#)
- On QFX5100 switches or a QFX3500 or QFX3600 Virtual Chassis, IP packet frames of 1500 bytes might drop when family mpls is configured on a logical interface. [PR1199919](#)
- On EX4600/QFX3500/QFX3600/QFX5100 Series switches, traffic received from the MPLS core at the PHP node might not get forwarded to the egress ECMP IPv4 next hop. [PR1212519](#)

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 sfcsmnpd 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 sfcsmnpd process from the Director. [PR1165565](#)
- On QFabric system, SNMP does not work due to dead-lock on sfcsmnpd threads in a rare condition. [PR1192627](#)

Platform and Infrastructure

- In customer setup ingress node is RSNG and egress node is NNG, since NNG has all vlan information this fix will work. Say if ingress node is NNG and egress node is RSNG, this fix may not work because egress node RSNG may not have the incoming vlan information. This fix will work only if incoming vlan information is available in egress node. As per design for server node groups vlan information will be selectively downloaded as per configuration so this fix may not work in case were incoming vlan is not available in egress node. [PR1103274](#)
- On QFX5100 with VXLAN feature, all encap traffic would be dropped due to remote VTEP is pointing to failure next hop in PFE. Then you may see the syslog message below. BRCM-VIRTUAL,brcm_virtual_venh_install(),1479:VENH installation failed .. nh-id(14987) [PR1136540](#)
- On QFX Series switches, if a VLAN tag (e.g. VLAN 10) is assigned to a VLAN hardware token 4, when VLAN 10 carries the routing protocol traffic (for example OSPF/ISIS/BGP), that traffic will be put into a wrong queue (Queue 22) which has rate limit (100pps), then it might cause protocol flap. [PR1146722](#)
- 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. To work around this issue, you can explicitly specify the interface name for which to clear ARP entries, as follows: clear arp interface <interface name>. ARP entries are properly cleared when using this form of the command. [PR1159447](#)
- If DHCP packets with MPLS tags are sent to the CPU on a QFX5100 node acting as a PHP node, the logical interfaces index on the packet notification might not be set correctly, and the DHCP packets might be dropped. [PR1164675](#)
- When the system log severity level 7 debug level is set, this debug message is printed on a per-packet basis---/kernel: setsocketopts: setting SO_RTBL_INDEX to 1. [PR1187508](#)

QFabric Systems

- In QFabric scenario, if one power supply unit is removed and inserted back to DG (director group), the alarm message is not very clear. [PR1165890](#)
- 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](#)

Routing Protocols

- On QFX Series switches, when a neighbor device sends a flood of Link Layer Discovery Protocol (LLDP) traffic bigger than 1,000 pps to the QFX, Link Aggregation Control Protocol (LACP) flaps might be seen on unrelated interfaces. [PR1058565](#)
- On QFX5100 and EX4600 switches, if you use the Network Configuration Protocol (NETCONF) to add or delete firewall filters on an integrated bridging and routing (IRB) interface, the Packet Forwarding Engine Manager (fxpc) might generate a core file. [PR1155692](#)
- FXPC crash may happen during an ECMP route delete from LPM table. This might have happened due to large scale route change operation. SDK vendor provided a fix as a resolution. [PR1158517](#)
- Loopback filter not working due to higher priority system dynamic filter. Implicit DHCPv6/v4_l3_tag filter installed is conflicting with the configured loopback filter [PR1159024](#)
- On QFX5100 and EX4600 switches, when a limit traffic filter is configured with TTL=1 packets accepted on the loopback interface, the host-bound unicast packets with TTL=1 (for example, OSPF packets) might be dropped. [PR1161936](#)
- On a QFX3500 switch, if you configure one interface with PIM and the interface sends hello packets, and then you change its PIM hello-interval from non-zero to 0, the interface sends hello packets continuously. [PR1166236](#)
- On EX4600/QFX Series switches with logical interface, if family mpls is configured first and then other families like inet/inet6 are configured on the logical interface, then the other families configuration might not be programmed correctly in PFE, which can result in traffic not getting forwarded on the newly configured families. [PR1166595](#)
- On QFX5100 switches, if you apply a firewall filter on the loopback interface with the match condition for packets with TTL 0/1 and with "policer" set as the action, the term does not catch the packets. [PR1166936](#)
- ALL traffic destined for leaked route are forwarded to CPU. The traffic expected to be treated as transit. Ping between the default routing instance and routing-instance which leaking routes via the rib-group takes around 40-90 ms. juniper@abc:~\$ ping 10.1.1.1 PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data. 64 bytes from 10.1.1.1: icmp_seq=1 ttl=62 time=17.9 ms 64 bytes from 10.1.1.1: icmp_seq=2 ttl=62 time=42.3 ms 64 bytes from 10.1.1.1: icmp_seq=3 ttl=62 time=15.3 ms 64 bytes from 10.1.1.1: icmp_seq=4 ttl=62 time=18.0 ms 64 bytes from 10.1.1.1: icmp_seq=5 ttl=62 time=21.1 ms [PR1167156](#)
- On QFX3500 or QFX5100 switches, when parity errors occur on interfaces, they might affect the memory management unit (MMU) memories. MMU counters can be

corrupted, the interface buffers might be stuck, and there might be interface flaps and traffic loss on the affected ports. As a workaround (restoration only), reboot the system.

[PR1169700](#)

- On EX4600/QFX5100 series switches with virtual chassis scenario, if configure primary and secondary RTG (redundant-trunk-group) links on fpc0/master and fpc1/backup respectively and then perform NSSU. During a NSSU upgrade, when the original master/fpc0 goes down, secondary RTG links on original backup/fpc1 become active and they forward traffic for about roughly 17 seconds and thereafter stop forwarding any traffic across. Traffic starts flowing again once the primary RTG links take over the control. [PR1170258](#)
- Currently, on QFX5100/EX4600 with filter based GRE (for a filter with decapsulate gre action), first, only one prefix is supported per filter. If the prefix is a destination address, it should be a /32 prefix. Second, the filter only supports one term. And last, filter change may don't take effect (the workaround is to unbind and bind the filter). With the fix each term having a decapsulate gre action in a filter can have multiple destination prefixes (max of 100 per term) with prefix length 32 and one source prefix (any prefix length/wild card). Filter change will take effect as expected. [PR1171053](#)
- On EX4600/QFX5100 switches, in rare cases, route insert failure in `_soc_alpm_128_write_pivot` function will lead to a loop in the code resulting in a watchdog timeout. This will result in the FPC crash and restart with a core dump. [PR1173980](#)
- On EX4600 and QFX5100 series switches, there are several profiles that allocate memory differently for MAC addresses and host addresses. These profiles can be configured as "l2-profile-one, l2-profile-two, l2-profile-three, l3-profile, lpm-profile". If multicast and unicast host entries reach the maximum number of the L3 host table in related profile, then multicast traffic will be dropped. [PR1177430](#)
- The static route cannot be configured with 'resolve' and 'retain' flags together and we have a check to ensure this. But if one of the flags is configured via 'set routing-options static defaults' and another flag is configured via static route then commit is accepted and this is causing rpd crash. [PR1178418](#)
- In some scenarios after ECMP route flapping on QFX switches traffic is blackholed. RIB programming is fine: `root@qfx> show route 172.16.2.1 inet.0: 843 destinations, 2707 routes (843 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 172.16.2.0/25 *[BGP/170] 6d 12:24:13, MED 106, localpref 100, from 10.1.32.2 AS path: ?, validation-state: unverified to 172.16.5.101 via ae2.0 > to 172.16.6.101 via ae4.0 [BGP/170] 6d 12:20:20, MED 106, localpref 100, from 10.1.32.4 AS path: ?, validation-state: unverified > to 172.16.6.101 via ae4.0 [BGP/170] 6d 12:20:24, MED 111, localpref 100, from 10.1.32.1 AS path: ?, validation-state: unverified > to 172.16.5.21 via ae1.0 [BGP/170] 6d 12:24:09, MED 111, localpref 100, from 10.1.32.3 AS path: ?, validation-state: unverified > to 172.16.6.21 via ae3.0 FIB programming is fine: root@qfx> show route forwarding-table destination 172.18.2.0/25 Routing table: default.inet Internet: Destination Type RtRef Next hop Type Index NhRef Netif 172.16.2.0/25 user 0 ulst 131157 4 indr 131098 19 172.16.5.101 ucst 2349 9 ae2.0 indr 131154 16 172.16.6.101 ucst 2397 9 ae4.0 But kernel nexthops point to discard nexthop in broadcom sdk. Hence traffic to destination is blackholed. PR1179610`

- On a QFX/EX4600 Virtual Chassis, the "local-bias" knob causes traffic loss on a direct link when an interface is changed from Layer 3 to AE and then back to Layer 3 due to the AE programming still resides in the Packet Forwarding Engine (PFE) and hardware. [PR1179960](#)
- The QFX5100 will exception (forward a copy of) transit IPv6 Neighbor Discovery traffic to the RE, allowing for a collateral partial local DoS attack. Refer to JSA10749 for more information. [PR1183115](#)
- In MC-LAG scenario with igmp-snooping configuration, when one link of MC-LAG is disabled, the IGMP report packet cannot be transferred correctly. It might cause impact for multiple traffic for IGMP report failing. [PR1183532](#)
- On EX4600/QFX Series switches, after native-vlan-id is configured and rolled back on a vlan-tagged sub-interface, ARP might not be resolved and traffic forwarding can be affected. [PR1184985](#)
- In a QFX5100 Virtual Chassis, if the master is halted or rebooted with some limited MAC persistence timer set, then in a specific sequence the IRB/inet MAC does not get programmed correctly in the BCM. [PR1188092](#)
- On QFX5100 switches with MPLS and LDP enabled, for packets with incoming labels that must perform a PHP (penultimate hop popping) operation on the QFX5100 switch, occasionally the packets are not processed and are dropped. [PR1190437](#)
- On QFX3500/QFX3600/QFX5100/EX4600 series switches, if a routing loop is created, the TTL of the packet does not reduce to 0 and eventually the packet is not dropped. [PR1196354](#)
- On QFX3500/QFX3600/QFX5100/EX4600 series switches, if disable IRB interface then reboot the switch. After the switch rebooted then enable IRB interface, after that IRB interface might not be reachable. [PR1196380](#)
- On QFX5100 switches, the DSCP action modifier of a family inet firewall filter does not properly modify or mark the DSCP bits on packets matching the firewall filter. [PR1205072](#)
- On a QFX5100 switch with VRF enabled, route leaking from the default routing table (inet.0 or inet6.0) to VRF might not work as expected. [PR1210620](#)
- On QFX5100 switches, port-range-optimize (both source and destination) might fail to be programmed into the hardware for an inet output filter. [PR1211576](#)
- On QFX5100 and EX4600 switches, in rare cases, the FXPC process might crash and restart with a core file generated upon LPM route install failure. After the switch restarts, services are restored. [PR1212685](#)
- On QFX5100 Series switches with Protocol Independent Multicast (PIM) configured, the system can go into swap due to low memory condition, and fxpc core dump might happen due to this out of memory condition. [PR1217343](#)
- A filter attached to the lo0 interface with terms containing either destination-port-range-optimize or source-port-range-optimize statements will unexpectedly discard traffic. [PR1228335](#)

Storage and Fibre Channel

- The FLOGI has come from the ENODE over a VF port on GW. The Proxy tries to find a N port for sending the FLOGI to FC switch. Due to some churn in the system the N port through which the WWNN is reachable has gone down or unavailable. This leads to the crash most likely in some race condition. [PR1152334](#)

User Interface and Configuration

- When a VRRP group is created on Qfabric, in race conditions, vrrpd will not find the VRRP group in its database and will not transition to master state if the VRRP route notification reaches vrrpd before the VRRP config push to Server Node Group(SNG). [PR1197443](#)

Virtual Chassis

- SDK can raise false alarms for parity error messages like "soc_mem_array_sbusdma_read" & "soc_ser_correction: mem write" on QFX5100. [PR1161821](#)

Resolved Issues: Release 14.1X53-D35

General Routing

- mgd-bsd and java high CPU issue fixed on release 14.1X53-D17.4 and 14.1X53-D35 [PR941833](#)
- Setting link speed to 100 Mbps does not work in the following situations: - When network interfaces are used on an EX4600 switch - When an EX4600-EM-8F expansion module is installed in a QFX5100-24Q switch or an EX4600 switch [PR1032257](#)
- On QFX series switches, the wrong source IP address is being used when the switch initiates traffic when em0 is configured with a 192.168.1.XXX/16 subnet. [PR1071517](#)
- On QFX3500 switches, if you remove 1- Gigabit Ethernet SFP transceivers from ports 0-5/42-47 and then insert 10-Gigabit Ethernet SFP+ transceivers in the same port, the 10GE SFP+ transceivers might not be detected. [PR1085634](#)
- On EX9200/EX4300/EX4600/QFX Series switches, if removing/inserting one QSFP, it might cause pfe process to crash. [PR1098385](#)
- On a QFX3500 switch with nonstop active routing (NSR) enabled, deleting a routing-instance or logical-system configuration might cause a soft assert of the rpd process. If NSR is not enabled, after you delete a routing-instance or logical-system configuration, executing "restart routing" might trigger this issue, too. This issue has no functional impact. [PR1102767](#)
- On a QFX5100 VCF in an auto-provisioned mode, when adding a new leaf device to the VCF, should zeroize device and reboot by "request system zeroize" if new leaf device has been configured any command. But the issue (interface still up) might be observed at the time of the reboot until the PFE re-initialized the interfaces. [PR1106194](#)
- On an EX Series or QFX Series switch configured as a DHCP client, the length of the DHCP Vendor ID is always 60 in DHCP discover packets when the vendor class ID is

configured, although the actual vendor-id name is less than 60. As per RFC 2132, the code for this option ("Vendor class identifier") is 60, and its minimum length is 1.

[PR1123111](#)

- Fix for this PR is in host OS. Host OS does not get upgraded if upgrade is done using ISSU. If customers upgrade to D35 using ISSU and they need fix for this issue. they need to copy 2 scripts to host OS. 2 Scripts are attached to this PR as attachments. Customer will have to download these scripts and copy then to host OS under the path /vmm/bin/qfx_setup_disk and /vmm/bin/qfx_mount_disk. Alternate method to get the fix is, upgrade to the image which has the fix using "request system software add <path to image> force-host" command. 'force-host' is mandatory to make sure host os upgrade will be done. [PR1127517](#)
- On QFX/EX4600 Series switches, in rare condition, the trunk interface may not get create due to data structure becomes out of sync between Packet Forwarding Engine (PFE) and control plane. [PR1128316](#)
- On QFX5100 Series switches with Open vSwitch Database (OVSDB) management protocol configuration that act as virtual tunnel endpoints (VTEPs), traffic being forwarded from ingress AE interface to egress tagged port may not be attach 802.1Q VLAN tag if both ports are located the same device. [PR1128507](#)
- On a QFX5100-48T switch, the 10G port is used to interconnect between QFX5100-48T and Intel X540 10G Ethernet NIC (Network Interface Card), the link speed has a chance appear to be listed as 1-Gigabit Ethernet if the 10 port on QFX5100-48T experiences a local fault. [PR1131392](#)
- On Juniper Networks devices that support OpenFlow, the openflowd process might crash after you issue the show openflow statistics tables command. [PR1131697](#)
- On QFX5100 switches with minimum-interval for a the Bidirectional Forwarding Detection (BFD) session configured to less than 1 second, the pre-ISSU check might be successful and continue to implement the ISSU, which causes the BFD session to flap. The expected behavior is that the pre-ISSU check for the BFD session should fail and ISSU would be aborted. [PR1132797](#)
- On QFX5100 Series switches, the Virtual Router Redundancy Protocol (VRRP) is configured on IRB interface associated with the private VLAN (PVLAN), traffic from the hosts on secondary VLANs (isolated VLAN or community VLAN) destined to VRRP MAC address might be dropped. [PR1135756](#)
- On QFX5100 Series switches, after disable the interface and hot swap 1G copper transceiver, link flap or link-up might be seen on SFP-T port though interface is configured as disabled explicitly. [PR1137204](#)
- On a QFabric system, the interface might not convert to 40-Gigabit Ethernet (xle) port after configuring a block of ports to operate as 40-Gigabit Ethernet (xle) ports in a QFX3600 node device. [PR1138444](#)
- In EVPN/VXLAN dual homed scenario with QFX5100 as leaf, after failure of a leaf which has switch or LAG interface with hold-time enabled, then some VLANs might not reconverge and traffic forwarding does not work as expected. [PR1140403](#)
- On QFX5100 switches, the openflowd process might generate a core file. [PR1142563](#)

- On QFX Series switches with Data Center Bridging and Exchange Capability (DCBX) enabled, when you are configuring a guaranteed minimum rate of transmission for a CoS traffic control profile, the Layer 2 Control Protocol daemon (l2cpd) might crash during the initial LACP setup. [PR1143216](#)
- On EX4600/QFX5100 switches, after performing command "show version detail", an error message **Error: abnormal communication termination with app-engine-management-service daemon** might be seen at the end of the output. [PR1144234](#)
- On QFX5100 and EX4600 switches, the Gigabit Ethernet (ge) interface might stop forwarding traffic when you hot-swap a transceiver from SFP-SX to SFP-T. [PR1144485](#)
- On EX4300 and QFX Series switches with PVLAN configured, if secondary VLANs (isolated VLANs or community VLANs) are configured with vlan-name, after binding or unbinding the isolated or community VLANs in the primary VLAN, packets loss might occur between existing VLANs. [PR1144667](#)
- On QFX5100 switches, if you delete an auto-negotiate configuration on a 10-gigabit interface (xe), the interface goes down as expected because the auto-negotiate setting is not matching with that on the peer interface. However, the interface might come up after the reboot even though auto-negotiate is still disabled. For release versions D37 and above, this situation will not be observed anymore. Also, to disable AN on Nirvana xe port, the speed of the interface must be set to 100M explicitly. [PR1144718](#)
- After the number of DHCP server IPs in the dhcp-relay configuration is modified (increased or decreased), messages log file will be filled with following error messages and eventually cause DHCP process (jdhcpd) to crash. jdhcpd:
%USER-3-DH_SVC_SENDMSG_FAILURE: sendmsg() from 10.161.102.1 to port 67 at 0.0.0.0 via interface 615 and routing instance VR08_v881_900_office_system failed: Can't assign requested address [PR1147831](#)
- On a QFX5100 Virtual Chassis, if you configure an aggregated Ethernet interface as an OVSDB interface with multiple subinterfaces that are configured under different VXLAN domains, removal of the last but one AE subinterface might reset VXLAN settings on the physical port that are part of the AE interface, resulting in packet drops. [PR1150467](#)
- On EX4600/QFX Series switches, in corner cases, the PFE manager (fxpc) might crash when an SFP-T transceiver is removed/inserted too quickly or the interface is deleted. [PR1152097](#)
- On EX4300/EX4600/QFX5100 Series switches, when an STP configuration is initially applied to an interface and the interface is down at that moment, executing "show/clear spanning-tree statistic interface" might cause the Layer 2 control protocol process (l2cpd) to crash. [PR1152396](#)
- On QFX5100 switches, a child member might drop the incoming Link Aggregation Control Protocol (LACP) frames when this child member is moved from an access-mode VXLAN LAG interface to a trunk-mode VXLAN LAG interface. [PR1153042](#)
- On a Qfabric system with the QFX3500/QFX3600 as a node device or the QFX3600-IC as an interconnect device, executing the "show snmp mib walk jnxMibs" command causes the chassis daemon (chassisd) process to crash. [PR1157857](#)

Class of Service (CoS)

- On a QFX5100-VC platform, when gr interface is configured, and then if it is deleted or deactivated, unicast traffic might not be forwarded well on the underlying L3 interface. [PR1154812](#)

EVPN

- On QFX5100 Series switches using EVPN with VXLAN, the Ethernet Segment Identifier (ESI) value of the most significant octet (type byte) must be 00 when manually configuring an ESI even though the switch accepts other configuration values. [PR1085837](#)
- storm-control: SC profile still shows up on PFE after it is removed from config [PR1099377](#)
- On QFX5100 Series switches with virtual extensible local area network (VXLAN) configured, the SIP/DIP (source IP/destination IP) to be 0.0.0.0 in VXLAN traffic after the device reboot due to sometimes VTEP Gateway daemon (vgd) might push remote MACs to the layer 2 learning daemon (l2ald) before the source VETP logical interface (IFL) is created. [PR1109838](#)
- On a QFX Series switch or Virtual Chassis which 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 (RE) reboots as part of the NSSU. [PR1126855](#)
- 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. [PR1126915](#)
- On QFX5100-96S standalone/VC/VCF mode, the packet forwarding engine manager daemon (fxpc) process might crash continuously when the SFP/SFP+ transceivers is removed and then inserted in the specific 10-gigabit ethernet (xe) interface (xe-*/0/95) which has extended-vlan-bridge/flexible-vlan-tagging configuration. [PR1159156](#)

High Availability (HA) and Resiliency

- On EX4300/EX4600 Series switches and 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. [PR1129004](#)

Interfaces and Chassis

- On a QFabric system with the IGMP snooping is enabled, every time the IGMP join/leave was allocated by sockaddr memory, but the memory is not freed accordingly. This might cause the fabric control protocol (rpdf) memory leak of 32-bytes and 48-bytes on the Network Node Group. When rpdf reaches its max memory limit, rpdf process crash will be seen. [PR1121875](#)
- On a Qfabric system without any config related to dot1x, a memory leak might occur in the dot1x daemon (dot1xd), this issue cause is dot1x is running on Redundant Server Node Group (RSNG) node despite it is not supported on Qfabric. [PR1131121](#)

- On QFX5100 switches, if an mc-ae member link is deleted and then re-added on an MC-LAG node, there could be a traffic loss of about 2 seconds. [PR1146206](#)

Layer 2 Features

- On QFX5100 switches, if you configure a PVLAN inter-switch-link on an existing working trunk port, normal VLAN traffic might break. [PR1118728](#)
- On EX4300, EX4600, and QFX Series switches, traffic received on the backup redundant trunk group (RTG) link might get forwarded to other interfaces following an RTG link failover. [PR1119654](#)
- If you reboot one FPC in a two-member Virtual Chassis, the traffic might not exit from the FPC after the FPC comes back online and rejoins the Virtual Chassis, and local registers might be incorrectly cleared, if the port number is the same on both the master and backup. [PR1124162](#)
- On QFX3500/5100 Series switches, while committing et interface inet plus mpls config with no-redirects knob having MTU setting, the protocol ARP might not be configured for the IFL in PFE. [PR1138310](#)
- On QFX VC/VCF, when firewall filter with "vlan" action is applied to the ingress interface of one member, traffic may not pass the inter-member to egress interface of another member. [PR1138714](#)
- On QFX5100 and EX4600 switches, after you delete one logical interface from one VLAN that is configured with multiple logical interfaces, the MAC address for other logical interfaces might not be learned again. [PR1149396](#)
- On an EX4300 switch in a VCF, if a Layer 3 AE interface is looped back with a Layer 2 port in the same VLAN, then traffic with the same destination MAC to the AE interface is dropped (for example, the ping address of the AE interface). [PR1157283](#)

Multiprotocol Label Switching (MPLS)

- On QFX/EX4600 Series switches, while receiving an IPv6 packet whose destination IPv6 address does not have an entry in the IPv6 neighbor table, they would fail to send out an IPv6 neighbor discovery packet and traffic to these IPv6 hosts might be dropped. [PR1134599](#)
- On QFX5100 switches, a ping from the CE to the PE (LHR) lo0 interface does not go through with explicit-null (RSVP). [PR1145437](#)
- On QFX Series switches, when action "load-balance" and match condition "rib mpls.0" are configured on two different terms of a policy, the commit operation might fail and produce an error message. [PR1147463](#)

Network Management and Monitoring

- On Junos Platform with private and internal interfaces used, whenever there is a software upgrade from any prior to 12.3 to any newer version, where kernel is holding older version value and mib2d comes with newer index value, mib2d might core and crash. There is no service impact. [PR1109009](#)

- On a QFabric system QFX3000-G/QFX3000-M, when big/large files in event capture directory /var/opt/dgscan/nodes/node/, the dgsnmp daemon might run at high CPU utilization. And at such times snmp polling does not work, and the sfcsmnpd and the dgsnmpd don't log any messages. As a workaround move or delete large files (>50mb) from /var/opt/dgscan/nodes/node/ to /var/tmp/. [PR1139852](#)

Platform and Infrastructure

- "show chassis forwarding-options" CLI output for 'l2-profile-three -> num-65-127-prefix' is incorrect (NONE) even if it is configured correctly. Configuration is applied as 'set chassis forwarding-options l2-profile-three num-65-127-prefix 3' but CLI command '> show chassis forwarding-options' still shows the output as 'NONE' for 'num-65-127-prefix' [PR1069535](#)
- On QFX5100 switches, adding or removing virtual routing and forwarding (VRF) instances that have many logical interfaces in the link aggregation group (LAG) might cause Link Aggregate Control Protocol (LACP) flapping. [PR1087615](#)
- 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 http://www.juniper.net/techpubs/en_US/junos14.1/topics/reference/configuration-statement/fabric-tree-root-virtual-chassis.html. [PR1093988](#)
- When we issue the PFE command "show brcmfm ifd all" there might be an FXPC core-dump on QFX5100 running 14.1X53-D12 [PR1119567](#)
- Multiple PFEMAN disconnects and reconnects between the master and backup within a short period of time can cause the backup to generate core files. [PR1123379](#)
- On QFX Series and EX4600 switches, if an AE interface is used as an ECMP next hop (load balance), traffic is not hashed evenly to all member interfaces correctly. [PR1141571](#)
- On MX Series routers, and EX Series and QFX Series switches, SSH authentication might fail due to improper file ownership. [PR1142992](#)

Routing Protocols

- On EX4600 and QFX5100 switches with Q-in-Q, if the native VLAN is configured on a Q-in-Q interface connected to a customer device (CE), the packets going out with the native VLAN ID (Customer-Vlan) are still tagged. [PR1105247](#)
- On QFX5100/EX4600 Series switches, when eRACL (Egress routing ACL filter) is applied to more than 64 interfaces, a memory corruption issue might occur, resulting in the Packet Forwarding Engine manager (fxpc) process to crash. [PR1123374](#)
- On QFX5100 series switches configuring gre interface over irb interface, then gre interface can become up but not able to ping IP address of gre interface in remote end. [PR1124149](#)
- On QFX5100 switches, you might see the "soc_mem_read: invalid index -1 for memory EGR_L3_INTF" log message. You can ignore the message; there is no functional impact on the switch. [PR1126035](#)

- This PR changed the behavior when using flexible vlan tagging and native-vlan-id to egress packets untagged for packets that are part of the native vlan. Previous these packets would egress tagged. [PR1130192](#)
- On a Qfabric system, the DHCPv6 packets are getting dropped in Network Node Group (NNG) due to internal filters. [PR1132341](#)
- Configuring analyzers might lead to sub-optimal use of allocated TCAM space. When this happens, the following logs might be displayed: [Sat Nov 21 08:45:18 2015 LOG: Err] PFE: Unknown next-hop (nh_id 2532) for sampling [Sat Nov 21 08:45:19 2015 LOG: Err] PFE: Unknown next-hop (nh_id 2532) for sampling [PR1136837](#)
- On QFX/EX4600 Series switches with dual-stacked interface, if the interface is configured to be part of a non default routing-instance and input IPv4 Filter Based Forwarding (FBF) with no matching condition is applied, the IPv6 packets received might be dropped. [PR1145667](#)
- On QFX5100 and EX4600 switches, if you use the Network Configuration Protocol (NETCONF) to add or delete firewall filters on an integrated bridging and routing (IRB) interface, the Packet Forwarding Engine Manager (fxpc) might generate a core file. [PR1155692](#)
- Loopback filter not working due to higher priority system dynamic filter. Implicit DHCPv6/v4_l3_tag filter installed is conflicting with the configured loopback filter [PR1159024](#)

Storage and Fibre Channel

- On EX4500 and QFX Series switches or a QFabric Fabric system with DCBX enabled, when the DCBX neighbor is up and then receives a normal LLDP packet (without DCBX TLVs) on the same port as the DCBX packets, the device might ignore the DCBX packets, causing session timeouts and a reset of the priority-based flow control (PFC) settings. [PR1095265](#)

Resolved Issues: Release 14.1X53-D30

General Routing

- On EX Series switches with integrated routing and bridging (IRB) interface configured, if the JSRV interface is created prior to the IRB interface after restarting the device or chassis daemon (chassisd), it might cause all IRB interfaces to be disappeared. [PR965097](#)
- On a Virtual Chassis Fabric (VCF), a small amount of Layer 3 unicast packet loss (for example, 0.2 - 0.3 sec) might be seen when a leaf node that is not in the traffic path is rebooted. [PR976080](#)
- On a QFX5100 platform, when to upgrade junos by topology-independent in-service software upgrade(TISSU), during "FPC Warm Booting" period of TISSU, a few packets drop might be seen on an SFP-T interface, this issue not seen with SFP-SX interface. [PR1027336](#)
- On EX4600 and QFX5100 switches, the Link Aggregation Control Protocol (LACP) in either slow mode or fast mode might go down and then come back up, causing a

timeout and a service outage during an In-Service Software Upgrade (ISSU) or a Nonstop Software Upgrade (NSSU). In addition, after the master Routing Engine is rebooted, the switches might experience intermittent traffic loss on non-LAG interfaces, and redundant trunk group (RTG) convergence times might be long. [PR1031338](#)

- On any EX/QFX Series switches with support to Media Access Control Security (MACsec) it might generate the error message following as below: "dot1xd[1634]: knl_ifcheck_chunk: Starting interface state recovery". This is a cosmetic issue, it has no functional impact. [PR1045144](#)
- Inconsistent/Incorrect AE IFD stats because of incorrect handling the child IFD stat flags. As AE stats is an aggregate of the child IFD stats, these requests are processed differently as compared to stand alone interfaces thereby introducing inconsistencies in the next poll cycle. [PR1048276](#)
- EDITED MP 8/28 On a QFX5100-48T switch, interfaces numbered 0 to 23 are sometimes not turned down during device reboot. This issue might be seen when a peer device is using 1G link speed. [PR1059876](#)
- In certain environments - with certain narrow operating temperatures or changing operating temperatures, there is a statistical probability of the QFX5100 850W AC power supply shutting itself down due to a bug in the power supply firmware logic comparing measured fan speed versus target speed at temperature. [PR1062224](#)
- On a QFabric system, if configuring management address for LLDP on the Network Node group (NW-NG) interfaces, the Link Layer Discovery Protocol daemon (lldpd) might be continuous crashed. [PR1062445](#)
- On QFX5100 switches, enabling error-correcting code (ECC) ELV. [PR1064567](#)
- On QFX5100 switches that are configured with the "include-option-82 nak" option so that Dynamic Host Configuration Protocol (DHCP) servers include option 82 information in NAK messages, two copies of option-82 might be appended to DHCP ACK packets. [PR1064969](#)
- On EX4600 switches and QFX Series switches, when a pair of devices configured for multichassis link aggregation (MC-LAG) are both using active mode when rebooting, traffic can drop for a while on one of the switches. [PR1069644](#)
- The SNMP walk for the dot1dBasePortIfIndex object might return a value of 0, which is not a valid SNMP ifindex for an interface. [PR1070532](#)
- On a QFabric system, if configuring "system accounting events", the device creates audit process(auditd) child processes for every accounting events, but multiple child processes may not terminate, which result in high CPU utilization of the auditd. [PR1070701](#)
- On EX Series and QFX Series switches, issuing the "show interfaces extensive" command or polling SNMP OID ifOutDiscards provides a drop count of zero. [PR1071379](#)
- On a QFX5100-24Q-AA switch, in few of the cases, after the switch reboot, the guest virtual machine (VM) may not get the field-programmable gate array (FPGA) devices for use. Consequently any application or utility trying to use the FPGA device will fail. [PR1073076](#)

- On a QFX5100-24Q-AA switch, if the PFE manager (FXPC) restarts due to any reason (crash or planned restart), then the guest virtual machine (VM) will lose its PCIe devices. Consequently, any utility or application using those devices will lose the access to them. This may result in failures of the utilities and/or applications. [PR1073084](#)
- On QFX5100 Series switches, 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](#)
- After powercycling QFX5100 in QFabric chassis status LEDs are going off [PR1074310](#)
- On QFX5100 Series switches, the SFP management interfaces might fail to come online. [PR1075001](#)
- QFX Series: Insufficient entropy on QFX systems (CVE-2016-1273); Refer to <https://kb.juniper.net/JSA10746> for more information. [PR1075067](#)
- On QFX5100 switches, if more than 1K virtual extensible LAN network identifiers (VNIs) are created by Open vSwitch Database (OVSDB), the VTEP gateway daemon (vgd) might generate a core file. [PR1075189](#)
- On a QFX5100 Virtual Chassis, the log messages as "fpc0 vccpd irt socket connect failed (no route to host)" are seen continuously, it is harmless. [PR1075437](#)
- A QFX5100 switch with a BIOS version older than V18.7 does not have error-correcting code (ECC) memory enabled by default. This might cause issues because it limits correction of memory corruption. [PR1075915](#)
- On QFX5100/EX9200 Series switches, when configuring the VLAN name and Logical Switch (LS) for OVSDB, if the VLAN name or LS using the UUID format, the configuration would not to commit. [PR1075919](#)
- On QFX5100 Series switches, if you configure both Q-in-Q tunneling and IGMP snooping, IGMP reports do not egress. As a result, multicast traffic is flooded instead of being sent to requested receivers. [PR1076324](#)
- On a QFabric system, if configuring Internet Group Management Protocol (IGMP) snooping, the Virtual Router Redundancy Protocol (VRRP) multicast packets might be dropped. [PR1077085](#)
- On a QFabric system, the sfid-bcm memory might be leaked with a core file generated during multicast data packet handling. [PR1077678](#)
- On QFX/OCX1100 Series switches, when the Encapsulated Remote Switched Port Analyzer (ERSPAN) output IP address is reachable via more than one route, the analyzer goes down. [PR1077700](#)
- On EX4600, QFX5100, QFX3500, and QFX3600 Series switches, when the device acts a transit router between the DHCP server and the DHCP relay agent, and DHCP server/relay is not configured, the device might not forward the DHCP ack packets to the destination address. Instead, packets are sent to Routing Engine (RE) if the packets' destination port was 68. [PR1079826](#)

- On EX9200 and QFX5100 switches, if you configure DHCP relay with the DHCP server and the DHCP client in separate routing instances, unicast DHCP reply packets (for example, a DHCP ACK in response to a DHCP RENEW) might be dropped. [PR1079980](#)
- On QFX5100-48T Series switches, if rebooting the device, the 10gbase-T interfaces do not go down until after the software has reloaded. It might cause the peer device service impact due to failover invalid. [PR1081105](#)
- On QFX5100 switches, the maximum number of LAGs is now 1000. [PR1082043](#)
- On QFX5100 Series switches, if Class of Service (CoS) configuration is changed on a physical interface while running traffic, the host inbound packets might be affected and cannot be processed, and the PFE manager (fxpc) process crash with a core file generated, which result in Aggregated Ethernet (AE) interface goes down due to LACP time out. [PR1082224](#)
- On EX4600 switches and QFX Series switches, you must use the -C and -S option with a DHCP request - if you do not, the client might not receive the DHCP ack packets. [PR1082473](#)
- On QFX5100 Series switches, if installing license for VCF feature and with Junos OS release 14.1X53-D25, the device might raise an error information "license not valid for this product" and fail to install. [PR1084235](#)
- On a QFabric system, if using CLI command "request system reboot all" or switchover Director Group (DG) mastership, it might cause the Packet Forwarding Engine manager(pafxpc) crash and generates the core file. [PR1087420](#)
- On a QFabric system with Junos OS release 14.1X53-D15 only, the device could not forward the DHCP unicast packets in VLAN. [PR1088393](#)
- On QFX Series switches, when a large number of small form factor pluggables (SFPs) with Digital Optical Monitor (DOM) support are inserted, the CPU utilization of the of the PFE manager daemon (fxpc) might increase (maximum value 50%) due to a large number of iterations of SFP diagnostics polling. The thread that causes the high CPU to have a low priority might not cause any problems to the functionality. As a workaround, if DOM statistics are not important, disable the diagnostics by issuing the <vty> test sfp diagnostics disable command, or increase the diagnostic-interval (default is 3 seconds) to bring down the CPU utilization by issuing the <vty> test sfp periodic diagnostic-interval 10 command. [PR1091512](#)
- On a Virtual Chassis Fabric, if configuring VCF in autoprovisioned or preprovisioned mode, and enable LLDP on em0 interface to connect other VCF members. When the VCP interface flap, it might cause the Virtual Chassis Control Protocol Daemon(vccpd) to crash and generates the core file. [PR1095199](#)
- On QFX5100 Series switches, in a corner case, the BIOS upgrade is getting updated with soft reboot might cause device got stuck at "RE-FPGA-DRV: Please standby while rebooting" message. [PR1097318](#)
- On QFX5100 Series switches, when the Open vSwitch Database(OVSDB) controller is configured, and then changing the inactivity-probe-duration and/or SSL port might cause the controller port to get overwritten with default values. As a workaround,

configured the controller port number after configuring inactivity-probe-duration. [PR1098869](#)

- On EX4300/EX4600/QFX Series switches, when VLANs name contains "-vlan" and then add the interface to this VLAN, it might cause VLAN does not work. As a workaround, change the name of VLANs to another. [PR1100609](#)
- On QFX5100 Series switches, the unsigned Python scripts might not execute successfully due to no executable permissions, which result in the Zero Touch Provisioning (ZTP) process fails. As a workaround, use chmod command to change the permissions of Python scripts file. [PR1101680](#)
- On the QFX5100 with the maximum-ecmp 16, the ECMP scale will be 256 groups though it should be 1k groups. [PR1105851](#)
- On a QFX Series Virtual Chassis Fabric (VCF) or Virtual Chassis with GRES enabled, the backup Routing Engine might continuously reboot after you configure "forward-and-send-to-re" or "forward-only" under the [edit interface interface-name unit unit-number family inet targeted-broadcast] hierarchy. [PR1106151](#)
- On a QFX5100 VCF in an auto-provisioned mode, when adding a new leaf device to the VCF, should zeroize device and reboot by "request system zeroize" if new leaf device has been configured any command. But the issue (interface still up) might be observed at the time of the reboot until the PFE re-initialized the interfaces. [PR1106194](#)
- On a QFX5100 Virtual Chassis, the MAC address is not learned on an AE interface configured as a VXLAN Layer 2 port and with the interface mode configured as access. The issue is observed only with AE interfaces that span multiple Virtual Chassis members and when the member node is rebooted or power cycled. [PR1112790](#)
- On QFX5100 Series switches in the Open vSwitch Database (OVSDB) scenario with VxLAN configured, MAC learning might not work well across the interface which is dynamic changed, and the interface is bounded on the link from routers between Bare-metal Server (BMS) to router connected locally and also between BMS to router connected through spine. [PR1115546](#)
- On QFX3000-G or QFX3000-M QFabric System with a 14.1X53 release, configure few VLANs as needed. When deleting this configuration and then reconfigure them. The flood traffic for a specific VLAN might not reach the interfaces within the same VLAN on another Node. [PR1116817](#)
- On QFX/EX4600 Series switches, in rare condition, the trunk interface may not get create due to data structure becomes out of sync between Packet Forwarding Engine (PFE) and control plane. [PR1128316](#)
- On QFX5100 Series switches with Open vSwitch Database (OVSDB) management protocol configuration that act as virtual tunnel endpoints (VTEPs), traffic being forwarded from ingress AE interface to egress tagged port may not be attach 802.1Q VLAN tag if both ports are located the same device. [PR1128507](#)
- On QFX5100 Series switches, after disable the interface and hot swap 1G copper transceiver, link flap or link-up might be seen on SFP-T port though interface is configured as disabled explicitly. [PR1137204](#)

- From 14.1X53-D36, there is a commit check added to prevent more than one IFL per physical interface(IFD) assigned to one single VLAN. [PR1144123](#)
- On EX4300/EX4600/QFX3500/QFX3600/QFX5100 series switches, if you insert bad SFP or SFP+ optic in a port and replace it with a good optic, then the good optic might not come up. [PR1144190](#)
- On a QFX5100 Virtual Chassis, if you configure an aggregated Ethernet interface as an OVSTDB interface with multiple subinterfaces that are configured under different VXLAN domains, removal of the last but one AE subinterface might reset VXLAN settings on the physical port that are part of the AE interface, resulting in packet drops. [PR1150467](#)
- On QFX5100 switches, a child member might drop the incoming Link Aggregation Control Protocol (LACP) frames when this child member is moved from an access-mode VXLAN LAG interface to a trunk-mode VXLAN LAG interface. [PR1153042](#)
- On Junos based platforms, if they are configured as DHCP client, DHCP offer packet which giaddress is not zero might be dropped. [PR1191452](#)
- On QFX5100-96S with 850W AC power supply inserted, in certain environments, because of a software defect, there is a statistical probability of the QFX5100 850W AC power supply shutting itself down. [PR1203591](#)

Class of Service (CoS)

- On QFX Series switches, 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 of a LAG interface. When there is a mix of LAG interface bundles and interfaces that are not LAG members on a device, do not use * wildcard to apply the CoS configuration globally to all device interfaces. [PR1001605](#)
- On EX4600/QFX Series switches, when applying fixed classifier to the ingress port, the IEEE802.1p CoS values of the egress packets are incorrectly, which result in the peer device handle the packets with the wrong way. [PR1099187](#)
- On QFX5100 and EX4600 switches, if you channelize a 40-Gigabit Ethernet QSFP+ interface into four 10-Gigabit Ethernet ports and try to apply the CoS configuration to one of the specific channels, multicast traffic might get dropped. [PR1108103](#)

EVPN

- On a QFX5100-VC platform, while rebooting a spine node which has active route to reach other Vxlan Tunnel End Points (VTEPs), Packet Forwarding Engine manager(fxpc) might create core files and crash. [PR1088992](#)
- On QFX5100 Series switches VXLAN ports, while receiving DHCP discover packets, there will be incorrect/additional headers on the VXLAN encapsulated DHCP frames, and when these frames are sent by PFE to Kernel, the kernel might drop these incorrect VXLAN udp frames. [PR1107793](#)
- On QFX5100 Series switches with virtual extensible local area network (VXLAN) configured, the SIP/DIP (source IP/destination IP) to be 0.0.0.0 in VXLAN traffic after the device reboot due to sometimes VTEP Gateway daemon (vgd) might push remote

MACs to the layer 2 learning daemon (l2ald) before the source VETP logical interface (IFL) is created. [PR1109838](#)

- On QFX5100 Series switches, when the logging action is set on VXLAN port in firewall filter, the forwarding traffic may get duplicated due to the device encapsulated the packet in VXLAN header using both multicast and unicast destination address. [PR1110818](#)
- On a QFX Series switch or Virtual Chassis which 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 (RE) reboots as part of the NSSU. [PR1126855](#)
- 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. [PR1126915](#)
- On QFX5100 switches, if a trunk interface is a VXLAN port, tagged frames matching the native VLAN ID might be sent out with the native VLAN tagged. [PR1164850](#)

Interfaces and Chassis

- The log message "DCD_CONFIG_WRITE_FAILED" repeatedly appears in the log file. [PR1088577](#)
- On EX4600/QFX Series switches with MC-LAG Inter-chassis Link (ICL) configured, when multiple servers are connected to the MC-LAG peers, if the server side configures LACP behind the EX4600/QFX switches and only negotiates LACP on one of the interfaces, it might lead to MC-LAG link failure. [PR1113903](#)

Junos Fusion Provider Edge

- On a Junos Fusion topology, if a QFX5100 switch is running Junos OS Release 14.1X53-D16 with Enhanced Automation, and you try to autoconvert the switch into a satellite device from the aggregation device, the conversion might fail. As a workaround, install the regular version of Junos OS Release 14.1X53-D16 on the switch prior to the conversion. [PR1072806](#)

Layer 2 Features

- In a mixed QFX3500 and EX4300 Virtual Chassis that has configured persistent MAC and MAC limiting, traffic is not received on Aggregated Ethernet (AE) interfaces on EX4300 switches when the EX4300 switches are acting as the linecard role. [PR1033618](#)
- On EX4600 switches and QFX Series switches, if the extended-vlan-bridge statement is configured for an interface and igmp-snooping is enabled, the interface might drop multicast traffic. [PR1071436](#)
- On QFX5100 Series switches, if the device managed by Open vSwitch Database (OVSDb) with large scale (i.e 2k vni, 4k sub-interface, 40k MAC), it might cause the PFE manager (fxpc) process to crash with a core file generated. [PR1078118](#)
- On EX4600 switches and QFX Series switches, the PFE manager process (FXPC) might crash, with a core file generated, under either of two circumstances - when an

interface is flapping or when you issue the CLI command "clear ether-switch table"
[PR1080132](#)

- On EX4600 switches and QFX Series switches, when an interface without spanning tree protocol (STP) configured receives a VSTP or PVST+ packet where the frame is tagged with a VLAN that is not configured on the device, the switch might change the packet's VLAN ID to a wrong VLAN ID for the VSTP/PVST+ frame and forward it (rather than dropping the frame). [PR1081275](#)
- On QFX5100 Series switches, when device configured VXLAN, at VXLAN I2-side, the egress ports are always selected based on layer2-headers of the inner packets instead of default layer2-payload. If the VXLAN traffic with inner MAC address are fixed. It might cause VXLAN decapsulated packets can not be load-balanced at AE interfaces. [PR1084591](#)
- On a QFX5100 Virtual Chassis, when device were part of an OVSDB-managed VXLAN, and if configuring multiple LAG interface on different switch member over Equal-cost multipath (ECMP) for Layer 3 VXLAN interfaces, the load balancing of the LAG member interface does not work. [PR1090791](#)
- On a EX4300/EX4600/QFX VC/VCF except EX4300 VC, when configuring Protocol Independent Multicast (PIM) on the integrated bridging and routing (IRB) interface and enable IGMP-snooping on related VLAN, if the multicast send and receive interface both on the non-master Flexible PIC Concentrator (FPC), then failover the Routing Engine (RE) mastership might cause multicast traffic to drop. [PR1091645](#)
- On EX4600/QFX5100 Series switches with L3VPN scenario, ping packets sent from CE to remote CE may not work for back to back PE connection. [PR1096698](#)
- On QFX5100 and EX4600 switches running under Junos OS Release 14.1X53-D10 or later, when DHCPv6 solicitation packets go through the device with Q-in-Q configured, the packets might be dropped by peers due to the S-tag not being added. [PR1103793](#)
- On EX4300, EX4600, and QFX Series switches, traffic received on the backup redundant trunk group (RTG) link might get forwarded to other interfaces following an RTG link failover. [PR1119654](#)
- If you reboot one FPC in a two-member Virtual Chassis, the traffic might not exit from the FPC after the FPC comes back online and rejoins the Virtual Chassis, and local registers might be incorrectly cleared, if the port number is the same on both the master and backup. [PR1124162](#)
- On QFX3500/5100 Series switches, while committing et interface inet plus mpls config with no-redirects knob having MTU setting, the protocol ARP might not be configured for the IFL in PFE. [PR1138310](#)
- On QFX5100/QFX3500, buffer is corrupted on port 0 (*/*/0) and error message MACDRAINTIMEOUT and dc bcm_check_stuck_buffers are observed, which could eventually lead to port 0 (*/*/0) flapping. [PR1162947](#)
- On QFX5100 switch, syslog may contain repeated messages like so: fpc12 Unit: 0 port 47 start error detected. [PR1164096](#)

Multiprotocol Label Switching (MPLS)

- On EX4600/QFX Series switches, when configuring the "labeled-unicast" in BGP, the incoming labeled packets might be dropped. [PR1080528](#)
- On QFX5100 Series switches with MPLS and ECMP enabled. By default, the ECMP policy effects for all protocol families that running on the box include MPLS. Because the QFX5100 does not support ECMP MPLS, so the box will install a UNILIST (ECMP route) as multiple UNICAST (a common route). In that case, Packet Forwarding Engine will install multiple copies of MPLS swap label in the egress MPLS table when multiple egress layer 3 nexthops pointing to the same swap label. When plenty of such MPLS routes are installed in the PFE, it might cause MPLS routing table exceedings its scale limit. This will result in a new MPLS route cannot be installed on the PFE. It might affect traffic when this issue happens. This optimization ensures Junos installs a unique entry in the egress MPLS swap table in the PFE when multiple egress layer 3 next hops are pointing to the same MPLS swap label. [PR1087476](#)
- When QFX/EX4600 Series switches are acted as Provider Edge (PE) devices with multiple L3VPNs configured, while pushing 3 labels either through VPN/LDP/RSVP or VPN/BGP/LDP, they might apply the incorrect bottom labels. [PR1089648](#)
- On EX4600/QFX5100 Series switches, when the device configured Ethernet-over-MPLS(L2 circuit) with high scale routes, if restarting the Routing Protocol Daemon (rpd) many times in continuously, it might cause the L2 circuit drops the forwarding traffic. [PR1091867](#)
- On EX4600/QFX Series switches, if LDP/MPLS explicit-null is set on egress PE Devices, packets with label value of 0 are not hitting the IPv4 firewall filter, which is configured under the core-facing interfaces (PE-P). [PR1099334](#)
- QFX5100 don't support the ECMP load balancing for mpls, But no commit errors when configured the ECMP to match on rib table mpls.0 on the code 14.1X53-D12 [PR1102230](#)
- Ping over LSP shows different behavior in regards to HLIM. [PR1179518](#)
- For 2 label PUSH cases, both labels are consuming entries in the same label table. This might result in instabilities of MPLS tunnels and packets drop when add/delete routes. Correct behavior should be that tunnel label goes in one table and VRF label should go in another table. [PR1185550](#)

Network Management and Monitoring

- On a QFabric system, when configuring SNMP to communicate with SNMP server, the device might stop responding to SNMP requests and SNMP polling is not working at random intervals. [PR1061518](#)
- On a QFabric system, if replacing old RSNG node to new RSNG node, it might cause the status of Director Group (DG) still showing the old node as "CONNECTED". [PR1071067](#)
- On a QFabric system, the SNMP process (snmpd) may restart and generate a core file when clients send excessive queries to Juniper Networks enterprise-specific Class-of-Service (CoS) MIB (mib-jnx-cos). [PR1078596](#)

Platform and Infrastructure

- The CPU utilization value is incorrect in the Cloud Analytics Engine probe response statistics. [PR1024840](#)
- QFX Series: PFE panic while processing VXLAN packets (CVE-2016-1274); Refer to <https://kb.juniper.net/JSA10747> for more information. [PR1074501](#)
- On EX4600 and QFX Series switches, MAC addresses on one VLAN might be installed in the hardware but missing from the Ethernet-switching table if the following steps were taken: 1. Configured "vlan-id-list" for a VLAN range "A" with commit 2. Deleted the VLAN range "A" and re-added the VLAN range "B" in the same commit 3. If A + B >= 4096 [PR1074919](#)
- On EX4300, EX4600, and QFX5100 Series switches, when Multiple Spanning-Tree Protocol (MSTP) is used for a VLAN and the link aggregation group (LAG) interface belongs to the VLAN but the LAG interface is not part of MSTP, then that VLAN traffic does not pass on the LAG interface. [PR1084616](#)
- On a QFX5100 Virtual Chassis, frequent MAC move events can put the system into an inconsistent state, which results in a PFE manager (FXPC) process crash with a core file generated. [PR1086108](#)
- On a QFX5100-24Q/QFX5100-24Q-AA switch, if configuring flexible-vlan-tagging and encapsulation on the expansion module(eg. QFX-EM-4), it might cause multicast traffic loss which sent to the interface on the expansion module. [PR1087014](#)
- On QFX5100 switches, adding or removing virtual routing and forwarding (VRF) instances that have many logical interfaces in the link aggregation group (LAG) might cause Link Aggregate Control Protocol (LACP) flapping. [PR1087615](#)
- On EX4600 and QFX5100 switches, when Spanning Tree Protocol (STP) is enabled on an S-VLAN, that S-VLAN's spanning tree protocol (STP) bridge data protocol unit (BPDU) packets might be dropped by the S-VLAN interface if the S-VLAN interface is an aggregated Ethernet (AE) interface. [PR1089331](#)
- On EX4300, EX4600, and QFX Series switches with a firewall filter configured, BGP sessions can go down under certain circumstances. When a BGP traffic term with accept action is configured in the firewall filter, and a log action is configured in the firewall filter with a discard/reject action in another term, BGP sessions might go down when this firewall filter is applied to the lo0 (loopback) interface. [PR1089360](#)
- On a Virtual Chassis Fabric(VCF) with Junos OS release 14.1X53-D25 onwards, when the switch member of VCF rebooting, Broadcast/Unknown/Multicast(BUM) traffic which pass through the Virtual Chassis Port(VCP) will be dropped until rebooted member joins back. [PR1093606](#)
- On QFX5100 Series switches, when device installs large scale route entries but not exceed than the max limitation(16k), the multicast route entry might not be added in the Packet Forwarding Engine(PFE) with "Table full" log messages. [PR1093665](#)
- On QFX5100 Series switches with VXLAN configured, after delete/add the VXLAN Network Identifier (VNI), the traffic is not getting load-balanced across layer 3 links to remote vxlan-tunnel-end-point(VTEP). [PR1094547](#)

- On EX4600 and QFX5100 switches, when flow control is configured on an interface, and pause frames are sent to this interface, the interface might go down. [PR1098055](#)
- On EX4300, EX4600, and QFX Series switches, while creating trunk interfaces that carry a large number for VLAN members which include a VLAN of IRB, multicast or broadcast traffic such as OSPF and ARP that are sent through the VLAN might be dropped, thereby impacting the protocol adjacency [PR1100001](#)
- On QFX5100 Series switches, when VLAN interface mac limit is configured, mac limit is not applied on VXLAN/OVSDB interfaces. [PR1101203](#)
- On EX4300, EX4600, and QFX5100 switches, when you configure a Layer 3 link aggregation group with Link Aggregation Control Protocol (LACP) enabled, and an aggregated Ethernet (AE) interface goes down due to LACP failures, the AE interface still accepts and forwards traffic. [PR1101273](#)
- On EX4300/EX4600/QFX Series switches, when configuring preemptive-cutover timer for a redundant trunk group (RTG), when the primary goes down, is replaced by the secondary link, if the secondary link goes down within the preemptive cutover time (by default, it is 120 seconds), even at this moment the primary link is up, the primary link is still in the blocked state. [PR1101678](#)
- When we issue the PFE command "show brcmfm ifd all" there might be an FXPC core-dump on QFX5100 running 14.1X53-D12 [PR1119567](#)
- Multiple PFEMAN disconnects and reconnects between the master and backup within a short period of time can cause the backup to generate core files. [PR1123379](#)
- If DHCP packets with MPLS tags are sent to the CPU on a QFX5100 node acting as a PHP node, the logical interfaces index on the packet notification might not be set correctly, and the DHCP packets might be dropped. [PR1164675](#)

QFabric Systems

- On a QFabric system, if configuring "remote-debug-permission", the Director Group (DG) should allow a login without providing the password to the component. However, it was observed that the DG have to prompted password to login to a node device. [PR1068276](#)

Routing Protocols

- On EX4600/QFX Series switches, if configuring a filter term to permit the VSTP BPDU packets, it might not work to match the packets. [PR1016394](#)
- On EX4600 and QFX Series switches, if filter-based forwarding (FBF) is configured on an IRB interface that is enabled for Virtual Router Redundancy Protocol (VRRP) also, when the host uses the VIP address as the gateway, the switch will not forward packets from that host to the destination routing instance via FBF. This is expected behavior based on the implementation of family inet filters. As a workaround, configure the hosts to use the physical IP address of the IRB interface, rather than the VRRP VIP address, as the gateway. [PR1025312](#)
- On EX4300 switches, EX4600 switches, and QFX Series switches, after you configure a hold-time timer for an interface member of a multichassis link aggregated Ethernet

(MC-AE), and then reboot the active node device, a loop can occur with the hold timer.
[PR1077019](#)

- On QFX5100 Series switches, if a link aggregation group (LAG) member is added or removed from a LAG port that is bound to a filter-based forwarding (FBF) filter, packets hit this filter may be not forwarded to the right destination. [PR1078195](#)
- On EX/QFX Series switches, when IGMP snooping for IGMPv3 is configured, IGMP snooping may not correctly while receiving an IGMPv3 report with "to exclude" followed by another IGMPv3 report with "to exclude {null}"/"ALLOW_NEW_SOURCES".
[PR1081093](#)
- On EX4600 and QFX Series switches, you might not be able to commit the configuration when the arp-type match condition is configured in a firewall filter. [PR1084579](#)
- On an EX4600 or QFX Virtual Chassis, if you reboot the master routing engine (RE), traffic might be lost due to an RE failover delay of around 15-20 seconds. [PR1085148](#)
- On a standalone QFX Series switch, if you configure a nested firewall filter and then attempt to commit the configuration, the firewall compiler process (dfwc) might crash and generate a core file, leading to commit failure. [PR1094428](#)
- On a QFX VCF, if the switch works as part of a target subnet, while receiving the targeted broadcast traffic, the packets might be forwarded to the destination with the switch's MAC address as the destination MAC address, where it should be converted into a Layer 2 broadcast frame with destination MAC address FFFF.FFFF.FFFF. [PR1114717](#)
- On QFX5100/EX4600 Series switches, when eRACL (Egress routing ACL filter) is applied to more than 64 interfaces, a memory corruption issue might occur, resulting in the Packet Forwarding Engine manager (fxpc) process to crash. [PR1123374](#)
- On QFX5100 series switches configuring gre interface over irb interface, then gre interface can become up but not able to ping IP address of gre interface in remote end.
[PR1124149](#)
- This PR changed the behavior when using flexible vlan tagging and native-vlan-id to egress packets untagged for packets that are part of the native vlan. Previous these packets would egress tagged. [PR1130192](#)
- On QFX5100 and EX4600 switches, if you use the Network Configuration Protocol (NETCONF) to add or delete firewall filters on an integrated bridging and routing (IRB) interface, the Packet Forwarding Engine Manager (fxpc) might generate a core file.
[PR1155692](#)
- FXPC crash may happen during an ECMP route delete from LPM table. This might have happened due to large scale route change operation. SDK vendor provided a fix as a resolution. [PR1158517](#)
- On QFX3500 or QFX5100 switches, when parity errors occur on interfaces, they might affect the memory management unit (MMU) memories. MMU counters can be corrupted, the interface buffers might be stuck, and there might be interface flaps and traffic loss on the affected ports. As a workaround (restoration only), reboot the system.
[PR1169700](#)

- On EX4600/QFX5100 switches, in rare cases, route insert failure in `_soc_alpm_128_write_pivot` function will lead to a loop in the code resulting in a watchdog timeout. This will result in the FPC crash and restart with a core dump. [PR1173980](#)
- On EX4600 and QFX5100 series switches, there are several profiles that allocate memory differently for MAC addresses and host addresses. These profiles can be configured as "l2-profile-one, l2-profile-two, l2-profile-three, l3-profile, lpm-profile". If multicast and unicast host entries reach the maximum number of the L3 host table in related profile, then multicast traffic will be dropped. [PR1177430](#)
- In some scenarios after ECMP route flapping on QFX switches traffic is blackholed. RIB programming is fine: `root@qfx> show route 172.16.2.1 inet.0: 843 destinations, 2707 routes (843 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 172.16.2.0/25 * [BGP/170] 6d 12:24:13, MED 106, localpref 100, from 10.1.32.2 AS path: ?, validation-state: unverified to 172.16.5.101 via ae2.0 > to 172.16.6.101 via ae4.0 [BGP/170] 6d 12:20:20, MED 106, localpref 100, from 10.1.32.4 AS path: ?, validation-state: unverified > to 172.16.6.101 via ae4.0 [BGP/170] 6d 12:20:24, MED 111, localpref 100, from 10.1.32.1 AS path: ?, validation-state: unverified > to 172.16.5.21 via ae1.0 [BGP/170] 6d 12:24:09, MED 111, localpref 100, from 10.1.32.3 AS path: ?, validation-state: unverified > to 172.16.6.21 via ae3.0 FIB programming is fine: root@qfx> show route forwarding-table destination 172.18.2.0/25 Routing table: default.inet Internet: Destination Type RtRef Next hop Type Index NhRef Netif 172.16.2.0/25 user 0 ulst 131157 4 indr 131098 19 172.16.5.101 ucst 2349 9 ae2.0 indr 131154 16 172.16.6.101 ucst 2397 9 ae4.0 But kernel nexthops point to discard nexthop in sdk. Hence traffic to destination is blackholed. PR1179610`
- The QFX5100 will exception (forward a copy of) transit IPv6 Neighbor Discovery traffic to the RE, allowing for a collateral partial local DoS attack. Refer to JSA10749 for more information. [PR1183115](#)
- On QFX3500/QFX3600/QFX5100/EX4600 series switches, if a routing loop is created, the TTL of the packet does not reduce to 0 and eventually the packet is not dropped. [PR1196354](#)

User Interface and Configuration

- On EX4300/EX4600/EX9200/QFX Series switches, when configuring an interface range, if the interface range includes large-scale physical interfaces, and is configured with the "family" option set to "ethernet-switching", committing the configuration might take a long time to complete. [PR1072147](#)

Resolved Issues: Release 14.1X53-D27

General Routing

- On EX Series switches with integrated routing and bridging (IRB) interface configured, if the JSRV interface is created prior to the IRB interface after restarting the device or chassis daemon (chassisd), it might cause all IRB interfaces to be disappeared. [PR965097](#)
- On a QFX5100 platform, when to upgrade junos by topology-independent in-service software upgrade (TISSU), during "FPC Warm Booting" period of TISSU, a few packets

drop might be seen on an SFP-T interface, this issue not seen with SFP-SX interface. [PR1027336](#)

- On QFX5100 switches, enabling error-correcting code (ECC) ELV. [PR1064567](#)
- On a QFX5100 Virtual Chassis, the MAC address is not learned on an AE interface configured as a VXLAN Layer 2 port and with the interface mode configured as access. The issue is observed only with AE interfaces that span multiple Virtual Chassis members and when the member node is rebooted or power cycled. [PR1112790](#)

EVPN

- On a QFX5100-VC platform, while rebooting a spine node which has active route to reach other Vxlan Tunnel End Points (VTEPs), Packet Forwarding Engine manager (fxpc) might create core files and crash. [PR1088992](#)
- On QFX5100 Series switches VXLAN ports, while receiving DHCP discover packets, there will be incorrect/additional headers on the VXLAN encapsulated DHCP frames, and when these frames are sent by PFE to Kernel, the kernel might drop these incorrect VXLAN udp frames. [PR1107793](#)
- On QFX5100 Series switches, when the logging action is set on VXLAN port in firewall filter, the forwarding traffic may get duplicated due to the device encapsulated the packet in VXLAN header using both multicast and unicast destination address. [PR1110818](#)

Layer 2 Features

- On a QFX5100 Virtual Chassis, when device were part of an OVSDB-managed VXLAN, and if configuring multiple LAG interface on different switch member over Equal-cost multipath (ECMP) for Layer 3 VXLAN interfaces, the load balancing of the LAG member interface does not work. [PR1090791](#)
- On a EX4300/EX4600/QFX VC/VCF except EX4300 VC, when configuring Protocol Independent Multicast (PIM) on the integrated bridging and routing (IRB) interface and enable IGMP-snooping on related VLAN, if the multicast send and receive interface both on the non-master Flexible PIC Concentrator (FPC), then failover the Routing Engine (RE) mastership might cause multicast traffic to drop. [PR1091645](#)

Multiprotocol Label Switching (MPLS)

- When QFX/EX4600 Series switches are acted as Provider Edge (PE) devices with multiple L3VPNs configured, while pushing 3 labels either through VPN/LDP/RSVP or VPN/BGP/LDP, they might apply the incorrect bottom labels. [PR1089648](#)
- On EX4600/QFX5100 Series switches, when the device configured Ethernet-over-MPLS (L2 circuit) with high scale routes, if restarting the Routing Protocol Daemon (rpd) many times in continuously, it might cause the L2 circuit drops the forwarding traffic. [PR1091867](#)

Platform and Infrastructure

- On a QFX5100-24Q/QFX5100-24Q-AA switch, if configuring flexible-vlan-tagging and encapsulation on the expansion module (eg. QFX-EM-4), it might cause multicast traffic loss which sent to the interface on the expansion module. [PR1087014](#)

- On QFX5100 Series switches with VXLAN configured, after delete/add the VXLAN Network Identifier (VNI), the traffic is not getting load-balanced across layer 3 links to remote vxlan-tunnel-end-point (VTEP). [PR1094547](#)

Resolved Issues: Release 14.1X53-D26

General Routing

- On a QFX5100 platform, when to upgrade junos by topology-independent in-service software upgrade (TISSU), during "FPC Warm Booting" period of TISSU, a few packets drop might be seen on an SFP-T interface, this issue not seen with SFP-SX interface. [PR1027336](#)
- On a QFX5100-48T switch, interfaces numbered 0 to 23 are sometimes not turned down during device reboot. This issue might be seen when a peer device is using 1G link speed. [PR1059876](#)
- In certain environments - with certain narrow operating temperatures or changing operating temperatures, there is a statistical probability of the QFX5100 850W AC power supply shutting itself down due to a bug in the power supply firmware logic comparing measured fan speed versus target speed at temperature. [PR1062224](#)
- On QFX5100 switches, enabling error-correcting code (ECC) ELV. [PR1064567](#)
- On EX4600 switches and QFX Series switches, when a pair of devices configured for multichassis link aggregation (MC-LAG) are both using active mode when rebooting, traffic can drop for a while on one of the switches. [PR1069644](#)
- The SNMP walk for the dot1dBasePortIfIndex object might return a value of 0, which is not a valid SNMP ifindex for an interface. [PR1070532](#)
- On a QFX5100-24Q-AA switch, in few of the cases, after the switch reboot, the guest virtual machine (VM) may not get the field-programmable gate array (FPGA) devices for use. Consequently any application or utility trying to use the FPGA device will fail. [PR1073076](#)
- On a QFX5100-24Q-AA switch, if the PFE manager (FXPC) restarts due to any reason (crash or planned restart), then the guest virtual machine (VM) will lose its PCIe devices. Consequently, any utility or application using those devices will lose the access to them. This may result in failures of the utilities and/or applications. [PR1073084](#)
- On QFX5100 Series switches, 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](#)
- On QFX5100 switches, if more than 1K virtual extensible LAN network identifiers (VNIs) are created by Open vSwitch Database (OVSDb), the VTEP gateway daemon (vgd) might generate a core file. [PR1075189](#)
- On a QFX5100 Virtual Chassis, the log messages as "fpc0 vccpd irt socket connect failed (no route to host)" are seen continuously, it is harmless. [PR1075437](#)

- A QFX5100 switch with a BIOS version older than V18.7 does not have error-correcting code (ECC) memory enabled by default. This might cause issues because it limits correction of memory corruption. [PR1075915](#)
- On QFX5100/EX9200 Series switches, when configuring the VLAN name and Logical Switch (LS) for OVSDB, if the VLAN name or LS using the UUID format, the configuration would not to commit. [PR1075919](#)
- On QFX5100 Series switches, if you configure both Q-in-Q tunneling and IGMP snooping, IGMP reports do not egress. As a result, multicast traffic is flooded instead of being sent to requested receivers. [PR1076324](#)
- On EX4600, QFX5100, QFX3500, and QFX3600 Series switches, when the device acts a transit router between the DHCP server and the DHCP relay agent, and DHCP server/relay is not configured, the device might not forward the DHCP ack packets to the destination address. Instead, packets are sent to Routing Engine (RE) if the packets' destination port was 68. [PR1079826](#)
- On EX9200 and QFX5100 switches, if you configure DHCP relay with the DHCP server and the DHCP client in separate routing instances, unicast DHCP reply packets (for example, a DHCP ACK in response to a DHCP RENEW) might be dropped. [PR1079980](#)
- On QFX5100 Series switches, if Class of Service (CoS) configuration is changed on a physical interface while running traffic, the host inbound packets might be affected and cannot be processed, and the PFE manager (fxpc) process crash with a core file generated, which result in Aggregated Ethernet (AE) interface goes down due to LACP time out. [PR1082224](#)
- On EX4600 switches and QFX Series switches, you must use the -C and -S option with a DHCP request - if you do not, the client might not receive the DHCP ack packets. [PR1082473](#)
- On QFX5100 Series switches, if installing license for VCF feature and with Junos OS release 14.1X53-D25, the device might raise an error information "license not valid for this product" and fail to install. [PR1084235](#)

Layer 2 Features

- On QFX5100 Series switches, if the device manage by Open vSwitch Database (OVSDB) with large scale (i.e 2k vni, 4k sub-interface, 40k MAC), it might cause the PFE manager (fxpc) process to crash with a core file generated. [PR1078118](#)
- On EX4600 switches and QFX Series switches, the PFE manager process (FXPC) might crash, with a core file generated, under either of two circumstances - when an interface is flapping or when you issue the CLI command "clear ether-switch table" [PR1080132](#)
- On QFX5100 Series switches, when device configured VXLAN, at VXLAN I2-side, the egress ports are always selected based on layer2-headers of the inner packets instead of default layer2-payload. If the VXLAN traffic with inner MAC address are fixed. It might cause VXLAN decapsulated packets can not be load-balanced at AE interfaces. [PR1084591](#)

Multiprotocol Label Switching (MPLS)

- On EX4600/QFX Series switches, when configuring the "labeled-unicast" in BGP, the incoming labeled packets might be dropped. [PR1080528](#)

Platform and Infrastructure

- On EX4600 and QFX Series switches, MAC addresses on one VLAN might be installed in the hardware but missing from the Ethernet-switching table if the following steps were taken: 1. Configured "vlan-id-list" for a VLAN range "A" with commit 2. Deleted the VLAN range "A" and re-added the VLAN range "B" in the same commit 3. If A + B >= 4096 [PR1074919](#)
- On EX4300, EX4600, and QFX5100 Series switches, when Multiple Spanning-Tree Protocol (MSTP) is used for a VLAN and the link aggregation group (LAG) interface belongs to the VLAN but the LAG interface is not part of MSTP, then that VLAN traffic does not pass on the LAG interface. [PR1084616](#)
- On EX4300, EX4600, and QFX Series switches with a firewall filter configured, BGP sessions can go down under certain circumstances. When a BGP traffic term with accept action is configured in the firewall filter, and a log action is configured in the firewall filter with a discard/reject action in another term, BGP sessions might go down when this firewall filter is applied to the lo0 (loopback) interface. [PR1089360](#)

Routing Protocols

- On EX4600/QFX Series switches, if you configure a multiple user-vlan-id term in a firewall filter and then apply it, only the first VLAN uses the term entry. [PR1065060](#)
- On QFX5100 Series switches, if a link aggregation group (LAG) member is added or removed from a LAG port that is bound to a filter-based forwarding (FBF) filter, packets hit this filter may be not forwarded to the right destination. [PR1078195](#)

User Interface and Configuration

- On EX4300/EX4600/EX9200/QFX Series switches, when configuring an interface range, if the interface range includes large-scale physical interfaces, and is configured with the "family" option set to "ethernet-switching", committing the configuration might take a long time to complete. [PR1072147](#)

Resolved Issues: Release 14.1X53-D25

General Routing

- On an QFX3500 configured for Layer 2 Protocol Tunneling, if the customer facing ports are configured as LAG interfaces then the LLDP packets are not tunneled across the switch. [PR871079](#)
- In case you are using QFX5100-48T-6Q, **show chassis hardware** displays QFX5100-48C-6Q like below. ----- **root@host> show chassis hardware**
Hardware inventory: Item Version Part number Serial number Description
Chassis TR0214999999 QFX5100-48C-6Q -----, [PR1006271](#)

- On QFabric node devices, interface flaps and resulting traffic drops can occur as a result of a Network Time Protocol (NTP) update. When this problem occurs, the string "SCHED_SLIP" appears in the log files. [PR1008869](#)
- On EX4600 and QFX5100 switches, the Link Aggregation Control Protocol (LACP) in either slow mode or fast mode might go down and then come back up, causing a timeout and a service outage during an In-Service Software Upgrade (ISSU) or a Nonstop Software Upgrade (NSSU). In addition, after the master Routing Engine is rebooted, the switches might experience intermittent traffic loss on non-LAG interfaces, and redundant trunk group (RTG) convergence times might be long. [PR1031338](#)
- To avoid a traffic loop, an ingress check is implemented on the vcp port for ingress traffic coming from a fpc which has been disconnected from VC or VCF. [PR1041995](#)
- 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 EX4300, EX9200, QFX Series, and MX Series platforms, naming a VLAN "vlan-rewrite" causes an error when you commit the configuration. [PR1054996](#)
- On a QFabric system, the Terminal Access Controller Access Control System (TACACS+) authentication fails to work in Junos OS release 14.1X53-D15. Other platform or other release is not affected. [PR1055775](#)
- SNMP polling may not work in QFabric with Junos OS Release 14.1X53-D15.2 code. [PR1058886](#)
- On a QFX5100-48T switch, interfaces numbered 0 to 23 are sometimes not turned down during device reboot. This issue might be seen when a peer device is using 1G link speed. [PR1059876](#)
- On a QFabric system, if any QSFP+ optics on 40-gigabit data plane (fte) uplink port is removed or inserted in a QFX3600 node device, it might cause other fte port and 40-Gigabit Ethernet(xle) port get detached. As a workaround, remove and re-insert the detached optics. [PR1060463](#)
- When you use the SNMP GET request to poll jnxOperatingState for FPCs that are not present on a Virtual Chassis Fabric (VCF) or an EX Series Virtual Chassis, incorrect results are displayed. Non-existent FPCs might be reported to be UP and RUNNING. This issue does not affect SNMP walks. [PR1061960](#)
- On QFX5100 switches, when a Gigabit Ethernet interface on a fiber Small Form-factor Pluggable (SFP) is configured with the speed of 1G, and full duplex and no auto-negotiation are enabled, the interface goes down. [PR1063118](#)
- On EX4600/QFX Series switches Virtual Chassis(VC) or Virtual Chassis Fabric(VCF) mode, when Redundant Trunk Groups(RTG) link failover, Media Access Control(MAC) refresh packets will be sent out from non RTG interface which belong the same Virtual Local Area Network(VLAN) with the RTG interface, it might cause the traffic drop because of MAC flapping. [PR1063202](#)
- When a Redundant Trunk Group's (RTG) primary link is down and the backup link is an active link, when the primary link comes back online to once again become the active link, other interfaces using that RTG can drop MAC addresses. This applies to

EX4600 and QFX5100 switches, and QFX3500 and QFX3600 switches using Virtual Chassis (VC) or Virtual Chassis Fabric (VCF). [PR1063226](#)

- On QFX5100-48T Series switches, wrong description is shown in "show chassis hardware", description for PIC 0 is displayed 48x10BaseT-6x40G, but it should be 48x10GBaseT-6x40G. [PR1071557](#)
- On QFX5100 switches, if more than 1K virtual extensible LAN network identifiers (VNIs) are created by Open vSwitch Database (OVSDb), the VTEP gateway daemon (vgd) might generate a core file. [PR1075189](#)

EVPN

- On a QFX5100 switch, traceroute does not work as expected when troubleshooting VXLAN packets if ECMP is enabled because the traceroute packets are not forwarded to the same interfaces as the data packets. [PR1035730](#)

Interfaces and Chassis

- On EX4300 switches, EX4600 switches, and QFX switches with Spanning Tree Protocol (STP) enabled, if you have configured an interface as an edge port when spanning-tree interface mode is configured as point-to-point, enabling Bridge Protocol Data Unit (BPDU) protection on those edge ports might not work as expected. This is a typical configuration for multichassis link aggregation (MC-LAG) interfaces. [PR1063847](#)

Junos Fusion Provider Edge

- On a Junos Fusion topology, if a QFX5100 switch is running Junos OS Release 14.1X53-D16 with Enhanced Automation, and you try to autoconvert the switch into a satellite device from the aggregation device, the conversion might fail. As a workaround, install the regular version of Junos OS Release 14.1X53-D16 on the switch prior to the conversion. [PR1072806](#)

Layer 2 Features

- On QFX/EX4300/EX4600 Series switches, traffic flooding or forwarding might cease completely whenever the administrator change the vlan-id for PVLAN to vlan-id-list with range of vlan-ids. [PR1046792](#)
- sfid-bcm memory leak will be seen on RSNG side when sflow setting is existing even if that Node does not have route for collector [PR1053813](#)
- On EX4600/QFX switches, with LAG interface enabled. In some rare scenarios, certain AE member might not inherit STP FORWARD state from its parent AE interface, resulting in the member interface STP state staying in DISABLE. consequently, data packets going through the affected member interface will get dropped. Disable and enable the fault member link will restore AE member interface with correct STP state. [PR1059718](#)
- On EX4600 switches and QFX Series switches, if the extended-vlan-bridge statement is configured for an interface and igmp-snooping is enabled, the interface might drop multicast traffic. [PR1071436](#)

Multiprotocol Label Switching (MPLS)

- MPLS auto-bandwidth does not reset MAX Avg Bandwidth when overflow or underflow threshold limit is configured. It may lead to wrong bandwidth reservations occasionally. [PR954663](#)
- On QFX5100/EX4600, DHCP Relay packets having MPLS tag are getting dropped on RE (Routing Engine), so DHCP client cannot obtain a valid address from DHCP server. RE expect packets with pure IP and not MPLS. After the fix, strip the MPLS tag from the DHCP Relay packet first, and then send pure IP DHCP Relay packet to RE for further process. [PR1060988](#)

Platform and Infrastructure

- On a Virtual Chassis Fabric (VCF), when the master routing engine (RE) is rebooting, traffic passing through the Virtual Chassis Port (VCP) will be dropped. This applies to broadcast traffic, unknown traffic, and multicast (BUM) traffic. [PR1006753](#)
- In situations where QFX Series Switches are expected to generate ICMP redirects, they will also duplicate the incoming packet, causing duplicate responses by the end device. Configuring no-redirects will stop the generation of ICMP redirect packets, however it will not stop the duplication of the packet. To stop the duplication of the packets, ICMP redirects need to be turned off at the Packet Forwarding Engine (PFE) level. [PR1022354](#)
- On QFX Series, EX4300, or EX4600 switches or Virtual Chassis, if you delete aggregated Ethernet (AE) interfaces to which many VLANs are associated, the CPU usage of the Packet Forwarding Engine manager (fxpc/pfex) process might become high. The duration of the high CPU utilization is proportional to the number of AE interfaces deleted. [PR1035669](#)
- On QFX5100 platform in standalone/VC/VCF scenario, the packet forwarding engine manager daemon (fxpc) may crash occasionally. This issue might be caused by multiple events (eg. the fxpc process does not handle signals properly or change the configuration of VC/VCF or after the NSSU etc). However, the issue is more likely to happen if there are any QFX-SFP-1GE-T plugged in. [PR1055331](#)
- On QFX/EX4600 Series switches, when Dynamic Host Configuration Protocol (DHCP) packets with double tag going through the trunk interface which configured Virtual Local Area Network (VLAN) members was bound Layer 3 (IRB) interface, it might be cause the DHCP packets dropped. [PR1059557](#)
- On QFX5100/EX4600 Series switches, installing routes beyond maximum limit might cause the PFE manager (fxpc) process crash and generates the core file. [PR1062349](#)
- On EX4300/EX4600/QFX Series switch, traffic might be flooded out of an interface where the destination MAC address is present in MAC table. [PR1066405](#)

QFabric Systems

- On a QFabric system, CLI command "show interface descriptions" may provide incomplete output or may not provide any output. [PR1057104](#)
- The SSH sessions are flapping between Junos Space and QFABRIC when it is being managed by SPACE/ND. [PR1062750](#)

Routing Protocols

- If a QFX Series switch with per-packet load balancing enabled has multiple Equal Cost Multiple Paths (ECMP) next hops and these also have multiple ECMP next hops, ECMP entries might be installed twice if they have overlapping members. The duplicate entries result in those links carrying twice the traffic of the other links in the ECMP group.
[PR936707](#)
- On QFX Series switches, if configure 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](#)
- On QFX5100 Series switches with a large number of firewall terms configured, if an In-Service Software Upgrade (ISSU) is performed from versions 13.2X51-D25 and below to versions 13.2X51-D26 and above, firewall filters configured after this upgrade method will not be programmed. [PR1051779](#)
- In a rare condition, the routing protocol daemon (rpd) might crash and create a core file if there is internal BGP (IBGP) route churn and BGP next hop fails to update.
[PR1060133](#)
- On QFX and EX4600 Series switches, moving the integrated routing and bridging (IRB) interface to other routing instance, it might cause the traffic drop because of the Address Resolution Protocol (ARP) resolve fail. [PR1063949](#)
- On EX4600/QFX Series switches, if you configure a multiple user-vlan-id term in a firewall filter and then apply it, only the first VLAN uses the term entry. [PR1065060](#)

VPNs

- "ESI TLV not received for ifd" seen very often in the logs. There is no service impact.
[PR1060609](#)

Resolved Issues: Release 14.1X53-D16

General Routing

- In case you are using QFX5100-48T-6Q, "show chassis hardware" displays QFX5100-48C-6Q like below. ----- root@QFX5100-48T>
show chassis hardware Hardware inventory: Item Version Part number Serial number
Description Chassis TR0214999999 QFX5100-48C-6Q -----.
[PR1006271](#)
- 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 EX4600 and QFX5100 switches, the Link Aggregation Control Protocol (LACP) in either slow mode or fast mode might go down and then come back up, causing a timeout and a service outage during an In-Service Software Upgrade (ISSU) or a Nonstop Software Upgrade (NSSU). In addition, after the master Routing Engine is rebooted, the switches might experience intermittent traffic loss on non-LAG interfaces, and redundant trunk group (RTG) convergence times might be long. [PR1031338](#)

- To avoid a traffic loop, an ingress check is implemented on the vcp port for ingress traffic coming from a fpc which has been disconnected from VC or VCF. [PR1041995](#)
- On a QFX5100 switch, issuing the request system reboot command might not shut down the SFP-T interfaces. [PR1050650](#)
- 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](#)
- 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](#)
- Packets are not mirrored when mirror IP address is configured on remote device. [PR1052028](#)
- On QFX Series or EX4600 switches with a primary link as an aggregated Ethernet (AE) interface and a secondary link on a redundant trunk group, if the primary link fails, the secondary link might not take over. [PR1052977](#)
- On EX4300, EX9200, QFX Series, and MX Series platforms, naming a VLAN "vlan-rewrite" causes an error when you commit the configuration. [PR1054996](#)
- On a QFabric system, the Terminal Access Controller Access Control System (TACACS+) authentication fails to work in Junos OS release 14.1X53-D15. Other platform or other release is not affected. [PR1055775](#)
- SNMP polling may not work in QFabric with 14.1X53-D15.2 code [PR1058886](#)
- On a QFabric system, if any QSFP+ optics on 40-gigabit data plane (fte) uplink port is removed or inserted in a QFX3600 node device, it might cause other fte port and 40-Gigabit Ethernet(xle) port get detached. As a workaround, remove and re-insert the detached optics. [PR1060463](#)
- When you use the SNMP GET request to poll jnxOperatingState for FPCs that are not present on a Virtual Chassis Fabric (VCF) or an EX Series Virtual Chassis, incorrect results are displayed. Non-existent FPCs might be reported to be UP and RUNNING. This issue does not affect SNMP walks. [PR1061960](#)
- On QFX5100 switches, when a Gigabit Ethernet interface on a fiber Small Form-factor Pluggable (SFP) is configured with the speed of 1G, and full duplex and no auto-negotiation are enabled, the interface goes down. [PR1063118](#)
- When a Redundant Trunk Group's (RTG) primary link is down and the backup link is an active link, when the primary link comes back online to once again become the active link, other interfaces using that RTG can drop MAC addresses. This applies to EX4600 and QFX5100 switches, and QFX3500 and QFX3600 switches using Virtual Chassis (VC) or Virtual Chassis Fabric (VCF). [PR1063226](#)

Layer 2 Features

- On QFX Series switches, when if a routed VLAN interface is configured with family ISO, the ISO maximum transmission unit (MTU) of the interface is reduced from 1500 (default) to 1497 bytes. Any transit ISO traffic larger than 1497 bytes might be sent to the CPU and cause latency issues. [PR955710](#)
- On QFX Series switches, adding or deleting a subinterface from an aggregated Ethernet (AE) interface might cause momentary packet loss when class of service (CoS) is applied on AE interfaces, even though the traffic is not on this particular AE interface. [PR1045466](#)
- sfid-bcm memory leak will be seen on RSNG side when sflow setting is existing even if that Node does not have route for collector [PR1053813](#)
- On EX4600/QFX switches, with LAG interface enabled. In some rare scenarios, certain AE member might not inherit STP FORWARD state from its parent AE interface, resulting in the member interface STP state staying in DISABLE. consequently, data packets going through the affected member interface will get dropped. Disable and enable the fault member link will restore AE member interface with correct STP state. [PR1059718](#)

Multiprotocol Label Switching (MPLS)

- On QFX5100/EX4600, DHCP Relay packets having MPLS tag are getting dropped on RE (Routing Engine), so DHCP client cannot obtain a valid address from DHCP server. RE expect packets with pure IP and not MPLS. After the fix, strip the MPLS tag from the DHCP Relay packet first, and then send pure IP DHCP Relay packet to RE for further process. [PR1060988](#)

Platform and Infrastructure

- The CPU utilization value is incorrect in the Cloud Analytics Engine probe response statistics. [PR1024840](#)
- The commit synchronize command fails because the kernel socket gets stuck. [PR1027898](#)
- On QFX Series, EX4300, or EX4600 switches or Virtual Chassis, if you delete aggregated Ethernet (AE) interfaces to which many VLANs are associated, the CPU usage of the Packet Forwarding Engine manager (fxpc/pfex) process might become high. The duration of the high CPU utilization is proportional to the number of AE interfaces deleted. [PR1035669](#)
- On QFX/EX4600 Series switches, when the device receive a lot of Address Resolution Protocol (ARP) request packets with high rate, ARP reply packets loss might be seen. [PR1041195](#)
- On QFX5100 and EX4600 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 QFX5100 platform in standalone/VC/VCF scenario, the packet forwarding engine manager daemon (fxpc) may crash occasionally. This issue might be caused by multiple events (eg. the fxpc process does not handle signals properly or change the

configuration of VC/VCF or after the NSSU etc). However, the issue is more likely to happen if there are any QFX-SFP-1GE-T plugged in. [PR1055331](#)

- On QFX/EX4600 Series switches, when Dynamic Host Configuration Protocol(DHCP) packets with double tag going through the trunk interface which configured Virtual Local Area Network(VLAN) members was bound Layer 3 (IRB) interface, it might be cause the DHCP packets dropped. [PR1059557](#)

QFabric Systems

- On EX4300/EX4600/QFX Series switches with Junos OS release 14.1X53-D10 onwards, the multicast routes aging might not work, it would the stale multicast route entries to remain. [PR1053316](#)
- The SSH sessions are flapping between Junos Space and QFABRIC when it is being managed by SPACE/ND. [PR1062750](#)

Routing Protocols

- On QFX Series switches, if configure 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](#)
- In a rare condition, the routing protocol daemon (rpd) might crash and create a core file if there is internal BGP (IBGP) route churn and BGP next hop fails to update. [PR1060133](#)
- On QFX and EX4600 Series switches, moving the integrated routing and bridging (IRB) interface to other routing instance, it might cause the traffic drop because of the Address Resolution Protocol (ARP) resolve fail. [PR1063949](#)

Resolved Issues: Resolved Before Release 14.1X53-D16

- [Interfaces and Chassis](#)
- [Layer 3 Protocols](#)
- [OVSDDB](#)
- [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](#)

- See Also**
- [New and Changed Features on page 91](#)
 - [Changes in Behavior and Syntax on page 122](#)
 - [Known Behavior on page 127](#)
 - [Known Issues on page 136](#)
 - [Documentation Updates on page 206](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 207](#)
 - [Product Compatibility on page 213](#)

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.

- See Also**
- [New and Changed Features on page 91](#)
 - [Changes in Behavior and Syntax on page 122](#)
 - [Known Behavior on page 127](#)
 - [Known Issues on page 136](#)
 - [Resolved Issues on page 144](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 207](#)
 - [Product Compatibility on page 213](#)

[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 208](#)
- [Upgrading Software on QFX5100 Standalone Switches on page 208](#)
- [Performing an In-Service Software Upgrade \(ISSU\) on page 211](#)
- [Preparing the Switch for Software Installation on page 211](#)
- [Upgrading the Software Using ISSU on page 211](#)

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 *Installing Software Packages on QFX Devices* 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 <https://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 211](#)
- [Upgrading the Software Using ISSU on page 211](#)

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 *Installing Software Packages on QFX Devices*.
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 might 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
```

- See Also**
- [New and Changed Features on page 91](#)
 - [Changes in Behavior and Syntax on page 122](#)
 - [Known Behavior on page 127](#)
 - [Known Issues on page 136](#)
 - [Resolved Issues on page 144](#)
 - [Documentation Updates on page 206](#)
 - [Product Compatibility on page 213](#)

Product Compatibility

- [Hardware Compatibility on page 213](#)

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:

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

- See Also**
- [New and Changed Features on page 91](#)
 - [Changes in Behavior and Syntax on page 122](#)

- [Known Behavior on page 127](#)
- [Known Issues on page 136](#)
- [Resolved Issues on page 144](#)
- [Documentation Updates on page 206](#)
- [Migration, Upgrade, and Downgrade Instructions on page 207](#)

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:
<https://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:
<https://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:
<https://www.juniper.net/documentation/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, 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

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

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

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

Revision History

13 November, 2018—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D48

20 August, 2018—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D47

17 May, 2018—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D47

28 February, 2018—Revision 6, Junos OS for the EX Series and QFX Series, Release 14.1X53-D46—Addition to QFX Series “Documentation Updates”

15 February, 2018—Revision 5, Junos OS for the EX Series and QFX Series, Release 14.1X53-D46—Moved PR1240845 from Resolved Issues to Known Issues

17 January, 2018—Revision 4, Junos OS for the EX Series and QFX Series, Release 14.1X53-D46—Addition to EX Series “Changes in Behavior and Syntax”

12 December, 2017—Revision 3, Junos OS for the EX Series and QFX Series, Release 14.1X53-D46—Addition to QFX Series “Changes in Behavior and Syntax”

6 December, 2017—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D46—Addition to QFX Series “New and Changed Features”

29 November, 2017—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D46

14 September, 2017—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D45—Addition to QFX Series “Known Behavior”

1 August, 2017—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D45

21 June, 2017—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D44—Updates to EX Series “Changes in Behavior and Syntax” and “Known Behavior”

14 June, 2017—Revision 1, Junos OS for the EX Series and QFX Series, Release 14.1X53-D44

13 June, 2017—Revision 4, Junos OS for the EX Series and QFX Series, Release 14.1X53-D43—Addition to QFX Series “Known Behavior”

8 June, 2017—Revision 3, Junos OS for the EX Series and QFX Series, Release 14.1X53-D43—PR link fix

6 June, 2017—Revision 2, Junos OS for the EX Series and QFX Series, Release 14.1X53-D43—Addition to EX Series “Known Behavior”

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 © 2018 Juniper Networks, Inc. All rights reserved.

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

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